

CA Service Desk Manager

Administration Guide

Release 12.7.00



This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2012 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Technologies Product References

This document references the following CA Technologies products:

- CA Asset Portfolio Management (CA APM)
- CA CMDB
- CA Business Intelligence
- CA Business Service Insight (CA BSI)
- CA Configuration Automation (formerly known as CA Cohesion ACM)
- CA Embedded Entitlements Manager (CA EEM)
- CA Enterprise Workload Automation (CA EWA)
- CA Process Automation (formerly known as CA IT PAM)
- CA Management Database (CA MDB)
- CA Management Portal
- CA Network and Systems Management (CA NSM)
- CA Portal
- CA Remote Control Manager (CA RCM)
- CA Service Desk Manager (CA SDM)
- CA Service Management
- CA Siteminder
- CA Software Delivery
- CA Spectrum® Infrastructure Manager (CA Spectrum)
- CA Wily
- CA Workflow

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At <http://ca.com/support>, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at <http://ca.com/docs>.

Contents

Chapter 1: Introduction 27

Audience	27
What You Need to Know	27
Service Management Processes and Best Practices.....	28

Chapter 2: Managing Servers 29

Number of Servers	29
Server Configuration Tasks.....	30
Configuring TCP/IP	31
Change a Server Configuration.....	32
ITIL Configuration	32
ITIL Service Disciplines	33
Employee and Guest Interface Options	34
Activity Log Security	36
PDA Interface	37
Tablet Support.....	38
REST Sample Mobile User Interface	39
Deploy the REST Mobile Sample User Interface	40
Enable the REST Mobile Sample UI	42
Example: Analyst Interface Home Page	43
Example: Analyst Searches for an Incident	44
Example: Analyst Creates an Incident.....	44
Example: Analyst Logs an Internal-Only Comment.....	45
Example: Employee Interface Home Page	46
Example: Employee Creates an Incident.....	47
Example: Employee Edits an Incident	47
Tomcat Logging	48
Servlet Defaults	48
REST Logging	49
Enable CXF Logging	50
Start a Secondary Server	50
Start the Primary Server.....	51
Server Status	53
Configure SSL on Tomcat.....	53
Configure a Secondary Server for Web Screen Painter.....	55
How to Deploy CMDBf Web Services.....	55

Stop a Server (Windows)	55
Stop a Server (UNIX)	56

Chapter 3: Defining Business Structure 57

How to Define the Business Structure	57
Contacts	58
Groups	58
Sites	58
Locations	59
Organizations	59
Define the Business Infrastructure	60
Object Definition Order	60
Families and Classes	60
Manufacturer and Models	61
Service Status	61
Vendor Types and Vendors	61
Configuration Items	62
External Asset Management Tools	62
Multi-Tenancy	63
Service Provider	63
How Multi-Tenancy Works	64
User Interface Impact	70
Support Automation Impact	73
Knowledge Management Impact	74
How to Use Multi-Tenancy	75

Chapter 4: Implementing Policy 85

Policy Implementation	85
Notifications	85
Activity Associations	86
Activity Notifications	87
Object Contact Notifications	88
Notification Methods	89
Email Notifications	92
Notification Rules	96
Example: Create a Message Template	105
Artifacts in Messages	107
Notification Codes and Phrases	107
Manual Notification Recipients List	111
Previous Assignee Notifications	112
Configuration Item Notifications	114

Survey Notifications	115
How to Add URL Hyperlinks to Notifications	116
Notification Log Reader.....	116
Email Administration	120
Mailboxes	121
Mailbox Rules.....	121
Mailbox Policies	128
How to Implement Mailboxes.....	129
Multiple Mailboxes	135
How to Configure Email Replies	136
Artifacts.....	138
Artifact Use Considerations	139
Artifact Protection and Security.....	140
Service Level Agreements	141
SLA Usage	141
Classic SLA Processing	142
Service Types and Events	142
How to Implement Service Types	143
Predefined Service Types	144
Events Setup.....	145
How to Create Service Targets.....	146
Service Contracts.....	147
Service Contracts Migration.....	148
Time to Violation.....	149
Time Zones and Workshifts.....	150
Security.....	153
CA EEM and CA Workflow User Base Configurations	153
Security Considerations	162
CA EEM Authentication for CA Process Automation.....	162
User Authentication	163
How CA SDM Authenticates Users.....	163
External Authentication	164
Validation Types.....	165
Logged In and Licensed User Counts.....	166
Internal Logs.....	166
CA SDM Integration.....	166
Data Partition Associations	167
Data Partitions	167
Data Partitions Setup	168
Constraint Specifications.....	168
Constraint Types	170
Create a Data Partition Constraint for CAB Assignments.....	171

Configure Knowledge Management Data Partition Constraints for Role-Based Permissions	172
Surveys	173
Configure Your System for Surveys	173
Prepare a Survey	173
Define Survey Notifications.....	174
Survey Reporting.....	175
Managed Survey	175
Web Services	176

Chapter 5: Configuring User Accounts 177

Contacts	177
Contact Definitions.....	177
Groups	179
Contact Types.....	179
Determine Behavior Based on Contact Type	180
Notification Setup Based on Contact Type	180
Select Contacts Based on Contact Type	180
Special Handling Types.....	180
How to Configure Special Handling Contacts.....	181
Associate a Contact to a Special Handling Type.....	182
LDAP Directory Data	183
Configure LDAP Options.....	184
Verify LDAP Integration.....	186
Create a Contact Automatically	188
Access Type Assignments From LDAP Groups	188
Batch Import Contacts Using LDAP Data.....	189
Batch Update Contacts Using LDAP Data	192
LDAP Authentication	194
Transport Layer Security	194
Attribute Mapping	195
Troubleshooting	196

Chapter 6: Managing Roles 207

Roles	207
Predefined Roles	207
Role-Based Security.....	210
How Access Types Work	210
Role Records	213
Functional Access Areas	214
Data Partitions	219
Role-Based Navigation	219

Tabs	219
Predefined Tabs	221
Web Forms	223
Form Groups	223
Menu Trees	224
Menu Tree Resources	225
Menu Bars	225
Toolbars	227
Go Resources	227
Help Sets	228
How to Implement a Custom Role	228
How to Implement a Custom Menu Tree	229
Create a Role Record	231
Create a Tab Record	232
Create a Menu Bar Record	233
Create a Web Form Record	234
Copy a Menu Tree	235
Create and Customize a Menu Tree	236
Create and Publish a Help Set	237
Switch Roles	239

Chapter 7: Establishing Support Structure 241

Support Structure	241
Models	242
Internal Model	242
External Model	243
Combined Model	244
CA Workflow	245
Workflow at Runtime	246
Select a Workflow Process Definition	246
Workflow Tasks	247
CA Process Automation Workflow Integration	247
CA Process Automation Components	248
CA Process Automation Integration with CA SDM at Run Time	249
How to Create a Process Definition	250
Create a Start Request Form	251
Attach a CA Process Automation Process Definition	252
Shared Codes	253
Priority Codes	254
Severity Codes	255
Impact Codes	255

Urgency Codes	255
Status Codes	256
Request Status Codes.....	256
Change Order Status Codes	258
Issue Status Codes.....	259
Task Status Codes.....	260
Task Types	261
Incident Tracking	262
Install Incident Tracking	262
Request/Incident/Problem Areas.....	263
Request/Incident/Problem Area Properties	264
Define Request/Incident/Problem Areas for Self-Service.....	266
Change Order and Issue Categories	266
Predefined Change Categories.....	267
Predefined Issue Categories.....	268
Rules for Changing Categories on a Ticket	268
Category Properties	269
Define Change and Issue Categories for Self-Service.....	271
Automatic Closure of Tickets	272
How to Define Auto Close Ticket Settings.....	272
How to Define an Auto Close Activity Notification	273
Related Ticket Activities	274
How to Define Activity Notifications for Related Tickets	275
How to Define Related Ticket Activity Notifications	275
Priority Calculation	276
How Priority Calculation Manages Ticket Values	277
Priority Calculation Generates Urgency Value After Saving Self-Service Tickets	279
How to Set Priority Calculation	281
Status Transitions and Dependent Attribute Controls	291
Work with Status Transitions and Dependent Attribute Controls	292
Configure Status Transitions	292
Configure Dependent Attribute Controls.....	294
Web Services Methods	295
Predefined Transition Flows	295
Best Practice: Predefined Status Transitions	299
Status Transitions for Self-Service.....	300
How Transitions for Self-Service Work.....	301
How to Create or Update Transition Types for Transitions	302
How to Link Transition Types to Transitions	302
Activate Predefined Transition Types	303
Timers.....	304
Time Zones	305

Service Type Event Triggers	306
Time Zone Event Triggers.....	306
Time Zone Rules	307
File Attachments	308
File Attachment Upload and Download	308
Repositories	309
Announcements	312
Internal Announcement Visibility	312
Specify Announcement Urgency	313
Stored Queries Setup	313
Sequence Numbers	314
Audit Log Use	315
Integration with CA Network and Systems Management.....	315

Chapter 8: Controlling System Behavior 317

Options Manager Usage.....	317
How to Modify the System Environment	318
Events.....	319
Macros.....	319
Macro Types.....	320
Use Macros with Events.....	320
Use Macros in the Behavior for Issue and Change Categories.....	321
Use Macros with Multiple Notification	321
Use Macros with a Site-Defined Condition	322

Chapter 9: Configuring the Web Interface 323

Web Interface Configuration.....	323
How the Web Interface Works.....	323
Web Director and Web Load Distribution.....	324
Web.cfg and CA SDM	325
Assign a Web Engine to a Web Director	325
Define the Web Engine Role by Web Director Parameters	326
Non-SSL Environment with Basic Load-Balancing.....	327
Global SSL Environment with Basic Load-Balancing.....	327
Targeted Login in a Non-SSL Environment with Optional Load-Balancing.....	327
Targeted SSL Login in a Mixed Environment with Optional Load-Balancing.....	328
Web Directors and Web Engines Configuration.....	328
How to Implement an SSL Login Environment with a Primary Server Only	330
How to Create a System with a Primary Server Only	330
Verify the Web Director Parameter Values	333
SSL Login Environment Verification	334

Secondary Servers and Web Capabilities	334
How to Prepare Web Capabilities	334
How to Create a System with Secondary Servers	335
Server Configuration	337
Make Changes Using pdm_edit.pl	348
Start Web Director	353
How Web Director Handles User Sessions	354
Improve Performance With Browser Caching.....	355
Configure the Microsoft Internet Information Server	355
Configure Apache.....	356
Clear the Cache	357
Record Locking Behavior in the Web Interface	358
Print CA SDM Web Pages	358
Configuration File Modification	359
SchedExpMaximum.....	367
SellListCacheExclude	367
SellListCacheMax	368

Chapter 10: Configuring Auto Assignment 373

Auto Assignment	373
Auto Assignment Relationships.....	374
Auto Assignment Methods.....	374
How to Begin Implementing Auto Assignment	375
Areas and Categories	375
Analyst Groups	375
Analysts	376
How to Auto Assign Tickets to a Group and Not to Contacts Within the Group	376
Auto Assignment by Location	377
Auto Assignment by Workshift	378
Default Group and Assignee.....	380
Auto Assignment Enablement.....	380
Auto Assignment Override	381
Assignment Controls	382
Manual Assignment	382
Assignee_set Option	382
Iss assignee_set.....	382
Area_Defaults	382
Require Assignee and Group Options	383
Templates.....	384
CA Network and Systems Management Interface	384
Activity Logging	384

Enable Activity Logging for Additional Attributes	385
Auto Assignment Tracing	385
Stored Queries	386
How Auto Assignment Assigns Tickets	386
How Auto Assignment Assigns Workflow Tasks.....	392
Configuration Item-Based Auto Assignments	395
How Configuration Item-Based Auto Assignment Works	395
Enable Configuration Item-Based Auto Assignments	399

Chapter 11: Managing Your Database 401

Database Management Utilities.....	401
Select and Configure your Database	401
Database Loader	402
How to Create and Use an Input File	403
Drop and Restore Constraints	404
Database Backup.....	404
Database Restore	405
Database Table Replacement.....	405
Data Extraction.....	406
Use the Data Extractor on UNIX.....	406
Data Selection for Extraction	406
Data Dereferencing	407
How to Use pdm_deref Example	408
Use the Dbadmin Mode	410
Archive and Purge Rules.....	411
Run Archive and Purge Rules	412
View Archive and Purge Rules.....	412
Start Archive and Purge Using a Third-Party Scheduler	413
Archive and Purge Rule Definitions.....	413
Archive and Purge History.....	416
Attachment Handling (arcpur)	418
How to Restore Archived Data.....	419
Archiving and Purging KPI Data.....	420
Archive and Purge Knowledge Management Forums.....	421

Chapter 12: Using the Text API 423

Text API	423
Command Line Interface.....	424
CA Network and Systems Management Interface	424
Input Format	425
Keywords.....	426

Keyword Input Conventions	428
Format an Email Message To Update a Ticket	428
Start and End Email Message Delimiters	429
How the Text API Uses Artifacts.....	429
How to Set Up Notification Replies to Update Tickets.....	430
Conversion Methods	433
The Configuration File	435
Options.....	436
Defaults	436
Ignore Incoming	437
Example Input	438

Chapter 13: Managing Version Control 439

How Version Control Works	439
Version Control Files	440
Primary Server Control Files.....	440
Secondary Server and Client Control Files	441
Version Control for Installation Customizations	441
Version Control Server Modes	442
Version Control File Syntax	443
Version Control Parameters.....	444
Uncontrol a Component.....	448

Chapter 14: Managing Configuration Items 449

Using the Web Interface	449
View Configuration Items.....	450
Create a Configuration Item.....	450
Update a Configuration Item	451
Associate a Maintenance Window with a CI	451
View Associated Change Windows	452
View Configuration Item History.....	452
Inactivate a Configuration Item	452
Reactivate a Configuration Item	453
Contact, Location, and Organization CIs	453
Create a CI from a Base Object	454
Select a Base Object for a CI.....	454
Edit CI Details for a Base Object	455
Edit CI Attributes for a Base Object.....	455
Create a Base Object CI Using GRLoader	456
CI Relationships	456
CI Relationship Types	457

Create a Relationship Type	457
Manage a CI Relationship.....	458
Create a CI Relationship	459
View Relationships for a CI.....	459
Inactivate a CI Relationship	460
Reactivate a CI Relationship	460
Inactivate CI Relationships (Edit in List)	461
Inactivate a CI Relationship Using GRLoader	461
Reactivate a CI Relationship Using GRLoader	462
Delete a CI Relationship from the Database	463
CI Relationship History and Comparison.....	463
Versioning	464
Uses of Versioning.....	465
Shared Asset and CI Audit Trail Records	466
Versioning Terminology	466
Sources of Versioning Data	469
CA SDM Change Management Integration	470
CA APM Integration.....	471
CI Versioning Management.....	472
CA SDM Change Management	488
View CI Attributes in Other CA Products.....	489
Using the CMDBf Viewer	490
CMDB Visualizer	490
Perform Root Cause Analysis	492
Visualizer Administration	492
Add a Discovered Asset	493
Asset and CI Flags	494
CI Reconciliation	495
MDR-Based Reconciliation	496
How to Identify and Resolve Ambiguous CIs	498
Review and Modify Inbound Data Using Transaction Work Area (TWA)	510
Manage Staged Transactions	514
Transaction Work Area	515
Populating the TWA	517
How to Use the Web Interface to Update Data in the TWA	528
Manage Relationship Transactions	531
How To Load Transactions into the CMDB.....	532
TWA Administration.....	534
CA CMDB Data Maintenance	538
CA CMDB Family/Class Structure	538
Change Family/Class of a Single CI	539
Change the Family/Class of a List of CIs	540

Change CI Family/Class Using GRLoader	540
Extending CA CMDB	541
Configuration Audit and Control Facility (CACF)	551
CACF Administration and Policy Definition	553
Managed Attributes	566
Managed Change States	566
Change Specifications	569
How Change Verification Occurs.....	574
How to Archive and Purge Audit Data	578
Implement a Change Verification Strategy	579
Planning and Implementing Change Verification.....	583
Change Verification Best Practices.....	589
Verify a CI Attribute Value Update Manually.....	593
Example: Allow Rogue Updates Only From a Specific Location	597
Example: Upgrade Laptops in Your Organization	598
Example: Lock Down Nonverified Change Orders	599
Example: Allow a CI Update If No Matching Change Order Exists	600
Example: Defer All Updates from CA Configuration Automation to the TWA	601
Example: Only Log the Policy Results as a Test	601
Example: Reject a CI Update	602
Example: Allow Change Orders Created Without Specifications	602
Example: Do Not Allow Change Orders Created Without Specifications.....	603
Example: Allow Rogue Inserts from Selected Sources	603
Example: Allow a Rogue Update for a Nonproduction CI	604

Chapter 15: Administering MDRs 605

What is an MDR?.....	605
MDR Classes and MDR Names	606
How does an MDR Complement CA SDM?	606
MDR to CA SDM Definition	607
MDR Launcher	607
Define a URL to Launch an MDR	607
Set Up a CA APM MDR Provider.....	609
Launch in Context from CA CMDB to CA APM	610
CI Properties that Support MDR Federation	610
Federated Asset ID	611
MDR Name	611
MDR Class.....	611
MDR Definition with CA Cohesion ACM Installation	612
CA Cohesion ACM MDRs	613
How to Associate an MDR to a CI Manually.....	614

CA Cohesion Automatic Import.....	614
CI to MDR Mapping.....	615
MDR Definition Administration.....	616
CA Cohesion ACM Report.....	616
Using GRLoader	617
CI Naming Conventions and Restrictions	617
system_name Naming Convention	618
Using the CMDBf Viewer	620
How to Update Metadata Files for CMDBf Mapping	620
How To Display MDR Attribute Values With CA CMDB Attribute Names	622
How To Hide MDR Provider Attributes	623
How To Define MDR Attributes Without CA CMDB Equivalents	623
Define CMDBf Data Provider Metadata	624

Chapter 16: Managing Change 625

Change Management in CA SDM	625
Change Management Components	626
View the Change Calendar	627
CAB Responsibilities	627
How the CAB Process Works	628
Assign Members to the CAB Group.....	629
Change Manager Responsibilities	629
How the Change Manager Role Works	630
Define Tasks for the Change Manager Role	631
Change Categories, Status, and Risk Levels.....	632
View the Change Order Scoreboard.....	633
Define a Change Order Stored Query.....	633
Configure Change Manager Options	634
Change Calendar	635
Add Schedule Information to a Change Order	635
iCalendar Event Templates	636
Export Schedules to iCalendar	636
Schedule Views	637
Scheduling View Configuration	639
How to Schedule Change Orders.....	647
Use the Change Scheduler Example.....	648
How to Schedule Change Windows	652
View Change Windows.....	653
Associate a CI with a Maintenance Window	653
View Associated CIs.....	654
Create a Blackout Window Example	655

Create a Global Maintenance Window	656
Conflict Analysis and Collision Detection	656
CA Workflow Visualization	656
How to Visualize the Workflow	657
Change Management Process Definition for CA Workflow	658
Change Management Process Definition Components	659
How to Set Up the Change Management Process Definition	659
How the Change Management Process Definition Works	665
ActivityNode Actor not found: Update Object -Service Desk r12	675
Change Order Does Not Close	676
CAB Console and Reporting.....	676
Manage CAB Groups	677
Assign CAB Groups	678
CAB Approvals	679
Change CAB Console Properties.....	679
Change Management Reporting	680
Risk Assessment	681
How to Implement the Risk Survey	682
How to Access a Risk Survey Directly from a URL	684
Impact Explorer	685
Launch Impact Explorer	686
Explore Attached CIs	686
View a CI in Impact Explorer	687
Add a Related CI to a Change Order.....	687
Display the CI Descendants List.....	688
Launch CMDB Visualizer from Impact Explorer	688
Configuring Impact Explorer	688

Chapter 17: Managing Reports 691

CA Business Intelligence Reports	691
Reporting Scenarios	692
Reporting Components	692
Reporting Data Flow Diagram	695
Display Web-Based Reports in InfoView	696
Security and Authorization.....	697
Groups and Users.....	697
CA SDM Data Partitions in Infoview	698
Universe and Universe Connections	698
Report Folders.....	699
Access Levels	701
How to Point an Existing CA Business Intelligence Server to a CA SDM Server	702

Create an ODBC Data Source	702
Configure the Universe	703
Export the Universe	704
How to Set Up Data Partitions Security for Reporting	704
Add the CA SDM Privileged User to CMC	705
Define Universe Database Credentials	705
Establish Data Partitions	706
Replicated Database for Offline Reporting	706
Role-Based Reports	706
Define Role-Based Reports for the Role	707
Display New Reports on the Reports Tab	708
Web-Based Reports	715
BusinessObjects InfoView Interface	715
Navigate to Reports	715
InfoView Preferences	716
Schedule Reports	716
Data Analysis Setup	717
Publish and Distribute Reports	717
Key Performance Indicators	717
KPI Types	718
Predefined KPIs	718
KPI Daemon	719
System KPIs	719
Stored Query KPIs	721
SQL KPIs	721
KPI Fields	722
Retrieve Ticket Data	724
Troubleshooting	726
Ad Hoc Reports	728
Web Intelligence Interface	728
How Ad Hoc Reporting Works	729
Example Ad Hoc Reports	729
Example: View All Open Priority 1 and 2 Requests for All Users	730
Example: View All Open Requests that Do Not Include a Status of Work In Progress	731
Example: View All Closed Requests in the Last 30 Days for Users Whose Last Name Begins with "C"	732
Dashboard Reports	734
View Dashboards and Reports in InfoView	735
Write CA Business Intelligence Reports	735
CA SDM ODBC Driver	736
Write SQL for BusinessObjects Reports	737
PDM Functions	737
Attribute Aliases	740

pdm_isql Interactive SQL	740

Chapter 18: Generating Reports in CA SDM **741**

Generate Reports	741
Database Views	741
Basic Views	742
Advanced Views	744
Reports Method Setup	745
Reports Formatting	745
Column Sort Order Modification	746
Summary and Detail Reports	746
Analysis Reports	747
Generate Request or Issue Reports	747
Generate Request Area or Issue Category Reports	747
Generate Request Area Priority or Issue Category Priority Reports	748

Chapter 19: Managing Knowledge **749**

Knowledge Management	749
Find Procedures for Knowledge Management	750
Knowledge and Multi-Tenancy	750
How to Set Up a Knowledge Base	751
Import Sample Knowledge Data	751
Knowledge Base Monitoring	751
Knowledge Base Re-Index	752
Index and De-Index Queue Settings for Batch and Instant Processing	752
How to Use Documents in the Knowledge Base	753
Knowledge Submission from CA SDM	755
Knowledge Submission from Self-Service	755
Document Attributes	755
Document Permissions	755
Resolution Editing	756
Document Publication Preparation	756
Document Publication	757
Version Documents	757
Document Expiration	758
Document Archive and Purge	758
Knowledge Search	758
Forums	759
Knowledge Tree Documents	759
Knowledge Documents Schedule	760
Knowledge Schedule Filter	760

Knowledge Schedule Views.....	762
Scheduling View Configuration	763
Access Export/Import.....	769
How to Export/Import Knowledge	770
Export/Import Documents	772
Export/Import Packages.....	773
View Export/Import Templates.....	774
pdm_ket Utility—Knowledge Export Tool.....	780
pdm_kit Utility—Knowledge Import Tool	781
Allow Users to Export and Import Knowledge	782
Web Services	782

Chapter 20: Administering Knowledge Management 783

Knowledge Administration	783
Find Procedures for Knowledge Administration	784
Knowledge Management Roles and Functions	784
Knowledge Management User Interfaces.....	785
Knowledge Management Configuration and Management Functions	786
Self-Service Knowledge Options	786
Documents and Users	791
How to Manage Role Privileges and Document Visibility	794
Action Content	795
View Action Content	796
Create Action Content (Action URL).....	797
Create Action Content (Internal HTML)	797
Edit Action Content	798
Search for Action Content.....	799
Document Approval Process	799
Approval Process Manager	800
Define an Approval Process for Document Editing	801
Create an Approval Process Template	802
Document Status Definitions	804
Automated Policies	806
View Automated Policies	807
How to Set Up Automated Policies	808
Create an Automated Policy	808
Edit an Automated Policy.....	809
Schedule Automated Policies.....	809
View Document Lifecycle Policy Reports	810
Knowledge Document Control	810
Comment Types	811

Document Templates	816
How to Create Knowledge Document Links	823
Knowledge Categories	825
Create a Knowledge Category	826
Modify a Category	829
Delete a Category	830
Move a Category	831
Copy a Category with Document Links	831
Copy a Category without Document Links	832
Manage Category Permissions	832
Reports and Metrics	835
Knowledge Report Card	835
Web-Based Reports	836
Role-Based Report Web Forms	837
Search	837
KT Search Engine	837
Recommended Documents	847
Set Up Default Search Settings	851
CA SDM Integration Options	852
Define Field Mapping	853
Define Issue Search Configuration	855
Define Request/Incident/Problem Search Configuration	856
Knowledge Suggestions	857
Define Issue Categories	858
Define Request/Incident/Problem Areas	859
Configure Self-Service Policies	860
Solution Survey	861
Define FAQ Settings	861
Define Solution Survey Settings	863
Knowledge Management System Settings	864
Configure General Settings	864

Chapter 21: Administering Support Automation 867

Automating Support in Your Environment	867
Live Assistance	867
Support Automation Analyst Administration	870
How Analysts Launch Live Assistance	871
How to Configure Live Assistance for Analysts	872
How End Users Join Assistance Sessions	873
How Analysts Automate Support for End Users	874
How Analysts Provide Live Assistance	874

Support Automation User Administration	875
How to Configure Support Automation Role Permissions	876
Support Automation Anonymous and Registered Users	876
How to Configure Support Automation for Guest Users	877
Support Automation Access Level Administration	878
Support Automation Activity Notification Administration	879
Support Automation Page Adaptations	880
Branding Administration	881
Localization Administration	881
Page Layout Configuration	882
Support Automation System Properties	882
Support Automation Queue Administration	883
Queue Management	883
How to Manage Queue Summaries	884
How to Manage the Queue Hours	884
Ticket Template Management	885
Administration Settings	885
How to Configure Support Automation Settings	885
How to Customize Support Automation Tools	887
Automated Tasks	887
Chat Preset Administration	892
Default Credentials	893
Disclaimers	893
Session Log Administration	893
View the Session Log	894
Support Automation Reports	894
Receive a Ticket Request	896
Invite the End User to an Assistance Session	896
Resolve the Ticket with a Chat Session	897
Provide Live Assistance	897
End the Assistance Session and Close the Ticket	898

Appendix A: View Field Descriptions 899

View Field Descriptions	899
View_Act_Log	900
View_Audit_Assignee	901
View_Audit_Group	902
View_Audit_Priority	902
View_Audit_Status	903
View_Change_Act_Log	903
View_Change	905

View_Change_to_Assets	910
View_Change_to_Change_Act_Log	911
View_Change_to_Change_WF	912
View_Change_to_Properties	914
View_Contact_Full	916
View_Contact_to_Environment	919
View_Group	920
View_Group_to_Contact	920
View_Issue	921
View_Issue_Act_Log	926
View_Issue_to_Assets	927
View_Issue_to_Issue_Act_Log	928
View_Change_to_Request	929
View_Issue_to_Issue_WF	933
View_Issue_to_Properties	935
View_Request	936
View_Request_to_Act_Log	942
View_Request_to_Properties	942

Appendix B: RFC 2251 LDAP Result Codes 945

LDAP Return Codes	945
LDAP Server Return Codes	945
LDAP Client Return Codes	950
LDAP-Associated RFC Standards	952

Appendix C: Reference Commands 955

bop_sinfo--Display System Information	956
dbmonitor_nxd--Database Monitoring Daemon	957
pdm_backup--Write Database to ASCII File	958
pdm_cache_refresh--Refresh Database	960
pdm_configure--Open the Configuration Window	961
pdm_d_refresh--Start Failed Daemons	961
pdm_deref--Dereference ASCII Data	962
pdm_discimp -- Discovered Asset Import	966
pdm_discupd -- Discovered Asset Update	967
pdm_edit--Configure Server Processes	968
pdm_extract--Extract Data from Database	969
pdm_halt--Terminate Daemons or Stop Services	972
pdm_init--Start Daemons	973
pdm_key_refresh--Refresh Cached Key Information	973
pdm_lexutil--Modify CA SDM Lexicons	974

pdm_k_reindex—Knowledge Re-Index Utility	975
When to Use pdm_k_reindex	976
Re-Index Tracking	976
Import and Re-Indexing	977
Index and De-Index Queue Settings for Batch and Instant Processing	977
pdm_listconn--List Active Connections	978
pdm_load--Add, Update, and Delete Database Records	980
pdm_logfile--Change stdlog Cutover Size	982
pdm_log4j_config Utility--Modify the log4j properties File	983
Utility Usage Examples	984
Modify the Log File Refresh Interval Manually	986
Modify the jsrvr.log Appender	987
Modify the jstd.log Appender	987
pdm_proctor_init--Start Proctor on Secondary Servers	987
pdm_replace--Replace a Database Table	988
pdm_rest_util--Manage the CA SDM RESTful Web Services Application	989
Undeploy the REST Web Services Application	990
pdm_restore--Restore a Database	990
pdm_status--Show Status of Daemons or Processes	992
pdm_task--Set Environment Variables	992
pdm_text_cmd--Text API Command Line Interface	993
Input Examples	995
pdm_uconv--Convert Local Charset to UTF-8	996
pdm_userload--Add, Update, and Delete Database Records	998
pdm_webstat--Return Web Usage Statistics	1001
report--Generate Reports	1005
rpt_srv--Generate Reports	1006
uniconv--Start UNIX CA NSM Event Converter Daemon	1007

Appendix D: Form Groups 1009

Customer Forms Group	1009
Employee Forms Group	1010
Analyst Forms Group	1012

Chapter 22: Troubleshooting Performance Problems 1035

Identifying Performance Problems in CA SDM	1035
Define the Performance Problem	1036
Use the CA Diagnostic Report Tool	1036
Execute the Interval Logging Script	1041
Gather Database Server Environment Details	1042
Review General Tuning Recommendations	1043

Chapter 1: Introduction

This section contains the following topics:

[Audience](#) (see page 27)

[What You Need to Know](#) (see page 27)

[Service Management Processes and Best Practices](#) (see page 28)

Audience

This guide is intended for the CA SDM administrator, the person who has overall responsibility for the administration of the product. Some tasks that the administrator performs include the following:

- Starting and stopping the services necessary to the CA SDM server.
- Setting up various system components.
- Determining the available options to users.
- Generating reports based on the service desk data.

The purpose of this guide is to help you to use CA SDM to implement, administer, and enforce service delivery. This guide helps you understand how the product solves the challenge of fully automating and managing service delivery from problem inception to resolution.

What You Need to Know

To successfully administer CA SDM, you should be familiar with the following:

- The operating environment where CA SDM is installed
- The operation of your web server
- Basic administration tasks

This guide assumes that the product has been installed successfully, based on the information in the *Implementation Guide*.

Service Management Processes and Best Practices

Implementing standardized processes and best practices directly impacts the effectiveness, productivity, and cost of the Service Support environment. CA provides a library of recommended processes and best practices for Service Management that aligns with industry standards and recognized best-practice frameworks including ITIL, CobIT, BS15000, and more. The processes described for CA SDM include the following:

- Incident Management
- Problem Management
- Change Management
- Request Management
- Configuration Management
- Release Management
- Knowledge Management
- Support Automation

Note: Information about the Best Practices library is available online. You can learn about the Service Management Best Practices of CA, including white papers and other collateral at <http://www.ca.com/sm/bp> <http://www.ca.com/sm/bp>. The strategic process expert partners of CA can help personalize the best-practice library for your organization.

Chapter 2: Managing Servers

This section contains the following topics:

- [Number of Servers](#) (see page 29)
- [Server Configuration Tasks](#) (see page 30)
- [Configuring TCP/IP](#) (see page 31)
- [Change a Server Configuration](#) (see page 32)
- [ITIL Configuration](#) (see page 32)
- [PDA Interface](#) (see page 37)
- [Tablet Support](#) (see page 38)
- [REST Sample Mobile User Interface](#) (see page 39)
- [Tomcat Logging](#) (see page 48)
- [Start a Secondary Server](#) (see page 50)
- [Start the Primary Server](#) (see page 51)
- [Server Status](#) (see page 53)
- [Configure SSL on Tomcat](#) (see page 53)
- [Configure a Secondary Server for Web Screen Painter](#) (see page 55)
- [How to Deploy CMDBf Web Services](#) (see page 55)
- [Stop a Server \(Windows\)](#) (see page 55)
- [Stop a Server \(UNIX\)](#) (see page 56)

Number of Servers

Your CA SDM installation has one or more server components that as an administrator you manage. The number of servers depends on your enterprise. Every installation has one primary server that manages overall CA SDM functionality. Installations can optionally have one or more secondary servers to manage specific features. After you install CA SDM, configure each computer that runs product components.

For example, your particular configuration can be a single computer acting as the primary server with all the necessary services installed. Or your installation can be spread across several computers, each serving a specific function. For example, one computer can act as the primary server and the database server, while a separate web server runs the web interface, and yet another computer authenticates users.

You can run the server configuration as part of the installation process, or you can run it later. You change the configuration if changes in your environment occur.

Note: For more information about the architecture of this product and the configuration options, see the *Implementation Guide*.

More information:

[Start a Secondary Server](#) (see page 50)

[Start the Primary Server](#) (see page 51)

Server Configuration Tasks

We recommend that you configure each computer that is designated to run CA SDM. As administrator, you configure the primary server, secondary server, and client installations. The initial configuration occurs as part of the CA SDM installation process.

In any CA SDM environment, you perform the following tasks:

- Configure servers
- Start, stop, and monitor the primary server
- Start, stop, and monitor a secondary server (if defined)

Note: CA SDM uses your database server to store and retrieve information. The database server runs as a separate service. For administrative information, refer to your database server documentation.

A change in the system environment can require changes to the server configuration. For example, you can change the CA SDM server configuration for the following reasons:

- Changes in the database management system
- Changes made to the CA SDM start-up template file by `pdm_edit.pl` (a feature used to add functionality and increase the performance of the enterprise system)
- Integration with a web server such as Tomcat
- Integration with CA EEM

More information:

[Start the Primary Server](#) (see page 51)

[Stop a Server \(Windows\)](#) (see page 55)

[Stop a Server \(UNIX\)](#) (see page 56)

[Change a Server Configuration](#) (see page 32)

Configuring TCP/IP

You can change the default TCP Internet Protocol (TCP/IP) setting on one or more servers. This setting cannot be forced on the client if it is not supported on the server.

The TCP/IP setting is controlled by using the `NX.env` file, which is found in the `$NX_ROOT` directory. Use a text editor such as WordPad to edit this file. The following option controls the TCP/IP setting:

```
NX_PROTOCOL_ONLY=mode
```

where *mode* can be one of the following values:

IPv4

In IPv4 mode, the system opens up sockets for slump processes that listen to IPv4 traffic.

IPv6

In IPv6 mode, the system opens up sockets for slump processes that listen to IPv6 traffic.

Mixed

In mixed mode, the system opens up sockets for slump processes that listen to both IPv4 and IPv6 traffic. Mixed mode is designed for clients with secondary servers that use a different Internet Protocol from the primary server (or each other).

Note: If IPv4 and IPv6 hosts coexist on the network, ensure that the appropriate transition strategies, tools, and mechanisms to support these technologies are in place before you change the server configuration.

Example

```
NX_PROTOCOL_ONLY=ipv4
```

Change a Server Configuration

You can use the configuration utility to make a change to a server configuration.

To change a server configuration

1. From the Windows Start menu, select Programs, CA, CA SDM, Configuration.

The CA SDM Configuration utility appears.

2. Complete the utility fields, and click Next.

The right panel changes to show the appropriate fields for the link highlighted in the navigation pane on the left.

3. Continue following the on-screen instructions to complete the installation, and click Finish.

The server configuration is changed.

ITIL Configuration

Information Technology Infrastructure Library (ITIL) is a collection of best practices in computer data center management. In addition to defining recommended processes, a major benefit of the ITIL framework is its precise definitions of commonly used data center terminology. Often in the IT world, the same word is used to mean different things; or different people use a particular word with a meaning that is individually nuanced. ITIL helps to avoid this problem.

The following ticket types are available:

- Request
- Change Order
- Issue
- Incident
- Problem

Important! CA SDM only supports an ITIL interface. The ITIL interface supports data objects that were not used in previous non-ITIL versions of the product, for example, problem and incident tickets.

ITIL does the following:

- Produces an ITIL interface for your CA SDM installation.
- Allows your CA SDM database, forms, and fields to differ from a standard installation, and to conform to ITIL conventions instead.

Important! When upgrading your existing system, clear the "Load default data" check box to retain the database tables and data; otherwise, all existing data is lost. For more information about migrating from a non-ITIL environment, see the *Implementation Guide*.

ITIL Service Disciplines

ITIL describes best practices for several disciplines. The Service Support and Service Delivery disciplines combine to provide the Service Management capability of an organization. Complex interrelationships among all ten of the Service Management disciplines interact to help ensure that IT infrastructure delivers a high level of service to businesses.

Service Support includes the following disciplines:

- Incident Management
- Problem Management
- Change Management
- Release Management
- Configuration Management

CA SDM specifically addresses Incident, Problem, Change, and Configuration Management.

Service Delivery includes the following disciplines:

- Service Management
- Availability Management
- Capacity Management
- Financial Management for IT Services
- IT Service Continuity Management

CA SDM specifically addresses Service Management.

Employee and Guest Interface Options

You can use CA SDM to configure separate interfaces for employees and guests. You configure these separate interfaces through the Options Manager in the Administration tab. The following values control these interfaces:

employee_intf_incident_support

Displays the following values:

- Request only
- Incident only
- Both Incident and Request

guest_intf_incident_support

Displays the following values:

- Request only
- Incident only
- Both Incident and Request

Important! If this is a new installation, ITIL is configured by default with the value set to *Incident only*. If you are migrating from a previous non-ITIL configuration, the options are installed, but the values are set to *Request only*.

Configure the Employee Interface

You can configure the employee interface to display incidents, requests, or both.

To configure the employee interface

1. Click the Administration tab.
The Administration console appears.
2. Click Options Manager, Request Mgr.
The Option List appears.
3. Click employee_intf_incident_support.
The Options Detail page appears.
4. Change the Option Value field to one of the following values:

Incident Only

(ITIL Default) Displays only Incident ticket types on the employee interface.

Request Only

Displays only Request ticket types on the employee interface.

Both Incident and Request

Displays both Incident and Request ticket types on the employee interface.

Click Save.

5. Click Refresh to confirm your selections.

The Options Detail is updated.

6. Close the Options Detail.

The Option List appears.

Configure the Guest Interface

You can configure the guest interface to display incidents, requests, or both.

To configure the guest interface

1. Click the Administration tab.

The Administration console appears.

2. Click Options Manager, Request Mgr.

The Option List appears.

3. Click guest_intf_incident_support.

The Options Detail page appears.

4. Change the Option Value field to one of the following values:

Incident Only

(Default) Displays only Incident ticket types on the guest interface.

Request Only

Displays only Request ticket types on the guest interface.

Both Incident and Request

Displays both Incident and Request ticket types on the guest interface.

Click Save.

5. Click Refresh to confirm your selections.

The Options Detail is updated.

6. Close the Options Detail.

The Option List appears.

Activity Log Security

The Activity Log security option prevents end users from updating fields on an activity log. You can select the internal option to prevent a customer from seeing the log.

Activity Log security affects activity logs from the following ticket types:

- Request
- Change order
- Issue
- Incident
- Problem

Enable Activity Log Security

You can enable Activity Log Security from the Options Manager in the Administration Tab.

To enable activity log security

1. Click the Administration tab.
The Administration console appears.
2. Click Options Manager, Request-Change-Issue.
The Option List appears.
3. Click activity_log_security.
The Options Detail page appears.
4. Click Edit, and select one of the following Option Values:

Editable

(Default) Allows all fields on the activity log to be editable through the web interface or web services.

Write Protected

Displays the activity log as read-only. If you select the internal option, only internal users can edit the log, and it cannot be viewed by the customer.

Note: If the security option is enabled, an error message displays that the activity log is read-only if you try to edit the log through the web interface or external sources such as web services.

Click Save.

5. Click Refresh to confirm your selections. Close Window.

Activity log security is enabled.

Important! The `activity_log_security` option cannot be uninstalled. You can only change the value of the option to Editable or Write Protected in Options Manager, Request-Change-Issue.

Impact on Web Screen Painter

The Activity Log Security feature, `$NX_ACTIVITY_LOG_SECURITY`, includes the following attributes (time_spent, time_stamp, and description) for the alg, chgalg, issalg objects in majic.

Example: `$NX_ACTIVITY_LOG_SECURITY` for Object alg in cm.maj

In this example, for object alg in cm.maj, `$NX_ACTIVITY_LOG_SECURITY` appears on the three attributes:

```
time_spent DURATION $NX_ACTIVITY_LOG_SECURITY {ON_POST_VAL update_cr_timespent( call_req_id )
50 ;
};
time_stamp DATE $NX_ACTIVITY_LOG_SECURITY { ON_NEW DEFAULT NOW ; };
descriptionÂÂÂ STRING ÂÂÂ $NX_ACTIVITY_LOG_SECURITY;
```

In Web Screen Painter, *Updatable only for new record* field is disabled when the value of the keyword evaluates to `WRITE_NEW`.

Note: For information about Web Screen Painter, see the *Implementation Guide*.

PDA Interface

CA SDM supports ITIL terminology for the Personal Digital Assistant (PDA) interface. This interface lets end users create Requests, Incidents, and Problems, and lets them search for the following ticket types:

- Incidents
- Problems
- Requests
- Change Orders
- Issues

Analysts can use the PDA interface. This interface honors the role present in the *PDA Role* field of the Access Type of the contact.

Displays the PDA Interface on all the unsupported devices and browsers.

Look up service for searching a contact works only on PDA browsers that support multiple tabs and multiple browser windows. Multi-tenancy works similar to the browser interface, but the drop-down list to select multi-tenancy is not provided. When you create the ticket, tenancy is based on the requestor, the affected enduser, and the analyst.

Automatic Priority Calculator (APC) helps in deciding priority while creating Incident/Problem, based on the attributes, urgency, impact, request area, and affected end user.

Note: An asterisk indicates a required field. For example, when you create a request, specify the Affected End User and Priority. The PDA Interface does not support attachments when you create tickets.

Tablet Support

The Apple iPad displays the Analyst Interface, Employee Self-Service, Customer Self-Service, and Vendor Self-Service interfaces with limited functionality only on its default browser. CA SDM does not support attachments on Tablets.

Note: All other Tablets display the PDA Analyst Interface. On the iPad, the Safari browser supports the full CA SDM User Interface. However, other browsers on iPad may display the full UI, but are not supported.

REST Sample Mobile User Interface

The REST sample mobile user interface provides two user interfaces: Analyst and Employee. This sample also supports the Administrator, Level 1 Analyst, and Employee roles. After you log in to CA SDM, the Administrator and Level 1 Analyst roles view the analyst interface. The Employee role only views the employee interface. For an invalid login, CA SDM rejects the login request and displays a message. If an administrator disabled Automatic Priority Calculation, the Priority field appears on create and edit pages.

Important! This functionality applies only to the sample program provided in CA SDM Release 12.7. It is not a design restricted in the REST Web Services application itself.

Fields that use autosuggest require a value selection from the auto suggestion list, such as the assignee and Incident Area. If you do not select a value, the user input for that field clears.

Note: We recommend that you navigate REST interfaces with the provided buttons, instead of using the browser back and forward options. For example, select Home or Cancel to return to a previous page.

You configure REST during the CA SDM configuration, but you enable the sample mobile interface [manually](#) (see page 41).

Analysts can perform the following tasks:

- View Announcements.
Default: Sorted by Posted Date
- View Incidents that CA SDM assigned to the Analyst.
- View Assigned and Unassigned Incidents.
- Sort Incidents.
Default: Sorted by Open Date.
- Search for an Incident.
- Create an Incident.
- Modify the Status, Urgency, and Impact of the Incident.
- View and update the Activity Log with a comment, or specify Callback or Research.
Default: Sorted by Activity Date.

Employees can perform the following tasks:

- View Announcements.
- View open and closed Incidents that an end user created.


- Sort Incidents, such as by Open Date.
- Create an Incident.
- Update the Activity Log with a comment.
- Search for an Incident.
- Modify the Incident status from Open to Closed, or Closed to Open.

Note: An asterisk indicates a required field.

Deploy the REST Mobile Sample User Interface

The REST mobile sample user interface lets a Service Desk Analyst manage Incidents on mobile devices. The System Administrator wants to enable REST Web Services and test the sample mobile user interface. The System Administrator enables REST Web Services during product configuration and copies files on the CA SDM computer to enable the REST mobile sample user interface. The Service Desk Analyst uses the mobile interface to manage Incidents that end users open in CA SDM.

The following diagram explains how a System Administrator enables REST Web Services and so that a Service Desk Analyst can manage Incidents in the queue:



1. [Configure REST Web Services](#) (see page 41).
2. [Enable the REST Mobile Sample User Interface](#) (see page 41).
3. [View the Sample Interface to Manage Incidents](#) (see page 42).

Configure REST Web Services

The System Administrator enables REST Web Services during the CA SDM configuration. By default, the CA SDM configuration does not enable REST Web Services.

Follow these steps:

1. Complete *one* of the following actions:
 - Install CA SDM and wait for the configuration dialog to appear.
 - If you already installed CA SDM and you previously did not configure REST Web Services, execute `pdm_configure` from the command-line interface. On Windows, click Start, Programs, CA, Service Desk Manager, Configure.

The configuration dialog appears.

2. Confirm the configuration information for the General Settings, System Accounts, Database, and Web Interface dialogs.

The REST Web Services dialog appears.

3. Enable the Configure REST Web Services option.
4. Specify the REST Tomcat Port.

Default: 8050

5. Specify the REST Tomcat Shutdown Port.

Default: 8055

6. Click Next and continue the CA SDM configuration.

The configuration completes.

Enable the REST Mobile Sample UI

CA SDM disables the REST Mobile Sample user interface by default to prevent undesired access to the MDB. The System Administrator enables this interface manually.

Follow these steps:

1. Locate the following directory on the CA SDM computer:

`$NX_ROOT/samples/sdk/rest/mobiledemo`

2. Copy this directory to the following location:

`$NX_ROOT/bopcfg/www/CATALINA_BASE_REST/webapps`

Note: You do not have to restart Tomcat.

The REST Mobile Sample user interface is enabled.

3. (Optional) If you want to disable the REST sample user interface, remove the `mobiledemo` directory from `/webapps`.

View the REST Mobile Sample Interface to Manage Incidents

The Service Desk Analyst views the REST mobile sample interface on their mobile device to manage Incidents in the queue. For example, the Service Desk Analyst opens Incident 30 to log an internal-only comment.

Follow these steps:

1. Open the following URL on your mobile device:
`http://hostname:REST-Tomcat-port/mobiledemo/login.html`
The sample REST mobile interface login page appears.
2. Log in to CA SDM using your credentials.
The REST mobile analyst home page appears.
3. Select Assigned Incidents.
4. Select an Incident, such as Incident 30.
5. Select Create Activity on the detail page.
6. Select Log Comment from the drop-down list.
7. Select Yes from the Internal option.
8. Enter a comment and select Save.

You have now successfully deployed the REST Mobile Sample UI so that Incidents can be managed in the queue.

Enable the REST Mobile Sample UI

CA SDM disables the REST Mobile Sample user interface by default to prevent undesired access to the MDB. The System Administrator enables this interface manually.

Follow these steps:

1. Locate the following directory on the CA SDM computer:
`$NX_ROOT/samples/sdk/rest/mobiledemo`
2. Copy this directory to the following location:

`$NX_ROOT/bopcfg/www/CATALINA_BASE_REST/webapps`


Note: You do not have to restart Tomcat.

The REST Mobile Sample user interface is enabled.

3. (Optional) If you want to disable the REST sample user interface, remove the mobiledemo directory from /webapps.

Example: Analyst Interface Home Page

The following image displays the REST sample analyst mobile home page:



Example: Analyst Searches for an Incident

In this example, an analyst logs in to the REST interface on the server to search for Incident 40.

Follow these steps:

1. The analyst logs in to CA SDM from the following URL on the mobile device:

`http://myservicedesk:8050/mobiledemo/login.html`

The Analyst Home Page appears.

2. The analyst selects Search Incidents.

The Incident Search page appears.

3. The analyst enters search criteria, such as the Incident number.

4. The analyst selects Submit.

The Incident search results appear.

5. The analyst selects Incident 40.

The Incident 40 details appear.

Example: Analyst Creates an Incident

In this example, an analyst logs in to the REST interface on the server. The analyst wants to create an Incident for a network problem.

Follow these steps:

1. The analyst logs in to the REST mobile interface.

2. The analyst selects Create New Incident.

3. The analyst completes the following information:

- Enters **Analyst** as the End User

Note: Entering the first three characters of a user name opens the auto-suggest list. You can use any CA SDM contact in this field.

- Enters **ServiceDesk** as the Assignee
- Selects Open from the Status drop-down list
- Selects 2-Soon from the Urgency drop-down list
- Selects 3-Single Group from the Impact drop-down list
- Enters a summary of the Incident

For example, the summary specifies that the analyst cannot access the company network on a laptop.

- Enters a description of the Incident

For example, the description specifies the name of the end-user laptop on the network and hardware details.

4. The analyst selects Save.

The Incident is created and a message displays the Incident number.

Example: Analyst Logs an Internal-Only Comment

In this example, an analyst opens an Incident to log an internal-only comment. Analysts can also specify Research or Callback for the Activity Log.

Follow these steps:

1. The analyst logs in to the REST mobile interface.
2. The analyst completes one of the following steps:
 - Selects My Incidents or Assigned Incidents
For example, the analyst opens an Incident 40 from the list.
 - Selects Search Incidents
For example, the analyst enters **40** to search for Incident 40.The detail page for Incident 40 appears.
3. The analyst selects Create Activity.
The Create New Activity Log page appears.
4. The analyst selects Log Comment from the Type drop-down list.
5. The analyst selects Yes from the Internal option.
6. The analyst indicates the amount of time spent on the Incident.


Note: If the analyst does not provide a time spent value, CA SDM displays the amount of time between opening the page and saving the activity log.

7. The analyst enters a description and selects Save.

The Incident 40 activity log is updated with an internal-only comment.

Example: Employee Interface Home Page

The following image displays the REST sample employee mobile home page:



Example: Employee Creates an Incident

In this example, an employee logs in to the REST interface on the server to create an Incident for a network problem on their computer.

Follow these steps:

1. The employee logs in to the REST mobile interface.
2. The employee selects Create a new Incident.
3. The employee completes the following information:
 - Enters their telephone number
 - Enters their email address
 - (Required) Selects 3-Quickly from the Urgency drop-down list
 - Enters **Network** as the Incident Area
 - Enters a summary of the Incident

For example, the summary specifies that the employee cannot access the company network on a laptop.
 - (Required) Enters a description of the Incident

For example, the description specifies the name of the end-user laptop on the network and hardware details.
4. The employee selects Save.

The Incident is created and a message displays the Incident number.

Example: Employee Edits an Incident

In this example, the employee wants to edit Incident 40.

Follow these steps:

1. The employee logs in to the REST mobile interface and selects My open incidents.

The Incident List page appears.
2. The employee selects Incident 40.

The Incident 40 details appear.
3. The employee selects Edit.
4. The employee updates the following information about the Incident:
 - (Required) Urgency
 - Incident Area

- (Required) Description
- 5. The employee selects Save.
The Incident is modified.
- 6. (Optional) The employee selects Close.
The Incident is closed.

Tomcat Logging

CA SDM uses a separate log4j.properties file for its web components such as Servlets and PDM_RPC. Support Automation, CMDB Visualizer, and REST also have a separate log4j.properties file. Starting or stopping Tomcat logging does not require you to recycle the Tomcat daemons.

The following list describes how the log4j.properties files of CA SDM components differ:

- CA SDM monitors the log4j.properties in the NX_ROOT\site\cfg and in the NX_ROOT\bopcfg\www\CATALINA_BASE\webapps\CAisd\WEB-INF for changes.
- Support Automation monitors the log4j.properties in the NX_ROOT\bopcfg\www\CATALINA_BASE_SA\webapps\SupportAutomation\WEB-INF.
- CMDB Visualizer monitors the cmdvisualizerlogging.properties in the NX_ROOT\bopcfg\www\CATALINA_BASE_VIZ\webapps\CMDBVisualizer\WEB-INF\classes.
- REST monitors the rest.log4j.properties in the NX_ROOT\site\cfg.

Use the pdm_log4j_config utility only for changing variables in the log4j.properties, cmdvisualizerlogging.properties, and rest.log4j.properties files. You cannot use the pdm_log4j_config utility to modify the polling interval.

CA SDM checks the log4j properties files for any changes periodically. Configure the time interval by modifying the NX_LOG4J_REFRESH_INTERVAL variable in the NX.env file.

Servlet Defaults

The following servlets log INFO level messages to the jsrvr.log file by default:

- PDMContextListener—Generates a log entry during the startup and shutdown of services.
- PDMweb—Generates a log entry from operations on the user interface.
- UploadServlet—Generates a log entry when you attach files to a ticket.
- pdmExport—Generates a log entry when you click Export on list forms.

- `pdm_report`—Generates a log entry when you click the Report menu on list forms.
- `page_cache`—Generates a log entry from operations on the user interface.
- `ProcessServlet`—Generates a log entry when you install CA Workflow and access CA Workflow tasks.

Note: The `pdm_rpc` daemon also generates a log entry from these actions.

- `BOServlet`—Generates a log entry when you configure CA Business Intelligence and clicks on the Reports tab.

REST Logging

REST uses the `pdm_rest_util.jar` and `rest-core.jar` packages, which contain log4j logging support. These components do *not* write any messages to the standard logs (`stdlog.*`), but through the log4j component. By default, these packages log INFO, ERROR, and WARN messages. Because the Java packages use the same log4j configuration properties file, each logs messages to the same output file. You view the *rest.log4j.properties* file in the `$NX_ROOT\site\cfg\` directory.

To increase the log level to trace or debug, update the following section in the *rest.log4j.properties* file:

```
log4j.rootCategory=debug, jrestlog
```

Locate the output file, as defined in the configuration properties file in the following directory:

```
$NX_ROOT\log\jrest.log
```

Enable CXF Logging

USDK> disables CXF logging by default because it can affect performance on a production environment. If your environment requires logging for debugging purposes, the administrator can modify the beans.xml file to enable CXF logging.

Follow these steps:

1. Locate the beans.xml file in the following directory:

NX_ROOT\bopcfg\www\CATALINA_BASE_REST\webapps\caisd-rest\WEB-INF

2. Locate the following section of the XML file:

```
<cxf:bus>
  <cxf:features>
  </cxf:features>
</cxf:bus>
```

3. Add <cxf:logging/> to the section, as shown in the following example:

```
<cxf:bus>
  <cxf:features>
  <cxf:logging/>
  </cxf:features>
</cxf:bus>
```

4. Save the XML file.

Start a Secondary Server

If your installation includes one or more secondary servers, you must start the secondary servers prior to starting the primary server.

In a UNIX environment, you start each CA SDM secondary server from the command line by using `pdm_proctor_init`.

In a Windows environment, you start each CA SDM secondary server by using the Remote Proctor server.

To start a secondary server in Windows

1. Select Services from the Control Panel.

The Control Panel appears.

2. Select the CA SDM Remote Proctor service and click Start.

The service is started.

More information:

[pdm_proctor_init--Start Proctor on Secondary Servers](#) (see page 987)

Start the Primary Server

Every CA SDM installation has one primary server that handles basic product functionality.

Important! If your installation includes one or more secondary servers, you must start the secondary servers before starting the primary server.

To start the CA SDM primary server in a Windows environment, from the Control Panel, select the CA SDM Server service, and click Start. You can start the service manually each time you need it, or you can configure it to start automatically like any other Windows service.

In a UNIX environment, start the CA SDM server from the command line using `pdm_init`.

The following table describes the processes that start automatically when you start the CA SDM server:

Process	Description
Daemon Agent (<code>pdm_proctor_nxd</code>)	The daemon agent responsible for starting the managed daemons
Daemon Monitor (<code>pdm_d_mgr</code>)	Monitors daemon processes
BOP Virtual DB (<code>bpvrtddb_srvr</code>)	BOP virtual database server
Data Dictionary (<code>ddictbuild</code>)	Rebuilds data dictionary each time the system is started—this runs and then goes away
KPI Daemon (<code>kpi_daemon</code>)	Manages the collection, organization, and storage of KPI data.
Oracle agent (<code>orcl_agent</code>)	Agent for Oracle database—many instances, depending upon load
Oracle DB (<code>orcl_prov_nxd</code>)	Oracle database provider
SQL agent (<code>sql_agent</code>) (runs only if you are using MS SQL)	Agent for SQL Server database—many instances, depending upon load
SQL DB (<code>sql_prov_nxd</code>) (runs only if you are using MS SQL)	Microsoft SQL Server database provider
Event Manager (<code>ehm_nxd</code>)	Event manager

Process	Description
License Manager (license_nxd)	Manages security
Message Dispatcher (sslump_nxd)	Dispatches messages
Notification Manager (Windows only) (bnotify_nxd)	Manages notifications in the Windows environment
Object Engine (domsrvr)	CA SDM server
Report Manager (pcrpt_nxd)	PC reporting
Software Version Control (pdm_ver_nxd)	Manages versions of specified system components
Method Engine (spel_svr)	Spell code interpretation server
Text API (pdm_text_nxd)	Text API daemon for email and CA NSM interfaces
Timed Events/Notifications (animator_nxd)	Timed events and notifications
User Validation (boplgm)	Request Management login
Web Engine (webengine)	Runs the engine for the web client
Archive Purge Daemon (arcpur_svr)	Manages the background Archive and Purge processing
BU Daemon (bu_daemon)	Handles FAQ Rating calculation for Knowledge Documents
DB Monitor (dbmonitor_nxd)	Monitors CA common tables for changes
EBR Daemon (bpebr_nxd)	Handles knowledge search requests
EBR Idx Daemon (bpeid_nxd)	Handles EBR re-indexing
KRC Daemon (krc_daemon)	Manages statistical calculations and notifications for the Knowledge Report Card
KT Daemon (kt_daemon)	Manages Knowledge Documents (KD approval process, permissions, notifications, and so on)
LDAP virtdb (ldap_virtdb)	Agent for communication with LDAP Servers
Mail Daemon (pdm_mail_nxd)	Handles outbound email notifications
Mail Eater (pdm_maileater_nxd)	Handles inbound email notifications
MDB Registration Server (mdb_registration_nxd)	Agent for handling MDB registration requests
PDM KTLC (pdm_ktlc)	Manages licensing for Knowledge Management.
PDM RPC (PDM_RPC)	Manages Web Services requests
Repository Daemon (rep_daemon)	Handles attachment repositories

Process	Description
Spell checker (lexagent_nxd)	Handles spell checking requests
Time-to-Violation (ttv_nxd)	SLA Violation forecaster
tomcat controller (pdm_tomcat_nxd)	Manages Tomcat services

Note: In this table, the Notification Manager process applies only to the Windows environment and the Default DB process applies only to the UNIX environment.

Server Status

You can use the `pdm_status` utility to display the status of a primary or secondary CA SDM server running on any operating environment.

More information:

[pdm_status--Show Status of Daemons or Processes](#) (see page 992)

Configure SSL on Tomcat

You can configure your system to use Secure Socket Layer (SSL) on a Tomcat web server.

To configure SSL on Tomcat

1. From the command line, change directories to the JRE install location, and enter the following command:

```
bin\keytool -genkey -alias tomcat -keyalg RSA
```

A `.keystore` file is created by default in the home directory of the logged in user. You can specify a different location during `.keystore` file generation. On UNIX, verify that the directory in which you generate the `.keystore` file has sufficient permissions for CA SDM access.

Note: For more about specifying a different `.keystore` file location, see your Tomcat documentation.

2. Answer the prompts appropriately. The default password is *changeit*.

Note: You can enter a password other than the default one. For more information, see your Tomcat documentation.

3. Edit the server.xml file located in the NX_ROOT\bopcfg\www\CATALINA_BASE\conf directory as follows:

```
<Connector port="8443" maxHttpHeaderSize="8192"
    maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
    enableLookups="false" disableUploadTimeout="true"
    acceptCount="100" scheme="https" secure="true"
    clientAuth="false" sslProtocol="TLS"
    keystoreFile="keystoreFile="C:\Documents and Settings\user\keystore"
    keystorePass="password"/>
```

4. Use the following commands to restart the Tomcat server:

```
pdm_tomcat_nxd -c stop
pdm_tomcat_nxd -c start
```

5. To access the CA SDM web interface with SSL, use `https://computer_name:8443/CAisd/pdmweb.exe`. You can also derive the web services URL in a similar fashion.

Note: You can specify a port other than 8443 in the server.xml file.

6. View and install the SSL certificate to access CA SDM.

The system is configured to use Secure Socket Layer (SSL) on a Tomcat web server.

7. (Optional) If you are accessing CA SDM with Internet Explorer, and the Windows Server 2003 is configured for SSL, the browser requires additional configuration and restart. On the Internet Options, Advanced Tab, clear the following options in the Security section:

- Check for server certificate revocation (requires restart)
- Do not save encrypted pages to disk

After you restart the browser, you can access the CA SDM using the SSL-enabled Tomcat server.

Configure a Secondary Server for Web Screen Painter

If you want a secondary server to use Web Screen Painter, add the @NX_WSP_CGI_URL to the nx.env file as follows:

1. Open the \$NX_ROOT/NX.env file for editing.
2. Locate the following line:
@NX_WEB_CGI_URL=http://hostname/CAisd/pdmweb.exe
3. Add the following parameter after this line:
@NX_WSP_CGI_URL=http://hostname/CAisd/pdmweb.exe
4. Save the changes.
5. Recycle the CA SDM services.

The secondary server can use Web Screen Painter.

More information:

[How to Modify the System Environment](#) (see page 318)

How to Deploy CMDBf Web Services

After you install CA SDM, you can deploy CMDBf web services.

To deploy CA CMDB web services

1. Verify that the web server is up and running.
2. Navigate to the CMDBHOME\sdk\websvc\CMDBF directory.
3. Run deploy_cmdbws.bat.

The CMDBf web services are deployed and started.

Note: For more information about CMDBf web services deployment, see the *CA CMDB Technical Reference Guide*.

Stop a Server (Windows)

You can stop a primary or secondary CA SDM server in a Windows environment.

To stop a CA SDM server in a Windows environment

1. Select Services from the Control Panel.

The Services window appears.

2. Select the service you want to stop, and click Stop:

CA SDM Server

Controls the primary server.

CA SDM Remote Proctor

Controls a secondary server.

The server is stopped.

Stop a Server (UNIX)

You can stop a primary or secondary CA SDM server in a UNIX environment.

To stop a CA SDM server in a UNIX environment

1. Close all CA SDM clients.
2. Stop the server by running the *pdm_halt* utility.

The server is stopped.

More information:

[pdm_halt--Terminate Daemons or Stop Services](#) (see page 972)

Chapter 3: Defining Business Structure

This section contains the following topics:

[How to Define the Business Structure](#) (see page 57)

[Define the Business Infrastructure](#) (see page 60)

[Multi-Tenancy](#) (see page 63)


How to Define the Business Structure

Business structure is a logical representation of the organization within which CA SDM operates. To customize your installation, you can set up the business structure objects to best suit your needs.

You set up the following objects to define the business structure:

- Sites
- Locations
- Organizations

The following diagram illustrates the relationship between the objects in the business structure:



Because an organization has (references) a location, which in turn has a site, you must define the objects in the following sequence:

1. Sites
2. Locations
3. Organizations

This way, when defining each object, you can select from existing objects at lower levels in the hierarchy.

Note: For more information about how to define each object, see the *Online Help*.

Contacts

You can define every person who uses your system on a regular basis, such as analysts and customers. These people are referred to as contacts.

Groups

A group is a collection of contacts that represent a specific area of responsibility within your service desk. Defining groups lets you assign responsibility for resolving a ticket when that responsibility is shared among several individuals.

For example, you might have a Dallas Human Resources group, which is responsible for dealing with personnel issues in the Dallas office of your organization. When an employee in that office has a problem, you can assign it to the Dallas Human Resources group for resolution.

You can associate request areas, locations, and a work shift with a group to be used in determining when the contacts in the group can accept automatic assignment of a request.

Sites

A *site* is a group of locations. An example of a site might be a city in which your enterprise has one or more physical locations, or a region in which you have a customer base that you support. Because locations reference sites, you should define sites before locations. This sequence lets you use the new site in the location.

Locations

Locations identify a specific physical place and let you track and administer areas that concern your enterprise, such as the address of a particular company, an office address city, a campus, a building, or even a floor of a building.

Locations are the only objects with which you can associate a physical address—all other objects, such as organizations, configuration items, and contacts, derive their addresses from their associated locations.

You should define locations before organizations; this sequence lets you use the new location in the organization.

You can use locations to assign requests automatically. To enable automatic assignment based on location, assign locations to request areas and groups. The request area locations match a request to a location based on the location of the request's configuration item or affected end user. The group locations are then used to select a group that is eligible for automatic assignment of requests in that location.

Note: On the Administration tab, tickets are referred to as request/incident/problem areas. For brevity, they are referred to here as request areas.

Organizations

Organizations describe internal departments and divisions or external companies that can be assigned to tickets, configuration item classes, and contacts (for example, they are used to identify the company with which a customer is associated).

You can assign a default service type to an organization that is assigned automatically to tickets when the organization is specified. This default assignment lets you associate a specific level of service to a ticket based on the assigned organization.

You can define configuration items for organizations. configuration item definitions let you specify the hardware, software, and services that are used by the organization. Linking configuration items to an organization (shared CIs) supplements the feature of linking CIs directly to a contact (private CIs). When you designate a contact as the affected end user of a request, you can then select a CI for the request from separate lists of shared and private CIs.

Define the Business Infrastructure

An important aspect of implementing your service desk using CA SDM is to define your business infrastructure by setting up the following:

- Configuration item families and classes
- Manufacturers and models
- Service statuses
- Vendors and vendor types

The following information provides a general description for each object and explains how the object is used in the product.

Note: For more information about how to define each object, see the *Online Help*.

Object Definition Order

You start at the bottom levels of the object hierarchy when you define objects. This way, when you define objects at higher levels, you can select from existing objects at lower levels in the hierarchy. For example, because a class has (references) a family, you define families first, followed by classes. Similarly, because configuration items are the top of the hierarchy, you define them last, after you define all the supporting objects. Therefore, you define data objects in the following orders:

Define First	Define Second	Define Third
Families	Classes	Configuration items
Manufacturers	Models	Configuration items
Service statuses	Configuration items	
Vendor types	Vendors	Configuration items

Families and Classes

Families classify configuration items by type and assign meaningful attributes for each one. *Classes* identify general categories of configuration items that your enterprise supports. Families are broad categories of configuration items such as hardware, software, and services. Classes are more specific categories within the broader family category. For example, the family hardware might contain classes such as, modem, router, repeater, and bridge.

Organizing your configuration items into families and classes makes it easier to manage your configuration items. For example, you can generate a list of configuration items that belong to a particular family or class.

Manufacturer and Models

Manufacturers identify the manufacturers of the various configuration items of concern to your enterprise. *Models* contain specific information about the products that a particular manufacturer provides to your enterprise. For example, you might define as a manufacturer a particular software company. Then, you would define as models each one of the applications that the company provides for your enterprise.

Defining manufacturers and models makes it easier to manage your configuration items. For example, you can generate a list of models provided by a particular manufacturer and generate a list of configuration items of a particular model.

Service Status

Service status identifies the readiness condition of configuration items. Examples of service status include: *in service*, *in repair*, or *discontinued*. Defining service status lets you track the availability and use of configuration items in your enterprise. For example, you can generate a list of configuration items that are currently in repair.

Vendor Types and Vendors

Vendor types are classifications for vendors that identify the type of company providing configuration items. For example, you might classify vendors that you lease configuration items from as leasers, while classifying vendors that provide service to you as providers.

Vendors identify the companies that supply your enterprise, including the type of company and a primary contact. Besides being referenced by configuration items, you can also reference a vendor in a user's contact record.

Defining vendor types and vendors gives you a convenient way to organize your configuration items. For example, you can generate a list of vendors that fall under a particular vendor type and generate a list of configuration items from a particular vendor.

Configuration Items

Configuration items are the devices, software, and services that make up your business infrastructure. The information associated with a configuration item uniquely identifies the configuration item and indicates its precise location. Configuration items can be associated with contacts (private configuration items) and organizations (shared configuration items). Configuration items let you identify, view, and specify the following:

- Identify configuration items by name, class, and family
- Specify inventory information
- Specify additional properties to define the configuration item
- Log and view comments associated with the configuration item
- Specify location information for the configuration item
- Specify service information, such as a service type, for the configuration item
- View and define contacts and organizations assigned to the configuration item
- Identify hierarchical and peer-to-peer relationships between configuration items
- View tickets associated with the configuration item

You are not required to define all this information for configuration items. However, if you define an optimal amount of configuration item information, you get a clearer picture when impact analysis is performed for the IT organization.

More information:

[Create a Configuration Item](#) (see page 450)

[Contact, Location, and Organization CIs](#) (see page 453)

[CI Relationships](#) (see page 456)

External Asset Management Tools

You can integrate your CA SDM installation with other asset management tools such as CA NSM, CA Asset Management, and CA APM. The asset management features of these tools from CA SDM include the following:

- CA NSM provides the `pdm_nsmimp` utility (for Windows only) to add asset information to CA SDM.
- CA Asset Management provides a complete set of hardware and software inventory functions. When viewing a CI in the CA SDM client, you can click the Asset Viewer button to display additional information that has been discovered about the asset.

Integration between CA SDM and CA Asset Management is activated when the products are installed on the same MDB.

- CA Asset Portfolio Management provides a complete set of asset lifecycle functions. When viewing a CI in the CA SDM client, you can click the Asset Viewer button to display additional financial information about the asset.

Integration between CA SDM and CA Asset Portfolio Management is activated when the products are installed on the same MDB.

Note: For more information about integrating your CA SDM installation with other asset management tools, see the *Implementation Guide*.

Multi-Tenancy

Multi-tenancy is the ability for multiple independent tenants (and their users) to share a single implementation of CA SDM. Tenant users only interact with each other in defined ways, as specified by their *roles* and *tenant hierarchies*. Typically, unless granted access by a role or hierarchy, each tenant views the CA SDM implementation as solely for its own use and cannot update or view another tenant's data.

Multi-tenancy allows tenants to share hardware and application support resources, which reduces the cost of both, while gaining many benefits of an independent implementation.

Note: For information about installing and configuring multi-tenancy, see the *Implementation Guide*.

Service Provider

The *service provider* is the primary tenant (owner) in a CA SDM multi-tenancy installation. The first tenant added to a CA SDM installation is always the service provider tenant. The service provider tenant cannot have a parent tenant.

CA SDM associates the privileged user (typically ServiceDesk on Windows, or srvcdesk on Linux/UNIX) with the service provider tenant.

Only the service provider tenant can perform any of the following CA SDM tasks:

- Set CA SDM options
- Set Knowledge Management options
- Set Support Automation options
- Create tables or columns

- Create, edit, or delete tenants
- Allow tenants to have subtenants
- Update public data

Note: An administrator can grant tenant users access to data other than their own. Non-Service Provider tenant analysts can only have access to their own tenant and sub-tenants, unless their function access is updated to include the tenant of the analyst. For example, a role definition can set separate read and write access to certain tenant groups for users within that role.

Important: The first created tenant is set to be the service provider; after that, the Service Provider check box and Record Status field are read-only.

More information:

[Service Provider Administration](#) (see page 64)

[Tenant Access](#) (see page 67)

[Create a Tenant](#) (see page 76)

Service Provider Administration

The service provider can permit tenants to administer their own settings. Tenant administrators have access to a subset of administration tasks that is the same for all tenants, so the ADMIN_TREE table is not itself tenanted. Instead, the default Tenant Administrator role defines the administration functions that are available to tenant administrators.

To designate a user as a tenant administrator, select the Tenant Administrator for that user's role.

How Multi-Tenancy Works

When multi-tenancy is active, you can grant each contact access to all tenants (public), a single tenant, or a tenant group (user-defined or product-maintained). A contact's role controls access, which specifies read and write access independently. Because tenant access is role-dependent and a contact can change roles during a session, contact tenant access also can change.

If multi-tenancy is installed, most CA SDM objects include a tenant attribute that specifies which tenant owns the object. Objects fall into three groups, depending on their tenant attribute and how it is used:

Untenanted

Defines objects without a tenant attribute. All data in these objects is public.

Examples: Priority and urgency.

Tenant Required

Defines objects with a tenant attribute that cannot be null (enforced by CA SDM, not the DBMS). All data in these objects is associated with individual tenants; there is no public data.

Examples: Ticket tables (Request, Issue, and Change Order).

Tenant Optional

Defines objects with a tenant attribute that can be null. Some of the data in these objects is public, and some is associated with specific tenants. Each tenant's view of the object is a merged view of the public data and their tenant-specific data.

Examples: Category and location.

When a user queries the database, CA SDM restricts the results to objects belonging to tenants that the user is authorized to access. This restriction applies in addition to any data partition restrictions that are in effect. This means that you never see data in tenant required and tenant optional tables except for the data that belonging to tenants that you are permitted to access.

When a tenant user asks to create or update a database object, CA SDM verifies that the object belongs to a tenant that the user's current role allows them to update, and that all foreign key (SREL) references from the object to other objects are to public (untenanted) objects; to objects from the same tenant; or to objects from tenants in the tenant hierarchy above the object's tenant. That is, a tenanted object is allowed to reference objects belonging to its parent tenant, to its parent's parent, and so on.

If a user that creates an object has update access to multiple tenants, the user must specify the tenant explicitly, either directly or indirectly.

Note: There is one exception to the SREL reference restriction. Certain SREL references (such as the assignee of an incident) are permitted to reference objects that belong to tenants in their containing object's tenant hierarchy. Such references are designated as `SERVICE_PROVIDER_ELIGIBLE` in the CA SDM object schema (the Majic). The `SERVICE_PROVIDER_ELIGIBLE` flag makes a difference only if the service provider tenant is not in the tenant hierarchy above the object's tenant; if the service provider tenant is in the hierarchy, tenant validation rules permit service provider references.

A service provider user asking to create or update an object is subject to the same restrictions as tenant users, except that service provider users can be authorized to create or update public objects. The active role of the service provider user controls this authorization.

Note: If CA SDM limits a user from updating tenant data, an error message can announce a data partition limitation. If you receive this error message, either data partition or multi-tenancy restrictions are in effect.

The Multi-Tenancy Option

You activate multi-tenancy by installing one of the following multi-tenancy options:

- **off**—Multi-tenancy is not in use. Multi-tenancy features are not available, and objects do not have a tenant attribute. This option is the default setting at a new CA SDM installation.
- **setup**—Multi-tenancy features are in effect for administrators, so that tenant-related objects and attributes are visible and editable. However, CA SDM does not enforce tenancy restrictions, and non-administrator users see no changes. This setting allows an administrator to prepare for multi-tenancy by performing such tasks as defining tenants or assigning objects to tenants without impacting normal use of CA SDM.
- **on**—Multi-tenancy is fully operational. All users see the UI changes appropriate to them, and CA SDM enforces tenancy restrictions.

Note: For more information about installing and implementing multi-tenancy, including additional options for its enforcement, see the *Implementation Guide*.

Tenant Information

You create and update tenants when you install multi-tenancy (in either setup or full enforcement mode). Information maintained for a tenant is similar to the data maintained for Organization, except for the following two attributes:

Logo

Provides a URL to an image file with the logo of the tenant. The logo is shown both on the Tenant Detail page itself, and as a substitute for the CA logo on web forms displayed by a tenant user or showing an object associated with the tenant.

Service Provider

Indicates whether the tenant is the service provider. The service provider tenant is always the first tenant added. When the administrator adds the first tenant, the following occurs:

- The first tenant becomes the service provider. This designation cannot be changed.
- The privileged user (usually ServiceDesk) and all system contacts (such as System_AHD_Generated) are set to belong to the new service provider tenant

Note: The system user "Administrator" is added in Windows only and is not assigned a tenant. The privileged user must assign a tenant to Administrator manually.

Tenant Access

The role of a CA SDM user governs both access authorization and the user interface. The set of roles available to users depends on their access type. Multi-tenancy lets you control the tenant or tenant group that a user can access within the role.

The Role Detail page provides Tenant Access and Tenant Write Access drop-down lists on its Authorization tab. Tenant Access is view-only, and Tenant Write Access also allows create and update.

You can assign the following associations to roles:

Same As Tenant Access (Tenant Write Access Only)

Sets Tenant Write Access to be the same as the Tenant Access setting. Default for Tenant Write Access and only valid for Tenant Write Access.

All Tenants

Removes tenant restrictions. CA SDM allows a user in a role with this access to view any object in the database (read access) or create and update (write access) any tenanted object in the database. When users with All Tenant access create an object, CA SDM requires that they select the tenant of the new object.

Single Tenant

Sets a role's tenant access to a named tenant. When this option is selected, a second field appears on the web UI that allows selection of a specific tenant. CA SDM restricts a user in a role with this access to view (read access) or create and update (write access) only those objects associated with the named tenant. This selection is valid for either Tenant Access or Tenant Write Access.

Tenant Group

Sets a role's tenant access to a user-defined or system-maintained tenant group. When the Tenant Group option is selected, a second field appears on the web UI that allows selection of a specific tenant group. CA SDM restricts a user with the role to view (read access) or create and update (write access) only those objects associated with one of the tenants in the group. When a user with tenant group access creates an object, CA SDM requires that they select the tenant for the new object. This selection is valid for either Tenant Access or Tenant Write Access.

Contact's Tenant

Sets a role's tenant access to the tenant of the contact using it. CA SDM restricts a user in a role with this access to view (read access) or create and update (write access) only those objects associated with their own tenant. This selection is valid for either Tenant Access or Tenant Write Access.

Contact's Tenant Group (Analyst Only)

Sets an analyst's role access to the tenant group that the analyst works with, as specified on the analyst's contact record. If the user with the role is not an analyst, this selection has the same effect as Contact's Tenant. It is valid for either Tenant Access or Tenant Write Access.

Contact's Subtenant Group

Sets a role's tenant access to the Subtenant group of the contact using it. CA SDM restricts a user in a role with this access to view (read access) or create and update (write access) only those objects associated with their own Subtenant group. This selection is valid for either Tenant Access or Tenant Write Access.

Contact's Supertenant Group

Sets a role's tenant access to the Supertenant group of the contact using it. CA SDM restricts a user in a role with this access to view (read access) or create and update (write access) only those objects associated with their own Supertenant group. This selection is valid for either Tenant Access or Tenant Write Access.

Contact's Related Tenant Group

Sets a role's tenant access to the Related Tenants Group of the contact using it. CA SDM restricts a user in a role with this access to view (read access) or create and update (write access) only those objects associated with their own Related Tenants Group. This selection is valid for either Tenant Access or Tenant Write Access.

All users can view public data, regardless of their current role's access rights. The Update Public check box controls whether a service provider user in the role has the authorization to create or update public data. Tenant users (users belonging to a tenant other than the service provider) cannot update public data, regardless of their role.

More information:

[Edit Tenant Access for a Role](#) (see page 69)

Edit Tenant Access for a Role

You can assign or edit tenant access for a role.

To edit tenant access for a role

1. Navigate to Security and Role Management, Role Management, Role List.
The Role List appears.
2. Click a role.
The Role Detail page appears.
3. Click Edit.
The Update Role page appears.
4. Select options for Tenant Access and Tenant Write Access.
Note: Exercise caution if you select different options for these settings.
5. Click Save.
The updated tenant access options are saved for the role.

Tenant Terms of Usage

A terms of usage statement presents the end user with an initial page statement when they log in to CA SDM. The statement reminds the user about the proper use of the product. The user must agree to the terms before they can continue to log in to CA SDM. Entries are written to the standard log and in the user event log after the attempted session login.

You can perform the following terms of usage actions:

- Create, update, and delete a terms of usage statement.
- Associate a terms of usage statement with a tenant.
Note: You must enable multi-tenancy and configure one or more tenants before you can associate a terms of usage statement with a tenant.
- Force the end user to accept the statement every time they log in.
- Let the end user ignore the initial statement by presenting a blank terms of usage statement.

More information:

[How to Configure Terms of Usage](#) (see page 70)

How to Configure Terms of Usage

The terms of usage statement presents the end user with an initial page statement when they log in to CA SDM. The statement reminds the user about the proper use of the product. The user must agree to the terms before they can continue to log in to CA SDM. If the end user selects Accept, CA SDM proceeds with the login and displays the main form. If the user selects Reject, CA SDM returns to the login. Entries are written to the standard log and in the user event log after the attempted session login.

Typically, you configure the contact tenant terms of usage statement. If the contact tenant is configured with an inactive terms of usage statement, the terms of usage is not configured, or if <empty> is selected in the Terms of Usage drop-down list, CA SDM displays the terms of usage statement for the tenants parent, grandparent, and so on. If no terms of usage statement is found at any level, CA SDM proceeds with the login. If you configure a tenant with a blank terms of usage statement, CA SDM proceeds with the login and displays the main form.

You can configure terms of usage as follows:

1. Enable multi-tenancy.
2. Configure one or more tenants.
3. Define a terms of usage statement.
4. Update a tenant to use the specific terms of usage statement.

Note: For detail information about creating and modifying terms of usage statements, see the *Online Help*.

User Interface Impact

Installing the multi-tenancy feature changes the user interface, depending on the authorization and tenant access associated with the user's role. The changes affect both tenant users and service provider users.

Tenant Users

If the role of a user is restricted to a single tenant and the user is not an administrator, you can substitute a custom tenant logo for the default CA Technologies logo on all pages. This substitution depends on whether a logo is defined on the Tenant Detail page for the tenant, and so is optional by tenant.

The only user interface change for non-administrator tenant users is that menu items or buttons allowing update or edit (the Edit button, or the Create New button on a list page) are suppressed for public objects, because tenant users are not authorized to update a public object.

Tenant Administrators

Tenant administrators with Single Tenant access who view tenant-optional objects see an additional interface change. List pages for these objects automatically include a Public column specifying whether the list row is public data. In addition, the first element in the search filter is a Public Data drop-down list containing the following selections:

- Include (default)
- Exclude
- Only

A tenant administrator with access to multiple tenants sees a Tenant column on list pages for any tenanted object. This column takes the place of the Public column in lists of tenant optional tables.

Users Who Can View a Tenant Group

If a user's role allows view access to multiple tenants, or a service provider user has Update Public authorization, CA SDM list forms change as follows:

Untenanted objects

Untenanted objects contain only public data. A service provider user is allowed to create or update an untenanted object only if their role has Update Public authorization. If not, the UI suppresses menu items or buttons allowing update or edit, such as the Edit button itself, or the Create New button on a list page. Tenant users cannot update public objects, and these users never see an Edit or Create New button on a list page for an untenanted object.

Tenant-required objects

Tenant-required objects contain only data associated with a particular tenant. List forms for these objects automatically include a Tenant column after the last link column. In addition, the search filter contains a tenant selector allowing the user to restrict a list to a single tenant.

Tenant-optional objects

Tenant-optional objects contain both public and tenant-specific data. List forms for these objects automatically include a Tenant column (a blank tenant indicates a public object). In addition, the search filter contains both a tenant selector and a Public Data drop-down list (the same one seen by tenant administrators).

Note: If tenant required tables incorrectly contain untenanted data in a multi-tenancy system, a public data drop-down list appears in tenant required tables and you get the following message: “AHD05358 There were nn untenanted active xxx objects at Service Desk startup.”

Users Who Can Update Multiple Tenants

If the user's role allows access to multiple tenants, or a service provider's user role has Update Public authorization (typical for an analyst that works for a service provider), detail pages change as follows:

Untenanted objects

Untenanted objects contain only public data. There are no changes to their detail pages for a service provider user with Update Public authorization. If the user is in a role without Update Public authorization, or does not belong to the service provider, read-only pages for untenanted objects have no Edit button.

Existing Tenant-required objects

Tenant-required objects contain only data associated with a particular tenant. Detail pages for existing tenant-required objects show the object's tenant as part of the standard page header.

Tenant-optional objects

Tenant-optional objects contain both public and tenant-specific data. The detail page for these objects depends on whether the user belongs to the service provider and is in a role with Update Public authorization:

- If a service provider user's role has Update Public authorization, the detail page is the same as that for tenant-required objects.
- If the user's role does not have Update Public authorization, or the user does not belong to the service provider, detail pages for public objects do not have an Edit button. Other detail pages are the same as those for tenant-required objects.

Support Automation Impact

The impact of multi-tenancy on your support environment depends on tenant and role restrictions placed on end users and analysts. The Service Provider manages read/write permissions for both public and tenant-specific data. For example, an analyst can handle assistance sessions from the public queue and a tenant-specific queue, but the analyst can only use Live Assistance tools enabled for each tenant.

For end users with Support Automation access, do not configure the end user to have write access to a tenant that is not a sub-tenant of the owning tenant of the end user, unless the foreign key (FK) group is altered to include the owning tenant. If the end user selects a login tenant, or is invited through a ticket in a tenant that does not meet this criterion, they receive an error when they try to access the Support Automation end user client. This restriction does not apply if the owning tenant of the end user is the service provider tenant.

For analysts with Support Automation access, do not configure the analyst to have write access to a tenant that is not a sub-tenant of the owning tenant of the analyst, unless the foreign key (FK) group is altered to include the owning tenant. If the analyst attempts to handle an end user, or invite an end user from a ticket, in a tenant that does not meet this criteria, the analyst sees an error.

Analysts and end users without read access to their tenant cannot launch the Support Automation client. For analysts, a warning message appears in CA SDM in this case, such as from the main Support Automation tab.

You can use the following roles to manage Support Automation users:

Support Automation Analyst

Provides end-user support using Live Assistance. The Service Provider determines the appropriate tenant access and can enable Live Assistance tools and read/write access to automated tasks.

Important! If a non-Service Provider analyst has write access to a parent, sibling, or unrelated tenant, function access must be updated for that tenant. Analysts without read access to their tenant cannot launch the Support Automation analyst client, and a warning message appears in CA SDM, such as from the main Support Automation tab or a ticket.

Support Automation Administrator

Configures the Support Automation environment for analysts and end users. The Service Provider determines your tenant access and lets you view a tenant drop-down list on List and Detail forms. These forms let you select specific tenants or public data when searching, creating, and modifying Support Automation data in a multi-tenancy environment.

Note: Objects such as queues, privacy levels, and chat presets are tenant optional.

Knowledge Management Impact

The impact of multi-tenancy on your knowledge environment depends on the tenant restrictions that are placed on users:

Tenant Users

Substitutes the tenant's logo for the default, if the role is restricted to a single tenant.

Tenant Administrators

Allows administrators view both public and tenant-specific data. List pages for these objects automatically include a Public column specifying whether or not the list row is public.

When searching for knowledge, the filter contains a Public Data drop-down list with selections of Include (default), Exclude, and Only.

Note: Objects such as approval process templates, categories, documents, files and forums are tenant optional.

Knowledge Categories and Documents

Knowledge Documents and Knowledge Categories are both tenant optional. Consider the following information for tenant optional and public objects:

- Public Knowledge Documents can only be added under public Knowledge Categories.
- Public Knowledge Categories can only be added under public categories.
- Tenant categories can be added under public categories and under a tenant's categories.
- Only tenant documents can be added under tenant categories.
- Public and tenanted documents can be added under public categories.

Note: The Cut/Copy/Paste category functionality is allowed only if the source and destination have the same tenant, or if the destination is public.

Repositories are defined as tenant optional, so the administrator can create different repositories for different tenants.

Embedded images are allowed only when the document and image are set to the same tenant. Attachment Folders, Attachments and Attachments to document links are also defined as tenant optional.

FAQ Rating

When viewing the FAQ rating for tenant users, consider the following information about public documents:

- Public documents are viewed by a larger audience than tenant users.
The FAQ rating of public documents is higher than a tenant's specific document.
- Each tenant has different needs, so usage patterns are different between tenants.

The Top Solutions on the CA SDM Self-Service home page displays the top five public documents, as well as a tenant's top five documents.

You can configure the Top Solutions by navigating to Knowledge, Solution Survey, FAQ Settings on the Administration tab.

Knowledge Report Card

The Knowledge Report Card allows analysts and administrators to view various metrics such as document creation, publication, hits, and votes. This report is for a predetermined time period, per analyst, category and organization.

When using the Knowledge Report Card to provide information for a role having single tenant access, the data is limited by the tenant criteria.

How to Use Multi-Tenancy

Use the following procedures to administer CA SDM multi-tenancy features.

View Tenants

You can view tenants from any Tenant List.

Note: This feature is available only if multi-tenancy is installed (in either on or setup mode).

To view a tenant

1. On the Administration tab, select Security and Role Management.
2. Click Tenants.
The Tenant List appears.
3. (Optional) Select a tenant from the Tenant List.
The Tenant Detail page appears.

Note: Tenant lists also are displayed on pages such as Tenant Group Detail and Tenants Affected by CI.

Create a Tenant

You can use the product to create a tenant.

To create a tenant

1. Select Security and Role Management, Tenants on the Administration tab.

The Tenant List page appears.

Note: The Security and Role Management, Tenants option is available only when multi-tenancy is installed (either on or setup).

2. Click Create New.

The Create New Tenant page appears.

3. Complete the editable fields if necessary:

Name

Displays the tenant name.

Service Provider

Identifies whether a tenant is the service provider. The first created tenant is always the Service Provider.

Tenant Number

(Information Only) Displays the tenant number. This field is not used by CA SDM.

Record Status

Sets the tenant to Active or Inactive.

Parent Tenant

Specifies another tenant above this tenant, making this tenant a *subtenant* in a tenant hierarchy.

Subtenants Allowed

Allows this tenant to have subtenants. The tenant cannot modify the setting.

Tenant Depth

(Information Only) Indicates the tenant depth of this tenant.

Supertenant Group

(Information Only) Identifies the system-maintained tenant group that contains this tenant and all tenants above it in the tenant hierarchy.

Subtenant Group

(Information Only) Identifies the system-maintained tenant group that contains this tenant and all tenants below it in the tenant hierarchy.

Foreign Key Group

(Information Only) Identifies the system-maintained tenant group that contains tenants that can be referenced from an SREL in data that belongs to this tenant. The foreign key group is the same as the supertenant group.

Related Tenant Group

(Information Only) Identifies the system-maintained tenant group consisting of both the supertenant and subtenant groups for this tenant.

Terms of Usage

Specifies the Terms of Usage statement for the tenant.

Logo

Specifies the URL for the tenant logo file, which can be any web image type.

Location

Displays the Location lookup page.

Contact

Displays the Contact lookup page.

Note: If no contact is associated with the respective tenant, the Email Address and Pager Email Address fields are inactive.

4. Click Save.
The tenant is created.
5. Close the window.
6. Right-click the Tenant list and click Refresh.
The Tenant List is updated and displays the created tenant.
7. (Optional) To assign this tenant to user-defined tenant groups, click Update Tenant Groups on the Tenant Groups tab.

Edit a Tenant

You can edit a tenant from the Administration tab.

Note: This feature is available only if multi-tenancy is installed (in either on or setup mode).

To edit a tenant

1. On the Administration tab, select Security and Role Management.
2. Click Tenants.
The Tenant List appears.

3. Right-click on a tenant, and click Edit.

The Update Tenant page appears.

Note: When you edit an existing tenant, the Subtenants and Configurations Items tabs appear.

4. Make the appropriate changes and click Save.

5. Close the Update Tenant page.

The Tenant List appears.

6. Right-click and select Refresh.

The updated Tenant List appears.

Tenant Hierarchies

A *tenant hierarchy* is a structured tenant group that is system-created or modified when you assign a *parent tenant* to a tenant. The tenant becomes a *subtenant* of the parent and higher tenants (if any) in that hierarchy.

Note: The service provider can create multiple unrelated hierarchies, or none. Even in a system with tenant hierarchies, you can define standalone tenants.

A subtenant typically represents a subdivision within its *supertenant*s. A subtenant can have its own business rules and data, and supertenant data is "pushed" to the subtenant automatically on a read-only basis.

CA SDM supports a tenant hierarchy of unlimited depth. However, the *service provider* can specify a limit on the total number of tenants and the depth of tenant hierarchies (default is four levels). The service provider also determines whether individual tenants can have subtenants.

Note: The service provider can participate in tenant hierarchies, but this is not required. The service provider cannot have a parent tenant.

Create a Subtenant

Subtenancy allows you to build and modify tenant hierarchies for organizational and data-sharing purposes. To place a tenant into a tenant hierarchy, you assign it a parent tenant.

To create a subtenant

1. On the Administration tab, select Security and Role Management, Tenants.

The Tenant List appears.

Note: The Security and Role Management, Tenants option is available only when multi-tenancy is enabled.

2. Click an existing tenant to Edit, or click Create New.

The Tenant Detail page appears. Enter any required data or changes.

3. Select a Parent Tenant.

Note: The Parent Tenant drop-down only displays tenants that are allowed to have subtenants.

4. Click Save.

The tenant is a subtenant of the parent tenant.

Note: When a tenant is a subtenant, it belongs to the Subtenant group of the parent tenant, as do the subtenants (if any) of that subtenant, and so on. The parent tenant joins the Supertenant group of the subtenant, as do the supertenant(s) (if any) of that supertenant, and so on. Each joins the Related Tenants group of the other.

System-Maintained Tenant Groups

CA SDM generates and maintains three tenant groups automatically for each tenant in a tenant hierarchy (*tenant* is the tenant name):

- *tenant_subtenants* (tenant, its child tenants, and their lower subtenants)
- *tenant_supertenant*s (tenant, its parent tenant and its higher supertenant(s))
- *tenant_relatedtenants* (entire single hierarchy)

System-maintained tenant groups can be used like user-defined tenant groups. However, only their names and descriptions can be modified.

View Tenant Groups

You can view tenant group information to show group members.

Note: This feature is available only if multi-tenancy is installed (in either on or setup mode).

To view the tenant group list

1. On the Administration tab, select Security and Role Management.
2. Click Tenant Groups.

The Tenant Group List appears.

Note: You can choose to view or hide system-maintained tenant groups.

3. (Optional) Select a tenant group from the list.

The tenant group information appears.

4. Modify the tenant group if necessary.

Create a Tenant Group

You can use the product to create a tenant group.

To create a tenant group

1. On the Administration tab, select Security and Role Management.
2. Click Tenant Groups.

The Tenant Group List appears.

Note: The Security and Role Management, Tenant Groups option is available only when multi-tenancy is installed (either on or setup).

3. Click Create New.

The Create New Tenant Group page appears.

4. Complete the following fields:

Tenant Group Name

Displays the name of the tenant group.

Record Status

Sets the tenant group as active or inactive.

Description

Displays a description of the tenant group.

5. Click Save.

The tenant group is created.

6. Close the window.

The Tenant Group List appears.

7. Right-click the Tenant List and select Refresh.

The Tenant Group List is updated.

8. Click Update Tenants on the Tenant Group Detail page to add tenant members to the group.

Edit a Tenant Group

You can edit a tenant group to manage its members and detail information.

Note: This feature is available only if multi-tenancy is installed (in either on or setup mode).

To edit a tenant group

1. On the Administration tab, select Security and Role Management.
2. Click Tenant Groups.
The Tenant Group List appears.
3. Right-click on a tenant group, and click Edit.
The Update Tenant Group page appears.
4. Make the appropriate changes and click Save.
5. Close the window.
The Tenant Group List appears.
6. Right-click in the window and select Refresh.
The updated Tenant Group List appears.

Tenant Data Assignments

CA SDM displays the tenant in the same format in both the View and Edit versions of a detail page for an existing object, because the tenant for an existing object cannot be changed from the web interface.

When you edit a tenanted object, drop-down lists on the edit page are automatically restricted to values that are public, owned by the same tenant as the base object or any tenants above it in the tenant hierarchy, or owned by the service provider (if the drop-down list applies to a SERVICE_PROVIDER_ELIGIBLE attribute).

There are no changes on the detail page for lookup fields associated with a tenanted object. If a user with access to multiple tenants clicks a lookup link to a tenanted table, the web engine automatically restricts the lookup to values appropriate for the attribute, and displays a banner message on the pop-up search or list page.

Note: Tenant restrictions are not displayed in the search filter, and they cannot be changed by the user.

Tenant becomes a selector (either a lookup or a drop-down list) when you ask to create a tenant-required object.

If the tenant field is empty, you can specify a tenant value directly by filling in the field, or indirectly by specifying a value for a tenant-implying attribute (such as Affected End User). The interface displays the following suffixes:

(T)

Indicates an attribute that is tenant-implying; that is a lookup to a tenant-required table.

(TO)

Indicates an attribute that is optionally tenant implying that is, a lookup to a tenant-optional table.

web.cfg properties control the text of these indicators.

Except for the tenant attribute itself, tenant-implying attributes are always shown as lookups, even if created with a dtlDropdown macro.

CA SDM automatically sets the tenant when you look up or autofill a tenanted value into any tenant-implying field (except that filling a SERVICE_PROVIDER_ELIGIBLE field with a reference to a service provider object does not set the tenant). After the tenant is set, lookups for tenant-implying fields are restricted in the same way as existing tenanted objects.

Note: Until you save the object, the tenant field remains editable, and you can change the tenant by directly updating it. When you change the tenant, CA SDM automatically clears any tenant-implying fields containing references to objects belonging to the previous tenant.

CA SDM typically initializes the Tenant selector to empty. You can change this behavior in several ways:

- Open the Create New page from a page such as the Quick Profile, which pre-populates a tenant-implying field
- Set the Retain Tenant user preference
This is a new user preference that initializes the tenant for new objects to the same tenant as the last detail page viewed or updated, or in the last list page search filter restriction.
- Open the page with a URL that explicitly specifies a tenant
This is not provided in any predefined URL, but is available to allow sites to create menu items or buttons that specify a tenant.

Note: If you create configuration items from another CA Technologies product (such as CA APM) or the command-line interface, then the object is Public.

More information:

[Create a Tenanted Object](#) (see page 82)

Create a Tenanted Object

The service provider can add tenant-specific data to objects such as issues, requests, change orders, and so on. You can add a tenant to a ticket (such as an incident) that is created from a Scoreboard tab.

To add tenant data to an incident

1. Click File, New Incident.
2. Complete any of the following steps:
 - a. Select the tenant from the Tenant drop-down.
 - b. Click Affected End User (or any other tenant-implying field).

The Contact Search page appears. Search for a user; you can filter the search by tenant.
 - c. Enter a name into the Affected End User field.

The tenant data completes automatically.
3. Continue to create the incident.

Activity Notifications

Activity Notifications control both the contents of notifications and which contacts receive notifications for various events in the history of a ticket.

In a multi-tenancy environment, the notification rule is a tenant-optional object. Public notification rules apply to all tickets; tenanted rules apply only to tickets with the same tenant as the rule, or to tenants in its subtenant hierarchy. The tenanting restriction is applied in addition to any condition specified with the rule itself.

More information:

[Copy Notification Rules](#) (see page 83)

Copy Notification Rules

Default Notification Rules are stored as public objects. If multi-tenancy is installed, you must create a copy of the Notification Rule for each of the tenants, otherwise the Update Contacts option is restricted.

To copy a notification rule

1. Navigate to Notifications, Notification Rules on the Administration tab.

The Notification Rules List appears.
2. Click a rule in the Symbol column.

The Notification Rule Detail page appears.
3. Click File, Copy.

Resume updating the Notification Rule.

Repositories

The repository (`doc_rep`) object is tenant-optional. Tenants can define their own repositories, and it is possible to define public repositories for objects such as attachments to public knowledge documents. Each tenant can have its own default repository, and you can specify a default public repository.

All attachments are either public or associated with a single tenant. If a tenant does not have its own default repository, the public repository is displayed as the default for its tenanted objects.

Chapter 4: Implementing Policy

This section contains the following topics:

[Policy Implementation](#) (see page 85)

[Notifications](#) (see page 85)

[Email Administration](#) (see page 120)

[Service Level Agreements](#) (see page 141)

[Security](#) (see page 153)

[User Authentication](#) (see page 163)

[Internal Logs](#) (see page 166)

[CA SDM Integration](#) (see page 166)

[Data Partition Associations](#) (see page 167)

[Surveys](#) (see page 173)

[Web Services](#) (see page 176)

Policy Implementation

A key component of configuring your service desk is implementing policies in a way that best matches your business process.

CA SDM provides a predefined policy implementation that is appropriate for some sites, and serves as a good starting point for others. Review the default implementation in all the policy definition areas to determine which parts might meet your needs and which parts you need to modify.

Notifications

With CA Service Desk Manager, you can automatically notify key personnel about ticket activities (researching, escalating) and events (opening a ticket, for example). You can also notify key personnel about Knowledge Report Card (KRC) and Support Automation Assistance Sessions. When a significant activity or event occurs, CA Service Desk Manager creates a notification message that does the following:

- Identifies the ticket activity or the notification event
- References the ticket
- Includes other optional information
- Can identify potential contacts

You can view a notification message for a ticket because of a system action. A system action includes opening, closing, or modifying a ticket through its history information.

Setting up automatic notifications involves the following tasks:

- Defining activity notifications that determine the types of activities that generate notifications.
- Defining object contact notifications that determine the object contacts that can be used to send notifications in an activity notification.
- Identifying the methods used to send messages.

More information:

[How Activity Notifications Work](#) (see page 88)

[Object Contact Notifications](#) (see page 88)

[Notification Methods](#) (see page 89)

Activity Associations

You can define activity types that you associate with an attribute of an object. Activity associations let you do the following:

- Track and report on changes to an object attribute.
- Mark an activity association as internal.

Internal activity associations restrict notifications to those contacts whose access type allows them to view internal logs. Contacts that are not authorized to view internal logs are not notified, even if they are included as recipients of the notification. All future activity logs created with the associated activity are marked as internal.

Note: An object attribute can have only one activity association.

Activity Notifications

An activity is an action that someone performs, such as resolving a ticket, sending a managed survey, running the Knowledge Report Card, and so on. Activity notifications identify the types of activities that can cause notification to occur automatically for tickets. You can define a notification to send for a specific activity. For example, you can send an email notification to an affected user for additional information. The user can reply to the email with the information requested through a device such as a PDA. Even daily activities such as returning a call, canceling or closing a record, increasing priority, or updating status can result in a notification being sent.

Typically, activity notifications signify the start or completion of an action about which you want to notify key personnel, such as when tickets are escalated, transferred, or monitored. When a ticket activity occurs, CA SDM checks to see if automatic notification is configured for the activity. If selected, the activity notification defines what message to send to the notification process.

You can also perform the following actions on an activity notification:

- Attach events.
- Create notification rules that set guidelines for the system to follow when sending notifications. For example, set up an activity notification that notifies the previous assignee (or group) when something changes in a ticket.
- Set up a manual activity notification that updates a ticket when the user replies.

Note: You cannot maintain different Manual Notify activity notifications per tenant, or copy a Manual Notify activity notification.

More information:

[Internal Logs](#) (see page 166)

[Notification Rules](#) (see page 96)

[Previous Assignee Notifications](#) (see page 112)

[Configuration Item Notifications](#) (see page 114)

[How Activity Notifications Work](#) (see page 88)

How Activity Notifications Work

You can define a notification that is sent for a specific activity. Activity notification can be defined for Requests, Incidents, Problems, Change Orders, Issues, Managed Surveys, Knowledge Documents, Knowledge Document Comments, Knowledge Report Card, Assistance Session, Contacts, Configuration Items.

When you define an activity notification, you can configure the following:

Events

Defines the event list for this activity notification. Events are procedures that an issue management system follows after a certain amount of time has elapsed.

Notification Rules

Defines new notification rules for requests/incidents/problems, change orders, issues. You can also modify the default rules that appear in the Notifications Rules list.

Survey

Defines a survey notification that allows the recipient of this activity notification to click a URL to view a survey. Surveys let you collect and analyze customer feedback.

Note: For more information about creating activity notifications for Requests, Change Orders, Issues, Managed Surveys, Knowledge Documents, Knowledge Document Comments, Knowledge Report Card, Assistance Session, Contacts, or Configuration Items, see the *Online Help*.

Object Contact Notifications

Object contact notifications let you notify recipients based on the current value of a field on the ticket. Instead of identifying a person to notify, as in a notification method, you identify an object. For example, you can identify the To field to ensure that notification goes to the person currently identified in the To field, even if the value has changed since the ticket was defined.

To create an object contact notification

1. On the Administration tab, select Notifications, Object Contact Notifications.

The Object Contact Notification List page appears.

2. Click Create New.

The Create New Object Contact Notification page appears.

3. Complete the following fields:

Symbol

Defines a unique identifier for the object contact notification.

Status

Specifies if the object contact notification is active or inactive.

Object Type

Displays the name of the object to which the attribute applies.

Object Attribute Name

Provides the name of the object contact notification (in the Symbol field) in Majic, which is internal CA code. The attribute name depends on the Object Type selection:

- If the Object Type is Issue or Workflow Task, the attribute name is assignee, requester, or group; these are the attribute names in the chg objects and they map to fields in the Change_Request tables.
- If the Object Type is an Issue Activity log, the attribute name must start with the attribute name in the activity log object that links it to a specific instantiation of the chg object. The attribute name could be change_id.group.

Description

Describes the object contact notification.

4. Click Save.

The new object contact notification appears in the Object Contact Notification List when you redisplay the list.

Note: For more information about configuring object contact notifications, see the *Online Help*.

Notification Methods

Notification methods describe how notification messages are delivered to users. Notification methods are scripts that are invoked based on activity notifications. The scripts use information supplied as variables to notify personnel or other systems of something that has occurred. For example, you can write a script that sends voice mail to the analyst assigned to a request to indicate that the request has been escalated.

Notification methods can be assigned for each contact record. The system looks up the notification method to use for specific contacts.

The standard notification methods for CA SDM are as follows:

- Email sends messages by electronic mail directly to the recipient through Simple Mail Transport Protocol (SMTP) mail. Messages are also sent to the recipient notification log.
- The notification function sends messages to the recipient notification log that the recipient can access when CA SDM is running.
- Pager email sends email to an address maintained by a paging system provider. The email text typically displays on an alphanumeric pager.

You can also create your own notification methods. For example, you can send a notification to a particular printer for periodic collection or by pager. To create a notification method, create a shell script that includes notification variables, and then enter the new notification method into CA SDM.

Note: For more information about notification methods, see the *Implementation Guide*.

More information:

[Notification Setup Based on Contact Type](#) (see page 180)

pdm_mail Utility--Send Email Information

The pdm_mail utility is used in notifications to send emails by sending email information to the pdm_mail_nxd process. The pdm_mail utility can also be used for commands, but not to both. If no parameters are used, then the default behavior of using the NX_NTF_xxxx variable to pass parameters is in effect.

For email, the utility is invoked as follows:

```
pdm_mail [-i [-s subject] [-e email_address] [-q]] [-p] [-M] [-F] [-T] [-B] [-H] [-N] [-R] [-h]
```

-i

Uses STDIN instead of NTF variables. The following parameters are used for STDIN email behavior only:

-e

Specifies the email address (for recipient).

-s

Specifies the subject for the email.

-q

Disables the display prompt for STDIN.

-p

Utilizes pager logic. This option includes using the pager email address instead of the regular email address. Only the plain text version of the notification is used (no HTML).

-M

Uses plain text only (no MIME) in body.

-F

Specifies the From address of the email.

-T

Specifies the Reply-To address of the email.

-B

Specifies the body charset. This might be useful for pagers that do not support the UTF-8.

-H

Specifies the header charset. This might be useful for pagers that do not support the UTF-8.

-N

Specifies the Delivery Status Notification (DSN) Notify option.

-R

Specifies the Delivery Status Notification (DSN) Return option.

-h

Displays help on the utility.

For commands, the utility is invoked as follows:

Command to mail server

```
pdm_mail -c option [parameter]
```

Command to mail eater

```
pdm_mail -x option [parameter]
```

check_interval

(-x only) Changes the check mail interval to a specified value (in seconds).

report_interval

Changes the report interval to a specified value (in seconds).

report_now

Forces a report to the logs. The counter is not reset.

send_q

(-c only) Sends local mail queue to remote mail server.

trace

Turns tracing on or off.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

Email Notifications

Email notifications help analysts communicate with employees and end users.

Employees and end users within an environment are sometimes mobile, and interact with CA SDM using email. You can set up CA SDM so that this type of user can do the following:

- Update a request/incident/problem/change order/issue through a reply to an email notification sent from CA SDM.
- Approve or reject a workflow task (classic or CA Workflow) using email.

Notifications let you communicate with stakeholders for a ticket as follows:

- Email notifications can show all recipients of the notification, advising the user of who the interested parties are.
- Analysts can send a manual notification to a temporary email address that is currently not associated with a contact. A temporary address is helpful, for example, when an end user is out of the office or having difficulty accessing their standard email account.
- Analysts can send attachments with a notification to communicate additional details to a user.

Temporary Email Addresses

A temporary email address is an address that is not associated with a contact in the system.

Temporary email addresses are useful in circumstances such as the following one:

1. An end user is out of the office or is having difficulty accessing their standard email account.
2. The analyst wants to use email to track interactions with the user.

3. The analyst sends a manual notification to a *temporary* email address for the user.
4. The analyst can view the activity log, which is updated with the manual notification.

Note: For information about how to set up temporary email addresses, see the *Online Help*.

Allow Temporary Email Addresses

You can let users send manual notifications to temporary email addresses. For example, analysts can send email to an address that is not associated with a contact record. Recipients cannot reply to temporary addresses when their email address is not associated with a contact record or does not have permission to update the ticket.

Note: Temporary email addresses are always SMTP email addresses, and are supported only when the Preferred Method supports SMTP.

To allow temporary email addresses

1. On the Administration tab, browse to Options Manager, Notifications.
The Options List appears.
2. Click notification_allow_temp_address.
The notification_allow_temp_address Options Detail page appears with default values set.
3. Click Edit.
The notification_allow_temp_address Options Update page appears with default values set, and you can edit the Description.
4. Click Install.
The notification_allow_temp_address Options Detail page appears.
5. Click Close Window.
6. Restart the CA SDM server.
Temporary email addresses are now allowed for manual notifications.

Disallow Temporary Email Addresses

By default, users cannot send manual activity notifications to temporary email addresses. For example, analysts cannot send email to an address that is not associated with a contact record. Administrators install the `notification_allow_temp_address` option to allow temporary email addresses. If the option is installed and you want to disallow manual notifications to temporary email addresses, you can uninstall the option.

Note: Temporary email addresses use SMTP email addresses and the Preferred Method notification method.

To disallow temporary email addresses

1. On the Administration tab, browse to Options Manager, Notifications.

The Options List appears.

2. Click `notification_allow_temp_address`.

The `notification_allow_temp_address` Options Detail page appears.

3. Click Edit.

The `notification_allow_temp_address` Update Options page appears.

4. Click Deinstall.

The `notification_allow_temp_address` Options Detail page appears.

5. Click Close Window.

6. Restart the CA SDM server.

The option is uninstalled and temporary email addresses are now disallowed for manual activity notifications.

Mailbox Polling

If an error occurs on the outgoing mail server, email notifications are not sent and queue in the `%NX_ROOT%\site\mail_queue` directory. When the mail server becomes active again, after an interval it processes and sends email. You can change the interval to recycle the email that was queued when the mail server was busy.

Notification email messages that the outgoing mail server fails to send are resubmitted until you do one of the following:

- Stop the Mail Daemon (`pdm_mail_nxd`) that handles outbound email notifications.
- Manually delete the messages from the `%NX_ROOT%\site\mail_queue` directory.

Set the Email Retry Interval Variable

You can define the time interval (in seconds) to retry failed attempts to send outgoing email to the mail server.

Note: CA SDM does not retry sending messages that the outgoing mail server accepts, but cannot be delivered. For these messages, the outgoing mail server retry capabilities and policies, if any, are in effect.

Retries are on a per-message basis. If the mail server is unavailable for a period, each message is retried when its own timer expires, rather than all the messages being sent at once. However, if you restart the outgoing mail daemon, all unsent messages attempt to be sent at that time, and if they all fail to be sent, their retry timers are all reset at the same time.

The setting (NX_EMAIL_RETRY_INTERVAL) in the NX.env file controls the retry interval. You can change the default retry interval setting on one or more servers.

To set the email retry interval

1. Navigate to the \$NX_ROOT directory on the server.
2. Use a text editor such as WordPad to open the NX.env file.
3. Modify the value of to the NX_EMAIL_RETRY_INTERVAL interval you want as follows:

```
NX_EMAIL_RETRY_INTERVAL=number_of_seconds
```

NX_EMAIL_RETRY_INTERVAL

Defines the time interval (in seconds) to retry failed email attempts.

number_of_seconds

Specifies the number of seconds for each email retry interval. The default time is 600 seconds (10 minutes). The minimum value that you can use is 20 seconds. If you set a value that is less than the minimum of 20 seconds or more than 2000000 seconds, the default value of 10 minutes is used.

4. Save and close the file.
5. Restart the CA SDM service.

The change takes effect.

Notification Rules

A notification rule sets guidelines for the system to follow when sending notifications. With notification rules, you can specify who is automatically notified, and under which circumstances.

You can define notification rules for requests, incidents, problems, change orders, issues, contacts, configuration items, and global and specific tenants. Each notification rule contains the following components:

Condition Macro

Configure to represent the desired values of one or more fields on a service desk ticket. The condition is evaluated at different points during processing to determine whether an action occurs.

Message Template

Contains a prototype of the message that is generated and sent to contacts when the notification is sent. A notification rule must contain a message template and they can be reused.

Object Contacts

Displays the object contact derived from ticket information, such as assignee, end-user, group, and so on. If the condition is met, these object contacts receive notification.

Contacts

If the condition is met, these contacts receive notification. Contacts are the database records that represent the users of the system.

Contact Types

Displays the contacts type that you want to receive the notification by default, for example, Analysts, Employee, Customer, and so on.

Related Activity Notifications

Displays a list of related activity notifications that use the notification rule.

Note: Users can receive multiple notifications for an activity notification that contains multiple rules for which they meet the conditions.

Notification Rules Configuration Options

You can specify the following options when you define a notification rule:

- **Default Notification Rule**—You can specify a default notification rule that automatically notifies all contacts identified in the referenced field on the ticket. To view the list of default notification rules you can modify, select the Administration tab, Notifications, Notification Rules.

- **Site-defined Condition Macro**—You can implement a *site-defined condition macro* to create your own condition that only notifies specific contacts under certain circumstances. To create a condition macro from the Administration tab, select Events and Macros, Macro Types, and then Site-defined Condition on the Macro Type list.
- **New Contacts**—When you define a notification rule, you can identify new contacts to receive notifications on the following tabs:
 - **Object Contacts**—Displays the available organizations, vendors, and configuration items for the selected Object type that receives notification about tickets.
 - **Contacts**—Displays the individuals who are added to the notification rule, regardless of their affiliation with the ticket.
 - **Contact Types**—Displays the users who are defined within the notification rule with the same classification, such as analyst or customer.

The Update button on each tab lets you search for new contact records to associate with the rule.

Note: Before you implement a notification rule, you must configure the appropriate contacts for your business structure.

- **Message Template**—Every notification rule must contain a message template. You can define your own message template or specify one of the default message templates available on the Message Template list page. To view this page, select the Administration Tab, Notifications, Message Templates.
- **Configuration Item Notifications**—When you define a notification rule, you can identify the person responsible for maintaining a configuration item to receive notifications when a ticket is created.
- **Previous Assignee Notifications**—When you define an activity notification, you can specify values in the notification rule that notify the previous assignee or group when an event activity occurs.

More information:

[How to Define the Business Structure](#) (see page 57)

[Example: Create a Message Template](#) (see page 105)

[Previous Assignee Notifications](#) (see page 112)

[Object Contact Notifications](#) (see page 88)

[Site-Defined Condition Macro Overview](#) (see page 98)

Site-Defined Condition Macro Overview

A manager wants to be informed about all ticket activities that occur for Priority 2 issues pertaining to their department. The administrator adds the manager as a specific contact on selected types of activity notifications, such as Initial and Log Comment. With this implementation, the manager receives all selected activity notifications for the entire system.

In a high-volume department, such as technical support, the unwanted e-mail traffic could be significant. To control their mailbox, the manager must implement mailbox rules to filter the notifications and hope that nothing important is missed.

The administrator defines a site-defined condition macro that only notifies the manager when the Priority is set to 2 and the Request Area is set to Applications (manager's department). The new condition is specified in the notification rule and the manager is added as the contact. A message template is defined that describes the ticket activity and the rule is added to the activity notification.

The activity occurs when the priority of a ticket is set to 2 and the request area is set to Applications. When this activity happens, the rule is implemented and the condition is evaluated. If the condition evaluates to True, a notification message describing the ticket activity is sent to the manager. Unwanted notifications are no longer sent to the manager.

Example Notification Rule

Notification rule examples provide introductory information to help you accomplish the following:

- Create a site-defined condition macro
- Create a notification rule
- Add the rule to the activity notification

When the activity occurs, the corresponding condition attached to the rule is implemented. If the condition is met (evaluates to *true*), a notification message that describes the ticket activity is sent to all contacts.

More information:

[Example: Create a Site-defined Condition Macro](#) (see page 99)

[Example: Create a Notification Rule](#) (see page 100)

[Example: Add the Rule to the Activity Notification](#) (see page 101)

Example: Create a Site-defined Condition Macro

In this example, you create a site-defined condition macro that checks for the specified condition.

1. On the Administration tab, browse to Events and Macros, Macros.
The Macro List page appears.
2. Click Create New.
The Create New Macro page appears.
3. Complete the following fields:
 - Symbol: Enter *Application P1*.
 - Macro Type: Select Site-defined Condition.
 - Object Type: Select Request.
4. Click Continue.
The Create New Macro Applications P1 page appears.
5. Complete the following fields:
 - Macro Description: Type *Macro=Request, Priority=1, Request Area=Applications*.
 - If all Conditions succeed return: Select TRUE.
 - Record Status: Select Active.
6. Click Save.
The Applications 1 Macro Detail page appears.
7. Click Add Condition.
The Create New Atomic Condition page appears.
8. Create an Atomic Condition to check for Priority=2 and complete the following fields:
 - Sequence: Type *10*.
 - Description: Type *Check the Priority*.
 - Select an Attribute: Select Priority.
 - Choose Operator: Select Equals.
 - Select Attribute or Data Value: Select Data Value.
9. Select 2 for the data value and click Save.
10. From the Macro Detail page, click Add Condition and complete the following fields:
 - Sequence: Type *20*.
 - Description: Type *Check Request Area=Applications*.

- Select an Attribute: Select Request Area.
- Choose Operator: Select Equals.
- Select Attribute or Data Value: Select Data Value.

11. Select Applications for the data value and click Save.

The new condition displays on the Conditions tab of the site-defined condition macro.

Example: Create a Notification Rule

In this example, you associate the site-defined condition macro with the notification rule, attach a message template, and add the Manager as a contact type.

1. On the Administration tab, browse to Notifications, Notification Rules.
The Notification Rules List page appears.
2. Click Create New.
The Create New Notification Rule page appears.
3. Complete the following fields:
 - Symbol: Type App1 Notify Managers.
 - Object Type: Select Request.
 - Description: Type *Notify Managers of Priority 1 Application Outages*.
4. Click Save & Continue.
5. Click the Condition link, and select the new Site-defined Condition Macro called Application P1.
6. Click Message Template.
The Message Template List page appears.
7. Select a message template from the list. Ensure that the Auto Notification option is enabled in the message template.
8. On the Contact Types tab, click the Update Contact Type button.
The Contact Type Search page appears.
9. Click Search.
The Notification Update Recipients list appears.
10. Select Manager from the list on the left, and click the contact selection button (>>).
11. When the contact is in the list on the right, click OK.
12. Save the notification rule.

The new notification rule appears in the Notification Rule List when you redisplay the page.

Example: Add the Rule to the Activity Notification

In this example, you enable the Auto Notification option in the message template, and add the notification rule to the activity notification.

1. On the Administration tab, browse to Notifications, Activity Notifications.
The Activity Notifications List page appears.
2. Select the Initial activity notification from the list.
The Initial Activity Notification Detail page appears.
3. Verify that the Object Type is set to Requests, so that the notebook appears in the configuration information for requests.
4. On the Notification Rules tab, click Update Notification Rules.
The Notification Rules Search page appears.
5. Type *App P1 Notify Manager* in the Symbol field and click Search.
The Notification Rules Assigned Update page appears.
6. Select *App P1 Notify Manager* from the list on the left, and click the contact selection button ($\geq$).
- The notification rule is in the list on the right.
7. Click OK.
8. Click the Condition link and select the new App P1 Notify Managers notification rule. (You can add additional rules by clicking the Update Notification Rules button.)
The Activity Notification page appears.
9. Ensure that the Auto Notification option is set in the message template as follows:
 - a. Click the Notification Rule link on the Notification Rules tab of the activity notification.
The Notification Rule page appears.
 - b. Click the Message Template link.
The Message Template Detail page appears.
 - c. Click Edit.
The Update Message Template page appears.
 - d. Verify that Auto Notification is enabled.
10. Click Save.

When the activity occurs, the corresponding condition attached to the rule is implemented. If the condition is met (evaluates to true), a notification message that describes the ticket activity is sent to all managers.

Example Default Notification Rules

CA SDM contains several default notification rules. You can view the default notification rules by navigating to Notifications, Notification Rules on the Administration tab.

More information:

[Example: Use the Request Priority 1 Rule](#) (see page 102)

[Example: Use the Notify Current and Previous Assignee of a Change Order Rule](#) (see page 103)

[Example: Use the True Priority Escalation Rule](#) (see page 104)

[Example: Use the Active request/incident/problem Rule](#) (see page 104)

Example: Use the Request Priority 1 Rule

The notification rules provide flexibility to the end user if they want to receive a notification for a request with priority 1. The *Request Priority 1* rule runs if the ticket type is a request and if the priority is set to 1. This notification is sent only when the condition attached to the rule is evaluated as true and when the message template attached to the rule is used.

Note: You can modify the condition and message template of a Notification Rule by opening the rule and modifying the Condition field with a macro. For more information about modifying conditions, see the *Online Help*.

The product also contains default Notification Rules named "Incident Priority 1" and "Problem Priority 1". The major differences between these rules and Request Priority 1 are conditions for the ticket types. For example, Request Priority 1 has the condition with a sequence "20" attached to it which limits the notification rule only to Requests

By default, no Object Contact or Contact or Contact Types are attached to the rule. The "Auto Notification" field on the Message Template attached to the rule is also set to No by default. To send a notification, you must add an Object Contact or Contact, or Contact Type to the rule and the Auto Notification field on the Message Template should be changed to Yes.

Example: Configure a rule to send a notification only for Requests with Priority set to 1

The rule should be connected to both the Initial and Escalate Activity Notifications. This helps ensure that the notification is received for a new priority 1 Request and that a Request is escalated to a priority 1.

1. Verify that either the default rule attached to the Initial and Escalate activity notification is either removed or the Auto Notification option is set to No.

This prevents the notification from being sent out for these default rules.

2. Open the Request Priority 1 rule.

3. Add an Object Contact or a Contact or a Contact Type to the rule.
4. Click Message Template and select the Auto Notification option.
5. Attach the Request Priority 1 rule to the Initial and Escalate activity notifications.
6. Do the following:
 - Create a new request and set the priority to 1.
Verify that the user receives the notification.
 - Create a new request and set the priority to a different priority, such as 2.
Verify that the user does not receive the notification.

Example: Use the Notify Current and Previous Assignee of a Change Order Rule

The *Notify Current and Previous Assignee of a Change Order* rule notifies both the current and previous assignees of the change order during a transfer. You can also create a similar rule for other ticket types.

Example: Notify both the current and previous assignees of a change order transfer

1. Verify that either the default rule attached to the Initial activity notification is either removed or the "Auto Notification" on the Message Template attached to the rule is set to No.

This prevents notifications from being sent out for these default rules.
2. Open the Notify Current and Previous Assignee of a Change Order rule.
3. Click the Message Template "Notify Current and Previous Assignee of a Change Order" and select the Auto Notification option.
4. Attach the "Notify Current and Previous Assignee of a Change Order" rule to the Transfer activity notification.
5. Do the following:
 - a. Create and save a change order with an assignee.
 - b. Transfer to a new assignee.Verify that both users receive the transfer notification.

Example: Use the True Priority Escalation Rule

The 'Escalate' Activity Notification occurs when the priority of a ticket changes, not just when the priority increases. Users are notified of escalations as well as de-escalations. The Notification Rules now provide a way to discriminate between a change from priority 2 to priority 1 and a change from priority 1 to priority 2

The *True Priority Escalation* rule evaluates the current and previous values of the Priority field. This rule notifies the attached Object Contact, Contact and Contact Types attached to the rule only when the priority of an Issue ticket changes from a lower priority to a higher priority.

Note: You can create a similar rule for other ticket types. The notification is only sent when the condition attached to the rule is evaluated as true and uses the message template that is attached to the rule. By default, no Object Contact, Contact or Contact Types are attached to the rule. The Auto Notification field on the Message Template attached to the rule is also set to No by default.

Example: Receive notification when priority changes from a lower to a higher priority for an Issue ticket type

1. Verify that either the default rule attached to the Escalate activity notification is removed or the Auto Notification option on the Message Template attached to the default is set to No.

This prevents notifications being sent out for these default rules.
2. Open the True Priority Escalation rule.

Add an Object Contact or a Contact or a Contact Type to the rule.

Verify that the Contact can receive notifications.
3. Click True Priority Escalation on the Message Template field and check the Auto Notification option.

Attach the rule to the Escalate activity notification.
4. Create a new issue with a priority, such as 5.
5. Change the priority to 4 and verify that the user received the notification.
6. Change the priority from 4 to 5 and verify that the user did not receive the notification.

Example: Use the Active request/incident/problem Rule

The *Active request/incident/problem* rule has a condition attached to it to evaluate to send the Initial notification only if the request/incident/problem is active.

This allows the customer to receive only the Close notification when an inactive request is created. This rule notifies the attached Object Contact/Contact and Contact Types attached to the rule only when an active request/incident/problem is created.

Note: You can create a similar rule for change orders and issue ticket types.

The notification is sent only when the condition attached to the rule is evaluated to true and uses the message template that is attached to the rule. By default, no Object Contact or Contact or Contact Types are attached to the rule. Also the "Auto Notification" field on the Message Template attached to the rule is defaulted to No.

Example: Receive a notification when you close an Incident

1. Verify that either the default rule attached to the Initial activity notification is either removed or the "Auto Notification" on the Message Template attached to the default rule which in turn is attached to the Escalate activity notifications is set to No.

This prevents notifications from being sent out for these default rules.

2. Open the Active request/incident/problem rule.
Add an Object Contact or a Contact or a Contact Type to the rule.
Verify that the Contact can receive notification.
3. Click the Message Template "Active request/incident/problem "attached to the rule and check the" Auto Notification" checkbox field.
4. Attach the Active request/incident/problem rule to the Initial activity notification.
5. Create a new Incident and set the status field to Closed.
6. Verify that the assigned user configured to receive the notification does not receive the Initial notification.

Example: Create a Message Template

Every notification rule must contain a message template. You can create templates that contain default values to use for notification messages.

Example: Create a Message Template

This example demonstrates how to create a message template.

To create a message template

1. On the Administration tab, browse to Notifications, Message Templates.
The Message Template List appears.
2. Click Create New.
The Create New Message Template window appears.

3. Complete the following fields:

Tenant

Specifies the tenant associated with this template.

Note: Selecting global (shared) allows you to use the template for global messages.

Symbol

Defines a unique identifier for this record.

Object Type

Specifies the object type associated with this template.

Record Status

Specifies either active or inactive for the status of this template.

Auto Notification

Specifies that the notification associated with this template is sent automatically when the activity occurs.

Notify Level

Indicates the relative importance of sending this notification. Select Emergency, High, Low, or Normal.

Notification Message Title

Specifies the summary title of the message. You can use variables to insert values from the ticket when a notification is sent. For example, you can include a variable that inserts the issue number into the message title.

Notification Message Body

Specifies the content of the message. You can use variables to insert values from the ticket when a notification is sent. For example, you can use a variable that inserts the analyst name, end user name, and description into the message.

HTML Message

Specifies a message in HTML to display when the recipient email program supports HTML messages. If the recipient receives the message on a pocket device, such as a cell phone or Blackberry, the message displays in plain text only.

Edit HTML Message

Opens the HTML Editor, which allows you to edit your HTML message.

Quick View

Displays the message as it appears to the recipient.

HTML Source

Displays the message in HTML source code.

4. Click Save.

The new message template appears in the Message Template List.

Artifacts in Messages

You can use the ARTIFACT keyword to specify how artifacts are handled in outbound messages, message templates, notifications, and auto-replies. The ARTIFACT keyword uses the following values:

- **NONE**—Specifies no validation of artifacts. This value is the same as not using the keyword.
- **PROTECTED**—Validates a ticket against the hash for confirmation. If confirmation fails, the artifact is considered invalid and filtering fails when filtering searching for an artifact ({{object_id}}).
- **SECURE**—Decrypts the ticket number. If the value is not a valid password, the artifact is considered invalid and filtering fails when filtering is searching for an artifact ({{object_id}}).

Notification Codes and Phrases

Notification phrases let you add a standardized piece of information or text to a number of different notification messages, without having to enter and maintain the text or formulae separately in each notification template. For example:

Reply to this notification to add additional information to the ticket

Phrases standardize text for use in multiple message templates. For example, you can maintain a common phrase such as a confidential notice in a single record and use it in multiple message templates. Notification phrases are also useful for message replies, such as a Reply Notice, or a web URL link. CA SDM provides phrases and you can create your own phrases. You can set a phrase as active or inactive for use in a message template globally. (Notification phrases are inactive by default.) When a phrase is inactive, the phrase is suppressed in all message templates that use the phrase.

Notification phrases can also be used in the automatic responses to incoming email messages. The processing context for this type of message is different; omit the prefix (change_id., issue_id., call_req_id.) used in certain macros such as ref_num and web_url for phrases that the message uses. As a result, you cannot share notification phrases between notification templates and email automatic responses.

For example, some of the phrases that CA SDM provides are as follows:

Symbol	Code	Phrase
Notify History - Change	notify_history_chg	Click the following URL to view the Notification List: @{change_id.web_url}+HTMPL=chg_lr.html+INSTANCE=@{id}
Notify History - Issue	notify_history_iss	Click the following URL to view the Notification List: @{issue_id.web_url}+HTMPL=iss_lr.html+INSTANCE=@{id}
Notify History - Request/Incident/Problem	notify_history_cr	Click the following URL to view the Notification List: @{call_req_id.web_url}+HTMPL=cr_lr.html+INSTANCE=@{id}

Example: New Phrases

The following phrases are examples of phrases that you can create:

Symbol	Code	Phrase
Notify Confidential	ConfidentialNotice	This email and any files transmitted with it are for the sole use of the intended recipient(s) and contain information that may be privileged and confidential. Any unauthorized review, use, disclosure or distribution is prohibited. If you are not the intended recipient of this email, please delete this email and any files transmitted with it and notify the sender immediately.
Notify Incident Reply	IncidentReply	To add a comment to your Incident, just reply to this email or include the line below (on a line by itself). %incident_id=@{call_req_id.ref_num}
Notify Incident URL	IncidentURL	Click the following URL to view: @{call_req_id.web_url}

Note: Use separate phrases for the plain-text and HTML versions of message templates or email auto-replies. HTML consolidates most whitespace into single spaces, so line- or paragraph-breaks must be specified with tags. HTML tags included in plain-text versions of messages are displayed rather than processed.

Notification Phrase Syntax

You insert notification phrases in message templates and email reply messages using the following macro syntax:

```
@{notification_phrase[code].phrase}
```

code

Specifies the unique Code value, such as ConfidentialNotice, of the Message Phrases (usp_notification_phrase) table.

Note: The usp_notification_phrase table lists common phrases that notification message templates can use. For information about the usp_notification_phrase table, see the *Technical Reference Guide*.

For example:

```
@{notification_phrase[IncidentURL1].phrase}
```

```
@{notification_phrase[RequestReply].phrase}
```

When CA SDM locates the macro, the phrase text from the usp_notification_phrase table replaces the syntax. If no matching code exists (or if it is inactive), an empty string replaces the macro. No errors are written to the standard log (STDLOG), instead a warning message is logged to help you resolve potential problems.

Note: Embedding phrases within other phrases is limited to a maximum depth value that you configure by setting the NX_MAX_EXPAND_DEPTH [environment variable](#) (see page 318). This limitation prevents problems which can occur when processing phrases that are accidentally self-referenced (embed themselves) or circular-referenced (when a phrase embeds a phrase into which it is embedded).

How Notification Phrases Appear in a Message Example

This example demonstrates how notification phrases appear in a notification message. The example uses the following codes and phrases:

Code	Phrase
ConfidentialNotice	This e-mail and any files transmitted with it are for the sole use of the intended recipient(s) and contain information that may be privileged and confidential. Any unauthorized review, use, disclosure or distribution is prohibited. If you are not the intended recipient of this e-mail, please delete this e-mail and any files transmitted with it and notify the sender immediately.

Code	Phrase
IncidentReply	In order to add a comment to your Incident, just reply to this email or include the line below (on a line by itself). %incident_id=@{call_req_id.ref_num}
IncidentURL	Click the following URL to view: @{call_req_id.web_url}

The following message template includes the notification phrases:

@{call_req_id.type.sym} @{call_req_id.ref_num} Closed.
Assigned to: @{call_req_id.assignee.combo_name}
Customer: @{call_req_id.customer.combo_name}
Description: @{call_req_id.description}

@{notification_phrase[IncidentURL].phrase}
@{notification_phrase[IncidentReply].phrase}

@{notification_phrase[ConfidentialNotice].phrase}

The phrases appear in a message as follows:

Incident 1234 Closed.
Assigned to: Analyst, Joe
Customer: Doe, John
Description: This is a description of my problem

Click on the following URL to view:
<http://helpdesk/CAisd/pdmweb.exe?OP=SEARCH+FACTORY=chg+SKIPLIST=1+QBE.EQ.id=400723>

In order to add a comment to your Incident, just reply to this email or include the line below (on a line by itself).
%incident_id=1234

This e-mail and any files transmitted with it are for the sole use of the intended recipient(s) and contain information that may be privileged and confidential. Any unauthorized review, use, disclosure or distribution is prohibited. If you are not the intended recipient of this e-mail, please delete this e-mail and any files transmitted with it and notify the sender immediately.

Manual Notification Recipients List

You can set up a default set of Available Recipients that the Manual Notification composition page presents for requests, incidents, and problems. Available Recipients streamline the manual notify process for analysts because you can set up a list of contact objects (for example, Affected End User) or individual contact names for easy use as recipients of manual notifications.

Adding an Object Contact Recipient adds the individual contact names that the Object Contact represents to the Recipients list (consolidating any duplicate entries). The same contact can be referenced multiple times for different Contact Objects such as Assignee and Affected End User. Some entries, such as the Stakeholders List contact object, can add multiple entries to the Recipients list.

Contacts and Contact Objects remain in the Available Recipients list after you add them to the list. This behavior lets you remove recipients without affecting the initial Available Recipients list.

Example: How the Available Recipients List Works

The following examples demonstrate how default recipients streamline the manual activity notification process.

Object Contact Recipients includes the following entries:

- UserA belongs to the "Assignee" and "Affected End User" Contact Objects.
- Stakeholders List includes multiple contact names. For example, UserB and UserC.

You do the following actions:

- Add both Contact Objects that refer UserA to the Recipients list.
Only one UserA entry is listed in the Recipients list.

- Accidentally remove UserA from the Recipients list.

You do not have to refer to the ticket to get the UserA name and search to add it back. You can quickly add UserA to the Recipients list because UserA is listed in the Available Recipients list.

- Accidentally remove one contact name, such as UserC, from the Recipients List which came from the Stakeholders List.

You can add the Stakeholders List from Object Contact Recipients to add the contact name again. Because duplicate entries are consolidated, other Contacts in the Stakeholders List who were not deleted from the Recipients List are not affected.

Previous Assignee Notifications

You can define Previous Assignee or Group values for an activity notification that detects changes to key fields when a ticket is saved. Previous values let you notify a previous assignee when a ticket is transferred, or notify both the current and previous groups when the priority of a ticket is escalated.

The Previous value fields of a ticket are local fields that exist only in memory and not in the database. The fields are populated during the save operation of a ticket and cleared at the completion of notification processing. A previous value field is associated with a particular activity type through an activity association.

You can define Previous values that detect changes to the following key fields of a ticket:

Field	Requests, Incidents, Problems	Change Orders	Issues
Status	Yes	Yes	Yes
Active	Yes	Yes	Yes
Assignee	Yes	Yes	Yes
Request Area/Category	Yes	Yes	Yes
Group	Yes	Yes	Yes
Impact	Yes	Yes	Yes
Priority	Yes	Yes	Yes
Urgency	Yes	No	No
Severity	Yes	No	No

There are several contacts that you can specify for each object type (request, incident, problem, change order, or issue), which notify the current and previous contacts when an activity occurs.

- **Assignee**—Person assigned to handle the ticket.
- **Assignee Previous**—Person previously assigned to handle the ticket.
- **Group**—Group assigned to handle the ticket.
- **Group Previous**—Group previously assigned to handle the ticket.

After the notification rule is saved, the Assignee Previous and Group Previous fields display on the Object Contact Notifications List page.

Example: Configure Current and Previous Values for Key Fields

The following usage example describes how an administrator configures current and previous values for key fields to help ensure that the previous support representative is notified when a request is transferred away from them.

1. **Situation**—A support representative is frustrated because a ticket was transferred away from them and they were never notified.
2. **Task**—The administrator adds the Assignee and Assignee Previous object contacts to the notification rule for the Transfer activity notification. They attach a message template and specify the current and previous assignees to notify on the request form.
3. **Action**—When the request is saved, the Assignee and Assignee Previous fields of the request are populated. When the activity occurs (ticket is transferred), the condition for the rule is evaluated.
4. **Result**—If the condition is met, a notification message that describes the ticket activity is sent to the current assignee and the previous assignee.

Notify Contacts When a Ticket is Transferred Example

You can notify both the current and previous contacts when a CA SDM ticket is transferred.

Example: Notify both the current and previous contacts when a ticket is transferred

1. On the Administration tab, browse to Notifications, Activity Notifications.
The Activity Notifications List page appears.
2. Select the Transfer activity notification.
The Transfer Activity Notification Detail page appears.
3. In the Object Type field, select Requests/Incidents/Problems.
4. On the Notification Rules tab, under Symbol, select the Default Transfer Notification Rule.
The Default Transfer Notification Rule page appears.
5. On the Object Contacts tab, click Update Object Contacts.
6. Click Search.
The Notification Rule Update Recipients page appears.
7. From the Object Contacts list, select Assignee and Assignee Previous from the list on the left, and click the contact selection button (≥>).

The selected item is added to the list on the right.

Note: Use the CTRL or SHIFT keys plus the left mouse button to select multiple object contacts.

8. Click OK.
9. Save the notification rule.

The Object Contacts list displays the selected object contact.

10. On the Default Transfer Notification Rule page, click Message Template. Select a template and ensure that the Auto Notification option is enabled.
11. Create a request, specify an Assignee, and click Save.
12. On the Request Detail page, select Activities, and Transfer from the File menu.
13. Specify a new Assignee, and click Save.

The notification is sent to the current and previous assignees when the transfer activity occurs.

Configuration Item Notifications

A configuration item (CI) notification lets you define an activity notification that is associated with a specific CI that is associated with a specific CA SDM ticket. This feature lets you track information about the users, organizations, and vendors of a CI.

You can specify the CI object contacts on the Notification Rules Update Recipients page, such as CI Maint Org, CI Primary Contact, and so on.

Notify the Primary Contact of a Configuration Item for an Issue Example

You can define an activity notification for a primary contact that are sent for a specific CI for a specific CA SDM ticket.

Example: Notify the primary contact of a configuration item for an issue

1. On the Administration tab, browse to Notifications, Activity Notifications.
The Activity Notification List page appears.
2. Select the Initial Activity Notification from the list.
The Initial Activity Notification Detail page appears.
3. Select the object type you want to use.
4. On the Notification Rules tab, select the Default Notification Rule link.
The Default Notification Rule page appears.

5. Select the Default Message Template link and ensure that the Auto Notification option is enabled.
6. Select the Object Contacts tab, and click Update Object Contacts.
The Object Contact Notification Search page appears.
7. Click Search. A list of object contacts appears.
8. Select CI's Primary Contact from the list on the left, and click the contact selection button ($\geq$).
- The selected item is added to the list on the right.
You can use the CTRL or SHIFT keys plus the left mouse button to select multiple object contacts. You can add one object for a request and multiple objects for a change order or issue.
The object contact is in the list on the right.
9. Click OK.
10. Save the notification rule.
The Object Contacts list displays the selected object contact.
11. Complete the following tasks:
 - On the Service Desk tab, create or update an existing CI.
 - Add the primary contact listed on the Object Contacts tab. The selected object contact appears on the Configuration Items Detail page.
 - Add the CI to the Issue.

When an activity event occurs, the rule is implemented and the condition is evaluated. If the criteria for the condition is met, a notification message that describes the ticket activity is sent to all contacts of the CI associated with this notification rule.

Survey Notifications

For an activity notification, you can define separate survey notification details for tickets, managed surveys, knowledge documents, and knowledge comments and send it to the customer who initiated the activity.

Using survey notifications, you can do the following:

- Specify a new survey or use one of the default survey's available with the product.
- Use notification methods that describe how notification messages are delivered to users.
- Define a survey message for the customer.

When a customer receives notification of a survey, the message body automatically includes a URL that they can access from their web browser to find and fill out the survey form.

An activity log is generated both when a survey notification is sent and when one is received back from a customer.

More information:

[Define Survey Notifications](#) (see page 174)

How to Add URL Hyperlinks to Notifications

The `web_url` field in the `Change_Request` and `Workflow_Task` tables holds a URL value that allows a user to access a particular change order or workflow task through the web interface. When used in email notifications, a user can click the URL and go to the web interface without any further querying.

Before you can implement URL hyperlinks in notifications, configure your system as follows:

1. Install and configure the CA SDM web interface.

Note: For more information about configuring the CA SDM web interface, see the *Implementation Guide*.

2. Using the Options Manager, configure and install the `web_cgi_url` option to specify the location of the CA SDM web engine, for example, `http://hostname/scripts/pdmcgi.exe`.

More information:

[Options Manager Usage](#) (see page 317)

Notification Log Reader

The Notification Log Reader displays the notifications received for the logged-in user by date, urgency, and status. With the Notification Log Reader, you can do the following:

- Change the sort order and set menu options to have the Notification Log Reader appear automatically when new messages are received.
- Double-click a notification message to request that CA SDM display the detail page for the ticket associated with the notification.

- Monitor notification messages by entering specific selection criteria to query the database for analysis or for selection of notification messages based on data entered in the fields. For example, you can list only those notification messages that have not been cleared by changing the Message Status field to Not Cleared.
- Clear notification messages to keep your list of notifications to a manageable size. Cleared notifications are not displayed when you first access the Notification Log Reader, although you can display them, if needed.

Set Notification Log Reader Options

You can set options for the Notification Log Reader to define how you are notified when new messages are received for an issue.

To set options for the Notification Log Reader

1. On the ServiceDesk tab, browse to View, Log Reader.
The Notification Log Reader page appears.
2. Use the check box to the left of each notification to set the following options. You can select items to perform operations such as Clear Selected or Delete Selected.

Header

Displays the header of the message, which usually contains the number of the ticket and the activity type.

Start Date

Displays the date and time the notification was sent to your Log Reader window.

Status

Displays the status of the notification.

Urgency

Defines the level of urgency for the notification (low, normal, high, or emergency), which indicates the relative importance of different activities. Urgency levels are predefined; however, the system administrator is responsible for setting up other aspects of notification, such as notification methods and activity associations. The system administrator also defines the method of notification used for contacts and groups for each urgency level.

Message Text

The full message text for the notification.

The Log Reader displays any changes.

3. Click Close.

The Notification Log Reader page closes and the options are set.

Personalized Responses

You can create personalized responses and attach them to requests, issues, and change order records when adding activities to the record. For example, you can append a personalized response on the Status Change or Log Comment windows available from the Activities menu.

Create a Personalized Response

You can create a personalized response to append to requests, issues, and change order records.

To create a personalized response

1. From the Administration tab, navigate to Service Desk, Personal Responses.

The Personal Response list page displays.

2. Click Create New.

The Create New Personalized Response page displays.

3. Complete the fields on the page:

Response Owner

Specifies the contact who owns the response. If this field is left blank, the response is available to all analysts.

Response

Specifies the text delivered to all those who receive this response. This field can be up to 1000 characters long.

You can use variables in this field, for example:

Ticket ref_num: @{call_req_id.ref_num}

Assignee: @{call_req_id.assignee.combo_name}

Customer: @{call_req_id.customer.combo_name}

Description: @{call_req_id.description}

4. Select the type of records for which you want this response available. Click Save.

A personalized response is created.

Personalized Response Variable Substitution

Variables can be embedded in the text of a Personal Response. These variables allow information to be substituted from the corresponding Request, Change Order, Issue, Problem or Incident. The syntax of the variables is the same as is used elsewhere in the CA SDM product, such as in the Activity Notification Message Templates and the Manual Notify Activity Message Text. The information can only be substituted from the corresponding Request, Change Order, Issue, Problem or Incident. Activity Notification Message Templates and the Manual Notify Activity Message Text allow information from the Activity Log Record to be included as well.

Check boxes for each object type (Requests, Change Orders, Issues, Incidents, and Problems) allow Responses to be filtered during selection. If the object type is not checked, the Response is not available for that object. For example, if only the Request box is checked, the Response is only presented in Activities for a Request.

A single Response can be used for all object types (Requests, Change Orders, Issues, Problems or Incidents). Because each object has different attributes, information that does not apply to the object is not substituted (for example, attempting to substitute the Request Number in a Response for an Issue).

A Response text example and the variable substitutions that occur for each object type follows:

This is Request # '@{call_req_id.ref_num}'
This is Change Order # '@{change_id.chg_ref_num}'
This is Issue # '@{issue_id.ref_num}'

For a *Request*, the following substitution occurs:

This is Request # 'cr_demo:11'
This is Change Order # "
This is Issue # "

For a *Change Order*, the following substitution occurs:

This is Request # "
This is Change Order # 'chg_demo:3'
This is Issue # "

For an *Issue*, the following substitution occurs:

This is Request # "
This is Change Order # "
This is Issue # 'iss_demo:6'

By using the "Display this Response for" check boxes, you can create different versions of a Response with the appropriate substitution variables for the corresponding object (Requests, Change Orders, Issues, Problems or Incidents).

The format of the substitution variables for the different objects is as follows.

Object	Variable Format
Request / Incident / Problem	@{call_req_id.attr}
Change Order	@{change_id.attr}
Issue	@{issue_id.attr}

The substitution occurs when the Response is copied to the User Description field. The Response is copied after it is selected from the Personalized Response drop-down list and the drop-down list loses focus.

Email Administration

Email lets you communicate with end users, such as employees or customers. End users can reply to email notifications to update or create tickets using email from a computer, or other device such as a mobile phone or smartphone. For example, a reply to an email notification sent to an end user for additional information can update a ticket with the information requested.

The following features handle email communication to and from the end user:

- **Email Notifications**—Handles outbound email.

Consider the following information about outbound email:

- The Mail Daemon (pdm_mail_nxd) sends outbound mail notifications, using Simple Mail Transfer Protocol (SMTP).
- You configure the mail server SMTP options using Options Manager, Email.
- You can use default email notification methods (or configure new methods) to send mail directly to the recipient through SMTP mail.

- **Mailboxes**—Handles inbound email.

Consider the following information about inbound email:

- The Mail Eater (pdm_maileater_nxd) retrieves inbound email for ticket creation and updates, using Post Office Protocol (POP3) or Internet Message Access Protocol (IMAP or IMAP4).
- CA SDM mailboxes define email accounts (inboxes).
- Mailbox rules define how each mailbox handles inbound mail.

More information:

[Email Notifications](#) (see page 92)

Mailboxes

CA SDM provides a default mailbox that connects to the mail server, and which requires additional configuration such as setting values for hostname, user name, password, and so on.

Some of the ways that you can define a mailbox to handle inbound mail are as follows:

- Configure a mailbox to use either the POP3 or IMAP protocol.
- Create multiple mailboxes to let individual tenants or organizations have different mailboxes, with separate behaviors for each one.
- Use mailbox rules to filter messages and set specific behavior, define automatic replies, and set specific actions (for example, to update tickets), based on the contents or properties of the message.
- Use mailbox rules to supply Text API defaults and establish email interfaces with other software. You can configure parameters (such as category, assignee, and so on) specifically for the interface.
- Use notification phrases in mailbox rules to create common elements for automatic replies across multiple rules.
- Use mailbox policies to help protect against mail-based abuse.

More information:

[Mailbox Polling](#) (see page 94)

Mailbox Rules

Mailbox rules let you configure any actions, replies, or both, that must occur for messages retrieved from a mailbox. The table (usp_mailbox_rule) maintains the rules which apply to the connection to each mail server account (usp_mailbox)—rules belong to specific mailboxes. You can delete rules, or make them inactive to disable them.

Note: When you create mailbox rules, you associate them with a specific mailbox—another mailbox cannot share those rules. If you want to use rules that belong to one mailbox for a different mailbox, recreate them for the other mailbox.

Mailbox rules also let you do the following:

- Specify various behaviors for messages from each mailbox, based on the contents of the messages.
- Supply Text API default parameters (such as category, assignee, and so on) for message processing and routing.
- Set message filters based on fixed keywords or regular expressions.

- Use a common notification phrase for the automatic responses for numerous mailbox rule definitions, as in the following example:

Thank you for submitting your request.
@{notification_phrase[RequestReply].phrase}

The phrase appears in replies to email delivered to any mailboxes that includes the rule.

More information:

[Mailbox Rule Actions](#) (see page 122)

[Pattern Matching in Mailbox Rules](#) (see page 123)

[Filter String Object Identifier Restrictions](#) (see page 124)

[Special Characters in Regular Expressions](#) (see page 125)

[Sample Text for Notification Phrases in a Mailbox Rule](#) (see page 126)

[How to Create a Mailbox Rule That Matches Every Inbound Message](#) (see page 127)

[How to Use the Mailbox Rules TextAPI Defaults and TextAPI Ignore Incoming Settings](#) (see page 127)

Mailbox Rule Actions

Mailbox rules let you perform any of the following email actions:

- **Ignore Email**—Does not process the email and does not reply.
This action is useful for system-level messages such as Out of Office or Mail Delivery errors.
- **Ignore Email and Reply**—Does not process the email, and replies to the sender. Response emails use the reply success messages and the reply failure messages are ignored.
- **Update Object**—Uses the filter string to determine the object identifier (for example, %Incident:{object_id}% in the email), and sends an update request to the Text API. If the object identifier is not found, the Text API does nothing.
This action typically handles email replies where the object identifier is embedded in the email. If no object identifier is present, the failure response message is typically sent.
- **Create/Update Object**—Uses the filter string to determine the object identifier (for example, %Incident:{object_id}% in the email), and sends an update request to the Text API. If the object identifier is found, the Text API updates a ticket. If the object identifier is not found, the Text API creates a ticket.
This action is the standard behavior of the mail daemon (Mail Eater) in which the email does or does not contain Text API keywords.

Note: For information about configuring mailbox rules, see the *Online Help*.

Pattern Matching in Mailbox Rules

Mailbox rules use regular expressions for pattern matching. Consider the following whitespace characters that apply to regular expressions in mailbox rules:

`\t`

Specifies a horizontal tab.

`\r`

Specifies a carriage return character.

`\n`

Specifies a line feed or new line character.

The characters that represent line breaks in text can vary with the operating system, mail server, and mail client, for example:

- UNIX uses a `\n`.
- Microsoft uses `\r\n`
- Macintosh uses `\r`
- MacOS X uses `\n`

In certain circumstances, the mail processing elements of CA SDM exchange or substitute one of these line break characters for another to establish or maintain a distinction between different text elements, such as message text and attached parameters. As a result, when you want to use a line or paragraph break, build your filters so that either `\r` and `\n` can be matched, whichever is found. If you want to indicate a line break between two keywords, build your filters so that a sequence of one or more `\r` and `\n` characters can be matched.

Line wrapping by the mail client that sends a message can cause unexpected line breaks to appear in the middle of the text which is expected to match your filter string, when the filter string searches the body of the message. A space can change to a carriage return, line feed, or both, or the carriage return, line feed, or both can be inserted after a space. If a message is composed in HTML, and contains bulleted or numbered lists or indented paragraphs, tabs can also be present after the mail client converts and sends the message. When you include spaces in the middle of a filter string, use a Regular Expression block which represents any whitespace of variable size. This block is `[ \t\r\n]+` (open-bracket, space, backslash, t, backslash, r, backslash, n, close-bracket, plus sign), and represents any sequence of one or more spaces, tabs, carriage returns, and line feeds.

Example: Use [\t\r\n] to Match a Keyword Exactly

This example demonstrates how you can use whitespace characters to match the keyword "request" and not match other possible keywords such as the following words:

requester
requesting
requested
orechestra

To match only the keyword request, precede and follow the keyword request by one or more whitespace characters as follows:

```
[ \t\r\n]request[ \t\r\n]
```

The Mail Eater matches only the word request or the word as part of a sentence, and not as part of another word such as requester.

Filter String Object Identifier Restrictions

Restrictions apply to the mailbox rule filter strings which determine the object identifier (for example, %Incident:{{object_id}}%) in an email. Text that surrounds an object identifier ({{object_id}}) must be unambiguous in both content and length; the text must clearly define the beginning and end of the ticket ID artifact value that is between the text.

The following restrictions apply to how the Mail Eater interprets the *start* of the ticket ID artifact value:

- The {{object_id}} placeholder must not be the first element in the filter string. At least one character is required, and generally a distinctive keyword, or a sequence of letters, numbers, and symbols must precede the object ID keyword.
- The character that immediately precedes the {{object_id}} placeholder must not be a repeatable or optional character (meaning, a character that a plus sign (+), an asterisk (*), or a question mark (?) follows) that can be part of a ticket ID artifact value. Repeatable or optional characters risk being ambiguous with the start of the ticket ID artifact value, unless they are whitespace characters. Whitespace characters (space, tab, carriage return, line feed) must not be part of a ticket ID artifact value.
- The character that immediately precedes the {{object_id}} placeholder must not be a repeatable or optional bracketed-set of characters which includes characters that can be part of a Ticket ID Artifact value.

The following restrictions apply to how the Mail Eater interprets the *length* of the ticket ID artifact value:

- The `{{object_id}}` placeholder must not be the last element in the Filter String. One or more characters must follow the `{{object_id}}` placeholder.
- The character that immediately follows the `{{object_id}}` placeholder must not be a repeatable or optional character (meaning, a character that a plus sign (+), an asterisk (*), or a question mark (?) follows) that can be part of a ticket ID artifact value. Repeatable or optional characters risk being ambiguous with the end of the ticket ID artifact value, unless they are whitespace characters. Whitespace characters (space, tab, carriage return, line feed) must not be part of a ticket ID artifact value.
- The first character after the `{{object_id}}` placeholder must not be a character that can be part of a ticket ID artifact value.
- Avoid the following characters immediately before and after the `{{object_id}}` placeholder:
 - All upper or lower-case letters
 - Numerals
 - The plus sign (+)
 - The slash (/)
 - The comma (,)
 - The period (.), because it can represent any character except for a line break, and thus can be any of the characters in this list.

because any of these characters can exist within the ticket ID artifact value. When a bracketed set (several characters between brackets, of which the filter check can match any one), precedes or follows the `{{object_id}}` placeholder, the bracketed set must not contain any of these characters.

Special Characters in Regular Expressions

Pattern matching for the filters in mailbox rules follows the rules for ASCII Regular Expressions with C-style special characters.

Important! We recommend that you are familiar with Regex syntax to use special characters in regular expressions.

For example, consider the following special characters for Regex patterns that apply to regular expressions in mailbox rules:

`\t`

Specifies a horizontal tab. In filter strings for mailbox rules, `\t` matches the beginning and end of text (and tabs), and is specific to the Mail Eater.

\r

Specifies a carriage return (return to the beginning of the current line).

Note: Do not use \r for searching message subjects or sender addresses.

\n

Specifies a newline (combination of line feed and carriage return).

Note: Do not use \n for searching message subjects or sender addresses.

\t, \r, and \n are the special characters that occur most often in regular expressions for mailbox rules.

Example: \t, \r, and \n Use

`[ \t]request[ \t]`

Searches text for the word request. The brackets match any one character from the set, including the space, so `[ \t]` matches a space or a tab.

`[\r\n]+critical[ \t\r\n]`

Searches text for the word critical at the start of a line, the start of the message, or as the entire contents of a line. The brackets match any one character from the set, and the + (plus sign) matches one or more of the previous, so `[\r\n]+` matches one or more carriage returns and newlines.

Sample Text for Notification Phrases in a Mailbox Rule

This example shows sample text that you can use to include notification phrases in a mailbox rule. You define separate versions of a notification phrase for plain text and for HTML when the phrase contains any line breaks or other formatting.

Use the following text in Success Text, Success HTML, or both fields on the Update Mailbox Rule page, Reply Success tab:

■ Success Text

Thank you for submitting your request.

@{notification_phrase[IncidentURL1].phrase}

■ Success HTML

Thank you for submitting your request.</p>

@{notification_phrase[IncidentURL1H].phrase}</p>

Note: For more information about Notification Phrases that is defined under Administration Tab, Notifications, Notification Phrases, see the *Online Help*.

More information:

[Notification Phrase Syntax](#) (see page 109)

[Notification Codes and Phrases](#) (see page 107)

How to Create a Mailbox Rule That Matches Every Inbound Message

You can create a mailbox rule that matches every inbound message that another mailbox rule does not filter.

To create this type of rule, set the filter as Subject Contains and the filter string as a period and an asterisk (".*").

- A period matches any character except the line break.
- An asterisk matches zero or more occurrences of the symbol immediately before it.

As a result, this combination matches zero or more characters that are not line breaks.

Example: A "Catch-All" Mailbox Rule

This example demonstrates how you can use a ".*" combination to match every inbound message:

Filter = "Subject contains"

Filter String = ".*"

How to Use the Mailbox Rules TextAPI Defaults and TextAPI Ignore Incoming Settings

The TextAPI Defaults and TextAPI Ignore Incoming fields let you specify default values for incoming mailbox rules, and specify TextAPI commands that should not be accepted in incoming emails. These fields work with the default values that are set in the [EMAIL_DEFAULTS] section and with the forbidden-commands list in the [EMAIL_IGNORE_INCOMING] section of the text_api.cfg file. When a conflict occurs between the definition in a mailbox rule and the definition in the text_api.cfg file, the value set in the mailbox rule applies.

The TextAPI Defaults field includes TextAPI keyword commands that are applied to a ticket when it is created from an email that matches a mailbox rule. The commands are not applied when the message affects an existing ticket.

To specify TextAPI Defaults commands, do the following:

1. Place each command on a separate line in the TextAPI Defaults field.
2. Format the commands as follows:

`OBJECT.FIELD=value`

Note: Do not include a leading percentage symbol (%), which is required only for corresponding commands that are embedded in the body of the email.

For example, format the commands as follows:

`REQUEST.PRIORITY=3`
`PROBLEM.CATEGORY=Facilities`
`INCIDENT.GROUP=Plumbing`

The TextAPI Ignore Incoming field lists TextAPI keyword commands that are not permitted to be used in the text of the incoming email message. Any commands that are listed in this field are ignored when they are found in an incoming email message.

To specify TextAPI Ignore Incoming commands, do the following:

1. Place each command on a separate line in the TextAPI Ignore Incoming field.
2. Format the commands as follows:

`OBJECT.FIELD`

Note: Do not include a leading percentage symbol (%), which is required only for the corresponding commands that are embedded in the body of the email.

For example, format the commands as follows:

`CHANGE.ASSIGNEE`
`PROBLEM.GROUP`
`REQUEST.EFFORT`

3. Define all commands used in either field in the [KEYWORDS] section of the `text_api.cfg` file. This file is located in the “site” subdirectory of the CA SDM installation directory.

Mailbox Policies

Mailbox policies protect your organization against certain types of email abuse. You can implement policies in the following ways:

- You can define email address inclusion and exclusion lists:
 - Inclusion lists limit email reception to only mail that is sent from particular email addresses (for example, `user@company.com`) or email domains (for example, `company.com`).
 - Exclusion lists sort out unwanted email addresses or email domains.

The default inclusion list for each mailbox is an asterisk (*). An asterisk by itself as a whole name in this list specifies World Domain, and represents all email domains not included in the exclusion list. This representation prevents ambiguity of whether the inclusion list is the complete list of permitted domains or the exceptions to domains in the exclusion list as follows:

- An inclusion list that includes World Domain specifies that all other entries in the inclusion list are exceptions to the exclusion list, and only domains or addresses in the exclusion list are blocked (with the exception of specific addresses in the inclusion list).
 - An inclusion list that does not contain World Domain specifies that only addresses and domains listed in the inclusion list are permitted to post, and only if the sender domain (if the specific address is not in the inclusion list) or address are not in the exclusion list.
- You can set the Maximum Messages Per Hour Per Address limit for Denial of Service (DoS) detection.

Maximum Messages limits the number of emails per hour that any one email address is permitted to send to that mailbox. If an email address exceeds the limit, the email address is added to the exclusion list. No more messages from that sender email address are processed for that mailbox until you remove the email address from the exclusion list.

Mailbox policies for inclusion and exclusion lists and the maximum number of emails allowed per hour pertain to the single associated mailbox only. An address which is added automatically to one mailbox exclusion list for violating the maximum number of emails restriction is not added automatically to the exclusion list for any other mailboxes, nor is the email count used for other mailboxes.

Email addresses that violate policy are logged to the STDLOG based upon the settings of log_policy_violation (None, Log First, or Log All).

How to Implement Mailboxes

CA SDM provides a default mailbox (named Default) that is active and that you can use in your organization. You can modify the default mailbox, create additional mailboxes, or both.

To implement mailboxes, do the following preliminary steps:

- Configure the mail server account that you want to use.
- Create contacts and specify their email addresses if you do not want to allow anonymous users to create and update tickets (Allow Anonymous option).

To configure the Default mailbox, do the following optional basic steps:

1. Edit the Default mailbox.
2. Update the mailbox rules and policies as appropriate.

The mailbox is active and begins polling. The first poll occurs after one second.

Configure the Mail Server

Notifications for an event (automatic and manual notify) are sent using a single mail server definition.

To configure the mail server

1. On the Administration tab, browse to Options Manager, Email.

The Option List appears.

2. Click the option that you want to install.

The Options Detail page appears.

3. Click Edit, complete the fields as appropriate, and click Install.

The mail server is configured to send notifications (outbound mail).

4. Repeat the procedure until all relevant Option List options are configured.

Note: You can use the `pdm_mail` command `-F` and `-T` parameters, which set the Sender and Reply-To addresses, respectively. By using these parameters in customized notification methods, replies to mail-received notifications can be directed to the appropriate mailbox.

More information:

[pdm_mail Utility--Send Email Information](#) (see page 90)

Email Options

The email interface sends email notifications and lets users create tickets from an email. The email options let you configure the outgoing mail protocols.

mail_from_address

Specifies the mail notification From: address. The address is in the format `Displayname<user@company.com>`.

mail_login_password

Specifies the SMTP server login password.

mail_login_userid

Specifies the SMTP server login userid.

mail_max_threads

Specifies the maximum number of concurrent SMTP connections that can attempt to communicate with the server.

mail_reply_to_address

Defines the reply to address for email notifications. This option is useful if emails are sent from one user account, but you want replies sent to another email address. The default value is the same as the from address.

mail_smtp_domain_name

Defines the domain name that some SMTP servers that understand the ESMTP protocols require. You can clear the domain name by setting the value to NONE.

mail_smtp_hosts

Specifies a space-separated list of SMTP host names for email notifications.

mail_smtp_host_port

Specifies an SMTP port to override the default SMTP port.

mail_smtp_security_level

Specifies the SMTP security level. Valid settings are: 0=no security, 1=basic authentication, 2=NTLM, 3=MD5 and 4=LOGIN. If you set this option to 1, set the mail_login_password and mail_login_userid options. Most SMTP servers do not require authentication.

Edit the Default Mailbox

CA SDM provides an active default mailbox that you can edit to suit the mail delivery needs of your organization.

To edit the default mailbox

1. On the Administration tab, browse to Email, Mailboxes.
The Mailbox List appears and lists active mailboxes.
2. Click Default in the Name column.
The Default Mailbox Detail page appears.
3. Click Edit.
4. Complete or update the other fields as appropriate:

Sequence

Specifies the sequence number of the rule. Messages are checked against rules in sequence number order from lowest to highest.

Mailbox

Specifies the mailbox to which this rule belongs.

Active

Sets the rule to active or inactive.

Filter

Specifies what part of the email to search for the filter pattern, for example, Subject contains.

Filter String

Specifies a regular expression string to match, for example, [\t\r\n]request[\t\r\n]. The placeholder {{object_id}} lets you specify the artifact value for the Text API to use for associating the message with a specific ticket.

Ignore Case

Specifies whether to consider letter case when matching patterns.

Action

Specifies the action to take when the filter criteria matches, for example, Create/Update Object.

Note: For information about rule actions, see the *Administration Guide*.

Action Object

Displays the type of ticket object to which message actions apply, for example, Request.

Minimum Artifact Type

Sets the minimum type of artifact checking that you want to allow:

- **NONE**—Specifies no validation of artifacts. This value is the same as not adding the keyword to the input file. Also accepts Text API ticket ID commands.
- **PROTECTED**—Validates a ticket against the hash for confirmation. If confirmation fails, the artifact is considered invalid and the email fails filtering where filtering is looking for an artifact ({{object_id}}).
- **SECURE**—Validates the ticket number from an encrypted data block. If the value is not a valid encrypted ticket number, the artifact is considered invalid and the email fails filtering where filtering is looking for an artifact ({{object_id}}).

Note: Types that are more secure than what is set are allowed. In other words, if you set the minimum type to PROTECTED, then both PROTECTED and SECURE are allowed, but NONE is not. Also, if PROTECTED or SECURE are selected, Text API ticket ID commands are not accepted.

Reply

Specifies a notification method to send an automatic response, for example, Email. If you do not set this option, no response is returned. The response indicates success or failure of the actions performed for the message, and is separate from any notifications.

Reply Subject

Specifies a subject line for the reply, for example, Service Desk Response.

Write to stdlog

Write email text to the standard log (stdlog) if the filter matches.

Log Entry Prefix

Specifies a prefix to add when writing email text to the standard log entries. Use this option to enable matching log entries to rules.

Add Subject Line

Adds the subject line from the original message to the message body before processing. You can append, prepend, or not add a subject line. The subject line is attached to either the ticket Description or a Log Comment, depending on the actions for the message.

Text API Defaults

Specifies additional default commands for the Text API when the filter matches. Combines with the contents of the [EMAIL_DEFAULTS] section of the text_api.cfg file.

Text API Ignore Incoming

Specifies additional Ignore Details for the Text API when the filter matches. Combines with the contents of the [EMAIL_IGNORE_INCOMING] section of the text_api.cfg file.

Details

Specifies information about the rule.

Success Text

Specifies the plain-text contents of a reply message when the message is processed successfully. For example:

Thank you for submitting an update to your request. A support technician will contact you within the next 24 hours.

Success HTML

Specifies HTML contents of a reply message when the message processes successfully. The following options let you edit and preview HTML text:

- **Edit Success HTML**—Opens the HTML Editor which you can use to format the HTML.
- **Quick View**—Previews the HTML.
- **HTML Source**—Shows the HTML source.

Failure Text

Specifies the plain-text contents of a reply message when the message does not process successfully.

Failure HTML

Specifies HTML contents of a reply message when the message does not process successfully. The following options let you edit and preview HTML text:

- **Edit Failure HTML**—Opens the HTML Editor which you can use to format the HTML.
- **Quick View**—Previews the HTML.
- **HTML Source**—Shows the HTML source.

Note: For more information about mailbox rules, see the *Administration Guide*.

5. (Optional) Create or update the mailbox rules and policies as appropriate.

Note: When you create mailbox rules, you associate them with a specific mailbox—another mailbox cannot share those rules. If you want to use rules that belong to one mailbox for a different mailbox, recreate them for the other mailbox.

Important! We recommend that you set the associated mailbox to Inactive before you configure a mailbox rule. Otherwise, any messages that the mail server retrieves between your first change and the last change are processed with whatever rules are in effect when the message is retrieved.

6. Click Save.

The changes to the default mailbox are saved and applied. The first poll occurs after one second.

Multiple Mailboxes

CA SDM can process and manage multiple mailboxes. Each mailbox can have its own definition, instead of using a single global set of definitions. You can define multiple mailboxes and use different templates or default values for each mailbox. Multiple definitions let individual tenants use separate mailboxes, or let an individual tenant or organization use different mailboxes, and have different behaviors for each mailbox. You can set up multiple mailboxes by using the Administration interface. Each mailbox uses the following tables:

- **usp_mailbox**—Defines the mailbox.
- **usp_mailbox_rule**—Specifies a set of rules for each mailbox.

Because mailbox rules supply Text API defaults, you can establish email interfaces with other software and parameters (such as category, assignee, and so on) that are configured specifically for the interface.

Note: IMAP servers support multiple mailboxes for a single account, but alternate mailboxes are not supported; only the default inbox is supported.

How Multiple Mailboxes Use Rules

The Mail Eater (`pdm_maileater_nxd`) component on the primary server uses mailbox connections and rules to read and process messages from one or more accounts on one or more mail servers. The Mail Eater processes mailboxes serially (only one mailbox is processed at a time), and processes rules in sequence number order.

Multiple mailboxes use rules as follows:

1. Upon primary server startup, the Mail Eater reads the following tables:

usp_mailbox

Represents a connection to a mail server.

usp_mailbox_rules

Represents the rules that apply to the connection (`usp_mailbox`).

Contact_Method

Represents the Contact Methods used for automatic replies.

Document_Repository

Represents the Document Repositories for storing attachments.

The Mail Eater automatically detects changes to the objects in any of these tables, including the addition of additional objects. If a change is made to `usp_mailbox` or `usp_mailbox_rule`, the polling interval for the affected mailbox is rescheduled to one second after the change is applied.

2. At the interval defined by each mailbox, the Mail Eater retrieves each email in the inbox for the associated account, and processes the email as follows:
 - a. Checks the email address for policy violations. When the Mail Eater finds a violation, processing stops, and the standard log is affected according to the mailbox definition.
 - b. Compares the email to each rule (`mailbox_rule`) that belongs to that mailbox.
 - c. If a matching rule is found, submits the message to the Text API for posting, and replies to the user as appropriate based on the specified action for the rule.

For reply emails, the filter string identifies the object and uses the Text API for processing. After processing is complete, the response goes either to the Reply to or the From address.
 - d. After the Mail Eater finds a matching rule, no other rules are checked, and the Mail Eater processes the next email in the inbox.
 - e. If no matching rule is found, the message is discarded.
3. After the Mail Eater processes all emails for an inbox, the processed and discarded messages are purged from the mail server, and the next processing interval is scheduled.

How to Configure Email Replies

Email notifications that you use in mailboxes are specific to replies that are sent to a contact in response to their emails. You can configure email so that when a contact clicks a reply link in an email notification, the reply email is directed to a mailbox.

Note: This setup differs from the regular email notifications.

You use CA SDM Administration to configure email as follows:

1. Create a mailbox to handle inbound email on a server.
2. Browse to Options Manager, Email, and configure the outbound mail server.
3. (Optional) Specify a notification email address in the contact definition.
4. Browse to Notifications, Notification Methods, Email, and create or update a notification method that includes the following settings:

- Supports SMTP—enabled
- `pdm_mail [-F from_email_address] [-T reply_to_email_address]`

from_email_address

Specifies the address that is used as the From address of the message. This address overrides the sender address in the outbound mail server configuration.

reply_to_email_address

Specifies the address to which replies are sent. This address overrides the reply-to address in the outbound mail server configuration.

When a reply-to address is set in the outbound mail server configuration and no reply-to address is specified in the Notification Method, the reply-to address in the outbound mail server configuration is used with that Notification Method.

- “\$(REPLY_FROM)” keyword

If the keyword “\$(REPLY_FROM)” is specified as either address, that address is constructed from the username and mail server hostname for the mailbox. This keyword is only valid when a mailbox rule uses the Notification Method; Notification Methods that use it must not be used for any other purpose. For example, user name=dev, server name=mail32.ca.com, \$(REPLY_FROM)=dev@mail32.ca.com. Only use this keyword if your mail server is configured to accept the mail server name as equivalent to the email domain name. Use this keyword with caution: If the hostname is not fully domain-extended in the mailbox configuration (for example, mailserver1 instead of mailserver1.customer7.com), it is not extended automatically by the field interpreter.

Note: The from_email_address and reply_to_email_address are the addresses that appear in the From and Reply-To headers of the message when the user reads it. If the addresses are identical, you can specify only the from_address.

5. Browse to Email, Mailbox Rules, open each applicable mailbox rule, and select the new reply method in the Reply drop-down list. Save each updated rule.

When the contact replies to the email notification, the reply is addressed by default to the specified mailbox.

More information:

[pdm_mail Utility--Send Email Information](#) (see page 90)

Alternate Sender Address Identification

You can use a -m parameter in the subject of the message so that CA SDM identifies the sender of the message using a different email address from the one that originally sent it. The -m keyword, followed by a space and by the email address that CA SDM recognizes, must be the last elements of the subject line. Consider the following information when you use the -m parameter in the subject:

- Both the actual From address and the alternate address that is specified with -m are verified in the Inclusion and Exclusion lists.
- The email address that is specified as the alternate address must contain only the address, and not the accompanying display name.
- If more than one word follows the -m parameter in the subject line, the alternate address is not recognized.

Artifacts

An email *artifact* refers to something that arises from the mail process, for example, an email address that is included in a forwarded email. The Text API uses artifacts that contain a ticket ID (such as a reference number) for reply support. When the ticket ID is found, an existing Text API keyword (such as %INCIDENT_ID) is set and added to the input for the Text API. The Mail Eater identifies that a reply is associated with a particular ticket by finding the artifact in the message.

Mailbox rules let you specify the artifact and the value that the Text API uses. For example, you can define a rule for incidents as Incident:{{object_id}}%. When a rule finds Incident:1234, the Text API uses %INCIDENT_ID=1234 (1234 is the ref_num for the Incident). Because the artifact must be unique in an email and easy to find, you can make the artifact more distinctive such as %Incident:{{object_id}}%. A percentage sign (%), whitespace, or some other character that cannot appear in an artifact value must follow {{object_id}}. Uppercase and lowercase letters, numbers, forward slashes, commas, and plus symbols are potentially part of a value. The secure artifacts are Base64-encoded after encryption. If you do not use Secure artifacts, the characters that follow the artifact must not be contained in the ticket ID suffix, if any, which has been configured for that type of ticket.

When using the filter string of the mailbox rules to identify the ticket ID Artifact, the keyword {{object_id}} represents the position in the filter string where the ticket ID artifact is expected. This keyword is case sensitive, even if the mailbox rule is not.

Example: Email Artifact Use

The following example shows an ARTIFACT format for use in a mailbox rule for a change request ticket.

Usage: %REQUEST=@{call_req_id.ref_num}%

Example: %REQUEST=1234%

Artifact Use Considerations

When you construct the filter string of the mailbox rule, consider the following:

- A clear boundary must exist between the ticket ID artifact and the keywords which precede and follow it. We strongly recommended that you include whitespace text in this boundary text.
- Do not end the portion of the filter string that precedes the {{object_id}} keyword in a repeatable or optional pattern that can match the beginning of the ticket ID Artifact, and do not end a pattern whose length is ambiguous. For example, the filter string must not contain request(er|ed|ing)?{{object_id}}, because this construction causes an ambiguity whether a trailing er, ed, or ing is the end of the leading text or part of the prefix of an unprotected ticket ID.
- The portion of the filter string that follows the {{object_id}} keyword must not begin in a repeatable or optional pattern that may match the end of the ticket ID artifact, must not begin with a pattern whose length is ambiguous, and must contain at least one element of whitespace. For example:
 - The filter string must not contain {{object_id}}[A-Z]?, because the [A-Z]? may match the last character of the ticket ID artifact.
 - The filter string must not end with {{object_id}}Item, because it is possible for Item to appear in the ticket ID artifact, either as the suffix of a ticket ID in a plain-text or protected artifact, or as characters within a secure artifact.
 - {{object_id}} Item is acceptable, because the space cannot be part of a ticket ID artifact, and is not part of a protected or plain ticket ID artifact. However, {{object_id}}[\t\r\n]+Item (open-bracket, space, backslash, t, backslash, r, backslash, n, close-bracket, plus sign, +Item) is better, because the [\t\r\n]+ compensates for the mail program inserting a line break after the {{object_id}}.
- When you construct the filter strings for different mailbox rules, avoid using a filter string that completely includes another mailbox rule filter string, or in which the portion before or after a {{object_id}} keyword completely includes that portion of another mailbox rule filter string. Depending on the order in which these filters are checked, a message match intended for one filter can match with another, with a portion of the ticket ID artifact matching the additional text that distinguishes between the two filter strings.

Artifact Protection and Security

You use artifacts to protect mail delivery to a mailbox in the following ways:

- Use the ARTIFACT keyword in notifications and notification phrases.

Note: The artifact security level must match the mailbox rule Minimum Artifact Type setting for the mailbox to which the email is sent. Types that are more secure than what is set are allowed. In other words, if you set ARTIFACT=PROTECTED, then both PROTECTED and SECURE are allowed, but NONE is not allowed.

- Set the Minimum Artifact Type option in the Mailbox Rule Detail.

You can embed an artifact in a notification phrase and specify the security level for the artifact, for example:

`%REQUEST=@{ARTIFACT=PROTECTED:call_req_id.ref_num}%`

Note: For more information about the tables for the Minimum Artifact Type option (filter_min_artifact_type) and the Mailbox Rule Detail (usp_mailbox_rule table), see the *Technical Reference Guide*.

Example: No Validation of Artifacts

The following example shows no ARTIFACT format. No validation of artifacts can be performed.

Usage: `%REQUEST=@{call_req_id.ref_num}%`

Example: `%REQUEST=1234%`

Example: No Validation of Artifacts

The following example shows the ARTIFACT=NONE format. ARTIFACT=NONE is the same as not adding the keyword; no validation of artifacts can be performed.

Usage: `%REQUEST=@{ARTIFACT=NONE:call_req_id.ref_num}%`

Example: `%REQUEST=1234%`

Example: Validate a Ticket Number Against the Hash for Confirmation

The following example shows the ARTIFACT=PROTECTED format. The PROTECTED format validates a ticket against the hash for confirmation. The letter "A" denotes the type of formatting. Commas separate "A", the hash code, and the ticket number.

Usage: `%REQUEST=@{ARTIFACT=PROTECTED:call_req_id.ref_num}%`

Example: `%REQUEST=A,12345678,1234%`

Example: Encrypt the Ticket Number

The following example shows the ARTIFACT=SECURE format. The SECURE format encrypts the ref_num, and then encodes the result in Base64 encoding so it can safely be included in the message text. The letter "B" denotes the type of formatting. Commas separate "B" and a password encrypted ticket number.

Usage: %REQUEST=@{ARTIFACT=SECURE:call_req_id.ref_num}%

Example:

%REQUEST=B,da1jhr+9U5GVfd2VGH4dsnx2+PaSvygDS2e3lqjpjtyNSDW2u/KNPX61nopDu/KB%

Service Level Agreements

A service level agreement (SLA) or service type is an agreement between a service desk and its customers and usually describes the level of service to be provided by the service desk. If this level of service is not provided, the service desk can be penalized. For example, a service desk that operates on a “pay per service” basis may not receive full payment for service that does not meet an agreed upon level of service. Thus, most service desks view service level agreements very seriously and make every effort to meet the type of service specified in these agreements.

In addition, most service desks keep meticulous records when service level agreements are met or violated. The service types defined with CA SDM are designed to help the service desk personnel meet their service level agreements and keep the records they need to verify that their service level agreements have been met.

SLA Usage

In CA SDM, you set up service level agreements (SLAs) using service types and events to do the following:

- Use events to monitor tickets.
- Use service types to track vendor and organization commitments and schedules as they relate to specific tickets.
- Establish date and time controls for processing events and service types.

CA SDM lets you design and use SLAs as follows:

- **SLAs for multiple organizations**—You can assign different Service Types for each reference field depending on the affected organization of the ticket. For example, the Service Type for a Priority 1 ticket may be different for Organization A than Organization B.
- **Multiple SLA Tracking**—A ticket can track several Service Types at once, ensuring all aspects of a negotiated SLA are visible.
- **Time to Violation Projections**—A subsystem that constantly monitors tickets in CA SDM, and records when their SLA enters violation based on its current state.
- **SLA Violation Cost**—This Service type field helps with metrics and determining the severity of open tickets.
- **Service Types on Tasks**—Service Types are available on [workflow tasks](#) (see page 247).
- **Service Targets**—Service Types may also have one or more Service Targets for measuring whether the required service or repair completed within the required time frame.

Classic SLA Processing

In "classic" SLA processing (enabled if the `classic_sla_processing` option is installed in Options Manager) only one service type can apply to a ticket at any given time. When different attributes on a ticket have different service types associated with them, the higher ranked service type is used. The rank of a service type is defined when the service type is created, with the highest ranking being 1, the next being 2, and so on. For example, assume that the issue has a service type of 12-hour resolution (ranking 2), was assigned a priority code of 1, which has a service type of 4-hour resolution (ranking 1). The higher ranked service type determines the service behavior for the associated issue. In this example, 4-hour resolution is ranked higher than 12-hour resolution, so the 4-hour resolution service type is applied to the issue.

Note: For more information about the `classic_sla_processing` option, see the *Online Help*.

Service Types and Events

To keep track of vendor and organization commitments and schedules as they relate to specific tickets, you can attach events to service types. Service types provide a flexible and efficient way of attaching events automatically to tickets.

Service types provide automated support for the service level agreements you have established with your customers. Using service types, you can monitor any aspect of a time or condition-related situation. For example, you can set a policy that defines a resolution time for a particular issue and associate an action with it. Service types provide the capability to manage service providers, vendors, and the service desk itself, enabling easy performance and compliance measurement.

Service types let you do the following:

- Identify the level of service required for specific tickets automatically
- Identify service level violation conditions
- Send warning notifications automatically prior to expiration of service level agreements
- Set up service level agreement violation reports

How to Implement Service Types

Service types let you define actions based on user-definable conditions. Service types can be applied to tickets. For example, you can define a service type that specifies a resolution time for a particular ticket, and then associate a behavior that occurs automatically when the action occurs, such as notification of a support person of a pending violation. Because service types can be comprised of multiple events, and events can be comprised of multiple actions (on true and on false), you can define behavior that is more complex by adding multiple events to service types.

To implement service types, do the following:

1. Associate a service type with a respective attribute. Any number of service types can be defined and ranked in relationship with each other.
2. Assign the service types to CIs, contacts, organizations, issue categories, change categories, request areas, and priorities. Each record can be assigned to only one service type.
3. Define specific conditions that determine when a service level agreement is met or violated.
4. Identify actions that need to occur to assist the service desk in meeting its service level agreements.
5. Apply the correct actions to enforce the service level agreements automatically.

Example: Associate a Service Type with a Contact

In this example, you do the following steps to associate a service type with a contact:

1. Assign a service type on the Contact Detail page.

When the contact is assigned to a ticket, the Service Type field is automatically populated with the service type assigned in the contact record.

2. Enter Alpha, Beta as the name in the Affected End User field on the Issue Detail page and save the issue.

The Service Type field is populated automatically with the value defined in this contact record. You can override the default by selecting other service types.

Predefined Service Types

CA SDM provides predefined service types. You can modify the predefined service types and create service types to better suit the needs of your enterprise. The following predefined service types are included with CA SDM:

- 4-hour resolution
- 12-hour resolution
- 48-hour resolution
- 72-hour resolution

The 4-hour resolution service type is associated with priority code 1 when you first install CA SDM, which means that every ticket created with a priority of 1 is treated as if it should be resolved within four hours. Several service type events are associated with the 4-hour resolution service type to ensure quick resolution:

- The person for whom the ticket is created is notified that they can expect resolution within four hours.
- The group manager is notified if the ticket is not assigned within 30 minutes.
- A warning notification is sent to the assignee and group manager if the ticket is not resolved in 3.5 hours.
- Finally, a violation notification is sent to the assignee and group manager if the ticket is not resolved in 4 hours, and the SLA violation flag for the ticket is set.

None of the other predefined service types are associated with a priority code, but they are defined with events similar to those described for the 4-hour resolution service type. As system administrator, you can associate these service types with priority codes if you like. For example, you might make the following associations:

- Priority 2—12-hour resolution
- Priority 3—48-hour resolution
- Priority 4—72-hour resolution

Alternatively, you can remove the association between priority code 1 and the 4-hour resolution service type if it does not suit your needs.

The `cr_sla`, `chg_sla`, and `iss_sla` customization options specify that all requests, change orders, and issues must be associated with a service type. These options select the best service type and apply it to the request, change order, or issue. These options are installed automatically; however, you can uninstall or reinstall them using the Options Manager.

Note: For descriptions about the customization options, see the *Online Help*. For information about how to define service types and modify priority codes, see the *Online Help*.

Events Setup

CA SDM uses conditions and actions defined as events to schedule, process, and track tickets and service types automatically. Each event has three generic behavior characteristics: conditions, actions on true, and actions on false. In addition, a delay time is associated with the condition (and thus the event) and measured from the time the ticket is saved. The delay time determines when the condition is evaluated. You can also create automatic events to cause a specific action to occur every time a specific event occurs. For example, you can cause notification to occur automatically for a transferred issue with a priority of P1 that remains open for more than one hour.

- *Conditions* identify the event for which CA SDM should be prepared.

A condition is any measurable state of a ticket. A ticket with a state of “unassigned” is an example of a condition. This condition can then be associated with an action that occurs when the condition is true or when the condition is false. For example, in the predefined 12-hour resolution service type, there is an associated event named 1 hr unassigned with “Unassigned Issue” as its condition. The action defined for this event is to notify the issue group manager when the condition is true, with a delay time of one hour. If 12-hour resolution was assigned to an issue, the behavior would be to “wake-up” one hour after the record was saved, verify if the issue was still unassigned, and notify the issue group manager if it had not been assigned.

- *Actions* identify the processes that occur automatically when the condition is true or false after a specified amount of time.

You can assign a single action or a list of actions to occur for each monitored event. When an event occurs that matches a condition, CA SDM automatically processes the actions associated with the condition. In some cases, actions automatically process the object or can cause notification to occur.

Some of the predefined events for issues are shown in the following table. There are similar predefined events for requests and change orders, and several others relating to the predefined service types:

Event	Description
P1 Active Issue Notify assignee	Notifies the assignee when the issue is active for more than 1 hour and the priority is 1.
P2 Active Issue Notify assignee	Notifies the assignee when the issue is active for more than 1 day and the priority is 2.
Unassigned Issue Transfer	Transfers issues when they remain unassigned for more than 1 day; notifies the previous assignee, the issue creator, and the new assignee.

After events are defined, they can be attached manually or automatically to tickets and service types, as is the case with predefined service types. Multiple events can be attached to each service type.

Important! Do not use the Options Manager to install automatic events until the events are created. Turning on automatic events before events exist seriously impairs performance. For more information about how to define events, and event options such as `auto_events`, `chg_auto_events`, and `iss_auto_events`, see the *Online Help*.

More information:

[Options Manager Usage](#) (see page 317)

How to Create Service Targets

To minimize SLA violations, you can create a set of service target templates to measure each stage of ticket resolution. Like Service Types, each service target contains a condition and estimate of completion time. However, service targets do not provide an action mechanism.

During ticket creation, a Service Type assigns one or more service targets to track each stage of the ticket resolution. Each time the ticket changes, the active service targets evaluate the condition. If the condition is met, the ticket and activity log show the actual completion time. When the time exceeds the estimated time, the ticket displays the amount of time by which the target was missed.

Service targets let you do the following:

- Verify that tickets of the same Service Type follow the same service targets.
- Monitor whether tickets are closed within the required time frames.
- View information such as the remaining number of minutes before a service target completes.

When creating service target templates, consider the following:

- Consider the required outcome for meeting the service target. Use an existing Condition or Site-Defined Condition Macro to evaluate ticket data. If necessary, customize or write a new macro to manage the service target.
- Calculate the projected violation and penalty costs based on SLA agreements.
- Assign at least one service target template to a Service Type.

To create a service target template, do the following:

1. Create a service target template to manage requests, incidents, problems, change orders, or issues.
2. Link the service target template to a Service Type.
3. Assign the service target template detail to a ticket category such as request, incident, problem, change order, or issue tickets.

At the time of ticket creation, the appropriate template automatically attaches to the ticket based on the Service Type. Each time a user creates a ticket, the status of the service target automatically displays in the Service Type tab.

Service Contracts

The SLA model includes the Service Contract. The Service Contract defines the SLA for a particular organization, including its Service Types, Request areas, and Issue or Change Categories. These definitions are referred to as *private* Categories and *private* Service Types.

Note: The *private* Categories and Service Types can only be used on tickets where the Service Contract is used.

The Service Contract applicable to a ticket is determined by the ticket's *affected organization*, which is always the Organization of the Affected End User on the ticket (this is the Organization field on a Contact record). Only the Areas or Categories listed on the Service Contract can be selected for the ticket. In addition, the only Service Types that can be applied are the *private* ones listed on the Contract. This helps ensure that the SLAs for one Organization are not accidentally mixed with another's.

A Service Contract also maps Service Types to common reference fields on a ticket, such as Priority and Asset. This mapping associates Service Types with attributes of a ticket. For example, an Organization's contract can assign Service Types to each of the five Priority objects. When a ticket is created with a certain priority, the mapped Service Type is applied.

Categories and Service Types can be defined outside of a Contract and are considered *public*. The public definitions are used when a ticket has no Service Contract. The lack of a Service Contract can occur if the end user has no organization, or the organization's Contract is inactive. A public definition is a helpful backup or default mechanism. Public Service Types are set directly on categories and other reference field objects.

All applicable Service Types are assigned to the ticket. This ensures that all aspects of an SLA are visible and enforced, for example:

- A Printer may have a Service Type that requires a technician to be dispatched within two days.
- A priority object's Service Type may require a callback within one hour.

With both Service Types applied, these required actions are enforced.

Tracking multiple Service Types also helps prioritize tickets. For example, a ticket concerning a broken keyboard is assigned a low priority Service Type. However, if the affected end user is in urgent need of the keyboard, the service priority can be increased.

Note: The SLA model is enforced by default. Past versions of CA SDM applied only a single Service Type to a ticket. The Service Type selection involved finding the highest *ranked* Type among all the possible Service Types. A model using the ranking scheme can still be used by installing the Option 'classic_sla_processing'.

Service Contracts Migration

If CA SDM is installed as a migration, the Option `classic_sla_processing` is turned ON by default, so your SLA processing can continue as before the migration. This continuation gives you time to create appropriate Service Contracts and eventually deactivate `classic_sla_processing`.

When building Service Contracts, you do not have to create Service Types, Request Areas, or Categories. You can use the Copy button found on the Service Contract detail to copy existing objects to the Contract.

If the previous installation marked the support_lev field as required for any ticket type, this restriction must be removed. The support_lev field still exists, but is not set in the current model so a required field error results with new tickets. This affects the objects Request (cr), Issue (iss) or Change Order (chg).

Time to Violation

When the SLA model is in use, CA SDM's Time to Violation (TTV) system can help you track and prioritize tickets according to their projected violation time. This system monitors all active tickets and sets the projected violation time for each Service Type. You can report and sort tickets based on their violation time and cost, helping you resolve the most urgent and costly issues first.

The TTV system monitors all active tickets and evaluates their SLA events silently, to determine which events set the SLA violation flag. The events are not executed, but the system looks at the outcome of each event based on the current state of the ticket. If the evaluation results in an action that sets the SLA violation flag, the ticket's attached Service Type is updated with a Time to Violation value; this value is the date/time when the event fires that sets the SLA violation.

Evaluation occurs whenever a ticket is inserted or updated. Because tickets are often updated in rapid succession, the evaluation is delayed for a short period. The delay interval is controlled by the ttv_evaluation_delay Option. After the delay expires, the TTV system evaluates all the SLA events that might set the SLA violation flag.

Each Event has an optional condition and a set of actions (Macros) that start based on the outcome of the condition. To ensure adequate performance, Event template information is cached by the TTV system and is refreshed periodically. Projections made by TTV involving recently updated Event templates may be inaccurate for several minutes.

Note: The TTV projections appear on the Service Type tab of each ticket. The TTV system is activated with the ttv_enabled Option.

Time Zones and Workshifts

To address the complex business demands of automated application execution, CA SDM lets you define as many time zones and workshifts as you want, and to record them for easy reference.

- Time zones identify the time zone where the user, CI, and so forth, are located.
- Workshifts define the period during which event monitoring or the work hours of a service type or SLA are in effect.

Being able to determine a course of action to perform based on when an event occurs can be critical to the proper handling of the event. The time zones and workshifts you define are available for use by any of the CA SDM functions.

More information:

[Time Zones](#) (see page 305)

[Workshifts Setup](#) (see page 152)

Time Zones Setup

Time zone codes define the time zone from which a user usually accesses the system (that is, the user's local time zone) or the time zone in which a CI is located. Business hours are always entered in the time zone of the CA SDM server. This means that business hours can always be compared uniformly.

Time zones are used to manage service types, escalations, and overall response to affected end users based on the ability of CA SDM to present the correct time across multiple time zones. CA SDM automatically adjusts the offset between the time zone where a user is logging in or a CI is located, and the time zone where the server is running. Time zones use the Greenwich Mean Time (GMT) format.

Note: For information about how to define time zones, see the *Online Help*.

How to Manage Multiple Time Zones for Service Types

CA SDM servers and users can be located in different time zones. The time difference can cause users to miss Service Type expiration dates and times.

To correct the time difference, you can configure CA SDM to display Service Type expiration times in the end-user time zone. If two users in different time zones view the same ticket, each user sees Expiration Time values based on the local computer time zone. However, the Start Time values always reflect the server time zone.

To configure for the end-user time zone, do the following:

1. Create a server code that identifies the server name and time zone.
2. Create or update a Service Type. Set the Use End User Timezone field to Yes.

A value of Yes causes the Expiration Time to display and update according to each user time zone.

Example: Show Service Type Expiration Dates in Any User Time Zone

In this example, you configure CA SDM to show Service Type expiration dates in any user time zone. The server and user are on separate computers and in different time zones.

To create a server code that identifies the server name and time zone

1. On the Administration Tab of the host server, select Service Desk, Application Code.
2. Click Codes, Servers.
The Server List displays.
3. Click the Local Host server.
The Server Name Detail page appears.

4. Set the Time Zone. For example, set the time zone to GMT (EU). The local host name must match the NX_LOCAL_HOST value stored in NX.env for the server
5. Click Save.

The Host Server uses the new time zone. When the user views a ticket, the Start Time reflects the server time zone.

To create a service type

1. On the Administration Tab, select Service Types.
The list of Service Types appears.
2. Click Priority 1 Resolution or another Service Type that manages Priority 1 Requests.
The Update Service Type page displays.
3. Select the Use End User Time Zone check box.
4. Click Save.

The Service Type record updates.

To view the time zones on the ticket

1. On another computer, open a Request ticket, and set the Priority to 1.
Note: If you are using Service Targets, set the values in the ticket to cause the target to use Priority 1 Resolution.
2. View the ticket and click the Service Type tab.
The Start Time reflects the server time zone. The Expiration Time reflects the time on the end-user local computer.
3. Close the page that displays the Request ticket.
4. Change the time zone on the end-user local computer.
5. View the ticket Service Type on the end-user local computer to see the corresponding Expiration Time values based on the user time zone.
The Expiration Time reflects the new time zone.

Note: For detailed information about Server Types or Server Code, see the *Online Help*.

Workshifts Setup

Workshifts identify the days, dates, and times when an event or schedule is in effect. You can specify days or dates, or days and dates. Specifying a time is optional.

Note: For information about how to define workshifts, see the *Online Help*.

When you are monitoring events for tickets, workshifts define when the event is monitored or, in other words, when the clock is running. For example, using the predefined event P1 Active Issue Notify assignee, if a priority 1 issue is opened at 4:45 PM and the workshift schedule is 9:00 AM to 5:00 PM, the monitored event automatically sends notification to the issue assignee at 9:45 AM the next day.

Note: Workshifts are also used for the purpose of automatically assigning tickets to contacts.

More information:

[Establishing Support Structure](#) (see page 241)

Security

Before you allow people to use CA SDM, it is important that you set up security to determine the following:

- Which users can access the system
- What level or levels of access users can have
- How users are authenticated when they log in

Note: For details about how to accomplish security administration tasks, see the *Online Help*.

CA EEM and CA Workflow User Base Configurations

CA EEM is a central repository of user information (identities). CA EEM defines user authentication and access to other applications. If you have several CA Technologies products installed, some of them can use CA EEM to store identities and access policies. CA SDM only uses CA EEM for authentication. CA EEM is not a CA SDM configuration option and must be installed separately.

The CA EEM repository of user records is *either* of the following sources:

- An external LDAP directory
- Its own internal tables in the MDB

CA EEM has an LDAP interface for use when it is configured to use the MDB.

Note: The MDB tables used by CA EEM are different from the ones used by CA SDM.

If your organization uses a directory server, such as Active Directory or eTrust Directory, consider configuring CA EEM to use the directory for its user base. This configuration makes the users in your directory accessible by any other application that uses CA EEM. Because CA EEM centralizes access management, it is typically installed on a single server.

Note: For more information about installation, see the *Implementation Guide*.

CA Workflow

During CA SDM installation, you can optionally install CA Workflow, an embedded workflow engine. Because CA Workflow uses CA EEM for all user information, CA EEM manages all authentication attempts and certain permissions. CA Workflow is unaware of the CA SDM contact records.

Note: For information about configuring access for CA Workflow, see the *Implementation Guide*.

CA SDM

CA SDM stores contact information in MDB tables. These tables have no relationship to CA EEM. CA SDM does not use CA EEM for access or identity management. CA SDM manages its own access and security with Access Types and Data Partitions.

CA SDM uses CA EEM only for authentication. If you want to use CA EEM to authenticate users in CA SDM, install CA EEM. If you integrate CA SDM with CA EEM, it replaces the CA SDM operating system authentication with CA EEM authentication.

Note: For information about installing CA EEM, see the *Implementation Guide*. To integrate CA EEM and CA SDM, you must set the *eiam_hostname*, *use_eiam_artifact*, and *use_eiam_authentication* options in Options Manager, Security. For more information about these options, see the *Online Help*.

The CA SDM user base is separate from that of CA EEM, and therefore separate from the CA Workflow user base. The advantage for CA Workflow is that workflow items may be assigned to and completed by individuals outside of CA SDM. This is useful when CA EEM points to your organization's LDAP server.

The challenge lies in the CA SDM integration. A workitem cannot be assigned directly to the actual CA SDM contact record. Workitem assignment to CA SDM contacts can be made only if both of the following conditions are met:

- The CA SDM contact has a corresponding record in CA EEM
- Those records have matching userids

To summarize:

- CA SDM uses the MDB to store Contact information. CA SDM also features an LDAP integration, which allows it to create new Contacts from an LDAP server and synchronize existing contacts with the directory.
- CA EEM is CA's solution to centralized user management. If you have several CA products installed, they all may be using CA EEM to store identities and access policies.
- CA EEM may be configured to either point to an external (LDAP) directory or use the MDB to store user information. CA EEM itself has an LDAP interface for use when it is configured to use the MDB.

Note: The tables used by CA EEM in the MDB are different from the ones used by CA SDM.

- CA Workflow always uses CA EEM for its user information and authentication. Workitems can be assigned only to users known by CA EEM, and only users defined by CA EEM may access CA Workflow's IDE and Worklist web UI.

The following information discusses the various configuration options for the products and how they affect CA SDM and CA Workflow integration. Select the configuration that best matches your use of CA Workflow and (optionally) an LDAP directory server.

More information:

[CA EEM and CA Workflow User Base Configurations](#) (see page 153)


[CA EEM as LDAP Configuration](#) (see page 157)

Default Configuration

CA SDM and CA Workflow store user information after a default installation as follows:

- CA SDM stores user information in the CA SDM Contact tables in the MDB.
- CA Workflow uses CA EEM to store user information in the CA EEM User tables in the MDB.

The following diagram illustrates where the products store user information:




Consider using the default configuration if your site makes little or no use of CA Workflow or an LDAP directory server. This configuration can present some challenges. Each CA Workflow user requires a CA EEM user record.

External LDAP Server Configuration

Administering contacts is simplified if CA SDM and CA Workflow use the same LDAP directory server. This is especially useful if your site already uses an LDAP server as the main directory for users. Using CA SDM's LDAP capabilities, the userids of contact records match the ones in the directory, allowing for smooth assignment of workitems to CA SDM analysts. The integrations become even more seamless if CA SDM uses CA EEM for authentication.

The following diagram shows a scenario in which all users in the LDAP directory automatically have access to CA Workflow Worklist web application. CA SDM still uses the contact tables in the MDB for its user records, but these records can be imported from and periodically synchronized with the LDAP server records.




More information:

[LDAP Directory Data](#) (see page 183)

CA EEM as LDAP Configuration

When CA EEM is configured to use MDB rather than an external directory to store user information, CA EEM exposes the user directory using an LDAP interface. If your site does not use an external LDAP server, you can still get the advantages of external LDAP configuration by configuring CA SDM to use CA EEM as an LDAP source. This configuration can be useful if your site does not use an LDAP server but you want to consolidate user management in CA EEM. Other CA products also use CA EEM, which can greatly simplify user management.



Note: This configuration is applicable only when CA EEM is configured to use the MDB. If CA EEM is configured to an external LDAP server, configure CA SDM to point to the same LDAP server, *not* CA EEM.

Upgrade

If you are upgrading from an earlier version of CA SDM or batch loading many contact records, you must manually add these users to CA EEM for them to gain access to CA Workflow.

Configure the CA EEM r8.4 SP4 CR05 User Store

You can configure CA EEM r8.4 SP4 CR05 to store user records in an external LDAP directory or in its own internal MDB tables. When CA EEM uses an external LDAP directory, it is a read-only interface; you cannot add or modify users through the CA EEM interface.

Follow these steps:

1. Click Start, Programs, CA, Embedded Entitlements Manager, EEM UI.
The CA EEM user interface appears.
2. Click the Configure Tab.
3. Click the EEM Server subtab.
4. On the left-hand pane, click the Global Users / Global Groups link.
5. On the right-hand pane, select one of the following options:

- Store in internal datastore
- Reference from an external directory
- Reference from CA SiteMinder

Note: If you select the Reference from an external directory option, you are prompted for the LDAP server details.

6. Click Save.

The user store configuration for CA EEM is complete.

Note: For more information about CA EEM, see the *CA EEM Online Help*.

Configure the CA EEM r12 CR02 User Store

You can configure CA EEM r12 CR02 to store user records in an external LDAP directory or in its own internal MDB tables. When CA EEM uses an external LDAP directory, it is a read-only interface; you cannot add or modify users through the CA EEM interface.

Follow these steps:

1. Click Start, Programs, CA, Embedded Entitlements Manager, Admin UI.
The CA EEM user interface appears.
2. Click the Configure Tab.
3. Click on User Store subtab.
4. On the left-hand pane, click the User Store link.

5. On the right-hand pane, select one of the following options:

- Store in internal user store
- Reference from an external LDAP Directory.
- Reference from CA SiteMinder

Note: If you select the Reference from an external directory option, you are prompted for the LDAP server details.

6. Click Save.

The user store configuration for CA EEM is complete.

Note: For more information about CA EEM, see the *CA EEM Online Help*.

Add Users and Groups

If CA EEM is configured to reference an external directory, you cannot add users using the CA EEM user interface. CA EEM is a read-only interface to the LDAP server. You must use whatever interface is provided with your particular LDAP server product to update user records.

To add a new user record

1. Click Start, Programs, CA, Embedded Entitlements Manager, Admin UI/EEM UI.
2. Log in using the CA EEM administrator user name and password. These are specified during the CA EEM installation. CA EEM must be installed separately and is not a configuration option for CA SDM.
3. Click the Manage Identities tab.
4. On the left-hand pane, click the Users tab to search for and update existing user records.

Note: To manage CA EEM groups, click the Groups tab.

5. Click the icon to the left of the Users folder.

The form for creating a user record appears.

6. Complete the form and click Save.

The new CA EEM user record is saved in the MDB.

Note: The steps to edit an existing user record and maintain group records are similar to these steps. For information, see the *CA EEM Online Help*.

Configure CA Workflow Access in CA EEM

All logins to CA Workflow are authenticated by CA EEM. A person must have a CA EEM user record in order to access the CA Workflow IDE or Worklist application. The CA Workflow administrator, who is specified during CA SDM configuration, has full access.

By default, this user is used by CA SDM for the Workflow integration. This user account is set by the `cawf_username` and `cawf_password` Options in Options Manager. You must make sure that the username and password set in these options are correct and the user has full access to CA Workflow resources within CA EEM.

CA Workflow also uses CA EEM to restrict access to specific CA Workflow functions. This access is controlled by two Resources Classes called IDE and Process:

- The IDE resource has a single action named *login*, which gives login access to the IDE. A user must have permission for this action to login to the CA Workflow IDE application.
- The Process resource has the single action named *start*, which gives the ability to start a process instance. A user must have permission for this action to start processes from within the Worklist web application.

Note: All users known to CA EEM have access to the CA Workflow Worklist application to view and perform workitem tasks. This permission is only available to start new instances from the Worklist. These resource classes are defined with the CA SDM application instance in CA EEM; when logging into the CA EEM Web user interface, you must specify the CA SDM application instance in order to see the resources, policies and groups discussed here.

Users who need to either login to the IDE or start process instances need an authorization grant to the resources and actions named above. The CA SDM configuration adds two policies to CA EEM that grant access to these resources. For convenience, the policies grant access to two groups in its application instance: Workflow Administrators and Workflow Process Initiators. You can simply add users to the Workflow Administrators group for them to gain access to the IDE. Adding users to the Workflow Process Initiators group allows them to start processes from the Worklist application.

To add or remove users from the groups mentioned above

1. Login to the CA EEM Web UI
2. On the login page, select the CA SDM application and specify the CA EEM administrator name and password.
3. On the main CA EEM page, select Manage Identities.
4. Select Users search, enter the search criteria and perform the search.
5. Select a user in the result list.
6. On the user details display, add or remove group membership in the Application Group Membership section.

Note: If this section is not displayed, you may need to press the Add Application User Details button.

7. When you have finished, press Save.

The user is added or removed as applicable.

CA Workflow Workitem Assignment

A CA Workflow workitem's Role may be assigned to human and non-human actors. Examples of non-human actors can include a custom Java object, a command-line process or another CA Workflow process instance. For human actors, the workitem role is set to Global User List, which is the CA EEM user repository. This must be a semi-colon separated list of userids. For example:

ServiceDesk; abeju01; My Group

assigns the workitem to users ServiceDesk, abeju01 and anyone who belongs to My Group. This means that any one of these users can complete the workitem.

Note: All these users and groups must be known to CA EEM. So the group My Group is a group in CA EEM, and not a group in CA SDM.

To assign a workitem dynamically to a single user, set the Role userlist to \$MyUser.

Note: Do NOT add double-quotes around the string.

Declare a string attribute named MyUser in the process definition. When the workitem is created, whatever value is in MyUser is used for the workitem assignment. This means that you must assign a valid value to MyUser, either a single username or list of user names separated by a semi-colon. This assignment must be done before it is used in a workitem.

An example of assigning userids to variables is in the Order PC definition demo. It assumes the userid of a CA SDM Contact record matches the userid of a corresponding user record in CA EEM. The Order PC demo shows how to retrieve userids from CA SDM using the web service interface. The userids can come from the ticket (such as the assignee, category assignee, etc.).

To summarize setting the assignment of a workitem

1. Launch the CA Workflow IDE.
2. Double-click name of the Process Definition you wish to edit.
The Process-Definition appears in the Process Designer window.
3. Select the Roles tab.
4. Add or Update a Role.
5. Select Global User List from the list of available Actors and click Edit.
6. Enter a list of CA EEM userids separated by a semi-colon. You may use the Browse button to browse and select users known to CA EEM.

Security Considerations

When you first install CA SDM, the system is set up to allow maximum access to any contact that does not have an explicit access type defined in the contact record. You may want to make additional modifications to the predefined security implementation. At a minimum, you should perform the following steps before you allow people to begin working with the application:

1. Review the predefined access types to determine a reasonable default for your system.

Administrator is set as the default access type, which is not a good choice for most sites. For example, some sites offer read-only CA access to most members of the IT organization. If you set CMDB User as the default access type, you do not have to set the access type of new users unless they need additional privileges. Similarly, if most users require the privilege to write configuration information, you can select CMDB Analyst as the default access type.

2. Assign the access types of remaining contacts explicitly.

For example, if you select CMDB User as the default access type, modify the contact records for your analyst contacts to assign an access type of analyst.

Note: For more information, see the *Online Help*.

CA EEM Authentication for CA Process Automation

CA SDM and CA Process Automation communicate using a web services exchange over HTTP. Although every measure is made to pass minimal amounts of sensitive information between the products, a malicious entity can access user names, passwords, and proprietary information. You can take deliberate steps to secure server communication.

For CA Process Automation authentication, consider the following recommendations:

- As an option, you can configure CA Process Automation to use CA EEM as an authentication server. CA Process Automation implements default groups and policies within CA EEM. You can modify the default groups and policies to meet the needs of your organization.

- Using CA EEM eliminates the need to pass plain text user names and passwords for authentication purposes. If you are using multi-tenancy, CA EEM is required for enabling multi-tenancy within CA Process Automation.

Note: To achieve authentication security in this integration, it is not necessary to have CA SDM configured to use CA EEM. However, CA EEM is required for CA Process Automation multi-tenancy implementation. For information about implementing multi-tenancy with CA Process Automation, see the CA Process Automation user guide documentation. For information about configuring CA Process Automation to use CA EEM as an authentication server, see CA Process Automation installation and configuration documentation.

- Configure CA Process Automation to communicate using secure communications over HTTPS. HTTPS URLs use SSL/TLS to eliminate plain text exchanges while protecting proprietary and other sensitive data from accidental or malicious disclosure.

Note: For information about configuring CA Process Automation to use HTTPS, see the CA Process Automation installation and configuration documentation.

User Authentication

CA SDM provides a user authentication solution that you can customize as part of the access type. The same authentication is used by all CA SDM interfaces and by other CA products.

Note: You can set up CA SDM user authentication on a separate computer if it suits your needs. See the *Implementation Guide* for more information.

Authentication is flexible, allowing you to take advantage of external authentication mechanisms, such as Windows, HTTPD user validation or LDAP authentication. You can also select from a variety of internal authentication options, including operating system password, PIN, guest user access, or no access at all.

How CA SDM Authenticates Users

CA SDM authenticates users based on the user ID defined in their contact record. The product also does the following when a user requests access to the system:

1. If an external user ID is available (from HTTPD or Windows validation), CA SDM looks up the contact by login ID. If the contact is found and has an access type that permits external authentication, the user is allowed into the product.

2. If there was no successful external authentication, CA SDM prompts the user for a user ID and password. The product looks up a contact record for the user ID, obtains the access type, and then authenticates the user as specified by the access type.

Many installations find the predefined access types define authentication that is reasonable for that type of user; however, in some cases you may need to modify the authentication information for a predefined access type or define a new access type to handle a different authentication method for some of your users. You should review the authentication settings for the predefined access types to determine if they meet your needs, or if you need to modify them, or define additional types.

External Authentication

CA SDM permits users to access the system without supplying a user ID if all of the following conditions are met:

- External authentication is set for the user.
- The user's externally authenticated user ID is associated with a contact in your contact table.
- The contact record has an access type whose authentication definition permits external authentication.

External authentication does not permit users to access the system in the following cases:

- A user attempts access through a nonsecure server.
- A user attempts access but is assigned to an access type that does not allow external authentication.

None of the predefined access types use external authentication. If you want to use external authentication for users, consider modifying the employee, analyst, and administrator access types to set external authentication. Your individual site requirements and different types of users determine whether to allow external authentication. When external authentication is used, the server configuration controls the access to files and directories. When you define authentication for an access type, you can decide the usage as follows:

- Do not use any external authentication that is already implemented, such as the user login on Windows or validation by the HTTPD server.
- Use the authentication that is implemented and allow or deny access based on it.

Note: If external authentication is not allowed, the user is authenticated based on the validation type that you specify.

Following are some examples of external authentication:

- If a user who has administrator access logs into a Windows computer, the user can perform administrative tasks without re-entering any login information.
- If a user who has HTTPD server validation, the user can access the web interface without re-entering any login information. Because the administrator access type specifies the analyst web user type, the appropriate web interface for the analyst is presented automatically.

Validation Types

Validation types authenticate users only under the following circumstances:

- The user access type does not permit external authentication.
- The user access type permits external authentication, but the user has not been validated externally (for example, the user has attempted access through a nonsecure server).

CA SDM provides you with the following validation options:

- **No Access**—Users of this type have no access unless external authentication is allowed and is valid.
- **Open**—Users of this type have access, with no additional authentication required.
 - **OS**—Users of this type enter their operating system password for access. The operating system used for validation is the one running User Validation Host. This option is the default validation type for the administrator, analyst, and employee access types.
 - **PIN**—Users of this type gain access by entering the correct value for the PIN field in their contact record as their password. You define the PIN field by entering the field attribute name when you select PIN as the validation type. PIN is the default validation type for the customer access type, which uses the value in the customer ID (contact_num) field as the PIN.

Note: For more information about User Validation Host, see the *Implementation Guide*.

Note: For a list of attribute names for the cnt object, which is the object defined for the Contact table, see the *Technical Reference Guide*.

Logged In and Licensed User Counts

The webLicenseCt KPI counts the number of logged in users that are using a CA SDM license. The Licensed? check box on the Access Type page determines whether this contact is a licensed access type. Contacts assigned to unlicensed access types can only view or update their own personal data. Additionally, the webSessionCt KPI counts the total number of web sessions—both licensed analyst users and end users logging in to the web self-service interface are all counted. If a user is logged in multiple times to CA SDM, either on their computer or on another computer, multiple web sessions occur. If the user is an analyst, multiple licenses are used. Multiple open pages for the logged in user do not affect the web session and web license count.

Note: For more information about access types, see the *Online Help*.

Example: Licensed and Logged In Users Counts

Three end users are logged into the web self-service interface and reviewing some announcements.

At the same time five licensed analysts (that have the check box ticket in their access type) are logged into the analyst interface and working on incidents.

The webLicenseCt KPI shows a count of five, meaning that five licenses are currently being used.

The webSessionCt KPI shows a count of eight, meaning that eight total users are logged into the web interface of CA SDM.

Internal Logs

You can define whether a particular access type is qualified to view internal logs. If allowed to view internal logs, contacts see a check box labeled Internal on each of the Log Activity windows, which they can select to mark the activity as internal. When activities are marked as internal, only contacts with an access type that is qualified to view internal logs sees the activity or is notified of it.

CA SDM Integration

The access type also identifies the user type that contacts with this access type when they use other CA products with CA SDM. When a contact with this access type uses the other product, their user rights to that product are determined by the values you specify for the access type.

Data Partition Associations

Data partitions can be assigned to individual contacts, but the preferred method is to assign data partitions based on access type. After associating data partitions with different access types, you can associate a contact to a particular access type, and define its data partition. The access type has a specific option to override the contact data partition.

To associate a data partition to an access type, you set up data partitions that are meaningful to your site and then select one of the data partitions when you define or modify an access type.

Important! Other CA SDM security settings can take precedence over data partition.

Data Partitions

A data partition is a subset of a CA SDM database that controls a user's access to tickets and other data records based on their content. For example, you can restrict a user's view of the database to only those configuration items assigned to the user's organization. In addition, you can restrict a user to updating only assigned tickets and allow read-only access to other configuration items.

A data partition consists of a set of constraints. The data partition constraints identify the table that is being controlled by the data partition and the constraint type. Constraint types specify what the user can do, such as create, delete, update, view, and so forth, within the data partition. After you assign a data partition to an access type, which you in turn assign to a user's contact record, the constraints and constraint types are what control the user's access to the records in the CA SDM database table.

You can also view constraints independent of the data partition using them. For example, you may want to view all constraints for a specific table, regardless of the data partition.

Note: For more information about how to use data partition constraints, see the *Online Help*.

Data Partitions Setup

You can define an unlimited number of data partitions. Each data partition consists of a set of constraints and validations on each database table restricted by the data partition. For each table in a data partition, you can specify independent authorizations to view, update, create, or delete records using criteria specified in a format similar to an SQL WHERE clause. You can base the restriction on any attribute in the record being accessed, combined with any data in the user's contact record. This allows considerable flexibility when defining data partitions. For example, using the Vendor field in the Contact table allows data partition restrictions to be placed on vendors permitted to access CA SDM directly.

For performance reasons, CA SDM does not allow a data partition constraint to contain a Cartesian join. A Cartesian join results from a constraint containing an "OR" that does fully restrict all joined tables on both sides of the OR. To ensure that your data partition constraint does not produce a Cartesian product join, enter the following command:

Windows

```
bop_cmd -f $NX_ROOT\bopcfg\interp\bop_diag.frg "check_queries()"
```

UNIX

```
bop_cmd -f $NX_ROOT/bopcfg/interp/bop_diag.frg "check_queries()"
```

Important! Any data partitions that are flagged by this program must be updated appropriately. For more information about how to define data partitions, constraints, and constraint types, see the *Online Help*.

Constraint Specifications

You specify constraints and validation tests in Majic using the object definition metalanguage.

Note: For information, see the *Technical Reference Guide*.

Constraints defined in Majic closely resemble an SQL WHERE clause, with the following exceptions:

- Attribute names in the constraint are object attribute names, not the database attribute name from the schema.

- You can refer to the value of an attribute in the contact record for the logged-in user with a name of the following form, where *att_name* is the Majic name of the desired attribute:

`@root.att_name`

For example, specifying `@root.location` refers to the location ID of the current contact.

You specify joins with a specification of the following form, where *foreign-key* is the Majic name of the SREL attribute in the table for which you are writing the data partition constraint, and *attribute-in-referenced-table* is the Majic name of the attribute in the table being joined:

`foreign-key.attribute-in-referenced-table`

For example, to refer to the maintenance vendor of the asset associated with an incident report, specify the following:

`resource.vendor_repair`

This specification is recursive. For example, you could refer to the name of the vendor with the following:

`resource.vendor_repair.name`

The following table contains examples of valid constraints to use for the `Change_Request` table, used to store change order information:

Constraint Type	Code and Description
View	<code>assignee.organization = @root.organization</code> Specifies the user can only view change orders where the assignee's organization is the same as the user's organization.
Pre-Update	<code>requestor = @root.id</code> Specifies the user can only update the change orders where he is the caller or requester.

However, you cannot write a constraint that uses joins on both sides of the expression, as shown in the following example:

`assignee.org = requestor.org`

Constraint Types

You can assign the following types of constraints for each controlled table in a data partition:

Create

Specifies criteria that must be met before creating a record. When a user in the data partition attempts to create a record that does not match the create test condition, CA SDM displays the error message associated with the constraint and does not save the record.

Defaults

Specifies one or more assignment statements, separated by semicolons, defining values to be assigned to empty fields in a new record at the time the record is stored. The syntax of each assignment statement is, where *att_name* is the name of a Majic attribute from the record, and *value* can be an integer, a string enclosed in quotes, or a reference in the form *@root.att_name* to a Majic attribute in the contact record for the current user:

att_name=value

For tables updated for tickets, default values are placed into the record at the time it is displayed and are shown on the initial display of a new record. You can assign a default value to a reference field (a Majic SREL) by coding it in the form of a persistent ID. A persistent ID is an object name followed by a colon and an integer ID. For example, you can set a default value for category by including the following in the defaults specification, where PCAT is the target of the SREL (as shown in the Majic file) and 12345 is the ID number of the desired category:

category='PCAT:12345'

You can list persistent IDs for an object using a command of the form:

bop_odump domsvr pcat "" sym

Delete

Specifies criteria that must be met to delete a record. When a user in the data partition attempts to delete a record that does not match the delete condition, CA SDM displays the error message associated with the constraint and does not delete the record.

Pre-Update

Specifies the records in the controlled table that a user can update in the data partition. When a user in the data partition requests a record that does not match the pre-update condition, CA SDM makes the record read-only and displays the error message defined with the constraint.

Update

Specifies criteria that must be met when a record is saved. When a user in the data partition attempts to save a record that does not match the update condition, CA SDM displays the error message associated with the constraint and does not save the record.

View

Specifies the records in the controlled table that a user can view in the data partition. This constraint is automatically applied to all lists selected by a user in this data partition, in addition to any selection criteria explicitly specified by the user.

View can include joins to other tables and references in the form `@root.att_name` to Majic attributes in the contact record for the current or logged-in user. Valid examples are:

```
requestor.organization = @root.organization
```

```
requestor.organization.name = 'MIS'
```

```
assignee = @root.id
```

```
assignee.organization = @root.organization
```

Note: The Create, Delete, Pre-Update, and Update constraint types now support joins to other tables. They can also include references in the form `@root.attribute` to attributes in the contact record for the current user.

Create a Data Partition Constraint for CAB Assignments

You can create a data partition constraint that lets users update only change orders assigned to a CAB to which the logged in user belongs.

To create a data partition constraint for CAB change order user assignments, assign the following constraint values for a Change_Request controlled table in a data partition:

- Constraint type: Pre-Update
- Constraint specification: `cab.[group]group_list.member IN (@root.id)`

The logged in user can only update change orders that are assigned to a CAB to which the user belongs.

More information:

[CAB Console and Reporting](#) (see page 676)

Configure Knowledge Management Data Partition Constraints for Role-Based Permissions

Knowledge Management data partitions are enabled to let you use group and role permissions by default in CA SDM. If you are upgrading from a previous release, the migration tool updates your data partition constraints.

If you used custom data partition constraints to manage knowledge permissions in a previous release, update the constraints manually for the `O_INDEXES` and `SKELETONS` tables. You can view the default data partition constraints and apply the changes, as appropriate to your environment.

To update data partition constraints

1. On the Administration tab, browse to Security and Role Management, Data Partitions, Data Partition Constraints.
The Data Partition Constraints List appears.
2. Click Show Filter.
The search filter appears.
3. Enter **Service Desk Analyst** in the Data Partition search.
4. Enter **O_INDEXES** in the Table search.
5. Click Search.
The search results appear.
6. Click the *View* constraint type.
The Data Partition Constraint Detail page appears.
7. Click Edit.
The Update Data Partition Constraint page appears.
8. Modify the Constraint tab to replace "READ_PGROU in @root.pgrouPs" as follows:
`READ_PGROU in @root.pgrouPs OR`
`READ_PGROU.[pgrouP]contained_roles.role IN @root.id`
Save the constraint.
9. Open the *Delete* and *Pre-Update* constraint types from the Data Partition Constraints List.
10. Modify the Constraint tab to replace "WRITE_PGROU in @root.pgrouPs" in *Delete* and *Pre-Update* as follows:
`WRITE_PGROU in @root.pgrouPs OR WRITE_PGROU.[pgrouP]contained_roles.role IN @root.role`
11. Save the constraints.
12. Repeat the steps to update the *View*, *Delete*, and *Pre-Update* constraints in the `SKELETONS` table in your data partition.

The data partition constraints are updated.

Surveys

Customer surveys let CA SDM administrators systematically collect and analyze customer feedback about service desk performance. You can customize surveys to suit the needs of your site.

Configure Your System for Surveys

Before you can use surveys, you need to configure your system properly, which involves two steps:

1. Install and configure the CA SDM web interface. When a user accesses the URL for a survey, the web interface formats the survey and populates the survey information. See the *Implementation Guide* for more information.
2. Using the Options Manager, configure and install the `web_cgi_url` option to specify the location of the CA SDM web engine. See the chapter "Controlling System Behavior" and the *Online Help* for details.

Prepare a Survey

You prepare surveys using the Customer Survey List, which is a typical list window. For example, you can use this window to view all surveys or a subset based on search criteria that you enter; you can create new surveys; you can view details on a particular survey; and you can report on the surveys currently listed.

Each survey has the following features that you can define:

- A name that you can use for searching and reporting purposes
- An introduction that you can use to explain the purpose of the survey to customers
- An ordered list of questions for the customer to answer, each of which includes a set of possible answers
- An optional area where the user can enter free-form comments
- A completion message to display after the user submits the survey

Note: For more information about how to create surveys, see the *Online Help*.

Define Survey Notifications

The Survey tab on the Update Activity Notification page allows you to define a survey notification for an activity notification. When the selected activity notification is triggered, the contact who initiated the activity receives the survey notification. An activity log is generated both when a survey notification is sent and when one is received back from a customer.

To set up a survey notification

1. On the Administration tab, browse to Notifications, Activity Notifications.

The Activity Notification List appears.

2. Select the desired activity notification.

The detail page displays.

3. Click the Edit button.

The Update Activity Notification page displays.

4. Edit the fields as appropriate.

5. Select the appropriate object type from the drop-down list.

6. Click the Survey tab.

This tab contains the following fields:

Send Survey

This check box allows you to activate or deactivate the survey. If selected, the survey is sent to the contact when the selected activity notification is triggered.

Default Survey

Specify a default survey using the search icon or specify your own in the text box.

Notification Method

Choose *one* of the following notification methods:

- Email
- Notification
- Pager_Email

Survey Message Title

Enter the title for the survey.

Survey Message Body

Enter a message for the contact. When a contact receives notification of a survey, the message body automatically includes a URL that they can access from their web browser to find and fill out the survey form.

7. Save the activity notification.

When the selected activity notification is triggered, the contact who initiated the activity receives the survey notification.

Survey Reporting

Within CA SDM you can report on surveys from within the administration tab of the web client in all of the usual ways. For example, from the Customer Survey List window, you can choose Reports from the File menu, and then choose a Summary or Detail report. You can also choose Print Form from the various detail windows to print the form data for your surveys, questions, and answers.

You can also fashion your own reports based on the survey data maintained in the CA SDM database.

Managed Survey

The Managed Survey lets the CA SDM administrator select a desired survey sample population and match it to a specific survey. The administrator can then distribute targeted requests for respondents to take the survey at a specific time. This gives the administrator the flexibility of creating open survey periods, while maintaining the ability to have activity-based and category-based surveys related to Requests, Change Orders, and Issues.

The purpose of Managed Surveys is to provide a mechanism for managing surveys. This function can be useful when for survey forms that need to be monitored from time to time (for example, surveys only used during a short period every year or surveys that have been offline too long).

Important! If you want to send the survey to a large number of contacts, set the default value of `NX_SURVEY_ILIMIT` in `NX.env` to a higher limit, such as 1073741824.

Note: For more information about how to create managed surveys, see the *Online Help*.

Web Services

Web Services conform to data exchange standards that do the following:

- Let applications communicate over HTTP to various servers regardless of operating environment.
- Let most applications access CA product functionality.
- Let Web Service clients create tickets, update assets, search the knowledge base, and more.

Note: For more information, see Managing Web Services in the *Implementation Guide*.

Chapter 5: Configuring User Accounts

This section contains the following topics:

[Contacts](#) (see page 177)

[Contact Definitions](#) (see page 177)

[Groups](#) (see page 179)

[Contact Types](#) (see page 179)

[Special Handling Types](#) (see page 180)

[LDAP Directory Data](#) (see page 183)

Contacts

An important part of establishing a working service desk is defining the users who are going to access it. In CA SDM, users are named *contacts*, and you can perform several tasks to set up and manage them:

- Set up contacts manually.
- Organize contacts into groups that define areas of responsibility.
- Establish contact types to organize your CA SDM contacts into logical groupings based on how they use the system.
- Import LDAP user information into a CA SDM contact record.
- Assign a contact to a role to define the accessible system functionality.
- Assign a special handling type, such as Very Important Person (VIP), to a contact.

Contact Definitions

Everyone who uses CA SDM must be defined as a contact. A user's contact record defines the user information that the system needs as follows:

Basic Identification

Defines basic identification, such as the user's name and contact type. A contact's name is used as the primary identifier when selecting a contact or filling in contact information in other contexts.

Login

Defines login information, such as the user ID and in some cases, a PIN field to use as a password that verifies the user upon login. The user ID is used to identify the user in the contact table for authentication purposes and for determining the access types assigned to the user. Depending on how the administrator has configured security, another field such as the contact ID, can be specified as the PIN field and the user can use that as the password for login.

Security

Defines the access type that is assigned in their contact record or by a default access type, depending on how you set up the security for your system. In addition, a user's access type can be assigned based on their membership in an LDAP Directory group.

A user's access type determines all aspects of their security, including how they are authenticated to the system, which web interface they see, and what product functionality they can access.

Security management is a feature of the web interface.

Service Type

Determines the level of service a user receives. A contact's service type defines the level of service a user should receive. Service Level Agreements (SLAs) are negotiated with CA SDM customers, and service types serve as the mechanism for CA SDM to implement SLAs. By associating a service type with a user's contact record, you can guarantee that when a ticket for which a user is identified as the affected end user is created, the service type for the ticket is at least as good as the contact's service type.

Setting up SLAs using service types is a feature that you, as the administrator, perform using the web interface.

Automatic Assignment

Defines automatic assignment information, such as work shift and availability (used for analyst contact types only). You can set up analyst contacts to determine if they are eligible for automatic assignment. Automatic assignment is valid only for requests, and is defined as part of the request area definition. It is also linked to the groups to which the analyst belongs.

How to Send Users Notification Messages

Defines the notification information of a contact, which includes the following:

- various email addresses and telephone numbers to be used for notification
- method to be used for the notifications with different urgency levels
- work shifts during which notifications will be received

The notification delay calculation takes the Contact Time Zone into account. If the Contact Time Zone is not set, the server time zone is used, instead. Using the server time zone may result in notifications firing at times, perceived to be outside the Work shift settings.

Organizational information (such as location, organization, and department) lets you group contacts based on the organization to which they belong. For example, associating a contact with a location links the contact to a physical address and also helps in determining automatic assignment. The organization can be assigned a service type, making it easier to manage SLAs by organization rather than by individual contacts.

Groups to which a User Belongs

Organizes contacts into groups that represent specific areas of responsibility within your service desk. You can set up and define contacts using the web interface.

Groups

A group is a collection of contacts that share a common area of responsibility. In CA SDM, groups are implemented using the predefined group contact type, making a group just a special type of contact. A group has the same basic information as a contact, with the important additional feature that groups are one of the keys to automatically assigning requests. You can associate request areas, locations, and a work shift with a group. These attributes are used to determine if and when the contacts in the group can accept automatic assignment of a request.

Note: For information about defining groups, see the *Online Help*.

Contact Types

Contact types are used to categorize CA SDM users into logical groupings based on how they use the system. For example, some of the many contact types that are predefined by the system are analyst, customer, and group. These predefined contact types meet the needs of most CA SDM implementations; however, if your circumstances require it, you can modify the predefined contact types and create contact types. When you define users as contacts, you can associate a contact type with each one.

Note: For more information about defining contact types, see the *Online Help*.

Determine Behavior Based on Contact Type

The contact *type* determines which contacts display (and have permission) in different situations. For example, when you manually assign any type of ticket, such as a request or an issue, the field to specify the assignee requires that the person you specify have a contact type of analyst. If you choose to select a contact from a selection list for this field, only contacts with the type of analyst display in the selection list. Entering a contact with a different type displays the analyst search screen only.

Note: An important feature of the contact type is in the implementation of groups of contacts through the predefined group contact type.

Notification Setup Based on Contact Type

You can base notification on contact type, which allows you to send a notification message to all contacts of a particular type.

More information:

[Notifications](#) (see page 85)

Select Contacts Based on Contact Type

You can select users based on contact type in various contexts. For example, most list and selection windows that display contacts have a search field where you can select a contact type as a search criterion.

Special Handling Types

You can define special handling types that identify contacts who require special attention. You can use the special handling types that CA SDM provides or create your own types. You can view and locate tickets that specify an affected end user who requires special attention. For example, analysts can browse the V.I.P. folder on the Scoreboard to identify tickets that specify a VIP as the affected end user.

The following examples are contacts that special handling types can identify:

- Very Important Persons (VIPs) such as executives
- Customers with support renewal in process
- Customers with disabilities requiring special handling or equipment
- Visitors
- Contacts suspected of misusing or abusing system resources

When one or more special handling types are assigned to a contact, tickets that specify the contact in the Affected End User field show an alert banner, icon, or both. You can use ticket fields and special handling types to track tickets, and distinguish between two related but possibly distinct contact types. For example, a VIP (Affected End User) has an assistant (Requestor) acting on their behalf. When the Affected End User is a contact assigned to a VIP special handling type, an analyst can prioritize tickets more accurately.

More information:

[How to Configure Special Handling Contacts](#) (see page 181)

[Associate a Contact to a Special Handling Type](#) (see page 182)

How to Configure Special Handling Contacts

To configure special handling contacts, do the following steps:

1. Create special handling types.
2. [Associate a contact to any number of special handling types](#) (see page 182).
Similarly, a special handling type can have many contacts.

A contact that is associated with one or more Special Handling types is visually distinguished in the Contact Detail form and the Quick Profile browser using a banner at the top of each page. This banner displays the alert icon and alert text for each Special Handling type assigned to the contact.

Additionally, any tickets that identify the contact as the Affected End User are indicated as follows:

- Alert Icons and Alert text appear in a banner at the top of the ticket detail form.
- Alert Icons appear in the ticket list.
- The Scoreboard includes a V.I.P. folder and subfolders for each ticket type. The V.I.P. subfolders include tickets for affected end users who are VIP special handling contacts.

Note: The V.I.P. Scoreboard folder displays for analyst roles.

More information

[Associate a Contact to a Special Handling Type](#) (see page 182)

Associate a Contact to a Special Handling Type

You can assign a special handling type to a contact to alert analysts about tickets that affect end users with special requirements, such as for a visually impaired person, a contact that poses a security risk, and so on.

To associate a contact to a special handling type

1. On the Contact Detail page, select the Special Handling tab.
The Associated Special Handling List tab lists special handling types that are associated with the contact.
2. Click the Update Contact's Special Handlings button.
The Search filter appears.
3. Search for the special handling type that you want to associate to the contact.
The Special Handlings Update page appears.
4. Select one or more handling types from the left column and use the move button (>>) to move the types to the right column. Click OK.

Note: You can remove an association from a contact by using the move button (<<) to move the type from the right column to the left column. You can click the search icon to search for the value you want.

The contact is associated to a handling type.

CA SDM displays any of the following, depending on the handling type, when a ticket specifies the contact in the Affected End User field:

- An alert banner appears on the Contact Detail for the affected end user on a ticket.
- Alert text appears as a banner at the top of the ticket detail page and in the Quick Profile.
- Ticket lists highlight the contact row and show an alert flag.
- A V.I.P. folder appears in the Scoreboard for analyst roles. The folder contains all tickets that are associated with contacts (Affected End Users) that have a VIP special handling type.

LDAP Directory Data

Lightweight Directory Access Protocol (LDAP) is a network communications protocol for querying and modifying directory services running over a TCP/IP network. An LDAP directory is a tree structure that contains entries for managing users, groups, computers, printers, and other entities on a network.

CA SDM can be configured to access an LDAP directory, which allows you to use LDAP data in several ways:

- Synchronize contacts with LDAP user records. Synchronization can occur in the following ways:
 - **At login**—When a user logs in to the product, if an LDAP record exists for that user but a corresponding contact record does not exist, a contact record is automatically created based on the LDAP information.
 - **New Contact**—When you manually create a contact record, you can select an LDAP record and merge its attribute values with their corresponding fields in the new contact record.
 - **Batch Update**—You can run batch jobs to automate the processes of importing and updating contact records with information from the corresponding LDAP records.

Note: Synchronization with LDAP is a one-way process. LDAP data can be used to create and update contacts, but the product does not support updates to the LDAP directory.

- Assign CA SDM access types based on LDAP group membership.
- Implement an alternative method of performing CA SDM authentication.

Note: The `ldap_virtb` component provides LDAP integration functionality, is installed with CA SDM by default, and can run on either a primary or secondary server, regardless of operating system type. For information, see the *Implementation Guide*. The `$NX_ROOT/bopcfg/majic/ldap.maj` file specifies the mapping between LDAP attributes and contact record attributes.

Important! CA SDM requires that LDAP records have an entry in the last name field in order to search, view, and import the LDAP data.

Important! CA SDM supports *paged searching*, which searches all records in your LDAP directory. Paged searching also enables you to import new contact records or sync existing contact records from any number of LDAP records. These capabilities are limited, however, if you are using Sun Java System Directory Server or Novell eDirectory because these LDAP servers do not support paged searching. In that case, you can only search, import, and sync with the number of LDAP records specified by `NX_LDAP_MAX_FETCH`. For more information, see [NX.env File](#) (see page 197).

Configure LDAP Options

You can configure CA SDM to access LDAP directory data.

To configure CA SDM to access LDAP directory data

1. Manually install LDAP options using the Web Interface Options Manager.

Note: The options necessary for basic LDAP integration are identified as required in the Description column in the following table. Options identified as optional are features you can add only if all the required options are installed. The values you specify when installing these options are written to the \$NX_ROOT/NX.env file. For more information about the LDAP options and instructions for installing them, see the *Online Help*.

2. Restart the CA SDM service.

The changes take effect.

Option	Default Value	Description
default_ldap_tenant		Required for multi-tenancy installation. Specifies the default tenant assignment for contacts imported from LDAP. You must use the tenant UUID when setting the Option Value field. Note: You can get the tenant UUID from a database query. For example, "SELECT * FROM ca_tenant".
ldap_enable	Yes	Required. Enables LDAP integration with CA SDM.
ldap_host		Required. Specifies the LDAP database server host name or IP address.
ldap_port	389	Required. Specifies the LDAP server port number.
ldap_dn		Required. Specifies the LDAP server logon distinguishedName. Example: CN=Joe, CN=Users, DC=KLAND, DC=AD, DC=com If the LDAP server supports anonymous binds, this value can be empty.
ldap_pwd		Required. Specifies the password for LDAP server logon distinguishedName. If the LDAP server supports anonymous binds, this value can be empty.

Option	Default Value	Description
ldap_search_base		Required. Specifies the starting point for searches in the LDAP schema tree: (UNIX) You must specify a starting container. For example: CN=Users, DC=KLAND, DC=AD, DC=com (Windows) You do not have to specify a container. You may start at the top of the schema tree. For example: DC=KLAND, DC=AD, DC=com
ldap_filter_prefix	(&(objectClass=user)	Specifies the prefix applied to an automatically generated filter when searching for LDAP users. Note: This variable has been superseded by the ldap_user_object_class option. It is not available in Options Manager, but can be set manually in the NX.env file.
ldap_filter_suffix	)	Specifies the suffix applied to an automatically generated filter when searching for LDAP users. Note: This variable has been superseded by the ldap_user_object_class option. It is not available in Options Manager, but can be set manually in the NX.env file.
ldap_user_object_class	person	Required. Specifies the value of the LDAP objectClass attribute applied to an automatically generated filter when searching for LDAP users.
ldap_enable_group	Yes	Optional. Enables CA SDM access type assignment based on LDAP group membership.
ldap_group_object_class	group	Required only if the ldap_enable_group is installed. Specifies the object name applied to an automatically generated filter when searching for groups.
ldap_group_filter_prefix	(&(objectClass=group)	Specifies the prefix applied to an automatically generated filter when searching for LDAP groups. Note: This variable has been superseded by the ldap_group_object_class option. It is not available in Options Manager, but can be set manually in the NX.env file.
ldap_group_filter_suffix	)	Specifies the suffix applied to an automatically generated filter when searching for LDAP groups. Note: This variable has been superseded by the ldap_group_object_class option. It is not available in Options Manager, but can be set manually in the NX.env file.

Option	Default Value	Description
ldap_enable_auto	Yes	Optional. Enables auto generation of contact records from LDAP data.
ldap_sync_on_null	Yes	Optional. Overwrites existing CA SDM contact attributes with null data if the corresponding LDAP user attribute contains a null value.
ldap_service_type	Active Directory	Optional. Use this option if the CA SDM operating environment is Windows and the LDAP directory is <i>not</i> Active Directory (for example, eTrust or Novell). Note: On UNIX operating environment, "Non AD" functionality is used only if this option is <i>not</i> installed. If it is installed, the service type is set to Active Directory.
ldap_enable_tls	No	Optional. Specifies whether Transport Layer Security (TLS) is enabled during LDAP processing.

Verify LDAP Integration

After you have installed the necessary LDAP options, CA SDM users can import LDAP data on a case-by-case basis, eliminating the need to fill in all the contact attribute fields manually.

To verify that LDAP integration is correctly configured, complete the following steps using the Web Interface. If you encounter problems, see [Troubleshooting](#) (see page 196).

To verify that you can search for and import LDAP records

1. Select File, New Contact from LDAP on the Service Desk tab.
The LDAP Directory Search window appears.
2. Specify filter criteria, and then click Search. For example, you could enter b% in the Last Name field to retrieve a list of the LDAP user entries with last names that begin with the letter B.

Note: If your LDAP directory contains thousands of entries and you do not filter your search, your request attempts to retrieve *all* of the LDAP user records. This can cause the request to time-out and return zero records.

Search results matching your filter criteria are displayed.

3. Select an entry.

The Create New Contact window appears, populated with imported LDAP attribute values.

4. Click Save.

The contact record is created.

To verify that you can update a contact using LDAP data

Note: Before performing this procedure, for test purposes you may want to use whatever LDAP editing tool you have available to change one or more attribute values in the entry you used for the previous procedure. You can verify that the contact is updated with the latest LDAP data.

1. Select Search, Contacts on the Service Desk tab.

The Contact Search window appears.

2. Specify filter criteria to search for a contact that has a corresponding LDAP user entry. For example, you could search for the contact you created in the previous procedure.

Search results matching your filter criteria are displayed.

3. Select the contact you want to update with LDAP data.

The Contact Detail page appears, populated with the CA SDM contact information.

4. Click Edit.

The Contact Update page appears.

5. Click Merge LDAP.

The LDAP Entry List page displays a list of any LDAP user entries that correspond with the selected CA SDM contact.

To search the LDAP directory for other entries, you can click Show Filter, specify filter criteria, and then click Search.

Note: If your LDAP directory contains thousands of entries and you do not filter your search, your request attempts to retrieve *all* of the LDAP user records. This may cause the request to time-out and return zero records.

6. Click the LDAP entry of interest.

The LDAP Detail page displays the attribute values for the selected entry. Verify that you have selected the correct entry for the contact you want to update, then click Close Window.

7. On the LDAP Entry List page, right-click the entry that best matches the contact you want to update, and then select Merge into Contact.

The Contact Update page reappears, populated with the current LDAP attribute values. If the LDAP data has changed since you created or last updated the contact, the changes are reflected in the contact attribute fields.

Note: If you have installed the `ldap_sync_on_null` option, and the LDAP entry contains null values for any attribute fields that correspond to contact attributes that currently contain values, the values in the contact record are overwritten with null values when you save the contact data.

8. Click Save on the Contact Update page.

The contact is updated with the corresponding LDAP data.

Create a Contact Automatically

You can configure CA SDM to create a contact automatically from a corresponding LDAP user record whenever a new user logs in to CA SDM.

To enable this feature, install all of the required LDAP options plus the `ldap_enable_auto` option.

The contact record is automatically created as follows:

1. If a user logging in to CA SDM does not yet have a contact record, but the user's login name exists in an LDAP record, the LDAP data is automatically imported and a contact record is created.
2. The automatically created contact record inherits the default access type security settings.
3. The contact can then be assigned an access type explicitly, or the access type can be assigned based on the user's membership in an LDAP Group.

This process is completely transparent to the user, appearing as any other login session.

Access Type Assignments From LDAP Groups

You can configure CA SDM to assign access type values to contacts automatically, based on LDAP group membership. With automatic access type assignment enabled, if an LDAP user record that was used to create a contact belongs to an LDAP group that is associated with one of the CA SDM access types, then the contact is automatically assigned that access type. Otherwise, the contact inherits the default access type.

To enable automatic access type assignment, you must install the `ldap_enable_group` and `ldap_group_object_class` options.

Note: For details about installing the necessary options and associating LDAP groups with access types, see the *Online Help*.

Batch Import Contacts Using LDAP Data

You can run the `pdm_ldap_import` command-line utility to create CA SDM contacts in batch mode using LDAP data.

Note: In addition to creating contacts, `pdm_ldap_import` updates existing contacts if they are not synchronized with their corresponding LDAP entries. You can use the `pdm_ldap_sync` batch process to update existing contacts, but not create ones.

`pdm_ldap_import` has the following syntax:

```
pdm_ldap_import -l "ldap_where_clause" [-c "contact_where_clause"] [-u "userid"]
```

-l "ldap_where_clause"

Specifies the userids of LDAP records to be searched. Replacement variables are indicated with the '?' character. For example, for `userid = ?`. The default value is `userid = ?`. In this special case, id is mapped to the contact attribute `ldap_dn`.

Note: Use the keywords as defined in the `ldap.maj` file. You can also search by using the `memberOf = 'group_dn'` syntax.

-c "contact_where_clause"

(Optional) Specifies how to determine whether the contact record already exists. If the contact record does not exist, a new contact record is inserted. If the contact record does exist and is not synchronized with the current LDAP data, the contact record is updated.

-u "userid"

(Optional) Specifies the login name under which the `pdm_ldap_import` program runs.

Note: You can use wildcards with `pdm_ldap_import` to specify multiple records.

Examples: Batch Imports Using LDAP Data

This example imports a single LDAP record for `userid jsmith11`:

```
pdm_ldap_import -l "userid = 'jsmith11'"
```

This example imports all LDAP records with a `userid` that begins with the letter C:

```
pdm_ldap_import -l "userid = 'c%'"
```

This example imports all LDAP User records in the directory:

```
pdm_ldap_import -l "userid = '%'"
```

More information:

[Batch Update Contacts Using LDAP Data](#) (see page 192)

Batch Import Contacts by Date and Time

You can configure the `pdm_ldap_import` utility to import LDAP records that were created before or after a specified date and time. To enable this functionality, create an `ldap.mod` file with the following content:

```
OBJECT ldap {  
  ATTRIBUTES LDAP_Entry {  
    whenCreated whenCreated STRING ;  
  };  
};
```

This adds the *whenCreated* attribute to the LDAP object.

The rules for filtering records using the *whenCreated* attribute are as follows:

- Use only the `>=` or `<=` operator.
- Specify *all* characters for the date/time value, including the Z. Place a 0 in any location you do not wish to explicitly state (for example, the time of day).
- Place the date/time specification at the beginning of the filter; do not use leading 0s at the beginning of the string.
- Do not include the leading century. For example, to specify the year 2008, use 08.

Note: Single quotation marks must surround the date/time value.

Example: Using the *whenCreated* Attribute to Import LDAP Entries

The following example uses the *whenCreated* attribute to import LDAP entries created after 3/11/2008.

```
Pdm_ldap_import -l "whenCreated >= '080312000000Z'"
```

Example: Using the whenCreated Attribute to Search for LDAP Records

The following example uses the whenCreated attribute with pdm_ldap_test to search for LDAP records created after 3/11/2008.

```
pdm_ldap_test.exe -f "whenCreated>=080312000000Z" -a whenCreated
Starting ldap_test.exe...
LDAP Directory Type      : active directory
Service Desk Platform    : windows
Search Base              : DC=kirklandsd,DC=ca,DC=com
Search Filter             : (&(objectClass=person)(whenCreated>=080312000000Z))
Administrator Username    : CN=Administrator,CN=Users,DC=kirklandsd,DC=ca,DC=com
Administrator Password    : *****
LDAP Host                 : gecko.kirklandsd.ca.com
LDAP Port                 : 389
LDAP API Version          : 3
DN: CN=aixmail,CN=Users,DC=kirklandsd,DC=ca,DC=com
   whenCreated(17)(0): 20080312035327.0Z
DN: CN=hpmail,CN=Users,DC=kirklandsd,DC=ca,DC=com
   whenCreated(17)(0): 20080312035425.0Z
DN: CN=sunmail,CN=Users,DC=kirklandsd,DC=ca,DC=com
   whenCreated(17)(0): 20080312035726.0Z
3 Total LDAP records found...
```

Batch Import Summary and Log Data

The pdm_ldap_import command maintains a detailed log of all activity of each run. The ldap_logging.0-n log file is located in the \$NX_ROOT/log directory.

The following is an example of the summary data pdm_ldap_import returns at the command line:

```
pdm_ldap_import Starting...
pdm_ldap_import Summary: Processed(21) Updated(1) No Matches(7) New Contacts(11) Multiple Matches(0)
Empty Filter(2) Errors(0)
pdm_ldap_import Complete...
```

The following table describes the summary data.

Status	Count	Description
Processed	21	The total number of LDAP entries found
Updated	1	The number of contact records that were updated because the corresponding LDAP entry contained different information
No Matches	7	The number of LDAP entries with no corresponding contact record

Status	Count	Description
New Contacts	11	The number of new contact records that were created based on corresponding LDAP entries
Multiple Matches	0	The number of LDAP entries with multiple matching contact records, as defined by the <code>ldap_search_base</code> option
Empty Filter	2	The number of LDAP entries that cannot be used to generate a valid search filter
Errors	0	The number of LDAP entries that encountered an error during processing. For example, LDAP records that do not contain a value in a field required by CA SDM (such as Last Name) are counted as failures and cannot be imported.

Batch Update Contacts Using LDAP Data

You run the `pdm_ldap_sync` utility to update contact records in batch mode using LDAP data.

Important! This utility overwrites the existing tenant of the LDAP contact defined in CA SDM. If you want to retain the tenant value, you must modify `NX.env` by adding the `NX_RETAIN_TENANT_VALUE` variable manually, and set it to "yes". If this variable is set to "no", missing, or not set properly, the utility overwrites the tenant information.

Note: The `pdm_ldap_sync` utility synchronizes existing contacts with corresponding LDAP entries, but does not create contacts. You can use the `pdm_ldap_import` batch process to create contacts.

`pdm_ldap_sync` has the following syntax:

```
pdm_ldap_sync -l "ldap_where_clause" [-c "contact_where_clause"] [-u "userid"]
```

-l "*ldap_where_clause*"

Determines how to search for matching LDAP records. Replacement variables are indicated with the '?' character. For example, for `userid = ?`, the default value is `id = ?`. In this special case, `id` is mapped to the Contact attribute `ldap_dn`.

-c "*contact_where_clause*"

(Optional) Determines which contacts are used when searching for matching LDAP records.

Default: "`ldap_dn IS NOT NULL`"

-u "userid"

(Optional) Specifies the userid under which pdm_ldap_sync runs.

Note: You can use wildcards with pdm_ldap_sync to specify multiple records.

Examples:

This example establishes a baseline of contact records that have a corresponding LDAP record:

```
pdm_ldap_sync -l "userid = ?" -c ""
```

This example uses the default parameters to update all contacts that have an LDAP distinguishedName:

```
pdm_ldap_sync
```

This example updates a single contact:

```
pdm_ldap_sync -l "userid = ?" -c "userid = 'jsmith11'"
```

More information:

[Batch Import Contacts Using LDAP Data](#) (see page 189)

Batch Update Summary and Log Data

The pdm_ldap_sync command maintains a detailed log of all activity for every run. The ldap_logging.0-n file is located in the \$NX_ROOT/log directory.

The following is an example of the summary data pdm_ldap_sync returns at the command line:

```
pdm_ldap_sync Starting...
pdm_ldap_sync Summary: Processed(21) Updated(1) No Matches(7) No Changes(11) Multiple Matches(0)
Empty Filter(2) Errors(0)
pdm_ldap_sync Complete...
```

The following table describes the summary data:

Status	Count	Description
Processed	21	The total number of LDAP entries found
Updated	1	The number of LDAP entries with information different from their corresponding CA SDM contact record
No Matches	7	The number of LDAP entries with no corresponding CA SDM contact record

Status	Count	Description
No Changes	11	The number of LDAP entries with information identical to their corresponding CA SDM contact record
Multiple Matches	0	The number of LDAP entries with multiple matching contact records in CA SDM, as defined by the ldap_search_base option
Empty Filter	2	The number of LDAP entries that cannot be used to generate a valid search filter
Errors	0	The number of LDAP entries that encountered an error during processing

LDAP Authentication

You can use LDAP to authenticate users logging in to CA SDM. LDAP authentication is available when the CA EEM authentication component is integrated with CA SDM, which replaces the default validation performed by the host operating system. LDAP authentication is only applicable when CA EEM is configured to use an external LDAP directory and you have selected OS authentication for a user's validation type in an access type record.

When an CA EEM feature is activated, login requests are checked with the CA EEM server. A login request is granted only if the following occurs:

- The specified user ID matches a CA SDM contact record.
- The user ID matches a user profile in CA EEM.
- The user ID and password combination is successfully validated by CA EEM.

Note: For more information about using CA EEM for authentication, see the *Implementation Guide*. For more information about access type setup, see the *Online Help*.

Transport Layer Security

You can configure CA SDM to use Transport Layer Security (TLS) during LDAP processing. TLS, a secure communications protocol, is the successor of Secure Socket Layer (SSL v3) security. You install the ldap_enable_tls option to enable TLS.

Important! If this feature is enabled, all communications between CA SDM and the LDAP server are encrypted. If this feature is *not* enabled, all data communications (including the administrative login and password used to access the LDAP server) are sent in clear text.

Note: For information about configuring TLS, refer to your LDAP server and operating system documentation. For information about using the `ldap_enable_tls` option, see the *Online Help*.

Attribute Mapping

CA SDM contact record attribute values are synchronized with LDAP user attribute values based on the attribute mapping definitions in the `$NX_ROOT/bopcfg/majic/ldap.maj` file.

The following excerpt from `ldap.maj` illustrates mapping. Attribute names in the left column (`id`) are the CA SDM contact attribute names. The center column (`distinguishedName`) contains the corresponding LDAP attribute names.

```
id      distinguishedName      STRING 512;
last_name  sn,pzLastName      STRING ;
first_name givenName,pzFirstName  STRING ;
middle_name initials,pzMiddleName  STRING ;
userid     uid,sAMAccountName,pzUserName  STRING ;
phone_number telephoneNumber,pzWorkPhoneNumber STRING ;
```

If an SREL (a single relationship or foreign key in another database table) exists on CA SDM, the contact attribute value is synchronized with the corresponding LDAP value. If the SREL does not exist, it is not created automatically during LDAP synchronization processing.

Note: By default, attribute mapping is configured for the Microsoft Active Directory LDAP schema. If necessary, you can change the mapping by using a mod file.

How to Modify Attribute Mapping

You can change the default attribute mapping.

To change the default attribute mapping, do the following steps:

1. Navigate to `$NX_ROOT/site/mods/majic` and open the mod file.
2. Use MODIFY statements in the mod file as follows.
 - MODIFY statements must always be stated first in the file.
 - Following the MODIFY statements, any additional fields that are not in the `ldap.maj` file should be stated using the syntax shown in the following example.
 - If you define a field that contains a hyphen character in the attribute name, you must enclose the name in single quotes; otherwise, when you build the mod file, the attribute fails with a syntax error. For example, the following attribute name must be enclosed in single quotes:

```
c_nx_string1 'swsd-secret-question' STRING ;
```

3. Save and close the mod file.
4. Restart the CA SDM service.

Important! The web engine does not start if there is a discrepancy in the syntax or case.

Your changes take effect.

Example: Use MODIFY Statements

The following example shows how to modify two existing fields and add a new field.

```
//  
// Map CA SDM userid attribute to ADAM Userid  
//  
MODIFY ldap userid cn ;  
MODIFY ldap middle_name middleName ;  
OBJECT ldap LDAP {  
  ATTRIBUTES LDAP_Entry{  
    contact_num employeeNumber STRING ;  
  };  
};
```

Troubleshooting

The primary consideration when troubleshooting communications with an LDAP server is that seldom are any two LDAP implementations identical. CA SDM utilities can verify that LDAP integration is working correctly.

Note: CA SDM is preconfigured for integration with Microsoft Active Directory, eTrust, and iPlanet only. Integrating with other LDAP servers often requires changes and accommodations on both sides.

Show Status of Daemons or Processes

The `ldap_virtldb` process manages interactions between CA SDM and the LDAP virtual database.

To show the status of all CA SDM daemons (UNIX) or processes

1. Execute `pdm_status` at the command line with no parameters:

```
pdm_status
```

The `pdm_status` command shows the status of all CA SDM daemons (UNIX) or processes (Windows) on the system where the command is executed, for example:

DAEMON		STATUS	HOST	PID	SLUMP	CONNECT TIME
Agent antfarm		Running	antfarm	455	Tue Feb 17 17:55:12	
Ddict_rd	(ddictrd)	Completed	antfarm			
Data Dictionary	(ddictbuild)	Completed	antfarm			
...						
User Validation	(boplgm)	Running	antfarm	456	Tue Feb 17 17:55:21	

2. Examine the command output for the status of the `ldap_virtldb` process.

sostat Command

Run the following command without parameters to verify that bopLDAP is connected:

```
sostat
```

Examine the command output to see the status of bopLDAP.

NX.env File

Review the `$NX_ROOT/NX.env` file to verify that the basic LDAP options are correctly installed.

Depending on which LDAP options are installed, the NX.env file should include lines similar to the following:

```
@NX_LDAP_DN=quser
@NX_LDAP_ENABLE=Yes
@NX_LDAP_ENABLE_AUTO=Yes
@NX_LDAP_HOST=myserver
@NX_LDAP_PORT=389
@NX_LDAP_PWD=OBUNQXo7CmgbThZlCiMKIwJlA3UXdVNAOjUpHjstfDt2LBIDPgwtWA==
@NX_LDAP_SEARCH_BASE=dc=mycontroller, dc=xyz, dc=com
@NX_LDAP_SERVICE_TYPE=Active Directory
@NX_LDAP_SYNC_ON_NULL=Yes
@NX_LDAP_USER_OBJECT_CLASS=person
```

Important! Because Sun Java System Directory Server and Novell eDirectory servers do not support paged searching, LDAP search, import, and sync are limited to the value of NX_LDAP_MAX_FETCH records per invocation. The default value is 100. If you are using either of those LDAP servers, you may want to add NX_LDAP_MAX_FETCH to your NX.env file to specify the maximum number of LDAP records. You can set NX_LDAP_MAX_FETCH to any value less than the value of LDAP_SIZELIMIT_EXCEEDED or LDAP_ADMINLIMIT_EXCEEDED on your LDAP server.

More information:

[Configure LDAP Options](#) (see page 184)

pdm_ldap_test

Use the pdm_ldap_test command-line utility to test the connection to an LDAP directory, ensure that the search options are correctly configured, and test the TLS configuration.

By default, pdm_ldap_test uses the parameter settings that are entered in the \$NX_ROOT/NX.env file when you install, edit, or uninstall LDAP options. To override the defaults, you can specify parameters at the pdm_ldap_test command line.

To see the available parameters for this command, enter the following command:

```
pdm_ldap_test -h
```

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

Verify Connection to LDAP Server

To verify the connection to the LDAP server, run `pdm_ldap_test` without parameters:

```
pdm_ldap_test
```

Successful Connection to LDAP Server

If the connection is successful, you receive output similar to the following:

```
Starting pdm_ldap_test...
LDAP service type=active directory
Service Desk platform=windows
Using search base=DC=mycontroller,DC=xyz,DC=com
Using filter=(&(objectCategory=person))
ldap_init(myserver.mycontroller.xyz.com,389): (Success)
ldap_bind_s(Administrator) (Success)
LDAP API Verion 3
```

Failed Connection: Server Down or Incorrect Name or Port

If the connection fails because the server is down or you specified an incorrect LDAP server name or port, you receive output similar to the following:

```
Starting pdm_ldap_test...
LDAP service type=active directory
Service Desk platform=windows
Using search base=DC=mycontroller,DC=xyz,DC=com
Using filter=(&(objectCategory=person))
ldap_init(junk,389): (Success)
ldap_bind_s(Administrator) (Server Down)
```

Failed Connection: Invalid LDAP_DN or LDAP_PWD

If the connection fails because you specified an incorrect LDAP_DN or LDAP_PWD, you receive output similar to the following:

```
Starting pdm_ldap_test...
LDAP service type=active directory
Service Desk platform=windows
Using search base=DC=mycontroller,DC=xyz,DC=com
Using filter=(&(objectCategory=person))
ldap_init(myserver.mycontroller.xyz.com,389): (Success)
ldap_bind_s(junk) (No Such Object or Invalid Credentials)
```

View Search Parameters

To verify that the search parameters are correctly configured, run `pdm_ldap_test` without parameters:

```
pdm_ldap_test
```

Successful Search

When your search is successful, you see output similar to the following:

```
DN: CN=John A. Smith,CN=Users,DC=COMPUTERTEST
    c(2)(0): US
    displayName(14)(0): John A. Smith
    mail(14)(0): account02@mycompany.com
    givenName(4)(0): John
    initials(1)(0): a
    distinguishedName(38)(0): CN=John a.

Smith,CN=Users,DC=COMPUTERTEST
    objectGUID(3)(0): 314738
    pager(12)(0): ###-111-1111
    postalCode(5)(0): 11111
    SAMAccountName(7)(0): account02
    sn(6)(0): Smith
    telephoneNumber(12)(0): ###-342-6265
    userPrincipalName(16)(0): account02@COMPUTERTEST

DN: CN=Mike Johnson,CN=Users,DC=COMPUTERTEST
    displayName(10)(0): Mike Johnson
    givenName(4)(0): Mike
    distinguishedName(34)(0): CN=Mike

Johnson,CN=Users,DC=COMPUTERTEST
    objectGUID(12)(0): 312328
    SAMAccountName(7)(0): account03
    sn(5)(0): Johnson
    userPrincipalName(16)(0): account03@COMPUTERTEST
```

Failed Search: Invalid SEARCH_BASE

If the search fails because you have an invalid SEARCH_BASE, you receive output similar to the following:

```
Starting pdm_ldap_test...
LDAP service type=edirectory
Service Desk platform=windows
Using search base=o=SmartLabsx
Using filter=(&(objectClass=InetOrgPerson)
ldap_init(155.35.173.110,15389): (Success)

ldap_bind_s() (Success)
LDAP API Verion 3
ldap_search_st() (No Such Object or Referral)
```


Failed Search: SIZELIMIT_EXCEEDED, TIMEOUT

The search can fail with a SIZELIMIT_EXCEEDED or TIMEOUT message if you specify a filter that does not sufficiently narrow the search. Most LDAP Servers limit the size of the result set returned from a search request. If you exceed this limit, you receive a SIZELIMIT_EXCEEDED message. If your search request takes longer than the default timeout of 20 seconds, the LDAP server times out your request and you receive a TIMEOUT error message similar to the following:

```
Starting pdm_ldap_test...
LDAP service type=edirectory
Service Desk platform=windows
Using search base=o=SmartLabsx
Using filter=(&(objectClass=inetOrgPerson)
ldap_init(155.35.173.110,15389): (Success)
ldap_bind_s() (Success)
LDAP API Verion 3
ldap_search_st() (TIMEOUT or SIZELIMIT_EXCEEDED)
```

Failed Search: 0 Records Returned

The search may fail because your default filter is not correct. If pdm_ldap_test returns zero (0) records, always check the Using filter line, which is the base filter generated by the LDAP_FILTER_PREFIX and LDAP_FILTER_SUFFIX, or the LDAP_OBJECT_CLASS options:

```
Starting pdm_ldap_test...
LDAP service type=edirectory
Service Desk platform=windows
Using search base=o=SmartLabs
Using filter=(&(objectClass=inetOrgPerson)
ldap_init(155.35.173.110,15389): (Success)
ldap_bind_s() (Success)
LDAP API Verion 3
ldap_search_st() 0 records
```

Narrow Your Search

Use the -f parameter with the pdm_ldap_test command to specify a filter to be added to the base filter for narrowing the search criteria. You must use appropriate LDAP syntax and LDAP schema attribute names in your filter. Surround your filter with double quotation marks and use parenthesis to clarify the order of operator precedence.

For example, use the following command to search for all records where sn=Account_10001:

```
pdm_ldap_test -f "(sn=Account_10001)"
```

The `pdm_ldap_test` utility supports the following equality operators:

Equality Operator	Description
=	equal to
<=	less than or equal to
>=	greater than or equal to
~=	like

The `pdm_ldap_test` utility supports the following Boolean operators:

Boolean Operator	Description
&	AND
	OR
!	NOT

The AND and OR operators affect each set of parenthesis () in the search filter. The NOT only affects the first set of parenthesis. Always place these operators *before* the search filters to be operated on, rather than between them. They can be applied to any number of filters, as shown in the following examples:

```
"(&(sn=Brown)(initials=A))"
```

```
"(|(sn=Brown)(sn=Smith))"
```

```
"(!sn=Brown)"
```

Determine which Attribute Names have Values

Use the `-a "*"` parameter and the `-f` parameter with the `pdm_ldap_test` command to determine which attributes are defined for LDAP User or Group records. This test is useful for seeing if there are LDAP attributes that you want to map to Contact attributes, and to verify that a particular attribute has a value and should be available when creating or updating Contact records.

The following example shows output from an iPlanet directory:

```
pdm_ldap_test -a "" -f sn=Account_1000001
```

2 LDAP records found..

```
DN: cn=Account_1000001,ou=200K_Plus,o=SmartLabs
    sn(15)(0): Account_1000001
    objectClass(13)(0): inetOrgPerson
    objectClass(20)(1): organizationalPerson
    objectClass(6)(2): Person
    objectClass(18)(3): ndsLoginProperties
    objectClass(3)(4): Top
```

```
DN: cn=Account_1000001,ou=2_Plus,o=SmartLabs
    mail(28)(0): ThisIsTheMailingAddressField
    uid(13)(0): Login_1000001
    givenName(17)(0): GivenNameOfPerson
    sn(15)(0): Account_1000001
    objectClass(13)(0): inetOrgPerson
    objectClass(20)(1): organizationalPerson
    objectClass(6)(2): Person
    objectClass(18)(3): ndsLoginProperties
    objectClass(3)(4): Top
```

The following example shows output from Active Directory:

```
Ldap_test -a "" -f (&(sn=Brown)(initials=A))"
```

1 LDAP records found..

```
DN: CN=John A. Smith,CN=Users,DC=mycontroller,DC=xyz,DC=com
    objectClass(3)(0): top
    objectClass(6)(1): person
    objectClass(20)(2): organizationalPerson
    objectClass(4)(3): user
    cn(16)(0): John A. Smith
    sn(5)(0): Brown
    givenName(7)(0): John
    initials(1)(0): A
    distinguishedName(55)(0): CN=John A. Smith,CN=Users,DC=mycontroller,DC=xyz,DC=com
    displayName(16)(0): John A. Smith
    memberOf(52)(0): CN=Domain Admins,CN=Users,DC=mycontroller,DC=xyz,DC=com
    sAMAccountName(7)(0): smijo04
    userPrincipalName(25)(0): smijo04@mydomain.xyz.com
    objectCategory(63)(0): CN=Person,CN=Schema,CN=Configuration,DC=mycontroller,DC=xyz,DC=com
```

Turn on LDAP Tracing

Use the `pdm_logstat` utility to turn on trace logging to monitor LDAP use within CA SDM.

The `pdm_logstat` command has the following syntax:

```
pdm_logstat -f ldap_virtldb.c 1000
```

The following stdlog messages help you understand the status of the connection process.

Determine if ldap_virtldb Process Has Started

The first line to look for when analyzing stdlogs for LDAP messages is the startup of the `ldap_virtldb` process. The CA SDM's LDAP awareness begins only when this process starts.

Note: Even if LDAP integration options are not installed or set up, this process still runs.

```
06/03 17:00:18.27 cpasd1 bopLDAP 1964 SIGNIFICANT ldap_virtldb.c
```

```
680 STARTUP of LDAP_virtldb
```

Determine if All Required Options are Installed

If any of the required LDAP options have not been defined, the stdlog shows those that are missing, as illustrated by the following example:

```
06/03 17:00:18.72 cpasd1 bopLDAP 1964 SEVERE_ERROR ldap_virtldb.c 1023 LDAP Server port id missing
```

```
06/03 17:00:18.78 cpasd1 bopLDAP 1964 SEVERE_ERROR ldap_virtldb.c 1023 LDAP Server distinguished  
name missing
```

```
06/03 17:00:18.78 cpasd1 bopLDAP 1964 SEVERE_ERROR ldap_virtldb.c 1023 LDAP Server distinguished  
name password missing
```

Determine if the LDAP Connection is Successful

You can identify whether the LDAP connection is successful by looking at the entries in the stdlog. The entries should indicate that a connection has been successfully established with the LDAP server, as illustrated by the following example:

```
06/05 12:35:10.41 cpasd1 bopLDAP 1912 SIGNIFICANT ldap_virtldb.c 958 LDAP_SRVR connecting to  
host(Francisco.us.danconia.net) port(389)
```

```
06/05 12:35:11.01 frisco bopLDAP 1912 SIGNIFICANT ldap_virtldb.c 1002 LDAP_SRVR binding with  
username(simon)
```

Determine if the LDAP Connection is not Available

If a connection cannot be made to the LDAP server for any reason, *LDAP Entries*, *Merge LDAP*, or any other LDAP functionality, becomes disconnected and returns no results. In such instances, the stdlog shows messages similar to the following examples when accessing those operations:

```
06/03 17:00:32.25 cpasd1 bopLDAP 1964 SIGNIFICANT ldap_virtldb.c 219 LDAP server not available;
'register_producer' not processed
```

```
06/05 10:52:57.63 cpasd1 bopLDAP 1896 SIGNIFICANT ldap_virtldb.c 219 LDAP server not available;
'select_full' not processed
```

```
06/05 10:52:57.66 cpasd1 web:local 1868 ERROR sel_data_cache. 611 Error in ldap Select_Cache method
got_initial_count: LDAP server not available; 'select_full' not processed
```

```
06/05 10:52:57.66 cpasd1 bopLDAP 1896 SIGNIFICANT ldap_virtldb.c 219 LDAP server not available;
'select_cancel' not processed
```

Determine Actual Filter Used

CA SDM fetches records from the LDAP Directory according to the search base and the filter defined in Options Manager, and the search criteria entered by the user. Look for the following message to determine the actual filter generated in the search request:

```
06/24 14:18:28.32 mcxxx04- bopLDAP 3844 TRACE ldap_virtldb.c 853 Starting select full:
base=DC=kirklandsd,DC=ca,DC=com; filter=(&(objectCategory=person)(|(sn=Jones)(pzLastName=Jones)));
attributes=(uid,sAMAccountName,pzUserName,distinguishedName)
```

Determine Attributes Fetched

CA SDM fetches records from the LDAP Directory according to the search base and the filter defined in Options Manager. Look for the following message to determine if the SEARCH_BASE and attribute mapping as defined in ldap.maj and ldap_group.maj are correct:

```
06/24 14:18:28.39 mcxxx04- bopLDAP 3844 TRACE ldap_virtldb.c 766 Starting select short:
base=CN=John D. Jones,CN=Users,DC=kirklandsd,DC=ca,DC=com; filter=(&(objectCategory=person));
attributes=(modifyTimestamp,sn,pzLastName,givenName,pzFirstName,initials,pzMiddleName,uid,sAMAccountN
ame,pzUserName,telephoneNumber,pzWorkPhoneNumber,mobile,pzMobilePhoneNumber,department,pzDepa
rtment,facsimileTelephoneNumber,pzFaxPhoneNumber,pager,mail,pzEmailAddress,streetAddress,pzAddress,l,p
zCity,st,pzState,postalCode,pzPostalCode,c,pzCountry,o,memberOf)
```

Determine Which LDAP Data is Available and Not Available

Assuming that CA SDM was able to map to the LDAP object's ID attribute successfully, the attributes as defined in \$NX_ROOT/bopcfg/majic/ldap.maj are retrieved for each entry, or a message is logged indicating that an attribute has not been defined. A sample of these messages is as follows:

```
06/24 14:18:28.41 mclda04- bopLDAP    3844 TRACE  ldap_virtdb.c  1396 Value not available for  
'modifyTimestamp'
```

```
06/24 14:18:28.41 mclda04- bopLDAP    3844 TRACE  ldap_virtdb.c  1396 Value not available for  
'telephoneNumber,pzWorkPhoneNumber'
```

Chapter 6: Managing Roles

Note: For information about other aspects of the web interface, see [Configuring the Web Interface](#) (see page 323).

This section contains the following topics:

[Roles](#) (see page 207)

[Predefined Roles](#) (see page 207)

[Role-Based Security](#) (see page 210)

[Role-Based Navigation](#) (see page 219)

[How to Implement a Custom Role](#) (see page 228)

[How to Implement a Custom Menu Tree](#) (see page 229)

[Create a Role Record](#) (see page 231)

[Create a Tab Record](#) (see page 232)

[Create a Menu Bar Record](#) (see page 233)

[Create a Web Form Record](#) (see page 234)

[Copy a Menu Tree](#) (see page 235)

[Create and Customize a Menu Tree](#) (see page 236)

[Create and Publish a Help Set](#) (see page 237)

[Switch Roles](#) (see page 239)

Roles

Roles are the primary records that control CA SDM security and user interface navigation. Each role defines a focused view of the system by exposing only the functionality necessary for users to perform the tasks typically assigned to the role they perform within their business organization.

A user's default role determines the system view that is presented upon login. Users with multiple role assignments can switch from one role to another to see different views of the system without having to log out and log back in again.

Predefined Roles

You can use the predefined roles in their default configuration, modify them to meet your business requirements, or create new roles.

The following table describes the predefined roles installed with CA SDM. These roles are designed to align with ITIL v3 best practices, and thereby reduce the amount of site-specific customization required to bring your IT organization into ITIL compliance.

CA SDM only supports ITIL, and the CA SDM documentation is ITIL-oriented. For more information, see [ITIL Configuration](#) (see page 32).

Role Type	Role Name	Description
End Users	Configuration Viewer	Performs basic CI viewing and research tasks from inside your organization.
	Customer	Performs basic self-service tasks from outside your organization.
	Employee	Performs basic self-service tasks from inside your organization.
Analysts	Configuration Analyst	Performs tasks within the configuration item life cycle process and second-line CMDB support within your organization.
	Customer Service Representative	Supports users external to your organization, most often customers.
	Knowledge Analyst	Performs tasks within the knowledge management life cycle process.
	Level 1 Analyst	Provides first-line support within your organization.
	Level 2 Analyst	Provides second-line support within your organization, which requires more advanced subject matter expertise.
	Support Automation Analyst	Provides first-line support within your live assistance environment.
	Vendor Analyst	Supports a limited segment of your IT environment from outside your organization, such as vendor-specific hardware.
Managers	Change Manager	Manages the change order process, but typically not the analysts who work on change order tickets.
	Customer Service Manager	Manages Customer Service Representatives and the external support process.
	Incident Manager	Manages the incident process, but typically not the analysts who work on incident tickets.


Role Type	Role Name	Description
	Knowledge Manager	Supervises Knowledge Analysts, knowledge document reassignments and escalations, and day-to-day knowledge administration.
	Problem Manager	Manages the problem process, but typically not the analysts who work on problem tickets.
	Service Desk Manager	Handles escalations and supervises Level 1 Analysts. Also may manage overall service desk operations.
Administrators	Administrator	Performs administrative tasks throughout your CA SDM and Knowledge Management implementation. This role typically installs, configures, and integrates the products.
	Configuration Administrator	Performs administrative tasks related to your CA CMDB implementation. This role typically administers CMDB and configuration item infrastructure and data structures.
	Knowledge Management Administrator	Configures and monitors knowledge management settings.
	Service Desk Administrator	Performs administrative tasks on data and processes, such as creating and updating categories, contacts, service types, root causes, and so on.
	Support Automation Administrator	Performs administrative tasks related to your Support Automation environment, such as configuring queues and analyst tool permissions.
	System Administrator	Performs administrative tasks related to your CA SDM implementation, configuration and adaptation, such as setting options, configuring integrations and modifying web forms.
	Tenant Administrator	Performs multi-tenancy administrative tasks specific to a particular tenant or supporting organization.

Role-Based Security

Access types and roles are the primary components you use to control CA SDM security.

The following diagram shows an overview of how roles interrelate with other system objects to provide role-based security.

Note: For more information about other aspects security, see [Security](#) (see page 153).



How Access Types Work

Each access type for a user controls the following aspects of system behavior:

- How CA SDM performs web authentication when the user logs in
- The access level for the user
- Whether the user can modify web forms or the database schema using Web Screen Painter
- Which roles are available to the user

You can associate an access type with a contact by selecting the access type while creating or updating the contact record.

The following table lists the predefined access types, identifies their linked roles, and gives a brief description.

Access Type	Linked Roles	Description
Administration	■ Administrator (default) ■ Configuration Administrator ■ Employee ■ Level 2 Analyst ■ Service Desk Administrator ■ System Administrator ■ Tenant Administrator	Provides the highest level of security access to all key administration roles. Used during implementation and ongoing administration. Note: The Administration access type is preconfigured to allow administrators to switch to any of the linked roles. For example, to see a different view of the system, administrators can switch to the Employee role without having to log out and log in again.
Customer	■ Customer	Provides highly restricted access to <i>external</i> customers who use the self-service view.
Employee	■ Employee	Provides highly restricted access to <i>internal</i> employees who use the self-service view. Used to create new incident and update incident pages.
IT Staff	■ Configuration Analyst ■ Employee ■ Level 2 Analyst (default) ■ Knowledge Analyst ■ Knowledge Management Administrator ■ Knowledge Manager	Provides analyst-oriented access to users who work within your IT organization but are not actual members of the support team. This access is designed specifically for users who need access to Knowledge Management.

Access Type	Linked Roles	Description
Knowledge Management	■ Configuration Administrator	Provides administrative access tailored to users who administer Knowledge Management features.
	■ Configuration Analyst	
	■ Configuration Viewer	
	■ Employee	
	■ Knowledge Analyst	
	■ Knowledge Management Administrator (default)	
	■ Knowledge Manager	
	■ Level 2 Analyst	
Process Management	■ Change Manager	Provides access tailored to users who perform key process management roles.
	■ Configuration Analyst	
	■ Employee	
	■ Incident Manager (default)	
	■ Level 2 Analyst	
	■ Problem Manager	
	■ Service Desk Manager	
Service Desk Management	■ Customer Service Manager	Provides access tailored to users who manage IT support or external customer support functions (typically front-line support supervisors).
	■ Configuration Analyst	
	■ Employee	
	■ Level 1 Analyst	
	■ Level 2 Analyst	
	■ Service Desk Manager (default)	

Access Type	Linked Roles	Description
Service Desk Staff	■ Configuration Analyst ■ Configuration Viewer ■ Customer Service Representative ■ Employee ■ Level 1 Analyst (default) ■ Level 2 Analyst	Provides access tailored to users who perform support tasks. Access is focused on those users that perform frontline support.
Support Automation Admin	■ Support Automation Administrator	Provides access to users that perform Support Automation administration.
Support Automation Analyst	■ Support Automation Analyst	Provides access to users that provide live assistance to end users.
Vendor Staff	■ Vendor Analyst	Provides highly restricted access to external vendors, who work only on items directly related to their product (for example, a particular brand of hardware).

Role Records

You can assign roles to an access type, or directly to a user contact record. If a role assignment conflict occurs, the contact role assignments take precedence.

Each role record must be configured with the following components:

- One form group
- One user interface type
- Function access settings
- One or more tabs
- One help set

The following optional components can also contribute to each role definition:

- Menu trees
- Scoreboards
- Menu bars

- Toolbars
- One data partition
- Knowledge Management access
- Support Automation access levels
- Report web forms
- Go resources

Functional Access Areas

Functional access areas define the role level access to ticket records and other system components. The `usp_functional_access_type` table defines the area and the `usp_functional_access_level` tables describe user access.

The following table shows the default functional access areas:

Name	Code Name	New
Administration	admin	No
Incident/Problem/Request	call_mgr	No
Change Order	change_mgr	No
Inventory	inventory	No
Issue	issue_mgr	No
Knowledge Document	kd	No
Notification	notify	No
Reference	reference	No
Security	security	No
Announcement	announcement	Yes
Incident/Problem/Request Reference	call_mgr_reference	Yes
Incident/Problem/Request Template	call_mgr_template	Yes
Change Order Template	change_mgr_template	Yes
Change Order Reference	change_reference	Yes
Configuration Item	ci	Yes
Configuration Item Common Readonly	ci_common_ro	Yes
Configuration Item Reference	ci_reference	Yes
Contact	contact	Yes

Group	group	Yes
Issue Template	issue_mgr_template	Yes
Issue Reference	issue_reference	Yes
Location	location	Yes
Multisite Administration	multisite_admin	Yes
Multisite Reference	multisite_reference	Yes
Notification Reference	notification_reference	Yes
Organization	organization	Yes
Prioritization	prioritization	Yes
Service Level	service_level	Yes
Site	site	Yes
Stored Query	stored_queries	Yes
Support Automation	sa	Yes
Survey	survey	Yes
Tenant Admin	tenant_admin	Yes
Timezone	timezone	Yes
Workflow Reference	workflow_reference	Yes
Workshift	workshifts	Yes

How to Add a Functional Access Area

When you add a functional access area, the existing roles automatically have Modify access. You can review and change the access levels to grant the appropriate authority.

To add a functional access area, do the following:

1. On the Administration tab, select Security and Role Management, Functional Access.
2. Click Create New.

The Create New Functional Access List page appears.

3. Complete the functional access area fields as appropriate.
4. Click Save.

The Functional Access Detail page appears.

5. Apply access levels to [one](#) (see page 217) or [more](#) (see page 216) roles.

Note: For detailed information about functional access areas, see the *Online Help*.

Apply Access Levels to Many Roles

For a functional access area, you can set the access levels for every role to save time. Instead of using Role Management, you update the functional access area with the appropriate role access levels.

To apply access levels to many roles, do the following:

1. On the Administration tab, select Security and Role Management, Functional Access.
2. Click the name of the function area.
3. On the Roles Tab, click Edit in List.
4. Review and update each role as appropriate. The following access levels are available:

None

Denies the role access to the function object.

View

Grants read-only capability to the function object.

Modify

Grants read/write access to the function object.

5. Continue selecting roles and choosing access levels.
6. (Optional) Click Change All.
7. Click Save.

The changes for the roles apply immediately.

Note: For information about roles and functional access levels, see the *Online Help*.

Example: Change the Access Levels for the Announcement

This example shows you can role access levels for the Announcement functional access area.

1. On the Administration tab, select Security and Role Management, Functional Access.
The Functional Access Detail page appears.
2. Click Announcement.
3. On the Roles Tab, click Edit in List.
4. Select Level 2 Analyst.
5. Select View from the Access Level.

The Access Level for the Level 2 Analyst role highlights with the value of View.

6. Select Configuration Analyst.
7. Select Modify from the Access Level.
The Access Level for the Configuration Analyst highlights with the value of Modify.
8. Click Save.
A message confirms the change.
9. Log in as a Level 2 Analyst role.
10. Select View Announcements.
The Announcements page appears.
11. Click an announcement.
A message reminds you that as a Level 2 Analyst, you can only view announcements.
12. Set the Role to Configuration Analyst.
13. Select View Announcements.
The Announcements page appears.
14. Click an announcement.
The Update Announcement page appears. As a Configuration Analyst, you can modify announcements.

Apply an Access Level to a Role

You can use Role Management to change the way users access the user interface. When you change the access levels for a role, the user interface shows only objects, pages, and menu items based on the access level. For example, if a role can no longer create contacts, the File menu omits New Contacts.

To apply an access level to a role, do the following:

1. On the Administration tab, select the Security and Role Management, Role Management, Role List.
2. On the Role List, right-click the role name and select Edit from the short-cut menu.
3. Click Edit in List on the Function Access tab.
4. Click a function name.
The row highlights.

5. Update the functional access areas with the following access levels as appropriate:

None

Denies the role access to the function object.

View

Grants read-only capability to the function object.

Modify

Grants read/write access to the function object.

6. Click Save.
A message confirms the change. The role can immediately use the functional access area at the specified access level.
7. Verify the access level, by logging in as the role and checking menus, page options, and buttons.

Example: Grant the Level 2 Analyst Role Modify Access to Tickets

This example shows how the user interface changes when you grant a Level 2 Analyst access to modify tickets.

1. On the Administration tab, select the Security and Role Management, Role Management, Role List.

The Role List appears.

2. Right-click Level 2 Analyst and select Edit from the short-cut menu.
3. Click Edit in List on the Function Access tab.
4. Select Incident/Problem/Request Reference.
5. Select Modify on the Access Level.

The access level updates to Modify.

6. Click Save and log out.
A message confirms the change.
7. Log in as a Level 2 Analyst role.
8. Select Search, Incidents.
9. Click Search and open an incident.

The Incident Detail page includes an Edit in List button. As a Level 2 Analyst, you can modify the ticket.

Data Partitions


Data partitions are subsets of the CA SDM database that enable you to control access at the record level. You can associate a data partition to a role to control access to tickets and other records accessible through the web interface.

For information about working with data partitions, see [Data Partition Associations](#) (see page 167).

Role-Based Navigation

Each user's view of the web interface is defined by a role. Users with multiple role assignments can switch to multiple web interface views.

The following diagram shows how roles interrelate with other objects to produce a role-based presentation of the user interface.



Tabs

A tab is a graphical display entity that links to a role in order to present features to the users of that role. When a user logs in to the system, the main window displays the tabs assigned to the user default role.

Tabs define major subdivisions in the web interface main window. Each tab is configured to expose an appropriate set of user interface features to the role or roles that use it.

All roles must have at least one tab. You can associate one or more tabs with a role. Each tab has a sequence number that controls its display order in the main window. If only one tab is associated with a role, the tab starting page is displayed and not the tab.

You can configure tabs to include the following kinds of display features:

- A starting page (default web form) that is displayed when a user selects the tab. The starting page is a required element of all tabs. You can assign only one starting page to each tab.

Note: You can configure a starting page to display graphical reports from Business Object, which enables users to generate reports at run time. For more information, see [CA Business Intelligence Reporting](#) (see page 691).

- A menu bar, which presents drop-down lists of commands, such as File, View, and Search commands. The menu bar is optional. You can assign only one menu bar to each tab.
- A toolbar, which presents tool buttons for easy access to frequently used menu commands. The toolbar is optional.

CA SDM provides several predefined tabs. You can assign the predefined tabs to a role, modify the predefined tabs, and create custom tabs.

Important! Do not assign more tabs than your browser window can display; doing so causes tabs with higher sequence numbers to extend beyond the window viewable display area and make them inaccessible to the user.

Important! Include only tabs that contain forms that are included in the form group assigned to the role you are creating or editing. For example, do not assign the Customer tab or the Employee tab to the Administrator role; doing so causes an error when users attempt to access that tab. The role form group is specified in the Customization Form Group field on the Role Detail page, and is also displayed in the Form Group column on the Role List page. For a list of the web forms in each form group, see [Form Groups](#) (see page 1009).

Predefined Tabs

The following table shows the predefined tabs assigned to each role. The tabs are listed in sequence number order, indicating their left-to-right position in the window.

Note: In many cases, there are multiple versions of tabs with the same display name. For example, the Service Desk tab for the Administrator role provides full access to CA SDM functionality, while the Service Desk tab for the Change Manager role is more focused on change orders.

Role	Tabs
Administrator	■ Service Desk tab with full menu and scoreboard ■ Knowledge tab ■ Administration tab with full menu tree ■ Reports tab - Administrator ■ Change Calendar tab ■ CA CMDB tab ■ Support Automation tab
Change Manager	■ Reports tab - Change Manager ■ Service Desk tab - Change Manager ■ Change Calendar tab
Configuration Administrator	■ CA CMDB tab ■ Administration tab - Configuration Administrator
Configuration Analyst	■ CA CMDB Scoreboard - Configuration Analyst
Configuration Viewer	■ CA CMDB Scoreboard - Configuration Viewer
Customer	■ Customer tab
Customer Service Manager	■ Service Desk tab - Customer Service Manager ■ Knowledge tab ■ Reports tab - Customer Service Manager
Customer Service Representative	■ Service Desk tab - Customer Service Rep ■ Quick Profile tab ■ Knowledge tab
Employee	■ Employee tab

Role	Tabs
Incident Manager	■ Reports tab - Incident Manager ■ Service Desk tab - Incident Manager ■ Knowledge tab
Knowledge Analyst	■ Service Desk tab - Knowledge Analyst ■ Knowledge tab ■ Knowledge Management Schedule ■ Knowledge Report Card tab ■ Reports tab - Knowledge Analyst
Knowledge Manager	■ Service Desk tab - Knowledge Manager ■ Knowledge tab ■ Knowledge Management Schedule ■ Administration tab - Knowledge Manager ■ Knowledge Report Card tab ■ Reports tab - Knowledge Manager
Knowledge Management Administrator	■ Administration tab - Knowledge Administrator
Level 1 Analyst	■ Service Desk tab - Level 1 Analyst ■ Quick Profile tab ■ Knowledge tab
Level 2 Analyst	■ Service Desk tab - Level 2 Analyst ■ Knowledge tab ■ Change Calendar tab
Problem Manager	■ Reports tab - Problem Manager ■ Service Desk tab - Problem Manager ■ Knowledge tab
Service Desk Administrator	■ Administration tab - Service Desk Administrator
Service Desk Manager	■ Reports tab - Service Desk Manager ■ Service Desk tab - Service Desk Manager ■ Knowledge tab ■ Change Calendar tab

Role	Tabs
Support Automation Administrator	■ Service Desk tab - Level 1 Analyst ■ Support Automation tab ■ Support Automation Admin ■ Quick Profile tab ■ Knowledge tab
Support Automation Analyst	■ Service Desk tab - Level 1 Analyst ■ Support Automation tab ■ Quick Profile tab ■ Knowledge tab
System Administrator	■ Administration tab - System Administrator
Tenant Administrator	■ Administration tab - Tenant Administrator
Vendor Analyst	■ Service Desk tab - Vendor Analyst

Web Forms

Web forms define the pages that appear in the CA SDM web interface.

There are four web form types:

- Business Object Report URL
- HTML page
- GO resource
- Custom (for example, a URL for a third-party web page)

For information about creating custom web forms, see [Create a Web Form Record](#) (see page 234).

Form Groups

Form groups define the sets of CA SDM web interface pages that are available to a role. Each role has one form group. Users can display only the web pages that are included in the form group assigned to their role.

Each interface type has an associated form group, a set of HTML files that define the pages users with that interface type can see.

CA SDM provides the following predefined form groups:

- Analyst
- Customer
- Employee
- ITIL

You can use the predefined form groups in their default configuration, modify the predefined form groups, and create new form groups using Web Screen Painter. You can view a listing of the HTML filenames included in each predefined [form group](#) (see page 1009).

Menu Trees

Menu trees are the hierarchical listings of nodes (menu tree resources) that are displayed in the navigation pane on the left-hand side of the main web interface window.

A role can have a menu tree, which provides nodes for access to many functional areas of the system. For example, the predefined Administrator role has a menu tree that includes nodes to the System and Role Management administration features, Service Desk administration features, and many others.

For roles that include a menu tree, the menu tree provides access to a specified set of resources that provide access to functional areas of the system.

CA SDM provides predefined menu trees for the following roles:

- Administrator (admin_tree)
- CA CMDB Administrator (cmdb_admin_tree)
- Knowledge Management Administrator (kt_admin_tree)
- Knowledge Manager (kt_mgr_tree)
- Support Automation Administrator (sa_admin_tree)
- Service Desk Administrator (sd_admin_tree)
- System Administrator (sys_admin_tree)
- Tenant Administrator (tn_admin_tree)

You can edit the Name, Record Status, and Description fields of the predefined menu tree records, but you cannot customize them by adding or removing their menu tree resources.

To produce a custom menu tree, you can create new menu tree record or copy and customize one of the predefined menu trees.

Note: The non-modifiable Internal field on each menu tree record indicates whether the menu tree can be customized. A value of YES in the Internal field indicates a predefined menu tree, which cannot be customized. A value of NO indicates a site-defined menu tree, which can be customized. The Customize Menu button appears only on menu tree detail records with an Internal field value of NO.

When you attach a menu tree to a tab, it becomes available for all roles that have access to that tab.

Menu Tree Resources

Menu tree resources define the items users can access from the menu tree.

A menu tree resource consists of a name, description, and a URL fragment or HTML filename used by the web engine that controls the web page displayed.

Menu Bars

A menu bar is a user interface element that displays a horizontal list of menus in the web interface main window. Each menu contains a drop-down list of options or commands. You can define custom menu bars for any custom roles you might create.

Menu bar records specify the HTML form that controls the menu items that the menu bar can access.

Note: To define the functionality of the menu bar, you must use the Web Screen Painter application. For information about configuring the functionality of a predefined or custom menu bar, see the *Web Screen Painter Online Help*.

The following table lists the predefined menu bars and identifies the predefined tabs that use them.

Menu Bar	Associated Tabs
Administration	■ Administration tab - Configuration Administrator ■ Administration tab - Knowledge Administrator ■ Administration tab - Knowledge Manager ■ Administration tab - Service Desk Administrator ■ Administration tab - System Administrator ■ Administration tab - Tenant Administrator ■ Administration tab with full menu tree
CA CMDB	■ CA CMDB tab with full menu and scoreboard
Change Calendar	■ Change Calendar tab
Knowledge	■ Knowledge tab ■ Knowledge Management Schedule
Service Desk	■ Service Desk tab with full menu and scoreboard
Service Desk-Change Manager	■ Service Desk tab - Change Manager
Service Desk-Cust Service Mgr	■ Service Desk tab - Customer Service Manager
Service Desk-Cust Service Rep	■ Service Desk tab - Customer Service Rep
Service Desk-Incident Manager	■ Service Desk tab - Incident Manager
Service Desk-Knowledge Analyst	■ Service Desk tab - Knowledge Analyst
Service Desk-Knowledge Manager	■ Service Desk tab - Knowledge Manager
Service Desk-Level 1 Analyst	■ Service Desk tab - Level 1 Analyst
Service Desk-Level 2 Analyst	■ Service Desk tab - Level 2 Analyst
Service Desk-Problem Manager	■ Service Desk tab - Problem Manager
Service Desk-Service Desk Mgr	■ Service Desk tab - Service Desk Manager
Service Desk-Vendor Analyst	■ Service Desk tab - Vendor Analyst

Menu Bar	Associated Tabs
Support Automation-Analyst	■ Support Automation tab - Support Automation Analyst

Toolbars

Toolbars extend the functionality of menu bars by adding the capability to display one or more tool buttons to the right of the menus.

Tool buttons appear as icons on the toolbar. Clicking a tool button gives the user easy access to frequently used menu options or commands.

Note: You use the Web Screen Painter application to define the functionality of the toolbar. For information about configuring the functionality of a predefined or custom toolbar, see the *Web Screen Painter Online Help*.

Go Resources

The *Go* button provides an easy means of locating a particular record.

Go resources are a type of web form. If a role has associated Go resources, when a user logs in with that role the Go button appears in the upper-right corner of the main CA SDM window and in all popup windows. The Go button has two associated fields in the user interface:

- A drop-down list for selecting the type of record to search for (for example, Change Order)
- A text box for entering a value to identify a particular record (for example, 135 to locate Change Order 135)

By assigning Go resources to a role, you can specify the kinds of records users in that role can search for. For example, the predefined Administrator role has the following Go resources:

- Change Order
- Document by ID
- Incident
- Issue
- Knowledge
- Problem
- Request

- User by ID
- User by Name
- User by Phone

Help Sets

Help sets are the collections of online help topics available to users depending on their role assignments and current role setting. If you log in using the Administrator role, for example, you can view the online help topics included in the Administrator help set. If you switch to the Employee role, you can view the Employee help set.

Each predefined role has a corresponding predefined help set. You can create custom help sets for any custom roles you might define.

Note: For information about working with help set definitions, see the *Online Help*.

How to Implement a Custom Role

For many sites, the predefined roles are sufficient. There may be situations, however, when you want to create a custom role and tailor it to meet site-specific business needs within your organization.

The following process outlines the tasks required in implementing a new role. The example shown here describes how you might implement a role for a small group of analysts tasked with reviewing and authorizing change order tickets.

To implement a custom role, perform the tasks described in the following example:

1. Create a new role record using the following field values:

Role Name

Change Analyst

Code

chg_anal

Customization Form Group

Analyst

Preferred Document

Incident

2. Select Service Desk Analyst in the Data Partition field on the Authorization tab.
3. Select the Modify in the Change Orders field on the Function Access tab.

4. Enter the following values on the Web Interface tab:

Web User Interface Type

Analyst

Help View

Change Analyst

5. Select the following tabs:
 - Reports tab - Change Analyst
 - Service Desk tab - Change Analyst
 - Change Calendar tab
6. Select the following reports on the Report Web Forms tab:
 - Active Change Orders Aging by Priority for Status
 - Active Change Orders at Weeks End
 - Change Orders by Failed Service Type for Change Categories
7. Add the Change Order resource on the Go Resources tab.
8. Create a custom help set named Change Analyst that includes all content appropriate for the new role.

For more information, see [Create and Publish a Help Set](#) (see page 237).
9. Create the following custom tabs using features appropriate for the new role:
 - Reports tab - Change Analyst
 - Service Desk tab - Change Analyst
10. Create a custom menu tree that includes all nodes appropriate for the new role.

For more information see [How to Implement a Custom Menu Tree](#) (see page 229).

How to Implement a Custom Menu Tree

For many sites, the predefined menu trees are sufficient. There can be situations, however, when you want to customize a role by implementing a custom menu tree for it.

In most cases, it is easier to start with a copy of a predefined menu tree and then add, remove, or reorganize nodes within the hierarchy. Alternatively, you can create a menu tree and construct an all new hierarchy of nodes.

You can use either of the following methods to make custom menu tree available to a role:

- Replace the menu tree in the web form (Start Page) for the tab that shows the original admin_tree.
- Create a web form and attach the new web form with the new menu tree to a tab.

To implement a custom menu, perform the following tasks:

1. Copy one of the predefined menu trees.

Note: Make a note of value you enter in the Code field.

2. Create a web form using the following field values:

- **Type:** HTML

- **Resource:**

```
$cgi?SID=$SESSION.SID+FID=123+OP=DISPLAY_FORM+HTML=admin_main_role.html+KEEP.tree_code=menu_tree_code
```

Note: Specify the value of the code for the menu tree you created in Step 1 for menu_tree_code. The admin_main_role.html code uses the value of the KEEP.tree_code variable as its menu tree.

3. Create a tab record using the following field values:

- **Starting Page:** The web form you created in Step 2

- **Menu Bar:** Administration

Note: Administration is a generic menu bar used by many roles; it is not role-specific.

4. Assign the tab you created in Step 3 to the role you want to have access to the custom menu tree.
5. Log out of CA SDM and log back in again.

The Administration tab displays your custom menu tree.

More information:

[Create a Web Form Record](#) (see page 234)

Create a Role Record

Administrators can create customized roles to meet site-specific business requirements.

To create a role

1. Select Security and Role Management, Role Management, Role List on the Administration tab.

The Role List page appears.

2. Click Create New.

The Create New Role page appears.

3. Complete the following fields:

Role Name

Specifies the name that identifies the role wherever it appears in the user interface.

Code

Specifies the code that identifies the role to the system.

Note: After you save the record, this field value cannot be changed.

Record Status

Indicates whether the role is Active or Inactive.

Default?

Indicates whether this role is the default role.

Customization Form Group

Specifies a predefined or custom form group.

Preferred Document

Specifies the document used by this role for entering tickets into the system.

Description

Describes the role purpose. This description appears on the Role List page and can facilitate the task of assigning users to appropriate roles.

Click Save.

The role definition is saved and the Role Detail page appears.

Create a Tab Record

You can create custom tabs to appear on the web interface main page. When you link a tab record to a role record, it becomes available to users assigned to that role.

To create a tab

1. Select Security and Role Management, Role Management, Tabs on the Administration tab.

The Tab List page appears.

2. Click Create New.

The Create New Tab page appears.

3. Complete the following fields:

Tab Name

Specifies the name that identifies the tab within the administrative interface. For example, the tab name appears on the Tab List page.

Code

Specifies the code that identifies this tab to the system.

Note: After the code is defined, it cannot be changed.

Record Status

Indicates whether the tab is active or inactive.

Display Name

Specifies the name that appears on the tab graphical presentation in the user interface.

Starting Page

Specifies the initial web form that appears in the main window when a user selects this tab.

Important! The starting page and menu bar must belong to the same form group. Defining a tab with a starting page and menu bar that belong to different [form groups](#) (see page 1009) causes an error when users access the tab.

Menu Bar

Specifies the menu bar that appears in the main window when a user selects this tab.

Click Save.

The tab is created.

Create a Menu Bar Record

You can create customized menu bars to control the access to system functionality for the user-defined roles.

To create a menu bar

1. From the Administration tab, navigate to Security and Role Management, Role Management, Menu Bars.

The Menu Bar List page displays.

2. Click Create New.

The Create New Menu Bar page displays.

3. Complete the following fields:

Menu Bar Name

(Required) Specifies the name that identifies the menu bar. You can use this field to help identify the functionality available from the menu bar.

Code

(Required) Specifies the code that identifies this menu bar to the product. After the code is defined, it cannot be changed.

Record Status

Indicates whether the menu bar is Active or Inactive.

HTMPL Name

Specifies the name of the HTMPL form that contains the menu bar definition. The menu bar is designed using the Web Screen Painter.

Description

Describes the menu bar. Use this description to further identify this menu bar and the roles that use it.

Click Save.

The menu bar definition is saved and the Menu Bar Detail page appears.

Create a Web Form Record

Administrators can create customized web forms to be the starting pages for tabs, reports to display on tabs, "go button" resources, or another URL.

To create a web form

1. From the Administration tab, navigate to Security and Role Management, Role Management, Web Forms.

The Web Forms List page displays.

2. Click Create New.

The Create New Web Form page displays.

3. Complete the following fields:

Web Form Name

(Required) Specifies the name that identifies the web form.

Record Status

Indicates whether this form is active or inactive.

Code

(Required) Specifies the code that identifies the web form to the system. After the code is defined, it cannot be changed.

Note: This field specifies the `web_form_name` on the Properties tab for a multiframe form in Web Screen Painter.

Type

Specifies one of the following types of web form that you are creating:

- **HTMPL page**—Displays a web page to use as the starting page for one of the custom tabs you create.
- **Report**—Specifies a CA SDM report that displays on any tab.
- **Go Resource**—Specifies a "Go Button" resource.
- **Other**—Accesses any other external Web page through URL.

Description

Describes the web form. Use this description to further identify this web form, where it displays, and its purpose.

Resource

Specifies the code that calls the web form. This code can be command line code or a URL.

Example: Open a simple html form "menu_tab_dflt.html":

```
$cgi?SID=$SESSION.SID+FID=123+OP=DISPLAY_FORM+HTMPL=menu_tab_dflt.html
```

Click Save.

Copy a Menu Tree

You can copy an existing menu tree to use as a starting point for a customized menu tree.

To copy a menu tree

1. Select Security and Role Management, Role Management, Menu Trees on the Administration tab.

The Menu Tree List page appears.

2. Click the Menu Tree to copy.

The Menu Tree Detail page appears.

3. Click File, Copy.

The Create New Menu Tree page appears.

4. Complete the following fields:

Menu Tree Name

(Required) Specifies the name you assign that identifies this menu tree.

Code

(Required) Specifies the code that identifies the menu tree to the system. After the code is defined, it cannot be changed.

Record Status

Indicates whether this menu tree is active or inactive.

Description

Describes the menu tree. The description can be used to give additional details about the menu tree and the roles that use it.

Click Save.

The Menu Tree Detail page for the new menu tree appears.

5. Click Customize Menu.

A copy of the original menu tree appears.

6. Customize the menu tree as you want.

Note: For details about adding, removing, or editing menu tree resources, see the *Online Help*.

Create and Customize a Menu Tree

You can create and customize menu trees, based on one of the default menu trees provided.

To create and customize a menu tree

1. Select Security and Role Management, Role Management, Menu Trees on the Administration tab.

The Menu Tree List page appears.

2. Click Create New.

The Create New Menu Tree page appears.

3. Complete the following fields:

Menu Tree Name

(Required) The name you assign that identifies this menu tree.

Code

(Required) The code that identifies the menu tree to the system. Once the code is defined, it cannot be changed.

Record Status

Indicates whether this menu tree is active or inactive.

Description

A description of the menu tree. The description can be used to give additional details about the menu tree and the roles that use it.

4. Click Save.

5. Click Customize Menu.

A form appears, allowing you to customize the menu tree. At this point, the menu tree contains only a top node with the text you entered as the menu tree name.

6. Right-click the node in the menu tree and select Create New Node.

The Create New Node page appears.

7. Complete the following fields:

Node Name

Enter a name for the node. This is the name that appears in the menu tree.

Description

Enter a description for the node. This description can be used to further define the purpose of the node.

Resource

Enter the resource name directly in the field, or click the search icon to select the resource from a list. The menu tree resource determines the action to perform when the user selects the node from the menu tree.

8. Repeat steps 6 and 7 as many times as necessary to create the set of nodes you want to appear in the menu tree.

Note: For more information about adding, removing, or editing menu tree resources, see the *Online Help*.

9. Click Save.

The menu tree definition is saved and the Menu Tree Detail page appears.

Create and Publish a Help Set

You can create custom help sets for any custom roles you might define.

To create, populate, and publish a help set

1. Select Security and Role Management, Role Management, Help Sets on the Administration tab.

The Help Set List page appears.

2. Click Create New.

The Create New Help Set page appears.

3. Complete the following fields:

Help Set Name

The unique name of this help set.

Interface Type

The interface type of the help set (such as Analyst, Employee, or Customer).

Record Status

Indicates whether the help set is Active or Inactive.

File Name Prefix

The prefix you want to attach to the help files generated for this help set. Do not include spaces in the name.

Note: Assign a prefix that allows you to identify the files belonging to this help set. For example, you may want to use ANAL for an analyst's help set prefix.

Internal

This is automatically set to NO for user-defined help sets. The value in this field cannot be changed.

4. Click Save.

The Contents tab appears.

5. Click Define Content.

The Selected Help Update window opens.

6. Select the content you want to include in the help set.

Important! Some topics are required, and are included in your new help set regardless of whether you select them. For example, the CA SDM home page and other front matter topics are always included. Also, nested topics are dependent on their container topics. Container topics are included automatically if you include any of their nested topics. For example, if you select the "Use the Scoreboard" topic, the container topic "Navigate CA SDM" is included when you publish the help set.

7. Click OK.

The Selected Help Update window closes and the content is listed on the Contents tab.

8. Click Publish.

This generates the help set by packaging the selected topics into a help system you can display in a web browser.

9. Wait a few moments for the publishing process to complete; then select View, Refresh on the menu bar.

The View Help button becomes active.

10. Click View Help.

The custom help set appears in your default web browser.

Switch Roles

Any user with multiple roles assigned to them can switch roles without logging out and back in the system. Roles are assigned to users on their Access Type or contact record.

Note: For information, see [Configuring User Accounts](#) (see page 177).

To switch roles

1. Select the desired role from the Role drop-down list in the upper right corner of the main CA SDM page.
2. Click Set Role.

The web interface and available functionality for the logged in user change to match the current role.

Chapter 7: Establishing Support Structure

This section contains the following topics:

- [Support Structure](#) (see page 241)
- [Models](#) (see page 242)
- [CA Workflow](#) (see page 245)
- [CA Process Automation Workflow Integration](#) (see page 247)
- [Shared Codes](#) (see page 253)
- [Status Codes](#) (see page 256)
- [Task Types](#) (see page 261)
- [Incident Tracking](#) (see page 262)
- [Request/Incident/Problem Areas](#) (see page 263)
- [Change Order and Issue Categories](#) (see page 266)
- [Automatic Closure of Tickets](#) (see page 272)
- [Related Ticket Activities](#) (see page 274)
- [Priority Calculation](#) (see page 276)
- [Status Transitions and Dependent Attribute Controls](#) (see page 291)
- [Status Transitions for Self-Service](#) (see page 300)
- [Timers](#) (see page 304)
- [Time Zones](#) (see page 305)
- [File Attachments](#) (see page 308)
- [Announcements](#) (see page 312)
- [Stored Queries Setup](#) (see page 313)
- [Sequence Numbers](#) (see page 314)
- [Audit Log Use](#) (see page 315)
- [Integration with CA Network and Systems Management](#) (see page 315)

Support Structure

The *support structure* of your service desk consists of the components that users need for resolving problems. You can set up the CA SDM support structure to match your support model (internal, external, or both).

Note: You can customize issues, requests, and change orders to best suit the needs of your site. For more information, see the *Implementation Guide* and the *Online Help*.

Models

CA SDM supports the following service desk models:

- Internal Model
- External Model
- Combined Model

Note: For details about using the web interface to implement your selected model, see the *Online Help*.

Internal Model

An internal service desk supports employees who work for a company and have questions or problems with the products and services provided to them by the company. In CA SDM, the request is the basic unit of support when operating an internal service desk as follows:

- Requests are tickets that handle the questions or problems of employees, and are oriented toward supporting an infrastructure owned and administered by the support organization.
- Change orders are tickets that manage changes to the supported business infrastructure. Internal service desks often use requests as the primary ticket, attaching change orders in cases where the request must result in a change to the infrastructure.

If you are operating an internal service desk, do the following:

- Review the employee access type using the administrative function of the web interface to see that it meets your needs. If most of your contacts are employees who use the service desk for support, you might want to set "employee" as the default access type. This can save you from having to set the access type for each employee contact who is supported by your service desk.
- Review the employee contact type using the administrative function of the web interface to see that it meets your needs.
- Ensure that your contacts are set to have the appropriate access type and contact type. For example, if you set employee as the default access type, you also need to set the analyst contacts as the analyst access type.

The contact type is usually assigned automatically based on how you create the contact, but in some cases, the contact type might not be defined. Employees using the service desk for support should have a contact type of employee, whereas employees who work as support desk analysts should have a contact type of analyst.

You can work with contacts using the administrative function of the web interface.

If you are operating an internal service desk in which you are supporting employees, your support structure consists of the requests and change orders that they create and the underlying supporting features of those requests and change orders. As the administrator, you set up the support structure.

External Model

An external service desk supports customers who buy products or services from your company and have questions or problems with those products or services. In CA SDM, the issue is the basic unit of support when operating an external service desk. Issues are tickets designed to handle customer questions or problems and are oriented toward supporting products and services purchased by the customer.

If you are operating an external service desk:

- Review the customer access type using the administrative function of the web interface to see that it meets your needs. If most of your contacts are customers, you might want to set "customer" as the default access type.
- Review the customer contact type using the administrative function of the web interface to see that it meets your needs.
- Ensure that your contacts are set to have the appropriate access type and contact type. For example, if you set customer as the default access type, you also need to set the analyst contacts to the analyst access type.

The contact type is usually assigned automatically based on how you create the contact, but in some cases, the contact type might not be defined. Customers who use your service desk for support should have a contact type of customer and analysts who operate the service desk should have a contact type of analyst.

You can work with contacts using the administrative function of the web interface.

If you are operating an external service desk for which you are supporting customers, your support structure consists mostly of the issues that they create and the underlying supporting features of those issues. As the administrator, you need to set up the support structure using the information in the remainder of this chapter. Each topic states whether the information applies to your model.

Combined Model

Some companies need to operate both internal and external service desk models. In this case, you can do either of the following:

- Separate internal and external service desk models with separate service desk installations.
- Set up CA SDM to support both models.

This setup is convenient if your service desk analysts are cross-trained to support both employees and customers and if the distinction between internal and external support is not always clear. For example, you may have employees who purchase products from your company or customers who have problems and questions with the infrastructure of your company.

As the administrator, you must set up the support structure for a combined internal/external service desk that consists of issues, requests, and change orders, and their underlying supporting features. If you decide to use a combined internal and external service desk, do the following before setting up a structure:

- Review the customer and employee access types using the administrative function of the web interface to see that it meets your needs. If most of your contacts are customers, you might want to set "customer" as the default access type. If most of your contacts are employees using the service desk for support, you might want to set "employee" as the default access type.
- Review the customer and employee contact types using the administrative function of the web interface to see if it meets your needs.
- Verify that your contacts are set to have the appropriate access type and contact type. For example, if you set customer as the default access type, you also need to set the analyst contacts to the analyst access type, and the employee contacts to the employee access type.

The contact type is usually assigned automatically based on how you create the contact as follows, but in some cases the contact type might not be defined.

- Customers who use your service desk for support should have a contact type of customer.
- Employees who use the service desk for support should have a contact type of employee.
- Employees who work as support desk analysts should have a contact type of analyst.

You can work with contacts using the administrative function of the web interface.

More information:

[Contact Types](#) (see page 179)

[How Access Types Work](#) (see page 210)

CA Workflow

A CA Workflow process definition can be associated with any CA SDM [ticket type](#) (see page 32). The process definition that is applied to individual tickets is determined by the change category, issue category, or request/incident/problem area the ticket is assigned to.

If a CA Workflow process definition is associated with a category or area, when a ticket is assigned to that category or area, CA Workflow creates a process instance from the definition. Progress of the process instance is displayed on ticket's Workflow Tasks tab.

You can optionally install CA Workflow during CA SDM installation, or you may direct CA SDM to use a different instance of CA Workflow by updating several options in Options Manager, including:

CAWF_USERNAME

Specifies the CA EEM user who has full access (grants to the “IDE” and “Process” resources) to CA Workflow.

CAWF_PASSWORD

Specifies the password for the cawf_username.

CAWF_PM_LOCATION

Specifies the location: http://<server>:8090/pm/

CAWF_PM_URL

Specifies the URL:
http://<wf_hostname>:<wf_tomcat_port>/pm/services/pmService2

CAWF_WL_LOCATION

Specifies the location: http://<server>:8090/wl/

CAWF_WL_URL

Specifies the URL: http://<server>:8090/wl/services/wlService

These options specify various end points to CA Workflow. In most cases, you simply update the server and port to match the location of CA Workflow.

Important! Changing the CA Workflow server makes existing categories, areas, or tickets linked to process definitions and process instances on the old server inoperative.

Workflow at Runtime

When a ticket is saved with a category or area that is linked to a CA Workflow process definition, an instance of the definition is created. The progress of the instance is displayed on the Workflow Tasks tab of the ticket. Because CA Workflow may include branches, only completed and pending work items are displayed.

Each completed and pending work item, and the status of the overall workflow is displayed on the Workflow Tasks tab. Clicking on the Activity Name link displays the CA Workflow Worklist web application, which is used to complete individual work items.

The data for the Workflow tab is fetched directly from the CA Workflow server. If the CA Workflow server is unavailable, an error message is displayed in the tab.

Select a Workflow Process Definition

A CA Workflow process definition must have a string attribute named *usd_persid* that links with a change or issue category or request/incident/problem area. This attribute must be marked as an input parameter. When CA SDM starts the CA Workflow instance, it sets *usd_persid* to the unique identifier (*persistent_id*) of the ticket. One use of *usd_persid* is to make web service calls from CA Workflow to CA SDM.

Important! Change and issue categories can use either CA SDM's internal (classic) workflow system or the CA Workflow application; the two cannot be mixed. Request/incident/problem areas can use only the CA Workflow application. The internal classic workflow system is not supported for request, incident, or problem tickets.

To select a CA Workflow process definition

1. Select a change or issue category. For example, to select a change category, select Service Desk, Change Orders, Categories on the Administration tab.

The Change Category List page appears.

2. Click the symbol of the change category you want to edit.

The Change Category Detail page appears.

3. Click the Use CA Workflow button on the Workflow tab.

Note: If you continue this procedure and save the category, any CA SDM style workflow that is attached to the category is deleted.

A selection page appears lists the available CA Workflow process definitions.

4. Select one CA Workflow process definition.

The category is populated with the selection.

5. Click Save.

The category is updated with the selected process definition.

6. (Optional) Change the linked CA Workflow process definition by editing the category and clicking the CA Workflow definition hyperlink

The process definition selection page appears.

Note: CA Workflow process definitions are displayed in the selection page only if the definition has a string attribute named `usd_persid`. When the definition is started, this attribute is set to the `persistent_id` value of the ticket. The CA Workflow process definition can optionally include a string input parameter named *label*. If a label is defined, CA SDM sets its value to the ticket number. The ticket number is displayed in the CA Workflow worklist and helps provide context for the workitem.

Workflow Tasks

Workflow tasks identify the activities that must be completed on tickets associated with a specific category or area. As with properties and other aspects of the category or area, the tasks are automatically added to tickets when the category is selected. Defining workflow tasks lets you do the following:

- Specify and describe the type of task to be completed.
- Assign a sequential number to each task.
- Specify the user to complete the task.
- Track the progress toward resolving tasks by assigning specific states that are valid for each task.

When the status of a task changes, certain behaviors can execute. Behaviors let you define the specific tasks or processes that are performed when the workflow task reaches a specific state. For example, by defining behaviors, you can send an email notification to a specific manager when an approval task enters the Pending state.

Note: For more information about how to create workflow tasks and define behaviors for them, see the *Online Help*.

CA Process Automation Workflow Integration

CA Process Automation is a stand-alone CA product with features for automating and tracking hardware and software administration tasks in enterprise IT environments. CA Process Automation automates tasks and manages user interactions, such as approvals and notifications for compliance and accuracy within production environments.

When you integrate CA SDM and CA Process Automation, you can leverage the benefits of CA Process Automation workflow capabilities from key points within CA SDM. An effective integration between CA SDM and CA Process Automation requires you to understand both products.

More information:

[CA Process Automation Components](#) (see page 248)

[CA Process Automation Integration with CA SDM at Run Time](#) (see page 249)

[How to Create a Process Definition](#) (see page 250)

[Create a Start Request Form](#) (see page 251)

[Attach a CA Process Automation Process Definition](#) (see page 252)

CA Process Automation Components

CA Process Automation offers multiple capabilities and structures that facilitate a wide range of activities as part of CA Process Automation process management. For the integration with CA SDM, however, only the following CA Process Automation components are critical for CA Process Automation integration:

- **Process Definition**—Identifies a collective series of tasks, steps, and conditions that are structured in a specific order to be initiated, and completed by various individuals or parties. This component is the central building block of all CA Process Automation content.
- **Start Request Form**—An object containing descriptive information for end users. The Start Request Form presents a process definition to users while hiding the technical details of the process definition.
- **Keywords**—A list of pre-defined words or phrases to attach to Start Request Forms.
- **Automation Library**—An area within CA Process Automation that stores and shows Process Definitions and Start Request Forms.
- **Library Path or Reference Path**—A folder structure that organizes and describes Process Definitions and Start Request Forms within the automation library.
- **Process Instance**—An active entity that executes the rules that are defined in a process definition. The process instance progresses until the process definition state is complete.
- **Process Instance Log Messages**—A configurable, running record that details the progression of activities of the process instance. Log message categories are useful to the CA SDM integration with CA Process Automation.

Note: The scope of the definitions of CA Process Automation components is limited to the usage within CA SDM. For information about CA Process Automation components and capabilities, see the CA Process Automation documentation.

CA Process Automation Integration with CA SDM at Run Time

When you enable the integration, CA SDM users experience the following:

- On a new Request, Change Order, or Issue, a CA Process Automation process instance initiates based on the ticket category or area. Summary information immediately appears on the Workflow Tasks tab.
- When a Request Area, Change Category, or Issue Category changes, an attached CA Process Automation process instance terminates and a new process instance initiates.
- When a CA SDM user attempts to close a Request, Change Order, or Issue where the CA Process Automation process instance is not yet complete, the user cannot close the ticket. Instead, the user must first cancel the ticket. The Cancel status terminates the CA Process Automation process instance before the ticket closes.
- When a user wants to understand the state of the process instance without navigating away from the ticket, the user can click the ticket Workflow Tasks tab. The Workflow Tasks tab shows the process instance start date, end date, current state, and a current audit trail of messages indicating the path of the process instance.
- When a user wants to see the current path of the process instance relative to the entire process, the user selects the View Process button on the Workflow Tasks tab. The View Process button launches a graphical snapshot of the entire process instance, and shows the current path.
- When a user wants to see CA Process Automation interaction request forms that are waiting for user action, the user can select any entry in the Workflow Tasks tab. The Workflow Tasks tab contains an audit trail of process instance messages that appear on the CA Process Automation task list.

Note: When a user selects the CA SDM View Process button or CA Process Automation process instance messages, the system prompts for a CA Process Automation user name and password for a single browser session. After the initial prompt, the system does not prompt the user again until the CA SDM browser closes.

How to Create a Process Definition

When you create the process definition, you populate CA Process Automation with content to appear in CA SDM. You use the CA Process Automation graphical process designer to create, test, and check in a process definition. From CA SDM, you can also create macros to initiate CA Process Automation processes.

Note: For information about creating CA Process Automation macros, see the *Online Help*.

To create a process definition in CA Process Automation, do the following:

1. Log in to the CA Process Automation client as an administrative user.
2. Use the CA Process Automation graphical process designer to create, test, and check in a process definition. When you work with the process definition, use the instructions in the CA Process Automation user documentation.

Note: If you fail to check in the process definition before attempting to use it, the workflow cannot operate properly in CA SDM.

The following process definition items are available in CA SDM:

Process Name

Appears on the Request Area, Change Category, and Issue Detail page. The process name also appears on the Workflow Tasks tab of a ticket. The process name describes the process definition to CA SDM to Analysts and other users who manage tickets.

Process Reference Path

Appears on the Request Area, Change Category, and Issue Detail pages. The Process Reference Path also appears on the ticket Workflow Tasks tab. The Process Reference Path can be useful to describe the purpose of the process to end users. For example, "/Processes/Approval" is not helpful. Instead, a reference path like "/Office Supplies/Approvals/Over-200-USD" describes the workflow to manage orders that exceed \$200 US dollars.

Process Log Messages

Appears on the CA SDM ticket Workflow Tasks tab. By default, a record of activities stores with the process instance. Process log messages have the Process category. A process designer can create custom messages to appear on the Workflow Tasks tab of a CA SDM ticket.

Note: For information about configuring the `caextwf_log_categories` option to manage log messages, see the *Implementation Guide*.

Create a Start Request Form

When you create a Start Request Form, you associate it with the process definition and check it in. You include the appropriate keywords in the Start Request Form properties. If the appropriate keyword is missing from the Start Request Form properties, the Start Request Form and its associated process definition fail to appear in CA SDM.

To create a Start Request Form

1. Log in to the CA Process Automation client as an administrative user.
2. Open the CA Process Automation Library and navigate to the path Start Request Form.
The Start Request Form appears in the right pane of CA Process Automation library.
3. Select the Start Request Form from the list.
A shortcut menu appears.
4. Select Properties.
The Library Object Properties page appears.
5. (Optional) Click the General tab and modify the description of the Start Request Form. Add a description that identifies the proper usage of the Start Request Form and the associated Process Definition to the CA SDM Administrator.
6. Click the Keywords tab.
The Keywords tab is active.
7. Click the ab+ icon.
A row adds to the empty list.
8. Click the row.
A blinking cursor highlights the row and indicates the row is ready for typing.
9. Enter one of the following values to associate a keyword to the appropriate ticket area or category. For example, to make a Start Request Form available for a CA SDM request area, enter the pcat keyword.

Ticket	Use Keyword
request area	pcat
change category	chgcac
issue Category	isscat

10. Add a row to the list for each applicable keyword. For example, to make the Start Request Form appear on both request areas and change categories, add one row for the chgcac keyword and another row for the pcat keyword.

11. Click OK.

CA Process Automation saves the keywords and description and closes the Library Object Properties dialog.

12. Check in the Start Request Form.

Note: If you fail to check in the Start Request Form, the form fails to appear in CA SDM.

The CA Process Automation Start Request Form information appears on the CA SDM Start Request Form List. The CA SDM administrator can associate the CA SDM Start Request Form with the Process Definition, on a Request Area, Change Category, or Issue Category Detail page.

The following Start Request Form items appear in CA SDM:

Start Request Form Name

Appears on the Request Areas, Change Categories, and Issue Categories list pages.

Start Request Form Reference Path

Appears on the Request Areas, Change Categories, and Issue Categories list pages.

Start Request Form Description

Appears on the Request Areas, Change Categories, and Issue Categories list pages. The text in this field describes how the Start Request Form is appropriate for selection on a particular Request Area, Change Category, or Issue Category.

Attach a CA Process Automation Process Definition

When you attach a CA Process Automation process definition to a CA SDM ticket area or category, you create a static connection between a CA SDM ticket area or category and a CA Process Automation process definition.

When a CA SDM user creates or edits a ticket and selects a ticket category, the associated CA Process Automation process definition launches into a process instance. Pertinent information about the process instance appears on the Workflow Tasks tab on the ticket.

To attach a process definition

1. On the Administration tab, select Service Desk.
2. Navigate to the ticket Areas or Categories.

The Area or Category List appears.

3. Create or edit a ticket area or category.

The Update page appears.

4. Click the Workflow tab.

If the CA Process Automation Options are installed in the CA SDM Options Manager, the Use CA IT PAM button is available on the Workflow tab.

5. Click the Use CA IT PAM,

The CA IT PAM Start Request Form List appears. Each row in the list is a CA Process Automation Start Request Form.

Note: You can use CA Process Automation, CA Workflow, or the CA SDM internal workflow to manage different ticket areas or categories. However, a single category can only use one workflow tool at a time.

6. Click the value in the Name column to select the process definition associated with this Start Request Form.

The CA IT PAM Start Request Form List closes. The process definition name and process definition reference path appear on the Workflow tab.

7. Click Save.

The system saves the process settings. The next ticket that a user creates in the specified ticket area or category automatically attaches the workflow and creates a process instance. The ticket Workflow Tasks tab shows a summary of the process instance information. Additionally, a user can access additional information about the process instance by clicking View Process on the Workflow Tasks tab.

Shared Codes

In CA SDM, different types of tickets share certain underlying codes, such as priority, severity, impact, and urgency codes. Requests and change orders share some codes, and all types of tickets share other codes.

Consider the following information about shared codes:

- By default, numeric values rank them.
- You can customize the codes.
- You cannot add or delete shared codes.
- You can use impact, priority, severity, and urgency codes.

Based on your service desk model, set up the following codes:

Shared Codes	Description
Priority	Must be set up for all service desk models.

Shared Codes	Description
Severity	Must be set up for internal and combined service desk models.
Impact	Must be set up for internal and combined service desk models.
Urgency	Must be set up for internal and combined service desk models.

Note: For details about how to customize these codes using the administrative function of the web interface, see the *Online Help*.

More information:

[Priority Codes](#) (see page 254)
[Severity Codes](#) (see page 255)
[Impact Codes](#) (see page 255)
[Urgency Codes](#) (see page 255)

Priority Codes

Priority codes indicate a ranking order by which the service desk should respond to tickets (that is, they specify the level of attention a ticket should receive). Priority codes are referenced in requests, change orders, and issues; therefore, they apply to all service desk models.

You can use priorities to escalate tickets manually or automatically by monitoring events. In many service desk installations, priority codes are used on the scoreboard to provide analysts with a real-time status of their requests and change orders.

You can assign a service type to a priority code, which is then automatically assigned to tickets when the priority code is specified. This lets you associate a specific level of service to a ticket based on the assigned priority. For example, the system-defined service type, 4-hour resolution, is automatically associated with priority 1. Tickets that are assigned a priority of 1, therefore, are automatically assigned this service type, including all the service type events that are associated with the 4-hour resolution service type.

More information:

[Service Level Agreements](#) (see page 141)
[How to Implement Service Types](#) (see page 143)

Severity Codes

Severity codes identify the extent of the damage to equipment affected by a request. Severity codes are referenced in requests only; therefore, they apply only to internal and combined service desk models.

Note: Severity is often used as a synonym for priority. Some sites use priority only, ignoring severity altogether. If you want to distinguish between how serious a problem is on a technical level (severity) and how quickly you want it handled (priority), you can use severity and priority codes.

Impact Code values helps to calculate the Incident Priority.

Impact Codes

Impact codes measure the significance of a ticket on the functioning of the system. For example, if a change order could affect the functioning of the entire system, it would be assigned a high impact. Impact codes are referenced in requests and change orders only; therefore, they apply only to internal and combined service desk models.

Note: Impact and [urgency codes](#) (see page 255) are similar, but they have distinct purposes.

Urgency Codes

Urgency codes measure the significance of the request for users of the system (that is, they indicate the importance of the request to the overall production environment). For example, if a request could jeopardize the mission of the enterprise, the Urgency code can be a value of 5-Immediate. Urgency codes are referenced in requests only; therefore, they apply only to internal and combined service desk models.

Urgency and impact codes serve distinct purposes, but are often confused because they coincide. For example, a request to report a fire in a critical data center can have a 3-Single Group Impact and 5-Immediate Urgency. These codes apply because the fire impacts more than one group but not necessarily the entire organization. Because the data center is critical for operations, the urgency requires immediate attention.

Status Codes

Status codes are used to track the status of an item. Separate status codes track requests, change orders, issues, and workflow tasks. In each case, there are predefined status codes that you can use if they suit your needs. Otherwise, you can modify the predefined status codes or define new ones that are specific to your site. Depending on your service desk model, you set up the following codes:

Status Codes	Description
Request	Must be set up for internal and combined service desk models.
Change Order	Must be set up for internal and combined service desk models.
Issue	Must be set up for external and combined service desk models.
Task	Must be set up for all service desk models.

Status codes let analysts sort and select information based on the status so that they can carefully track their progress. How carefully you define the status codes determines how accurate the analysts can be in describing the actual status of an item.

You can mark any status code as active or inactive. When you mark a status code as inactive, it is no longer available for analysts to use, but it remains available for future use (that is, it is not deleted from the database). If you decide later to use the status code, you need only to go back to it and mark it as active.

Note: The same area definitions are available for request, incident, and problem tickets. On the Administration tab, these areas are referred to as request/incident/problem areas. For brevity, they are referred to here simply as request areas.

Request Status Codes

The following table describes the predefined status codes for request tickets.

Request Status Code	Description
Acknowledged	The receipt of a request has been acknowledged.
Closed	A request has been completely resolved.
Closed—Unresolved	A request has been closed but still must be resolved.

Request Status Code	Description
Fix in Progress	A request is pending a fix.
Hold	The service type events for the request are on hold.
Open	A request has been defined and is being used to monitor and manage its completion.
Problem Closed	A problem request is completely closed.
Problem Fixed	A problem request is fixed, but not closed.
Problem Open	A request has been identified as a problem.
Researching	A request is open pending additional research and analysis.
Work in Progress	Work is being done to fix a request.

If your site uses other terminology for identifying the request status, you should define status codes that suit your needs and ignore the predefined status codes, or change the definitions to match your use. For example, you may want to define additional request status codes, such as the following codes:

Request Status Code	Description
Duplicate	Requests that have been opened but may be a duplicate of an existing request for another user.
Emergency	Critical requests that must be addressed immediately.
Report	Requests that are resolved and closed, but should be reported on at a management level.
Test	Requests that are resolved, but should be tested for one week before closing.

Change Order Status Codes

The following table describes the predefined status codes for change order tickets.

Change Order Status Code	Description
Approval in Progress	A change order is open, pending approval.
Approved	A change order has been approved.
Cancelled	A change order has been cancelled.
Closed	A change order has been completed.
Hold	The service type events for the change order are on hold.
Implementation in Progress	A change order is being implemented.
Open	A service order has been defined in a change order and the change order is being used to monitor and manage its completion.
Rejected	A change order has been rejected.
Resolved	A change order has been resolved.
RFC	A request for change has been submitted.
Suspended	Stops workflow tasks on a change order.
Verification in Progress	A change order is being verified.
Backed Out	Implemented change order was backed out.
Implemented	Change was implemented.

If your site uses other terminology for identifying the status of a change order, you should define status codes that suit your needs and ignore the predefined status codes, or change the definitions to match your use. For example, you may want to define additional change order status codes, such as those listed in the following table:

Customized Request Status Code	Description
Duplicate	Change orders that have been opened but may be a duplicate of an existing change order for another user.
Emergency	Critical change orders that must be addressed immediately.
Report	Change orders that are completed and closed, but should be reported on at a management level.

Closure Codes

Use closure codes to define the final outcome of change orders, such as successful or unsuccessful. Set closure codes manually or as part of the Update Status activity on a change order when the status is closed, finished or resolved.

Note: For more information about creating or defining closure codes and details about the *require_closure_code* and *force_closure_code* options, see the *Online Help*.

Issue Status Codes

The following table describes the predefined status codes for issue tickets.

Issue Status Code	Description
Approval in Progress	An issue is open, pending approval.
Cancelled	An issue has been cancelled.
Closed	An issue has been completed.
Hold	The service type events for the issue are on hold.
Implementation in Progress	An issue is being implemented.
Open	An issue has been defined and open so that it can be monitored and managed until it is resolved.
Suspended	Stops workflow tasks on an issue.
Transaction in Progress	A transaction with a customer concerning this issue is in progress.
Verification in Progress	An issue is being verified.

If your site uses other terminology for identifying the status of an issue, you should define status codes that suit your needs and ignore the predefined status codes, or change the definitions to match your use. For example, you may want to define additional issue status codes, such as the following:

Customized Issue Status Code	Description
Duplicate	Issues that have been opened but may be a duplicate of an existing issue for another user.

Customized Issue Status Code	Description
Emergency	Critical issues that must be addressed immediately.
Report	Issues that are completed and closed, but should be reported on at a management level.

Task Status Codes

Task status codes describe the different possible states of a workflow task. Each task in an issue or change order workflow has its own status, separate from the ticket status. Workflow tasks allow analysts to keep track of how long it takes to complete individual tasks within a ticket.

The following table describes the predefined status codes for workflow tasks.

Task Status Code	Description
Approve	Task approved.
Cancelled	Task is cancelled and no further updates are possible.
Complete	Task is complete.
Pending	Task is started.
Reject	Task is rejected.
Reopen	Task has been re-opened.
Reopen–Wait	A prior task was reopened.
Skip	Skip task.
Wait	Task has not started.

If your site uses other terminology for identifying the status of a workflow task, you should define status codes that suit your needs and ignore the predefined status codes, or change the definitions to match your use.

For each task status code, you can assign a type of behavior that occurs when the task reaches this state, which provides much more information about the progress toward completing the task. You can also use the accumulate function to track time and cost involved in completing the ticket.

Task Types

Task types help determine the behavior of specific workflow tasks and task status codes. To produce defining characteristics for each type of task, you can identify the task status codes or specific states that can be used.

Because both change orders and issues use workflow tasks, set up task types for all service desk models. The task status codes identify the behaviors associated with each task. For example, you can set the Approval task type to allow Approve, Reject, Pending, and Wait as available states. When the Approval task type enters the Pending state, you can send notification to a specific manager, analyst, and so forth.

You can mark any task type as active or inactive. When you mark a task type as inactive, it is no longer available for analysts to use, but it remains available for future use (it is not deleted from the database). If you decide later to use the task type, mark it as active again.

Note: You can view the predefined task types in the Task Type List page of the administrative function of the web interface.

Example: Task Status Codes

The following table contains some example status codes:

Task Status Code	Description
Approval	Approve or reject ticket
Group End Task	End of group tasks
Group Start Task	Start of group tasks
Start Approval	Approval to start the ticket

In this example, the Group Start Task and Group End Task types define a group of tasks in an issue or change category that must be completed. Tasks in the group can be executed in any order. After the Group Start Task is in the Pending state (started), all tasks in the group are also placed in this state.

Note: For more information about using task types, see the *Online Help*.

Incident Tracking

Incident tracking lets analysts track an incident by selecting one or more flags for the incident. The information that analysts specify provides your organization with metrics about incidents for reports. For example, analysts can indicate that an incident was assigned incorrectly. When a large percentage of incorrectly assigned tickets appears in a report, your organization is aware that assignments must be adjusted.

For example, analysts can specify information to help your organization do the following:

- Improve SLA responsiveness and closure at lower levels within the support organization.
- Identify tickets that are incorrectly assigned.
- Indicate that a remote control tool was used to resolve the ticket.

You install the `efficiency_tracking` Options Manager option, so that analysts can use tracking options that appear on the Efficiency Tracking tab of incident detail pages.

Note: For more information about incident tracking, see the *Online Help*.

Install Incident Tracking

You can install an option to let analysts track an incident by setting particular flags for the incident. The information that analysts specify provides your organization with metrics about incidents for reports. After you install the option, the tracking flags appear on the Efficiency Tracking tab of incident detail pages.

To install incident tracking

1. On the Administration tab, browse to Options Manager, Request Mgr.
The Options List appears.
2. Click `efficiency_tracking`.
The `efficiency_tracking` Options Detail page appears with default values set.
3. Click Edit.
The `efficiency_tracking` Options Update page appears with default values set, and you can edit the Description.
4. Click Install.
The `efficiency_tracking` Options Detail page appears.

5. Click Close Window.

Incident tracking is installed; however, the Efficiency Tracking tab does not appear on incident detail pages.

6. Restart CA SDM.

The Efficiency Tracking tab appears on incident detail pages.

Request/Incident/Problem Areas

Request areas define the logical groupings into which you can organize request, incident, and problem tickets. For example, tickets pertaining to an application can be assigned to the predefined Applications area. Whenever an analyst assigns a ticket to a request area, all the information you have associated with the request area is automatically entered on the ticket. For example, if you indicate a service type, it becomes associated with the ticket and all its associated service type events.

Note: The same area definitions are available for request, incident, and problem tickets. On the CA SDM Web Interface Administration tab, these areas are referred to as request/incident/problem areas. For brevity, they are referred to here simply as request areas.

You can set the status of any request area to active or inactive. When you make a request area inactive, it is no longer available for analysts to use, but it is not deleted from the database. If you decide later to use the request area, you need only change the status back to active.

You can use request areas to do the following:

- Specify default values for the group and assignee fields on tickets.
- Automatically associate a level of service to tickets by assigning a default service type to the request area.
- Associate a survey with a request area.
- Select and report on tickets by area by defining your own custom request areas. Eventually study request trends and analyze problem causes. Focusing your view to specific request areas can help make these studies more significant and revealing.

The following predefined request areas are installed with CA SDM:

- Applications
- Email
- Hardware

- Networks
- Printer
- Software

Software is subdivided into several request areas.

Note: For information about defining and editing request areas, see the *Online Help*.

More information:

[Request/Incident/Problem Area Properties](#) (see page 264)

[Define Request/Incident/Problem Areas for Self-Service](#) (see page 266)

Request/Incident/Problem Area Properties

You can use properties to add custom attributes or qualities to a specific request area. If you have added properties to a request area, when an analyst assigns a ticket to that area the associated properties automatically appear on the ticket's properties tab. For example, you can add properties to the predefined email request area to specify the email server or mailbox size.

As you define properties, you can specify whether a value is required or optional. For [properties with a required value](#) (see page 264), users must enter a value (or accept the default) before they can save the ticket.

The `keep_tasks` option determines what happens when you assign an existing ticket to a different request area:

- If `keep_tasks` is not installed, existing properties (and workflow tasks) are removed from the ticket, and any properties or tasks associated with the new request area are added.
- If `keep_tasks` is installed, existing properties and tasks are retained on the ticket, and any properties or tasks associated with the new request area are added.

Note: For detailed instructions about adding properties to a request area and defining property validation rules, see the *Online Help*.

Property Validation Rules

You can define validation rules to restrict the property values users can enter to a predefined set of selectable values. For example, if you have defined a property named Operating System, you can define the validation rule as a drop-down list containing the options Windows, UNIX, and Linux.

Properties defined without validation rules are presented to the user as freeform text boxes, which allow any text string to be entered. Validation rules make reporting on area and category values less complex and less error prone.

Note: Property validation rules are reusable. They are not specific to a particular property. You can apply any existing validation rule to properties defined for change categories, issue categories, or request/incident/problem areas.

Depending on which type of validation rule you have configured for a property, when the user assigns a ticket to the category or area to which that property is attached, one of the following controls appears on the ticket properties tab:

- Text Edit Box—No validation rule has been defined and no default value can be specified. Users are expected to enter values that follow the examples you provide.
- Check Box—A two-state check box appears on the Properties tab. By default, the check box does not contain a checkmark. The user can accept the default, or select the check box. Boxes that contain a checkmark are displayed on the ticket detail page with Yes in the value column. Cleared boxes are displayed with No in the value column.
- Drop-down List—A drop-down list containing a predefined set of options appears on the Properties tab. If you have defined a default value, it is selected when the drop-down list is first displayed. The option the user selects is shown in the value column of the ticket detail page.

Define Request/Incident/Problem Areas for Self-Service

The Self-Service Include option lets you define which request/incident/problem areas to include on tickets for self-service. You can also define different self-service symbols than those seen by the analyst. When the ticket is saved, the self-service symbol displays in the Request/Incident/Problem Area field. If the ticket displays in the analyst interface, the normal symbol for the area appears. The self-service user is not allowed to edit the symbol; it is read-only.

To define a request/incident/problem area for self-service

1. On the Administration tab, select Service Desk, Requests/Incidents/Problems, Areas.

The Requests/Incidents/Problems Area List appears.

2. Select Edit In List.

The top portion of the page displays the editable fields.

3. Open the desired incident/problem/request area for editing on the Symbol list (Hardware.pc.install, for example).
4. Complete the following fields:

Self-Service Include

Specifies whether this request/incident/problem area is shown in the self-service interface.

Self-Service Symbol

Specifies a unique identifier for this request/incident/problem area in the self-service interface.

Active

Specifies whether the request/incident/problem area is active or inactive.

The request/incident/problem area is defined for self-service.

5. Click Save.

The updated request/incident/problem area appears in the Request/Incident/Problem Area List when you redisplay the list.

Change Order and Issue Categories

Change categories and issue categories define the logical groupings into which change orders and issues can be organized.

Note: Unlike request/incident/problem areas, change order categories and issues categories are managed separately:

- Set up change categories for internal and combined CA SDM models.
- Set up issue categories for external and combined CA SDM models.

You can use categories to specify default values for certain fields in tickets. You can automatically associate a level of service to tickets by assigning a default service type to categories. You can also associate a survey with a category.

For each category, you can specify attributes or qualities to be associated with the ticket and create a workflow that identifies all the individual tasks required to fulfill the ticket. By defining behaviors that are associated with the workflow tasks, you can notify key personnel when the status of the task changes or as activities close the ticket.

Whenever an analyst assigns a ticket to a category, all the information you have associated with the category is automatically entered on the ticket. For example, if you indicate a service type, it becomes associated with the ticket, and its associated service type events.

Note: For more information about defining and editing categories, see the *Online Help*.

More information:

[Predefined Change Categories](#) (see page 267)

[Predefined Issue Categories](#) (see page 268)

[Rules for Changing Categories on a Ticket](#) (see page 268)

[Category Properties](#) (see page 269)

[Define Change and Issue Categories for Self-Service](#) (see page 271)

Predefined Change Categories

The following sets of predefined change categories are installed with CA SDM:

- Add
- Change
- Move
- Retire

Note: All these category sets are subdivided into more specific categories. For example, the Change category set includes categories for changing servers and workstations.

Predefined Issue Categories

The following predefined issue categories are installed with CA SDM:

- Hardware.pc.install
- Software.pc.install

You can set the status of any category to active or inactive. When you make a category inactive, it is no longer available for analysts to use, but it is not deleted from the database. If you decide at a later time to use the category, change the status back to active.

Rules for Changing Categories on a Ticket

The following rules affect only workflow tasks.

- If the previous category used CA Workflow or Classic Workflow and the new category uses CA Process Automation the CA Process Automation process definition links to a workflow on a CA Process Automation server.
- If both the old and new categories use CA SDM workflow, the rules from previous releases apply.
- If the new category uses CA Workflow or CA Process Automation and the old uses CA SDM workflow, the following occurs:
 - All incomplete and pending (those tasks that can be updated) workflow tasks of the CA SDM 6.0 style, are set to Cancelled status, regardless of the KEEP_TASKS option. Any completed workflow tasks remain, but they cannot be reopened.
 - All incomplete and nonactive tasks (such as tasks in Wait status) are deleted.
 - The CA Workflow or CA Process Automation definition instantiates.
- If the new category uses CA SDM workflow and the old category uses CA Workflow or CA Process Automation, the following occurs:
 - The CA Workflow or CA Process Automation instance is forcibly set to a status of Terminated.
 - The new category workflow tasks are added as normal.

- If both the old and new categories use CA Workflow, the following occurs:
 - If the old and new categories point to the same process definition, the running instance from the previous category continues.
 - If the old and new categories point to different process definitions, the old instance is set to Terminated and the definition from the new category instantiates.

In a CA SDM workflow, a ticket with a running CA Workflow or CA Process Automation instance cannot be closed. The Accumulate, Expedite, and Insert Task functions are disabled for tickets using CA Workflow.

Category Properties

Properties are used to add custom attributes or qualities to a specific category. If you have added properties to a category, when an analyst assigns a ticket to that category the associated properties automatically appear on the ticket Properties tab.

For example, you can add properties to the predefined Software.pc.install issue category to specify memory size, processor type, and so on.

As you define properties, you can specify whether a value is required or optional. For properties with a required value, users must enter a value (or accept the default) before they can save the ticket.

Category Properties Retention

The keep_tasks option determines what happens when you assign an existing ticket to a different category:

- If keep_tasks is not installed, existing properties and workflow tasks are removed from the ticket, and any properties or tasks associated with the new category are added.
- If keep_tasks is installed, existing properties and tasks are retained on the ticket, and any properties or tasks associated with the new category are added.

Note: For detailed instructions on adding properties to a change or issue category and defining property validation rules, see the *Online Help*.

Property Validation Rules

You can define validation rules to restrict the property values users can enter to a predefined set of selectable values. For example, if you have defined a property named Operating System, you can define the validation rule as a drop-down list containing the options Windows, UNIX, and Linux.

Properties defined without validation rules are presented to the user as freeform text boxes, which allow any text string to be entered. Validation rules make reporting on area and category values less complex and less error prone.

Note: Property validation rules are reusable. They are not specific to a particular property. You can apply any existing validation rule to properties defined for change categories, issue categories, or request/incident/problem areas.

Depending on which type of validation rule you have configured for a property, when the user assigns a ticket to the category or area to which that property is attached, one of the following controls appears on the ticket properties tab:

- **Text Edit Box**—No validation rule has been defined and no default value can be specified. Users are expected to enter values that follow the examples you provide.
- **Check Box**—A two-state check box appears on the Properties tab. By default, the check box does not contain a checkmark. The user can accept the default, or select the check box. Boxes that contain a checkmark are displayed on the ticket detail page with Yes in the value column. Cleared boxes are displayed with No in the value column.
- **Drop-down List**—A drop-down list containing a predefined set of options appears on the Properties tab. If you have defined a default value, it is selected when the drop-down list is first displayed. The option the user selects is shown in the value column of the ticket detail page.

Define Change and Issue Categories for Self-Service

You can use the Self-Service Include option to define which change and issue categories to include on tickets for self-service. You can also define different self-service symbols than those seen by the analyst. When the ticket is saved, the self-service symbol appears in the Change (or Issue) Category field. If the ticket displays in the analyst interface, the normal symbol for the category appears.

To define a change (or issue) category for self-service

1. On the Administration tab, select Service Desk, Change Orders (or Issues), Categories.

The Category List page appears.

2. Select Edit In List.

The top portion of the page displays the editable fields.

3. Select the desired category from the Symbol list.
4. Complete the following fields:

Self-Service Include

Specifies whether this category is shown in the self-service interface.

Default: Yes

Self-Service Symbol

Specifies a unique identifier for this category in the self-service interface.

Active

Specifies whether the category is active or inactive.

The category is defined for self-service.

5. Click Save.

The updated category appears in the Change (or Issue) Category List when you redisplay the list.

Automatic Closure of Tickets

You can use a configurable setting to allow automatic closure of tickets (requests/incidents/problems, change orders, or issues). When a ticket is set to a Resolved status, the ticket is automatically closed in the number of business hours specified. The Auto Close activity notification sent to the end user displays the number of business hours before the ticket is closed. The number of hours is configurable and tenant-specific. If the status is changed before the configurable number of hours ends, the ticket closure is canceled.

Administrators, can perform the following actions:

- Define an Auto Close ticket setting to control the number of business hours, for the end user, before the ticket is automatically closed.
- Set up an Auto Close activity notification to notify the appropriate contacts when automatic closure is scheduled for a ticket.

If you use multi-tenancy, consider the following:

- The system uses the default public Auto Close setting when a tenant-specific Auto Close setting is not found.
- There is one Auto Close setting for each tenant.

Note: For more detailed information about performing these procedures, see the Auto Close Settings information in the *Online Help*.

More information:

[How to Define Auto Close Ticket Settings](#) (see page 272)

[How to Define an Auto Close Activity Notification](#) (see page 273)

How to Define Auto Close Ticket Settings

You can define the number of business hours before a ticket is closed (all ticket types) as follows:

1. On the Administration tab, select Service Desk, Application Data, Codes, Auto Close.
2. Click Create New on the list page.
3. Complete the following fields on the detail page:

Symbol

Defines the system setting name.

Request/Incident/Problem/Change Order/Issue

Defines the number of business hours, after the ticket is set to Resolved status, before the ticket is closed. If the status is changed before the number of hours ends, the ticket closure is canceled. 0 (zero) hours indicates that automatic closure is not implemented for the ticket type.

Description

Provides a brief description of the record.

Status

Indicates if the record is active or inactive.

The auto close setting is defined.

4. Click Save, Close Window.

The new setting appears on the Auto Close List page when you redisplay the list.

How to Define an Auto Close Activity Notification

You can change the settings in the default Auto Close activity notification to notify the appropriate contacts when automatic closure is scheduled for a ticket. The activity is valid for all CA SDM ticket types, and include a default notification rule for requests/incidents/problems, change orders and issues.

To define an auto close activity notification, do the following:

1. On the Administration tab, select Notifications, Activity Notifications.
2. Select the Auto Close activity notification on the list page to open it.
3. Update the default Auto Close Notification Rule and specify contacts to receive notification.
4. Click Save, Close window.

The updated Auto Close activity notification appears in the Activity Notification List when you redisplay the list.

Related Ticket Activities

When an activity is generated for a CA SDM ticket, you can propagate the activity to one or more related tickets. For example, a Problem record created from an Incident can update the Incident record when the Problem is resolved. When the activity occurs, an activity log is generated for the related ticket that includes the following information:

- Related ticket activity type, for example, Update Status
- Contact name
- Parent ticket type and its reference number
- Activity log description, for example, status updated from Work in Progress to Open

Activity logs are propagated to related tickets based on the properties set within each activity notification. The attributes of the related tickets are not modified. The following relationships are propagated:

- Problems propagate to all active Incidents.
- Change Orders propagate to all active Problems and Incidents.

As a system administrator, you can perform the following actions:

- Set up activity notifications to propagate related ticket activities.
- Configure a Related Ticket activity notification to notify the appropriate contacts when the activity is propagated to related tickets.

Note: For more detailed information about performing these procedures, see the Activity Notifications information in the *Online Help*.

More information:

[Activity Logging](#) (see page 384)

[How to Define Activity Notifications for Related Tickets](#) (see page 275)

[How to Define Related Ticket Activity Notifications](#) (see page 275)

How to Define Activity Notifications for Related Tickets

You can propagate activity logs to related tickets based on the properties set in each activity notification. The attributes of the related tickets are not modified.

To define an activity notification for related ticket activities, do the following:

1. On the Administration tab, select Notifications, Activity Notifications.
2. Open the appropriate activity notification for editing.
3. On the detail page, click the Related Ticket Activity check box to mark it as active.

Note: If you use multi-tenancy, do the following:

- Specify the appropriate tenant type from the Related Ticket Activity drop-down list.
 - Enter the name of the tenant in the tenant field, or click the search icon to search for a tenant.
4. (Optional) Update the default Notification Rule and specify contacts to receive notification.
 5. Click Save, Close Window.

The updated activity notification appears in the Activity Notification List when you redisplay the list.

How to Define Related Ticket Activity Notifications

You can change the default settings in the Related Ticket activity notification to notify the appropriate contacts when the activity logs are propagated to related tickets (requests/incidents/problems, change orders, and issues).

To define a related ticket activity notification, do the following:

1. Open the Related Tickets activity notification on the Activity Notification List page.
2. Click Edit and change one or more of the fields as appropriate on the detail page.
3. (Optional) Update the default Notification Rule for and specify contacts to receive notification.
4. Click Save, Close Window.

The updated Related Ticket activity notification appears in the Activity Notification List when you redisplay the list.

Priority Calculation

Priority calculation is a predefined set of values that automatically set Priority, Urgency, and Impact fields on problems and incidents. Priority calculation helps you manage incidents and problems for your business needs and IT capabilities. ITIL recommends that you prioritize tickets by using a data calculation that is based on Urgency and Impact values. Support organizations define this calculation based on their unique processes, how this calculation determines Service Level Agreements (SLAs), and other key events in the system. This calculation can also include the criticality of the CI that is linked to the incident and problem. Prioritizing tickets effectively helps you accomplish the following:

- Allocate IT resources for tickets
- Better serve customers
- Reduce costs

The CA SDM solution for Priority Calculation includes the following components:

- A Priority Calculation Matrix based on the Urgency and Impact values
- Default values for Urgency and Impact
- Urgency and Impact adjustment options based on Affected Service, Open Date, Affected End User, and Incident or Problem Area
- A Capture Reason option for manually modifying Urgency or Impact
- An Enable for Templates option for creating a ticket from a Template

When you install CA SDM, a Default priority calculation automatically manages ticket values. You can modify the Default priority calculation settings, or create additional priority calculations to manage incidents or problems. In the priority calculation, you define the outcome based on business scenarios to make the level of importance and ticket handling more consistent. Users can override some settings, but they cannot set the Priority on the ticket because this value is data-driven. For multi-tenancy, you or the tenants can create additional priority calculations with specific settings for each tenant.

When an analyst opens an incident or problem, the system automatically uses an active priority calculation and ticket values to generate Priority, Urgency, and Impact settings. The settings are based on one or more of the following fields:

- Urgency
- Impact
- Affected End User
- Incident or problem Area
- Open Date
- Affected Service

Analysts can override Urgency and Impact values as necessary. Depending on how you configure Options Manager, employees can only override incident Urgency values when the *urgency_on_employee* option is installed. When the Capture Reason flag is enabled and users override Urgency or Impact values and click Save, the Escalate Detail page appears to let users describe a reason for the change.

All ticket priority calculations, manual overrides, and reason information appear in the New Activity Log. If no priority calculation adjustments occur, the system does not create an activity log entry.

Note: If you migrated from a previous version, priority calculation is disabled by default. For information about how to enable priority calculation or retain your customizations, see the *Implementation Guide*.

How Priority Calculation Manages Ticket Values

The system adjusts problem and incident values based on active priority calculation settings to assist Analysts in handling tickets more effectively. The following table summarizes how priority calculation changes fields based on the priority calculation and user actions for problems, incidents, web service, email, and the Text API:

Action	Automatic Field Changes	Description
User changes Affected Service	Impact Priority	The system evaluates the Service Impact value on the CI of type Service to calculate the new Impact value. The CIs of type Service are defined as CIs with their class defined in the Enterprise Service family. If the open date of the ticket is within the blackout window time frame, the system increments a new Impact value based on the Impact Increment field. The system only replaces the Impact value when the new value is greater than the initial Impact value.
User changes Incident Area	Urgency	The Urgency value changes only when the new value is greater than the default value.
User changes Incident Area <i>and</i> Affected End User	Urgency Priority	If the user sets the Incident Area field first, the Urgency value changes after the user sets the Affected End User. The priority calculation sets the Priority.

Action	Automatic Field Changes	Description
User changes Urgency <i>and</i> Impact	Priority Impact	The system evaluates the Service Impact value on the CI of type Service to calculate the new Impact value. If the open date of the ticket is within a time defined by a blackout window, the system increments a new Impact value based on the Impact Increment field. The system only replaces the Impact value when the new value is greater than the initial Impact value. If the Administrator sets Capture Reason, the user must provide a reason for the modification. If the user changes the Urgency or Impact values, these values remain the same throughout entire ticket creation or ticket update unless the user modifies them again. However, the system can update the overridden values for Urgency or Impact the next time the user updates the ticket. After the system adjusts Urgency and Impact, the priority calculation sets the Priority value.
User selects New Incident based on the Knowledge Document and system has overrides for Knowledge Documents (see page 285)	Impact Urgency	The system always uses the Knowledge Document or knowledge solution values irrespective of whether the values are greater than or lesser than the default priority calculation values. For example, if a priority calculation has an Impact value of 3-Single Group and Urgency value of 3-Quickly, and Knowledge Documents have an Impact value of 2-Multiple Groups and Urgency value of 4-Very Quickly, the system applies the values from the Knowledge Document to the incident. The priority value always derives from the priority calculation.
User accepts Knowledge Document as solution to problem or incident	Impact Urgency	The system uses the values from the Knowledge Document for Impact and Urgency. The system also uses the priority calculation to set the Priority value.
User derives incident from Knowledge Document by selecting New Incident	Impact Urgency Priority	The system uses the priority calculation values.
User derives incident from Knowledge Document without system overrides for Knowledge documents (see page 285)	Impact Urgency Priority	The system uses the priority calculation values regardless of how the user created the incident for the Knowledge Document or knowledge solution.

Action	Automatic Field Changes	Description
Ticket accepts Knowledge Document as solution to problem or incident and the system does not override for knowledge documents	Impact Urgency Priority	The system uses values from problem or incident. The Priority value originates from the priority calculation.
Ticket that is in Read Only Mode accepts Knowledge Document as solution to the problem or incident and system overrides for Knowledge Documents	Impact Urgency Priority	The system uses the Impact and Urgency values from the Knowledge Document when the values are not empty. If the Impact/Urgency value in the Knowledge Document is empty the system uses values from the problem or incident. The Priority value originates from the priority calculation.
Ticket that is in Edit Mode accepts Knowledge Document as solution to the problem or incident and system overrides for knowledge documents.	Impact Urgency Priority	The system uses values from the problem or incident. The Priority value originates from the priority calculation.

Priority Calculation Generates Urgency Value After Saving Self-Service Tickets

By design, priority calculation generates Urgency values only after self-service users save incidents. Self-service users, such as VIP employees, employees, and anonymous users, can view the generated value after saving a ticket.

For self-service users, priority calculation uses the following settings and values to generate Urgency values:

- Urgency_On_Employee is set to Yes in Options Manager
- Override Urgency value is enabled in the active priority calculation for incidents
- *Web.cfg* Urgency settings such as AnonymousUrg for anonymous users, ESCEmpUrg for VIP employees, and EmpUrg for all other employees
- Area Urgency values
- Manual user overrides

The following table summarizes how priority calculation sets the Urgency value for self-service incidents:

Self-Service User Action	Urgency Value
The user saves an incident with the default Urgency and an empty Incident Area.	The ticket shows the default Urgency value from the <i>web.cfg</i> .

Self-Service User Action	Urgency Value
The user saves an incident after overriding the Urgency value.	Regardless of the Area Urgency, <i>web.cfg</i> , or priority calculation settings, the ticket shows the Urgency value that the user selected.
The user saves an incident after selecting an Incident Area. The Incident Area does not have a predefined Area Urgency value.	The ticket shows the default Urgency value from the <i>web.cfg</i> .
The user saves an incident after selecting an Incident Area that has a predefined Area Urgency value. The Override Urgency option is also enabled on the active priority calculation for incidents.	If the Area Urgency value is greater than the Urgency in the <i>web.cfg</i> , the ticket shows the Area Urgency value. However, the updated Urgency field is not visible while the user is creating or editing the ticket. When the user saves and reopens the incident, the updated Urgency value appears on the incident.
The user saves an incident after selecting an Incident Area that has a predefined Area Urgency value. However, the Override Urgency option is disabled on the active priority calculation for incidents.	The ticket shows an Urgency value from the <i>web.cfg</i> .
The user edits an existing incident that has an Incident Area with a predefined Area Urgency value.	The Urgency drop-down list shows the Area Urgency value and all applicable <i>web.cfg</i> values.

Note: For information about setting Urgency values for self-service users, see the *Implementation Guide*.

How to Set Priority Calculation

By default, ticket values such as priority are based on a priority calculation. You can find and adjust the initial values for Priority and Urgency in the web.cfg. The web.cfg has separate settings for various users such as guest, VIP-user, and employee.

Note: If you migrated from a previous version, priority calculation is disabled by default. Customized Incident and Problem Detail pages required additional configuration to operate properly. For information about how to enable priority calculation or retain your customizations, see the *Implementation Guide*.

To set priority calculation, do the following:

1. On the Administration tab, select Service Desk, Request/Incidents/Problems, Priority Calculation.
The Priority Calculation List appears.
2. Right-click the default priority calculation and select Edit from the short-cut menu.
The default Priority Calculation Detail page shows default settings for incident and problem tickets.
3. Review default priority calculation and adjust the values accordingly. When you set priority calculation values, consider the following issues for your working environment:
 - **Issue escalation**—When tickets require escalation to a particular VIP, you can increase the value for Urgency.
 - **Critical CIs**—For critical CIs, you can configure Service Impact for each CI.
 - **Critical Service Uptime**—When CIs require high availability, add a blackout window.
 - **Blackout window**—When CI-related tickets use a particular blackout window, you can increase the Service Impact value on the priority calculation.
4. Use the Manual Override setting to allow users to change tickets settings as necessary.
5. If you want a separate priority calculation to manage, problems or incidents, [configure the ticket type](#) (see page 284).
6. Click Save.

On the next new or updated ticket or Knowledge Document, the fields update according to the values in the active priority calculation.

7. Consider creating additional priority calculations for each ticket type. For multi-tenancy, create and activate additional priority calculations to manage tickets for each tenant.

Note: For information about priority calculation, see the *Online Help*.

Multiple Priority Calculations

You can set up more than one priority calculation. However, only one active priority calculation handles problems, or incidents, or both. For example, you can have an active priority calculation for problems, and another for incidents. You can also save several inactive priority calculations for future use.

Priority Calculation Assignment for Multi-Tenancy

You or your tenants can create tenant-specific priority calculations to manage incidents and problems. When you assign priority calculations for multi-tenancy, consider the following:

- When a priority calculation has no assigned tenant, it is considered public. The status of a public priority calculation is either Active or Inactive. A priority calculation is no longer considered public when it is assigned to a tenant.
- If a tenant has no priority calculation assignment, the Default priority calculation or another active public priority calculation automatically manages problems and incidents.
- A priority calculation manages problems and incidents for one tenant. However, a separate tenant-specific priority calculation can handle each ticket type. For example, Company X has one priority calculation to handle incidents and another to manage problems.
- When tenants create their own priority calculations while public priority calculations are active, the tenant-specific priority calculation applies only to tickets for the respective tenant.

For example, if the Default priority calculation is active, Company X tenant can create a tenant-specific priority calculation named `new_priority_calculation`. The `new_priority_calculation` settings and configurations apply only to Company X incidents and problems.

- If the tenant inactivates a priority calculation, the system uses an active public priority calculation to manage tenant problems and incidents.

For example, Company X inactivates the tenant-specific priority calculation while there is still an active Default priority calculation. Priority calculation remains enabled for Company X because the system uses the Default priority calculation to manage incidents and problems for Company X.

Note: Because tenants can delete their own priority calculation records, we recommend that you inactivate the public priority calculations that manage incidents and problems. Instead, you or the tenants can create tenant-specific priority calculations.

- When you disable multi-tenancy and there is more than one active priority calculation that manages tenants, leave *only one* priority calculation to manage incidents and problems. For example, you can inactivate all priority calculations except one to manage incidents and another to handle problems.

Note: For information about multi-tenancy, see the *Implementation Guide*.

How to Assign a Tenant to a Priority Calculation

For multi-tenancy, you can assign a tenant to a priority calculation. First, you inactivate public priority calculations. Then, you assign the tenant to a priority calculation and activate it.

To assign a tenant to a priority calculation, do the following:

1. On the Administration tab, select Service Desk, Request/Incidents/Problems, Priority Calculation.

The Priority Calculation List appears.

2. Edit each public priority calculation, such as Default. Set the status to Inactive and click Save.

The system disables the public priority calculations.

3. For each tenant, create or edit a priority calculation with tenant-specific settings for Impact, Urgency, and Priority.

The Create Priority Calculation or Update Priority Calculation page appears.

4. In the Name field, specify the tenant.

5. In the Status field, select Active.

6. Click Save.

The system applies tenant-specific values for Impact, Urgency, and Priority on new incidents and problems.

Note: For information about creating and editing priority calculations, see the *Online Help*.

Ticket Type Considerations for Priority Calculation

When you configure ticket types for a priority calculation, consider the following:

- The default priority calculation lets you manage both incident and problem ticket types.
- If you are migrating from a previous release, you enable the default priority calculation or create a priority calculation to manage problems and incidents.
- Although you can have many priority calculations, only one active priority calculation can handle a particular ticket type. For example, one active priority calculation can manage problems and another can manage incidents.
- If you want to create a priority calculation and an active priority calculation already handles a particular ticket type, you first disable the ticket type on the active priority calculation. For example, if you want a priority calculation to manage problems, you disable the problem ticket type on the active priority calculation and create an active priority calculation that manages problems.

Note: For information about enabling priority calculation and setting ticket types during migration, see the *Implementation Guide*.

Configure Ticket Types for a Priority Calculation

You can specify the ticket types that the priority calculation manages. When the priority calculation is active, it manages Priority, Impact, and Urgency values on new tickets.

To configure ticket types for a priority calculation

1. On the Administration tab, select Service Desk, Request/Incidents/Problems, Priority Calculation.

The Priority Calculation List appears.

2. Right-click a priority calculation and select Edit.

The Update Priority Calculation page appears.

3. Select or clear one or more of the following options:

Incidents

Enables or disables the priority calculation to manage new incidents.

Problems

Enables or disables this priority calculation to manage new problem tickets.

4. Click Save.

CA SDM uses the settings in the priority calculation to manage ticket values for new incidents, problems or both.

Use Ticket Templates to Calculate Priority Values

If you want ticket templates to calculate priority, you configure the priority calculation with the *Enable for Templates* option. If you enable this option the Urgency and Impact values come from the template, but the priority field comes from priority calculation record. The priority value displays as read-only.

If you do not enable this option, the Urgency, Impact and Priority fields come from the template. You can edit the priority field and no priority calculation is done for the ticket until it is saved.

To use ticket templates to calculate priority values

1. On the Administration tab, select Service Desk, Requests/Incidents/Problems, Priority Calculation.

The Priority Calculation List appears.

2. Select the priority calculation that you want to use for calculating priority with templates.

You can also create a priority calculation to use ticket templates to calculate priority values.

3. Select *Enable for Templates* from the Priority Calculation Options list.

4. Click Save.

The option is enabled.

Use Knowledge Documents to Calculate Priority Values

If you want knowledge documents to calculate priority, you can update the field mapping. After you configure the field mapping, the analyst can create incident or problem tickets from knowledge documents. The knowledge document calculates Impact and Urgency values on the tickets. If the Impact or Urgency value is missing, the values originate from the priority calculation.

Important! If you modify and save a ticket that already contains Impact and Urgency values calculated by a knowledge document, the priority calculation overrides the values set by Knowledge Management. The audit log displays these activities.

To use knowledge documents to calculate priority values

1. On the Administration tab, select Knowledge, Service Desk Integration.
2. Select Field Mapping.

The Field Mapping page appears.

3. For Impact and Urgency, select the following check boxes as appropriate:

Populate Empty Service Desk values

Specifies whether to use information from Knowledge Management to populate fields in issues or requests.

Override Service Desk values

Identifies the fields in issues or requests that correspond to fields listed in the Knowledge Management column.

Note: When the Override Service Desk values field is not enabled but Populate Empty Service Desk values field is enabled for Impact and Urgency, the knowledge values for Impact and Urgency override the Incident values.

4. Click Save.

Incidents and problems are created, using the Impact and Urgency values from the knowledge document to calculate the Priority value. If the values are missing, the ticket obtains the values from an active priority calculation. If no priority calculation is active for the ticket type, the system clears the Priority, Urgency, and Impact fields.

Manually Override the Impact Value

When you manually override the Impact value on a problem or incident, the active priority calculation that manages the ticket type automatically adjusts the Priority value.

To manually override the Impact value

1. Open the details page for the problem or incident you want to change.
2. Change the Impact value.

If there is an active priority calculation that manages the ticket type, the Priority value automatically changes based on the settings in the priority calculation.
3. Save the incident.

The Activity Log on the Incident Detail page reflects the Impact values changes.

Example: Manually override the Impact value on a new incident

1. Create an incident.

By default, the Urgency value is 3-Quickly. The Impact value is 3-Single Group. The Priority value is 3.
2. Override the Impact value to 1-Entire Organization.

The Priority value automatically changes based on the values in the active priority calculation that manages incidents.
3. Save the incident.

The Activity Log on the Incident Detail page reflects the Impact value changes.

Manually Override the Urgency Value

When you manually override the Urgency value on a ticket, the active priority calculation that manages the ticket type automatically adjusts the Priority value.

To manually override the Urgency value

1. Open the details page for the incident you want to change.
2. Change the Urgency value.

If there is an active priority calculation that manages the ticket type, the Priority value automatically changes based on the settings in the priority calculation.

3. Save the incident.

The Activity Log on the Incident Detail page reflects the Urgency values changes.

Example: Manually Override the Urgency Value on a New Incident

1. Create an incident.

By default, the Urgency value is 3-Quickly. The Impact value is 3-Single Group. The Priority value is 3.

2. Override the Urgency value to 5-Immediate.

The Priority value automatically changes based on the values in the active priority calculation that manages incidents.

3. Save the incident.

The Activity Log on the Incident Detail page reflects the change in the Urgency value.

Automatically Adjust Impact for a Problem or Incident

For Configuration Items defined with a family of Enterprise Service, you can automatically adjust the Impact value for problems or incidents. When you select the Problem or Incident Area and select an Affected Service, the impact adjusts according to the CI Service Impact settings for Enterprise Service CI's and the priority calculation.

To automatically adjust Impact for a Problem or Incident

1. Create a problem or incident for an Enterprise Service type CI.
2. Select an Affected Service.
3. Select a Problem or Incident Area.

If there is an active priority calculation that manages the ticket type, the Impact value changes based on the Increment Impact value (used for Blackout Window impact assessment) in the priority calculation and the Service Impact value from the affected service.

If you are using the default priority calculation, with a Service Impact for the Enterprise Service CI set to 1-Entire Organization, and the Problem or Incident is not opened within a Blackout Window the Impact value in the Problem or Incident is set to 1 and the Priority value on the ticket raises to a 2.

4. Save the ticket.

The Activity Log on the Incident Detail page reflects the Impact value's changes.

Adjust Impact for a Problem or Incident Example

The following example shows you how to adjust impact for a Problem or Incident.

1. Create an Enterprise Service CI named CI-APC and set the class as one that comes under the family Enterprise Service.

For example, you can set the class as Other Service, Business Services, or Infrastructure Service.

2. In the Service tab within the CI Detail page, set the Service Impact field to 2-Multiple Groups.
3. On the Service Desk tab, create an incident and set the Affected Service to CI-APC.
4. Save the ticket.

The Impact field in the incident reflects the value from the Service Impact of the selected affected service (CI-APC). In this case, the Impact value is set to 2 -Multiple Groups.

Note: If you are using the default priority calculation and the current ticket is created during a Blackout Window period then the Impact value increments by 1 and in the case above the Impact value is then set to 1-Entire organization.

Note: For information about creating CIs, see the Online Help.

Automatically Adjust Urgency for a Problem or Incident

For problems and incidents, you can automatically adjust Urgency and Priority by specifying an affected end user that requires Special Handling or by specifying a Problem/Incident Area that has an Area Urgency value.

When you assign special handling, with the Escalate Urgency turned on, for a contact or set an Area Urgency value for a Problem/Incident area, the Urgency value within the Problem/Incident automatically adjusts according to the values in the priority calculation and the Area Urgency value for the affected end user.

To automatically adjust Urgency for a Problem or Incident

1. Create a problem or incident.
2. Select an Affected End user. For an elevated urgency, select an Affected End user that requires Special Handling that has the Escalate Urgency on.

If there is an active priority calculation that manages the ticket type, the Urgency value changes based on the Urgency Increment value in the active priority calculation.

3. Select a Problem or Incident Area.

If there is an active priority calculation that manages the ticket type, the Urgency value changes based on the Area Urgency value in the Problem/Incident Area definition.

4. Save the ticket.

A confirmation message reminds you that the ticket requires special handling. The Activity Log on the Problem/Incident Detail page reflects the changes of the Urgency value.

Adjust Urgency for a Problem or Incident Example

The following example shows you how to adjust urgency for a Problem or Incident:

1. On the Administrator tab, create a contact named Non-VIP.
2. Create a special handling contact named VIP and set the Escalate Urgency value on.
3. Create an area named Test Area and specify Area Urgency as 2-Very Quickly.
4. On the Service Desk tab, create an incident.
5. For the Affected End User, select Non-VIP.
6. In the Incident Area, select Test Area and save the ticket.

The Urgency field reflects the Area Urgency value from the Incident Area definition. In this case, the Urgency is set to 2-Very Quickly.

7. Change the Affected End User to VIP and save the ticket.

If the default Priority Calculation matrix is being used, the Urgency value is incremented by 1 and is set to 1-Immediate. A confirmation message appears that reminds you that the ticket requires special handling. The Activity Log on the Incident Detail page reflects the Urgency value's changes.

Note: For information about creating contacts and areas, see the *Online Help*.

Status Transitions and Dependent Attribute Controls

You can use the following configurable controls to restrict ticket status flows for change orders, issues, incidents/problems/requests, and determine which fields are shown or required for each ticket status:

Transitions

Controls how users select available statuses on the incident/problem/request, issue, or change order form. For example, a problem is in a status of Open, and the transition flow only allows the analyst to update the status to Closed. In this example, the analyst has no other status options, which reinforces the problem management process.

Transitions let you define a subset of the full status list and specify the default new (or next) status of a ticket based on the current status. You can define unique status transitions for each ticket type. Consider using transitions when you want to restrict status workflows for your end users.

Dependent Attributes

Controls how attributes are designated as required (must supply) or locked (cannot update) depending on ticket status. For example, the Change Manager can prevent an analyst from editing the Summary attribute after a change order is approved. Consider using attribute controls when you want to restrict certain attributes based on the status.

Note: You can specify how strictly the system enforces Status policies by configuring the Status Policy Violations option in Options Manager (General Options). This option only applies to automated system processes, such as integrations and macros.

More information:

[Work with Status Transitions and Dependent Attribute Controls](#) (see page 292)

[Configure Status Transitions](#) (see page 292)

[Best Practice: Predefined Status Transitions](#) (see page 299)

Work with Status Transitions and Dependent Attribute Controls

To work with status transitions and dependent attribute controls, do the following:

1. On the Administration tab, configure the appropriate tenants, contacts, and roles for your environment.
2. On the Service Desk node, specify a ticket type (Change Order, for example), and select Status.

The Status List page displays active status codes.

3. Edit the appropriate status code (Acknowledged, for example) and use the controls available on the tabs at the bottom of the ticket's status detail page to do the following:
 - [Configure Status Transitions](#) (see page 292)
 - [Configure Dependent Attribute Controls](#) (see page 294)

Note: You can configure unique transitions and dependent attribute controls for each ticket type.

Configure Status Transitions

You can configure a subset of the full status list and specify the default next status of a ticket based on the current status. Transitions are enforced when the status is changed on the ticket detail form.

To configure status transitions

1. Click the Transitions tab at the bottom of the ticket's status detail page.

The Transitions List page shows all valid transitions for the status.

Note: When configured, linked transition types appear on the Incident and Request Transition list.

2. Click the Update Transitions button.

The Update Ticket Status Transitions page appears.

3. Configure the following check boxes as appropriate:

Allowed

Specifies a valid transition for the status. Use this option to restrict status workflows.

Default

(Optional) Specifies the default status transition. CA SDM applies the default transition when a user clicks the default transition button on the ticket detail form, or when a user (including a web services user) updates the status to a <d> value. There is only one default transition for each status (or one for each tenant in a multi-tenanted system).

Must Comment

(Optional) Specifies that an activity log comment for the transition is required when changing the status on a ticket.

Note: This option applies to CA SDM tickets only. It does not apply to other areas, such as web services or the edit in list functionality.

4. (Optional) Select a status code in the Name column to update its details.
5. Click Save.

The list of transitions configured for the new status appears on the Transition list. When the analyst selects the Status drop-down on the ticket form, the new status list appears.

Note: For detailed procedural information about defining status transitions, see the *Online Help*.

Configure Dependent Attribute Controls

You can determine which fields are shown or required for the ticket status.

Note: Before you configure dependent attributes as "required" for the ticket status, be aware that the Edit in List option that appears on the ticket's list page may not display the attribute field values that are required. If the required attribute field value is not already part of the saved ticket, and if it is not presented in the Edit In List format, then the ticket is not saved. Consequently, the analyst must edit the required dependent attribute field values on the ticket's detail page instead of using the Edit in List option.

To configure a dependent attribute control

1. On the ticket's status detail page, select the Dependent Attribute Control tab at the bottom of the page.

The Attribute Control List appears.

2. Click Create.

The Update Status Dependent Attribute Control page appears.

3. Complete the following fields:

Tenant

(Optional) In a system with multi-tenancy installed, specifies an optional tenant name. If a tenant is specified, the dependency applies only to that tenant and to its sub-tenants.

Attribute

Specifies the name of the attribute you want to control, for example, Summary.

Locked

Specifies the attribute as locked. A locked attribute associated with a status cannot be updated in a ticket with the same status. The attribute is unlocked when the status is changed.

Required

Specifies the attribute as required. A required attribute for the status cannot use a null value in a ticket with the same status.

4. Click Save.

The new attribute control for the status appears in the Attribute Controls List when you redisplay the page. When a user updates the ticket status, the system retrieves the list of required attributes corresponding to the new status, and updates the ticket form as appropriate. An error message appears at the top of the ticket form when a user attempts to close the ticket without filling in a required field.

Note: For detailed procedural information about defining dependent attribute controls, see the *Online Help*.

Web Services Methods

You can configure the following status transition and dependent attribute control SOAP web services methods:

getValidTransitions

Lists the transitions for a ticket. This method is modeled on the existing `getValidTaskTransitions` method, except that the argument can be a ticket or a status.

getDependentAttrControls

Lists the locked and required attributes for an attribute of an object that persistent id specified. The Status attribute is supported at this time.

Note: For more information about SOAP web services methods, see the *Technical Reference Guide*.

Predefined Transition Flows

For each ticket type, you can use the predefined status transitions provided with the product and modify them to accommodate your desired transition flow.

Note: Since status transitions can be shared between integrations such as CA Workflow, do not inactivate predefined status transitions unless explicitly requested.

To view the list of predefined transitions, do the following:

On the Administration tab, expand the Service Desk node, and select one of the following:

- Change Order Transitions
- Issue Transitions
- Request/Incident/Problem Transitions

The Transitions List displays the predefined transitions that let you control how a ticket (incident/request/problem, change order, and issue) continues through its lifecycle.

Note: For detailed procedural information about defining and modifying transitions, see the *Online Help*.

Incident Transition Flow

The following table shows the predefined Incident transition flow.

Current Status	Default Transition	Available Next Statuses
Acknowledged	In Progress <d>	Avoided, Awaiting Vendor, Cancelled, Closed, Closed Unresolved, In Progress, Open, Pending Change, Resolved
Avoided		Acknowledged, Avoided, Awaiting End User Response, Closed, Closed Unresolved, In Progress, Reject Solution, Researching, Resolved
Awaiting End User Response	Researching <d>	Closed, Closed Unresolved, In Progress, Open, Researching, Resolved
Awaiting Vendor	Researching <d>	Acknowledged, Closed, Closed Unresolved, In Progress, Open, Pending Change, Researching, Resolved
Cancelled		Closed
Closed		Open
Closed Unresolved		Acknowledged
Closed Unresolved		Closed
Closed Unresolved		Open
Hold		Acknowledged, Closed, In Progress, Open, Pending Change, Resolved
In Progress	Researching <d>	Acknowledged, Awaiting End User Response, Awaiting Vendor, Closed, Closed Unresolved, Open, Pending Change, Researching, Resolved
Pending Change	Researching <d>	Acknowledged, Closed, In Progress, Open, Resolved
Researching	Resolved <d>	Closed, Open, Resolved
Resolved	Closed <d>	Awaiting End User Response, Closed, Open

Problem Transition Flow

The following table shows the predefined Problem transition flow.

Current Status	Default Transition	Available Next Statuses
Acknowledged	In Progress <d>	Acknowledged, Approved, Cancelled, Closed, Fix in Progress, Open, Rejected, Researching
Analysis Complete	Approved <d>	Acknowledged, Cancelled, Closed, Fix in Progress
Approved	Fix in Progress <d>	Closed, Fixed, Pending Change

Current Status	Default Transition	Available Next Statuses
Awaiting Vendor	Researching <d>	Acknowledged, Closed, Closed Unresolved, Fixed, In Progress, Open, Pending Change, Researching
Cancelled	Closed <d>	Closed, Closed Unresolved, Open
Closed Unresolved		Acknowledged, Closed, Open
Fix in Progress	Fixed <d>	Approved, Cancelled, Fixed, Fix in Progress, Researching, Rejected
Fixed	Closed <d>	Closed
Hold	Researching <d>	Acknowledged, Closed, Fixed, In Progress, Open, Pending Change, Researching
In Progress	Researching <d>	Acknowledged, Approved, Cancelled, Closed, Fix in Progress, Pending Change, Rejected, Researching
Known Error	Fix in Progress <d>	Closed, Fix in Progress, Fixed
Open	Acknowledged <d>	Acknowledged, Approved, Cancelled, Closed, Fix in Progress, In Progress, Rejected, Researching
Pending Change	Fixed <d>	Closed, Fixed, Researching
Rejected	Closed <d>	Closed, Closed Unresolved, Open
Researching	Analysis Complete <d>	Acknowledged, Analysis Complete, Approved, Cancelled, Closed, fix in Progress, Fixed, Rejected

Issue Transition Flow

The following table shows the predefined Issue transition flow.

Current Status	Default Transition	Available Next Statuses
Acknowledged	In Progress <d>	Awaiting End User Response, Awaiting Vendor, Closed, Closed Unresolved, In Progress, Open, Pending Change, Resolved
Awaiting End User Response	Researching <d>	Acknowledged, Closed, Closed Unresolved, In Progress, Open, Researching, Resolved
Awaiting Vendor	Researching <d>	Acknowledged, Closed, In Progress, Open, Pending Change, Researching, Resolved
Cancelled	Closed <d>	Closed
Closed		Acknowledged, Open
Closed Unresolved		Acknowledged, Closed, Open
Hold		Acknowledged, Closed, In Progress, Open, Pending Change, Resolved

Current Status	Default Transition	Available Next Statuses
In Progress	Researching <d>	Acknowledged, Awaiting End User Response, Awaiting Vendor, Closed, Closed Unresolved, Open, Pending Change, Researching, Resolved
Open	Acknowledged <d>	Acknowledged, Avoided by Self Service, Awaiting End User Response, Awaiting Vendor, Closed, Closed Unresolved, In Progress, Pending Change, Resolved
Pending Change	Researching <d>	Acknowledged, Closed, In Progress, Open, Researching, Resolved
Researching	Resolved <d>	Closed, Open, Resolved, Awaiting End User Response, Closed, Open

Change Order Transition Flow

The following table shows the predefined Change Order transition flow.

Current Status	Default Transition	Available Next Statuses
Approval in Progress	Approved <d>	Approved, Cancelled, Closed
Approved	Scheduled <d>	Cancelled, Closed, Implementation in Progress, Scheduled
Backed Out		Approval in Progress, Closed, Open, RFC
Cancelled		Closed
Customer Hold		Cancelled, Closed, Implementation in Progress, Rejected, Scheduled, Verification in Progress
Hold		Cancelled, Closed, Implementation in Progress, Scheduled
Implementation in Progress		Backed Out, Cancelled, Closed, Customer Hold, Rejected, Scheduled, Vendor Hold, Verification in Progress
Open	RFC <d>	Approval in Progress, Cancelled, Closed, Customer Hold, Implementation in Progress, Rejected, RFC, Scheduled, Vendor Hold
Rejected	Closed <d>	Approval in Progress, Cancelled, Closed
RFC	Approval in Progress <d>	Approval in Progress, Cancelled, Closed, Customer Hold, Implementation in Progress, Open, Rejected, Scheduled, Vendor Hold
Scheduled	Implementation in Progress <d>	Cancelled, Closed, Customer Hold, Implementation in Progress, Vendor Hold, Verification in Progress, Cancelled, Closed, Implementation in Progress, Scheduled, Backed Out, Closed

Best Practice: Predefined Status Transitions

The predefined status transitions delivered with the product are Active in a new installation and Inactive after the upgrade. For every status listed on the Transitions List page, there is a default status transition (or next status). The path taken by the default status transition reflects the best practice. The additional status transitions listed on the Transitions List page are provided to fulfill various ticket management workflows. However, only Active status transitions that use this best practice can ensure that the proper workflow for managing Requests, Incidents, Problems, and Change Orders occurs. This best practice helps promote the movement of tickets to resolution and closure within the IT environment.

For example, the following predefined incident transitions listed on the Incident Transition list page are set to Inactive to help promote the resolution and closure of incidents:

Status	New Status	Default	Status Description	Record Status
Acknowledged	Closed	No	Acknowledged to Closed Status Transition	Inactive
Acknowledged	Closed Unresolved	No	Acknowledged to Close Unresolved Status Transition	Inactive
Acknowledged	Open	Yes	Acknowledged to Open Status Transition	Inactive
Awaiting End User Response	Acknowledged	No	Awaiting End User Response to Acknowledged Status Transition	Inactive
Awaiting End User Response	Open	No	Awaiting End User Response to Open Status Transition	Inactive
Awaiting Vendor	Acknowledged	No	Awaiting Vendor to Acknowledged Status Transition	Inactive
Awaiting Vendor	Closed	No	Awaiting Vendor to Closed Status Transition	Inactive
Awaiting Vendor	Open	No	Awaiting Vendor to Open Status Transition	Inactive
Closed	Acknowledged	No	Closed to Acknowledged Status Transition	Inactive
Closed Unresolved	Acknowledged	No	Closed Unresolved to Acknowledged Status Transition	Inactive
Closed Unresolved	Closed	No	Closed Unresolved to Closed Status Transition	Inactive

Status	New Status	Default	Status Description	Record Status
Hold	Acknowledged	No	Hold to Acknowledged Status Transition	Inactive
Hold	Closed	No	Hold to Closed Status Transition	Inactive
Hold	Open	No	Hold to Open Status Transition	Inactive
In Progress	Acknowledged	No	In Progress to Acknowledged Status Transition	Inactive
In Progress	Closed	No	In Progress to Closed Status Transition	Inactive
In Progress	Open	No	In Progress to Open Status Transition	Inactive
Open	Closed	No	Open to Closed Status Transition	Inactive
Pending Change	Acknowledged	No	Pending Change to Acknowledged Status Transition	Inactive
Pending Change	Closed	No	Pending Change to Closed Status Transition	Inactive
Pending Change	Open	No	Pending Change to Open Status Transition	Inactive
Researching	Open	No	Researching to Open Status Transition	Inactive

Status Transitions for Self-Service

Status transitions let you control the movement of a ticket from one discrete state to another (for example, from Open to Closed). For employees using self-service, you can include buttons on the Incident and Request detail forms to represent any [status transition](#) (see page 291).

Status transition buttons for incident and request process workflows appear in the employee interface when incident or request transitions are linked to active transition types. A transition type defines the button text and controls the behavior of the ticket detail form. When buttons are defined, the legacy Close Incident (or Request) and Reopen Incident (or Request) buttons are not displayed on the ticket detail forms. Instead, the employee can only change the status of the Incident or Request using the status transition buttons configured by the administrator.

By default, all predefined transition types delivered with the product are inactive, so status transition buttons are not in effect. As a system administrator, you can activate and modify predefined transition types or create transition types to accommodate your status transition workflows. After you create or modify a transition type, you can link them to any incident or request status transition.

Note: For more information about transition types, see the *Online Help*.

More information:

[How Transitions for Self-Service Work](#) (see page 301)

[How to Create or Update Transition Types for Transitions](#) (see page 302)

[How to Link Transition Types to Transitions](#) (see page 302)

[Activate Predefined Transition Types](#) (see page 303)

How Transitions for Self-Service Work

Transition types and their corresponding statuses control when employees can close and reopen tickets as follows:

1. Active transition types are linked to incident (or request) status transitions by the administrator.
2. The employee creates an incident using self-service.
3. The analyst assigned to the incident finds a solution and moves the ticket to the Resolved status.
4. When the ticket is in a Resolved status, the employee detail form displays status transition buttons to Accept or Reject the resolution.
 - If the employee accepts the resolution, the Resolved to Closed transition occurs.
 - If the employee rejects the resolution, the Resolved to Open transition occurs.
5. After the employee clicks a button, they can add their remarks in the resolution form that appears.

How to Create or Update Transition Types for Transitions

As a system administrator, you can create new or update existing transition types for incident and request status transition workflows on the Transition Types List page.

To create a transition type for a status transition, do the following:

1. On the Administration tab, select Service Desk, Requests/Incidents/Problems, Transition Types.
2. Click Create on the list page.
3. Edit the fields as appropriate on the detail page.

The transition type for the status transition is created.

4. Click Save.

The new transition type appears in the Transitions Type List when you redisplay the page.

To update a transition type, do the following:

1. Open the desired transition type for editing on the Transition Types List page.
2. Edit the fields as appropriate.
3. Click Save.

The updated transition type appears on the Transition Type list.

How to Link Transition Types to Transitions

When status transitions are linked to transition types, the employee ticket detail form displays status transition buttons to Accept or Reject the resolution. To link a transition type with a status transition, do the following:

1. On the Administration tab, select Service Desk, Requests/Incidents/Problems, Incident (or Request) Transitions.
2. Open the desired status transition for editing on the Request or Incident Transition List page.
3. Specify the desired transition type in the Transition Type field.
4. Click Save.

The transition type is linked to the status transition.

Activate Predefined Transition Types

By default, all predefined transition types that are delivered with the product are Inactive, so status transition buttons are not in effect. You can activate and modify these transition types to accommodate your desired status transition flow.

To activate a predefined transition type

1. Select Show Filter on the Transition Type List page.
The top portion of the page reveals additional search fields.
2. Select Inactive in the Record Status field and click Search.
The Transition Type List displays all inactive transition types.
3. Right-click the desired transition type and select Edit from the menu.
4. Select Active in the Record Status field.
5. Click Save, Close window.
6. Click Search.
The Transition Type List displays the active transition type.

More information:

[Predefined Transition Types for Incident Status Transitions](#) (see page 303)

[Predefined Transition Types for Request Status Transitions](#) (see page 304)

Predefined Transition Types for Incident Status Transitions

The following table describes the predefined transition types for incident status transitions:

Symbol	Button Text	Form Header Text	Incident Status Transition
Accept incident resolution button	Accept	Accept Resolution	Resolved to Closed
Reject incident resolution button	Reject	Reject Resolution	Resolved to Open

Symbol	Button Text	Form Header Text	Incident Status Transition
Reject Closure incident button	Request Closure	Request Closure	■ To Closed Unresolved from Open ■ Awaiting End User Response ■ Awaiting Vendor ■ In Progress ■ Acknowledged

Predefined Transition Types for Request Status Transitions

The following table describes the predefined transition types for request status transitions:

Note: The Closed Requested status for Requests is the equivalent of the Resolved status for Incidents.

Symbol	Button Text	Form Header Text	Request Status Transition
Accept request resolution button	Accept	Accept Resolution	Close Requested to Closed
Reject request resolution button	Reject	Reject Resolution	Close Requested to Open
Request Closure request button	Request Closure	Request Closure	■ To Canceled from Open ■ Awaiting End User Response ■ Awaiting Vendor ■ Approval In Progress ■ Acknowledged
Close request button	Close	Close Requested	From In Progress to Closed

Timers

Timers act as a stopwatch with various thresholds that give the analyst indications of elapsed time. You can define the amount of time the timer remains at each threshold, and have the timer change color, beep, or display a reminder as it reaches each threshold. A service desk analyst cannot control the stopwatch; only the administrator can control it.

Requests are the only type of ticket that uses timers; therefore, set up timers for internal and combined service desk models.

The following threshold values are predefined for the timer:

Threshold Duration	Color
00:00:00	Green
00:01:00	Yellow
00:05:00	Red

The predefined values start the timer at green. After one minute, the timer changes to yellow. After five minutes, the timer changes to red. The elapsed time is displayed when the analyst views the request in detail and is reset each time a new request is selected. You can add steps to this process or change existing steps.

You can mark any timer as active or inactive. When you mark a timer as inactive, it is no longer part of the process, but it remains available for future use (that is, it is not deleted from the database). If you decide later to use the timer, back to it and mark it as active.

Time Zones

You can set up specific time zones for servers, service types, contacts, and locations in your CA SDM system. You can set up local service types that apply to a specific time zone and global service types that apply across the entire enterprise. This setup eliminates the need for the administrator to know the time zone of a server and manually adjust the work shift times to fit different time zones.

More information:

[Service Type Event Triggers](#) (see page 306)

[Time Zone Event Triggers](#) (see page 306)

[Time Zone Rules](#) (see page 307)

Service Type Event Triggers

Service types define the events that are triggered after a specified delay time.

Example: Event Triggers by Work Shift Schedule

In this example, a work shift schedule associated with the service type constrains the actual time an event is triggered, as follows:

- Work shift is 8:00 a.m. to 5:00 p.m.
- Event delay is three hours
- Current server time is 3:00 p.m.

The event is triggered tomorrow at 9:00 a.m. server time because of the following:

- According to the current server time, two hours remain in the work shift for today (Current time is 3:00 p.m. Work shift ends at 5:00 p.m.)
- The event delay is three hours. Two hours of this delay are spent in the current day's work shift (3:00 p.m. to 5:00 p.m.). One additional delay hour is carried over to the next day's work shift (8:00 a.m. to 9:00 a.m.)

Consequently, the event begins at 9:00 a.m. the following day.

Time Zone Event Triggers

A time zone associated with the service type can constrain event trigger times.

Example: Event Triggers by Time Zone

In this example, a time zone associated with the service type constrains the actual time an event is triggered as follows:

- Work shift is 8:00 a.m. to 5:00 p.m.
- Event delay is three hours
- Current server time is 3:00 p.m.
- Current time in time zone is 12:00 p.m.

The event is triggered today at 6:00 p.m. server time because of the following:

- According to the current server time, two hours remain in the work shift for today (Current time is 3:00 p.m. Work shift ends at 5:00 p.m.).
- According to the current time zone time, five hours are left in the work shift for today (Current time is 12:00 p.m. Work shift ends at 5:00 p.m.).
- The event delay is three hours (12:00 p.m. to 3:00 p.m. time zone time. 3:00 p.m. to 6:00 p.m. server time).

Consequently, the event begins at 6:00 p.m. server time or 3:00 p.m. time zone time.

Time Zone Rules

You can specify a time zone for a server, for a location, and for a service type. You can also tell CA SDM to use the time zone of the Affected End User of a ticket. CA SDM uses the following rules to determine which time zone triggers an event:

- **Affected End User time zone**—CA SDM uses this rule when the following conditions exist:
 - The Use End User Time Zone option is selected.
 - A time zone is specified for the affected end user of the ticket.
- **Affected End User location time zone**—CA SDM uses this rule when the following conditions exist:
 - The Use End User Time Zone option is selected.
 - No time zone is specified for the affected end user of the ticket.
 - A time zone is specified for the affected end-user location.
- **Service type time zone**—CA SDM uses this rule when the following conditions exist:
 - A time zone is specified for the service type.
 - The Use End User Time Zone option is not selected.
- **Server time zone**—CA SDM uses this rule when the following conditions exist:
 - A time zone is specified for the server.
 - The Use End User Time Zone option is not selected.
 - No time zone is specified for the service type.
- **No time zone support**—CA SDM uses this rule when the following conditions exist:
 - No time zone is specified for the service type.
 - The Use End User Time Zone option is not selected.
 - No server records exist.

File Attachments

Often, you want to associate a file with a ticket; examples include the following files:

- A screenshot of a particular error.
- Log files that led to a problem.
- An email that explains a problem.
- A WebLink pointing to a web page that needs repair.

You set up your service desk so that users can work with attachments. All types of tickets can have attachments; therefore, set up attachments for all service desk models.

CA SDM classifies attachments as follows:

- **Stored**—Attachments that are stored are uploaded and stored in a repository. When an analyst reviews a stored attachment, the file is retrieved using a web browser and is viewed locally. The stored file remains in the repository.

Stored attachments can be uploaded to a repository using a web server that uses the HTTP protocol. Utilizing a web server for file storage allows storage and retrieval from the user interface. Shared drive repositories cannot be accessed effectively from a web client and are, therefore, not supported.
- **Linked**—Stores only a link to the file in the database.

More information:

[File Attachment Upload and Download](#) (see page 308)
[Repositories](#) (see page 309)

File Attachment Upload and Download

All client interfaces can access existing repositories for upload and download of file attachments except as follows:

- Shared file repositories can be accessed only if the repository daemon is running on a computer that can access the shared file. The server name on the repository record (detail form) must be a Windows computer that has access to the share. A <USDK repository daemon must also be running on the computer.

Note: Shared drive repositories are not supported on WIN2003.

- Download of zip files is based on when they were uploaded. Attachments from earlier releases are downloaded without unzipping them. Unzip the file on the client computer; otherwise, attachments uploaded from a client interface, are also downloaded without unzipping them. That is, the server unzips the file before returning it, during a client download request.

Note: Legacy clients continue to function the same way, and have ready access to new uploads to existing repositories.

Note: For information about how to work with attachments, see the *Online Help*.

Repositories

A repository represents a disk directory on a local or a remote computer. You can have one large repository or several small ones. How you decide to structure repositories is entirely up to your business practices and needs.

Moving or combining repositories is simple because all the attributes about a repository are defined in the repository record. Only the repository and the file name are associated with each ticket.

Files are uploaded to a repository and remain there. When a file is requested, one of the following actions occurs:

- A web browser (HTTP) retrieves and presents a copy of the file to the user.
- The application that is associated with the file extension opens the file automatically.

Example: Change the Directory Structure on a Web Server

This example shows how to change the directory structure on a web server.

Change the Upload Path in the repository record.

All existing attachments correctly point to the new location automatically.

Repository Setup

You need to set up the repositories before CA SDM users can work with attachments.

The distributed architecture lets a site configure their repositories according to their needs. The servlet for a repository does not have to reside on the same server as the attached files. Sites can have a central servlet to access all their distributed repositories or have a dedicated servlet for each of their repository servers.

Consider the following setups when configuring repositories:

- **Repository server on a protected CA SDM server**—Sites that designate a repository server on a protected CA SDM server (typically a secondary server behind a firewall) do not need to expose the servlet on that computer. Rather, they can specify a servlet running on the primary server or other unrestricted CA SDM server and still upload and download successfully to that server. Depending on the network between the two CA SDM servers, some performance impact on upload and download can occur.

Note: A good way to set up a remote repository is to install a secondary server with the repository on the remote server, and set the upload path to a local path.

- **Servlet on the same server as the repository server**—Sites that want optimal performance for attachment upload and download, and if network performance is an issue and if very large files are being attached, should consider this setup. This approach requires exposing the Tomcat port (typically 8080) on the server and should be noted if the computer is behind a firewall.

System Requirements

By default, the repositories are located on the primary server, and no extra configuration is required to upload file attachments. However, having the repository on a remote computer requires installing a secondary server with the rep_daemon running on the primary server. The system requirements for the server are as follows:

- Secondary server (for remote) with repository daemon running.
- Tomcat and Upload servlet

To set up a repository on a remote computer

1. Open a command prompt and enter:
`cd $NX_ROOT/samples/pdmconf`
2. Run `pdm_edit.pl`.
You are asked two questions. Answer those questions to return to the main menu.
3. Select 8 and press Enter to select Repository Daemon.
4. Enter the name of the secondary server and press Enter.

Note: The name of the secondary server must match the `NX_LOCAL_HOST` on the secondary server.

The main menu appears.

5. Select X and press Enter to Save and Exit.

The saved file is the `pdm_startup.rmt` file. This file is a replacement for `$NX_ROOT/pdmconf/pdm_startup.tpl`.

6. Rename your current pdmconf/pdm_startup.tpl file.
7. Copy your new pdm_startup.rmt file to pdmconf/pdm_startup.tpl.
8. Install a secondary server and start the proctor on the secondary server.
Note: For Windows NT systems, the proctor is a service and on Linux it is pdm_proctor_nxd.
9. Start the primary server and use pdm_status to see if your new daemon appears.
10. Change the Server Name and the Upload path to point to the remote computer where the repository daemon runs.

The remote computer is set up to upload file attachments.

Define the Repository

You must define a repository using the CA SDM administrative function of the web interface to use file attachments.

To define the repository

1. Select Attachments, Repository from the Administration menu to display the Repository List page.

2. Select File, New.

The Repository Detail page appears.

3. Enter data in the following fields:

Symbol

Identifies the repository with a name such as Large Files Only, Permanent Files, and so on.

Default Repository

Selects the repository automatically for all attachments. Users can change to a different repository, if needed.

Server Name

Specifies the server name of the repository.

Click Save.

4. Click OK.

The detail and list page close and the main page appears.

Announcements

You can use CA SDM to post announcements to users. Announcements help decrease the number of incoming calls and promote increased productivity through proactive resolution of tickets and communication of important information to all affected users. Users can scroll through stored multiple announcements.

Note: Announcements apply to all service desk models.

You can add new announcements and update existing ones. Announcements are part of the reference data function of CA SDM, so by using the access type you can control which contacts can create announcements.

An announcement can specify either or both of the following:

Location

Specifies a physical location, for example, a city, building, or floor.

Organization

Specifies an organization ID. When Organization is set for an Announcement, only individuals assigned to that Organization can see the Announcement.

When Location or Organization is set for an announcement, only contacts at that location or organization receive it. Any contact can still see all announcements that are not restricted by location or organization. For example, if neither is set, contacts in every location and organization can see an announcement.

Note: You can specify announcements using the administrative function of the web interface. For information about how to define announcements, see the *Online Help*.

More information:

[Internal Announcement Visibility](#) (see page 312)

[Specify Announcement Urgency](#) (see page 313)

Internal Announcement Visibility

You can control whether an announcement is visible to internal users by using the Internal check box on the Create New Announcement page. The View Internal Logs setting of the access type of the logged in user controls whether a user can view items marked as internal.

Specify Announcement Urgency

You can specify the urgency of an announcement record.

To specify announcement urgency

1. Create an announcement or navigate to the Announcement Detail page to edit an existing announcement.
2. Select one of the following values from the Announcement Type drop-down list:

Routine

Displays in black text.

Advisory

Displays in orange text.

Emergency

Displays in red text.

Click Save.

The announcement record is color-coded in the Announcements List page.

Stored Queries Setup

CA SDM supports use of stored queries to produce two kinds of data:

- Scoreboard Queries let you customize the web interface scoreboard by adding counter fields for items of interest.
- KPI Queries let you retrieve historical information for a specified time period on the value of a counter for use in web-based reporting.

Stored queries apply to all service desk models. Stored queries are not intended to bring users to a new page.

Administrators define the stored queries that are available to users. You can use the predefined stored queries that are installed with CA SDM, or define your own. You can define stored queries by using the administrative function of the web interface.

Consider the following information about stored query definitions:

- A valid stored query clause is integrated into an appropriate SELECT statement and passed to the underlying DBMS engine for processing. To develop the SELECT statements, see the object definition files in the following directories:
 - Windows: installation-directory\bopcfg:majic directory
 - UNIX: \$NX_ROOT/bopcfg/majic directory

- Stored queries use the object and attributes statements for building the SQL WHERE clause instead of the field names at the database level.
- CA SDM does not support Cartesian product joins for stored queries. To help ensure that your stored query does not produce a Cartesian product join, enter the following command appropriate for your system.
 - Windows: `bop_cmd -f installation-directory\bopcfg\interp\bop_diag.frg "check_queries()"`
 - UNIX: `bop_cmd -f $NX_ROOT/bopcfg/interp/bop_diag.frg "check_queries()"`Update any queries that are flagged as a result to avoid Cartesian product joins.
- The URL stored query is an option offered while creating stored queries. The URL stored query runs a query and returns numerical results that only work with Knowledge Management.

Note: For more information about the object definition syntax, see the *Technical Reference Guide*. For more information about defining stored queries, see the *Online Help*.

Sequence Numbers

When a ticket is opened, it is automatically assigned the next available sequential number. For example, if the last request opened is 5, the next one is assigned the number 6.

Important! After you install a new version of CA SDM, the internal record ID for all tickets is reset to 1. To help ensure that duplicate record IDs are not created, do not create tickets before restoring any backed-up data.

You can customize how requests, change orders, and issues are numbered by including a unique prefix or suffix in the numbering scheme for each one. For example, if you want to track requests by month, you can add a month identifier as a prefix or suffix to the request numbering scheme.

Note: For information about how to customize the numbers assigned to tickets using the administrative function of the web interface, see the *Online Help*.

Because you define a separate numbering scheme for each type of ticket, set up numbering schemes for all service desk models. You can control the format of the numbering of requests, change orders, and issues by changing the Sequence Number settings. By default, new tickets are numbered using consecutive integers. Because the number field of a ticket is actually a string field and not numeric, you assign additional string values to use as prefixes or suffixes when the ticket number is generated for a new ticket. For example, you can specify r, c, and i: for prefixes for requests, change orders, and incidents. This setup lets users easily differentiate between the various ticket types and prevents confusion.

Incidents and problems share numbering schemes with requests, because incidents and problems are internally different types of requests.

Audit Log Use

CA SDM creates an audit log that records the following information:

- All changes to an issue (Issue)
- All changes to a request (Call_Req)
- All changes to a change order (Change_Request)
- All changes to a data partition (Domain) tables
- The login ID of the person associated with the change and an associated day/date/time stamp
- The before and after values of the update, insert, or delete
- The creation of contacts
- Creation and update activities to the nr object

Note: The audit log captures modifications to only the contact data partition field, but no other contact record updates.

The audit log feature is automatically installed, and you enable it by installing two Audit Log options with the Options Manager: `audit_ins` and `audit_upd`. After you have installed these options, you can access the audit log on the Administration Tab by selecting Service Desk, Audit Log List. You can search the audit log list with a built-in search tool, and use it to facilitate report generation.

Note: For information about how to install and set values for options, see the *Online Help*.

Integration with CA Network and Systems Management

You can define event management message records and associated message record actions in CA NSM to create requests automatically, attach comments to existing requests, or post announcements to CA SDM.

Note: For more information about the CA NSM integration, see the *Implementation Guide*.

Chapter 8: Controlling System Behavior

This section contains the following topics:

[Options Manager Usage](#) (see page 317)

[How to Modify the System Environment](#) (see page 318)

[Events](#) (see page 319)

[Macros](#) (see page 319)

Options Manager Usage

The CA SDM web client lets you use the Options Manager from the Administrative tab to do the following:

- Set options
- Obtain a list of all the available options
- View a summary of each option such as the application with which it is associated, a brief description, and its status
- View the details of any specific option

When you view the detailed information for any option, you can get a comprehensive description of its functionality from the *Online Help*. The *Online Help* description includes any special actions that you need to take when changing the option. For example, after installing or uninstalling some options, you must recycle the CA SDM server for the option to take effect. If this is the case, the *Online Help* description for that option indicates it.

- Review the status of all options at the summary level
- Uninstall any of the available options

Many of the options are preconfigured and installed during CA SDM installation. Using the Options Manager to install or uninstall options can alter some or all of the default settings.

- Install any defined option

Note: During release cycles, options are occasionally made available through CA SDM support in response to user requests. To learn more about these options, contact CA SDM support. For more information about how to use the Options Manager, see the *Online Help*.

How to Modify the System Environment

CA SDM uses environment variables specified in the environment template file (NX.env.tpl) to determine certain behaviors. You can use environment variables to modify some system behaviors. You typically use the Options Manager to control system behavior, but at times CA Technical Support instructs you to modify a particular environment variable directly.

Consider the following when editing the environment template file:

- Environment variables set in this file can be overridden by setting the environment variable in the process space in which a process runs. Although convenient in some limited cases, this setup is usually not wanted. Preceding a variable setting with an at symbol (@) prevents variables in the process space overriding the variable. Unless there is a specific reason for allowing an override, the @ symbol always precedes the variable name in the template file.
- The comment characters for this file are pound (#) and exclamation point (!). The exclamation point character is also used to disable an option.

Important! Modify the template file (NX.env.tpl) and allow the configuration process to apply the changes to the environment file. Never modify the environment file (NX.env) directly, including on your client or secondary server installations.

To modify environment template file

1. Back up of the environment template file (.tpl) that corresponds to your system environment:
 - UNIX—\$NX_ROOT/pdmconf/NX.env.tpl.
 - Windows—installation-directory\pdmconf\NX.env_nt.tpl.
2. Edit the environment template file on the primary server. You can view and modify this file using any text editor (Windows users use WordPad).
3. Make the changes as instructed by your support technician, and save the changes.
4. Run the configuration utility on your primary server installation to apply the changes you made to the environment template file to the actual environment file.

Note: For information about how to run the configuration utility, see the *Implementation Guide*.

5. Stop and restart the CA SDM primary server for the changes made to the environment file to take effect. To avoid shutting down your system, your support technician can instruct you to stop and restart only certain processes rather than recycling your entire CA SDM server.

Events

You can configure events that are attached to objects to execute configured actions. Events are procedures that execute after a certain amount of time has elapsed. For example, an event sends a message to a service desk analyst if a "priority 1" issue is not received within an hour. Other parts of the system use events, for example, Service Types.

You can define events for Requests, incidents, problems, change orders, issues, contacts, configuration items, and global and specific tenants. CA SDM schedules the events execution time based on the delay time and workshift.

Note: For more information about events, see the *Online Help*.

Macros

Macros are small scripts that define either conditions or actions. When Events or Behaviors execute, they can execute one or more Action macros. Before macros are executed, you can use a conditional macro to determine which set of Action macros to execute.

You can use macros in the following areas:

- Events
- Issue or Change Behavior templates
- Activity Notifications

The product includes several macros. Users can create additional types.

Note: Customers cannot add Action or Condition macros but can create simple conditionals with site-defined conditions. Site-defined conditions are noncomplex macros that can be created from GUI dialogs; they are not replacements for condition-type macros.

For each macro, you specify the object type that you want the macro to use. If you create a site-defined conditional to verify a Requests values, you set the type to Request. When you select macros for Events and Behaviors, CA SDM only displays macros with a type matching the type on the Event or Behavior.

Macro Types

Each of the following macros has a specific type that describes its use:

Action

Executes some type of action, such as increasing the priority on an issue.

Attach Event

Attaches an Event object to a Change, Issue, Request, Contact, or Configuration Item.

Condition

Evaluates some true or false conditionals.

Execute an CA IT PAM Action

Executes a CA IT PAM action. CA IT PAM is a workflow engine that enables managing and reporting on groups of tasks that can require electronic approvals.

Execute a CA Workflow Action

Executes a CA Workflow action. CA Workflow lets you manage workflow processes in your environment, such as tracking activities that are completed on tickets associated with a specific category or area.

Remote Reference

Launches an external program on the server.

Multiple Notification

Sends a notification to one or more contacts. This macro type is especially useful, because you can specify the message, the recipients, and the urgency level.

Site-Defined Condition

Creates events based on true or false conditionals, based on attributes of the ticket to enforce service levels. For example, you can create a Condition macro that evaluates to true if a Request is priority 1 and is not yet assigned. You can attach the condition to an Event that executes soon after a request is created.

Use Macros with Events

On an Event detail, you optionally specify a conditional macro for the Event Condition field. You can also specify any number of macros (except conditionals) in the Execute on True and False columns. When the event fires, the conditional macro is evaluated first. If it evaluates to true, all the macros in the True column are fired. If it evaluates to false, all the macros in the False column are fired. The conditional macro could also check to see if the Assignee is set, the priority level, service type violation, and so on. If no Event Condition macro is specified, the event automatically executes the True macros.

On an Event detail, you can select from almost any type of macro to notify contacts, set values on the context object, execute a remote reference, or attach another event.

Use Macros in the Behavior for Issue and Change Categories

You can also use macros in the Behavior for Issue and Change categories. On a Category detail, you can look at the Workflow tab and select a workflow status. On the workflow detail, click the Behavior button. Similar to an event, you can set a condition and resultant actions.

From the Behavior detail's Transition tab, you can specify a Conditional-type macro to determine if the task can transition out of that status. This is fired when a user tries to change the task status from that value to a different one. If the condition evaluates to false, the status update is not allowed.

Real workflow tasks are created from the templates in the category template. When the status value of a task is changed and saved, the behavior specified in the category template is fired; meaning the condition and resultant actions are executed.

Note: Use an Attach Event macro to attach an event to a workflow task.

Use Macros with Multiple Notification

Multiple Notification

The message body of the multiple notification points directly to the table of the object selected. Therefore, when referencing fields on the object table, (CR (Call_Req) for example,) you can reference the field attribute name. If you want to specify the description information of the Call Request ticket, you can enter the following in the message body:

Description: @{description}

A QREL (Query Relation) is a relation which contains a list of objects defined by an SQL type WHERE clause. You can add the QREL to the description of a multiple notification macro. If you wanted to put the QREL in the description for act_log QREL <-- alg, you can use the following in a notification macro:

@{act_log.0.description}

You can use replacement variables to make notification messages more relevant and dynamic. These replacement variables take the form of @{attribute_path_here} where attribute_path_here is an attribute of some CA SDM object. When the notification is sent, the variable is replaced with the attribute value specified.

A notification (from an activity log or a multiple notification) always has some base context (a point of reference), which is usually a ticket or workflow task. On multiple notification, the Notification macro type field specifies the base object. You use the @{} syntax to reference any attribute on that object.

Example: A Multiple Notification of type Request references any attribute in the 'cr' (Call_Req) object. To include the Request description, specify @{description} (dot-notation is used to follow references to other objects). For example, to include the Request assignee last name: @{assignee.last_name}.

Use Macros with a Site-Defined Condition

Site-Defined Condition

The Conditional macro is made up of one or more *atomic* conditions. Each atom tests the value of a single attribute. Therefore, if you want to define the condition based on the values for Assignee and Category, the Conditional macro probably uses two atoms, one for Assignee and one for Category. When you use a Conditional Macro in an event, the individual atoms connected with a Boolean operator, mainly 'AND' or 'OR'.

The Conditional Macro Detail shows a list of Atoms, which is read from left to right as follows:

Attribute	Operator	Value	Logical	Translation
Assignee	Equals	Jones	AND	The assignee is Jones and...
Cost	Less than	600	OR	Cost is an integer value less than 600 or...
Group	Empty/NULL		AND	The Group field is blank.

Note: Logical connectors (AND, OR) connect two atoms. When you have two atoms such as one and two, they connect by atom ones connector. Atom twos connector is not used. Furthermore, some Operator selections are not useful, such as when an attribute is an *Assignee*, and you use the *Less than* operator—it would not make sense. For more information about creating conditional macros, see the *Online Help*.

Chapter 9: Configuring the Web Interface

This section contains the following topics:

- [Web Interface Configuration](#) (see page 323)
- [How the Web Interface Works](#) (see page 323)
- [Web Director and Web Load Distribution](#) (see page 324)
- [Web Directors and Web Engines Configuration](#) (see page 328)
- [Improve Performance With Browser Caching](#) (see page 355)
- [Record Locking Behavior in the Web Interface](#) (see page 358)
- [Print CA SDM Web Pages](#) (see page 358)
- [Configuration File Modification](#) (see page 359)

Web Interface Configuration

The *web interface* provides CA SDM functionality through the World Wide Web. The web interface enables independent browsing of the knowledge base, reducing the number of calls to the service desk and speeding resolution times. The interface also lets you do the following:

- Open, process, and close tickets
- Display and post announcements
- Access supporting data

After you install and configure CA SDM according to the instructions in the *Implementation Guide*, you can administer and secure CA SDM usage, and configure it to suit your needs. You configure the web interface security through the use of roles and access types. You can tailor the interface to use it with web browsers.

How the Web Interface Works

The CA SDM web interface uses a standard HTTP server, such as Tomcat, Apache, or Microsoft Internet Information Server (IIS).

- With the Tomcat server, a Java Servlet, `pdmweb.jar`, processes web requests.
- With Apache or IIS server, a CGI interface, `pdmweb.exe`, processes web requests.

The CA SDM web interface works as follows:

1. When a user requests a CA SDM web page, the HTTP server invokes `pdmweb.jar` or `pdmweb.exe` which sets up a connection with a CA SDM daemon (or Windows service) named the web engine.

2. The web engine interprets the user request. Most requests require the web engine to look up an HTML template (HTMPL) file and translate it into standard HTML.
3. Usually, the translation process requires the web engine to communicate with a CA SDM server to read or update the database and include database information in the generated HTML.
4. After the HTML is complete, the web engine sends it to pdmweb.jar or pdmweb.exe, which in turn sends it back to the user's browser.
5. After the data is sent back to the browser, pdmweb.exe terminates but pdmweb.jar stays active.

Web Director and Web Load Distribution

As web usage grows at your site, it can be helpful to distribute the load among multiple web engines, multiple HTTP servers, or both. Distribution lets you scale up the capacity of the web server to accommodate the level of usage demanded by your organization. You can use Web Director, provided with CA SDM, to distribute web load. Web Director receives connection requests from users, selects a web engine to handle the request, and redirects the request to that web engine. Except for the change in URL, this process is transparent to the end user, who always accesses CA SDM with the same URL, regardless of the number of web engines configured.

You can use Web Director in the following ways:

- To configure CA SDM for efficient use of secure sockets (SSL).
The HTTPS protocol allows web transactions to be encrypted, providing maximum security for sensitive data, especially passwords. However, pages using SSL are ineligible for browser caching, which can have a negative impact on performance.
- To direct logins to a specific web engine (or web engines).
After a user is authenticated, Web Director can move the session to a different web engine, which can be on a different HTTP server. This setup lets you configure SSL for a login web engine, providing maximum protection for your passwords while using standard HTTP protocol for other transactions, and maximizing their performance.
- To have multiple Web Directors, each handling a different group of web engines.
This setup can be useful in geographically dispersed organizations that want to locate groups of web engines physically closer to their end users.

Web.cfg and CA SDM

You can use the supplied `pdm_edit.pl` script to assist in configuring CA SDM web engines. Without the script, a user adding a web engine must manually specify the name, and create a unique instance of a `web.cfg` file. The script does these tasks automatically.

When you use the `pdm_edit.pl` script to configure a web engine, the script automatically creates and names a script for you. The name is constructed as follows:

```
"<Host_Name>-web[#].cfg"
```

host name

Specifies the server as primary when the web engine resides on a primary server, or specifies the actual server host name when the web engine resides on a secondary server.

[#]

Specifies an integer value automatically assigned by `pdm_edit.pl` and incremented for each web engine on a server.

For the default web engine configuration file, the CA SDM installation delivers one default web engine configuration file, defined by the following file:

```
$NX_ROOT/bopcfg/www/web.cfg
```

After the default web engine is modified through `pdm_edit.pl`, the web engine configuration file that is generated follows this new naming scheme:

```
<Host_Name>-web[#].cfg
```

Thereafter, the original `web.cfg` file, and any of its customizations, is ignored. If the user wants to preserve any customizations from the original `web.cfg`, the user must verify that the `$NX_ROOT/bopcfg/www/web.cfg` is selected as the template file to use when `pdm_edit.pl` creates the `<Host_Name>-web[#].cfg` files.

Assign a Web Engine to a Web Director

A web engine is assigned to a web director by the web director parameters in the following web engines file:

```
'<Host_Name>-web[#].cfg'
```

Notice that while each web engine requires its own `<Host_Name>-web[#].cfg` file, web directors do not use a web configuration file. This web engine/web director assignment is initially accomplished through the following utility script:

```
'pdm_edit.pl'
```

Additional changes to the web engine's <Host_Name>-web[#].cfg 'webdirector' parameters may be necessary to further modify web engine redirection behavior. If a web engine is not assigned to a web director (the default), the 'UseDirector' parameter in the web engine's configuration file is set to 'No'. CA SDM then ignores the remaining 'webdirector' parameters for the web engine.

After a web engine is assigned to a web director, the following parameters define how the web engine interacts with its assigned web director:

- UseDirector Yes | No | AfterLogin | BeforeLogin
- WebdirectorSlumpName (slump name is automatically assigned)
- WillingnessValue ranging from '0' to '10'
- RedirectingURL [http/https]://<machineName>/<cgi directory>/<cgi name>

Define the Web Engine Role by Web Director Parameters

Once a web engine has been configured to use a web director, the web engine takes on a distinct role. This role defines which web client requests the web engine is to handle. The roles are as follows:

- Service only logins (redirect non-logins elsewhere)
- Service only non-login requests (redirect logins elsewhere)
- Service both login and non-login requests ('general purpose' web engine)

This role is determined by the 'WebDirector' parameter values in the web engine's '<Host_Name>-web[#].cfg' file. The following table illustrates the web engine role to web director parameter setting relationship:

Web Engine Role	'<Host_Name>-web[#].cfg' 'webdirector' Parameter Settings
Service login requests	'UseDirector AfterLogin'; 'Willingness 0'
Service non-login activity	'UseDirector BeforeLogin'; 'Willingness [1 thru 10]'
General purpose	'UseDirector Yes'; Willingness [1-10]'

Non-SSL Environment with Basic Load-Balancing

You can use web director in a non-SSL environment with basic load-balancing. Web director balances load across all web engines according to each web engine willingness value. Each web engine can service login requests and nonlogin requests. The HTTP protocol communicates between the web client and the web server.

For each web engine under web director control, set the web director parameters in the web engine '*<Host_Name>-web[#].cfg*' as follows:

- UseDirector Yes
- WebDirectorSlumpName (do not change this value)
- WillingnessValue [1 through 10]
- RedirectingURL (the preappended protocol value can be either missing or 'http')

Global SSL Environment with Basic Load-Balancing

You can use web director in a global SSL environment with basic load-balancing. Web director balances load across all web engines according to each web engine willingness value. Each web engine can service login and nonlogin requests. The HTTPS protocol must be used for all communications between web clients and the web server.

For each web engine under web director control, set the web director parameters in the web engine '*<Host_Name>-web[#].cfg*' as follows:

- UseDirector Yes
- WebDirectorSlumpName (do not change this value)
- WillingnessValue [1 through 10]
- RedirectingURL (the preappended protocol value must be 'https')

Targeted Login in a Non-SSL Environment with Optional Load-Balancing

You can use web director for a targeted login in a non-SSL environment with optional load-balancing. The login-only web engine services only the login requests. The remaining web engines under web director control, service all other requests. This configuration puts the entire burden of servicing login requests on the specified login-only web engines. The HTTP protocol is used for communications between the web client and the web server.

For the login-only web engines, set the web director parameters in the web engine '*<Host_Name>-web[#].cfg*' as follows:

- UseDirector AfterLogin
- WebDirectorSlumpName (do not change this value)
- WillingnessValue 0
- RedirectingURL (the preappended protocol value can either be missing or 'http')

For the nonlogin web engines, set the web director parameters in the web engine '*<Host_Name>-web[#].cfg*' as follows:

- UseDirector Before Login
- WebDirectorSlumpName (do not change this value)
- WillingnessValue [1 through 10]
- RedirectingURL (the preappended protocol value can be either missing or 'http')

Targeted SSL Login in a Mixed Environment with Optional Load-Balancing

You can use web director for a targeted SSL login in a mixed SSL/non-SSL web environment with optional load-balancing. Every web-login request is redirected to and serviced by the SSL web engines with all other requests being redirected to and serviced by the non-SSL web engines. The HTTPS protocol must be used for all communications between the web client and the SSL web engines.

Web Directors and Web Engines Configuration

Become familiar with the following information to add and configure web directors and web engines, based on server and component types:

SSL

Specifies the Secure Socket Layer Protocol.

pdm_edit.pl

Specifies the CA SDM Perl script used to add and modify web engines, web directors, object managers, and other components.

CGI I/F

Specifies the script name for web engines and web directors assigned by pdm_edit.pl. This value is the name of an actual CGI executable when IIS or Apache is used as the HTTP server; it is a servlet parameter when Tomcat is used as the HTTP server.

Examples: (web engines) pdmweb1, pdmweb2, (web directors) pdmweb_d1, and pdmweb_d2

Systems that use servlet servers like Tomcat do not need CGI/IF executables. The CGI/If is simulated by a servlet running on the servlet server, instead these systems require a web.xml file in the directory:

\$NX_ROOT/bopcfg/CATALINA_BASE/webapps/CAisd/WEB-INF

Pdm_edit.pl creates sample web.xml files that are named <hostname>-web.xml.tpl. You copy these files to the appropriate server in the WEB-INF directory so they replace the web.xml.tpl file on each server, and then reconfigure the secondary server.

SSL Login

Specifies the CA SDM configuration in which one or more web engines reside in an SSL-enforced virtual directory and exclusively handles all web-client login requests using encrypted HTTP communication. After user authentication is complete, a web director redirects the web client to an eligible web engine residing in a non-SSL-enforced virtual directory for the remainder of the client session, using unencrypted HTTP communication.

Secure Login Web Engine

Specifies the web engine in the SSL login environment that exclusively handles all CA SDM web-client user authentication requests.

Non Secure Web Engine

Specifies the web engine in the SSL login environment that handles all web client nonuser authentication requests after user authentication has taken place through the secure login web engine.

Web Engine Role

Specifies the duties performed by a web engine after it has been assigned to a web director. These roles include the following:

- Only Process web client login requests.
- Process everything but web client login requests.
- Process all web client requests.

More information:

[How to Implement an SSL Login Environment with a Primary Server Only](#) (see page 330)

[Server Configuration](#) (see page 337)

[Web Engines](#) (see page 344)

How to Implement an SSL Login Environment with a Primary Server Only

You can implement an SSL login environment that uses a primary server only.

Note: The primary server must be installed and configured to create the system.

To implement an SSL login environment, do the following:

1. Verify that the SSL-enabled web server has successfully imported an SSL certificate.
2. Create a copy (including subdirectories) of the directory '\$NX_ROOT/bopcfg/www/wwwroot', and assign it the following name:
'\$NX_ROOT/bopcfg/www/wwwrootsec'
3. Add a new virtual directory for the web server named *CAisdsec*. Point this virtual directory to the following physical directory:

'\$NX_ROOT/bopcfg/www/wwwrootsec'

Verify that the virtual directory permissions for *CAisdsec* match the *CAisd* virtual directory permissions for script execution. Enforce SSL for the *CAisdsec* virtual directory.

Note: In this example, *CAisdsec* is user-defined and can be renamed.

4. Create a system with a primary server only.

How to Create a System with a Primary Server Only

You can create a system that uses a primary server only.

Note: The primary server must be installed and configured to create the system.

To create a system with a primary server only, do the following:

1. Using `pdm_edit.pl` add, edit webdirector entries, or both, as you want. One web director can distribute load among several web engines. Assign all new web engines and web directors to the host named `primary`. This host is the primary server.

Note: Auto-Assignment of webdirector CGI I/F assigns names from within `pdm_edit.pl`:

- The first webdirector on the primary server is automatically assigned the name `pdmweb_d1.exe`.
- The second webdirector on the primary server would be named `pdmweb_d2.exe`.
- The third would be `pdmweb_d3.exe`, and so on.

A secondary server could also have webdirectors with CGI I/F values of `pdmweb_d1.exe`, `pdmweb_d2.exe`, and `pdmweb_d3.exe`. No naming conflicts exist because the webdirector names are unique for the local computer on which the webdirectors are running.

2. Read the following information that can apply depending on your CA SDM web engine environment:
 - Number of web engines to add—When implementing any web engine load-balancing scheme, SSL-Login, or both, at least two web engines must be assigned to the same webdirector.
 - SSL-login environment—When defining the Secure-Login web engine in `pdm_edit.pl`, this web engine must be under the control of a webdirector. Set the Secure-Login web engine's protocol parameter value to `https` so that the webdirector constructs the correct redirecting URL value for the web client.
3. Save changes and exit `pdm_edit.pl` script, specifying which file to use for a `web.cfg` template when prompted.

Note: Webdirectors do not use a '`<Host_Name>-web[#].cfg`' file. However, web engines require a unique '`<Host_Name>-web[#].cfg`' file. Sample `web.cfg` files are automatically generated using the `pdm_edit.pl` script when saving and exiting the script. The user is prompted to specify a `web.cfg` file to be used as a template in the creation of the web configuration files. The customizations in the original `web.cfg` can be imported to the new web configuration files by specifying the original `web.cfg` as the desired template file.

4. Copy and save the following files, because a backup of these files is useful whenever you decide to restore the original environment:
 - `$NX_ROOT/pdmconf/pdm_startup.tpl`
 - `$NX_ROOT/pdmconf/pdm_startup`
 - `$NX_ROOT/bopcfg/www/web.cfg` file

- Any existing primary-web[#].cfg files
 - \$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF/web.xml and web.xml.tpl
5. Move \$NX_ROOT/samples/pdmconf/pdm_startup.rmt to \$NX_ROOT/pdmconf and rename it to pdm_startup.tpl.
 6. For each web engine that has been assigned to a webdirector, ensure that the web engine's '<Host_Name>-web[#].cfg' "webdirector" parameters have been set correctly by examining the file in a text editor. If necessary, modify the 'webdirector' parameter values to reflect the web engine role you want. Then copy them to the directory: \$NX_ROOT/bopcfg/www.
 7. Move all \$NX_ROOT/samples/pdmconf/primary-web[#].cfg files to the \$NX_ROOT/bopcfg/www directory.
 8. If using a servlet server like Tomcat, the pdm_edit.pl utility created web.xml.tpl files that can replace the web.xml.tpl file on each server hosting a webengine. These files are name primary-web.xml.tpl. Rename the files and copy them to the directory: \$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF directory.

If using an HTTP server like IIS or Apache, create copies of 'pdmweb.exe' in the \$NX_ROOT/bopcfg/www/wwwroot directory, a 'pdmweb[#].exe' for each web engine, and a 'pdmweb_d[#].exe' for each webdirector that has been defined in pdm_edit.pl. Verify that the 'pdmweb[#].exe' and 'pdmweb_d[#].exe' are named according to the correct CGI I/F values from the pdm_edit.pl script definitions (for example: 'pdmweb1.exe', 'pdmweb2.exe', pdmweb_d1.exe', and so on).
 9. If you are using IIS and want to add Server Extensions for each CGI interface, you can take the file primary-site.dat file and copy it to your \$NX_ROOT/bopcfg/www directory as site.dat. When the system is reconfigured these sites is added to IIS.
 10. Reconfigure the primary server without reinitializing the database and start services.
 11. After reconfiguring verify that the current settings are valid. Start the CA SDM daemons. Verify that there are no errors in the stdlog files. Use pdm_status to view the daemons and their status. From a browser, you are able to access the system: <http://localhost:8080/CAisd/pdmweb.exe>.
 12. For Knowledge Management to CA SDM integration, if SSL has been enforced for CA SDM, the CA SDM URL protocol value must be changed.
 - a. From the Knowledge Tool Settings Manager, General, Integration, change the CA SDM URL protocol value from http to https.
 - b. Save and exit.
 13. Open a web browser to the CA SDM login page and verify that a user can log in and that the expected redirect/login behavior is observed.

Verify the Web Director Parameter Values

You can verify web director parameter values.

To verify parameter values

1. For Secure Login Web engines, edit the `<Host_Name>-web[#].cfg` as follows:
 - a. Change the CAisd parameter value from `/CAisd` to `/CAisdsec`.
 - b. Change the UseDirector parameter value from Yes to AfterLogin if the web director uses pass thru authentication.
 - c. Change the Willingness parameter value from 5 to 0.
 - d. Verify that the RedirectingURL value protocol is listed as https.
 - e. Change the RedirectingURL `<cgi directory>` value from `CAisd` to `CAisdsec`.
 - f. Save the changes.
2. For non-secure engines handling all other activity, edit the `<Host_Name>-web[#].cfg` files for the non-secure web engines that you want to handle the non-login activity. Verify that the CAisd parameter value is `/CAisd`.
 - a. Change the UseDirector parameter value from Yes to BeforeLogin.
 - b. Maintain the Willingness value of 5 or set it to any integer value from 1 to 10, depending on the particular loading weight desired.
 - c. Verify that the RedirectingURL value protocol is listed as http.
 - d. Verify that the RedirectingURL `<cgi directory>` value is `CAisd`.
 - e. Save the changes.

The web director parameter values are verified.

SSL Login Environment Verification

You can verify the SSL login environment for web engines as follows:

- The secure-login web engines must reside in the physical directory mapped to the SSL-enforced virtual directory (CAisdsec in this example).

For secure-login web engines, create instances of pdmweb.exe in the \$NX_ROOT/bopcfg/www/ wwwrootsec directory with the name of pdmweb[#].exe. The executable name must match the CGI I/F value for each secure-login web engine defined in pdm_edit.pl.

Example: If pdm_edit.pl is assigned the CGI I/F value of the secure-login web engine to pdmweb2, create a physical copy of pdmweb.exe, and rename it pdmweb2.exe.

- The non-secure web engines and web directors must reside in the physical directory mapped to the non-SSL virtual directory CAisd.

For non-secure web engines and web directors, create instances of pdmweb.exe in the \$NX_ROOT/bopcfg/www/wwwroot directory. A copy of pdmweb.exe must exist for each non-secure web engine and webdirector that has been defined in pdm_edit.pl. Rename the copies so that the new names of the executables match the CGI I/F values that have been defined for the web engines and web directors in pdm_edit.pl.

Example: If pdm_edit.pl is assigned the CGI I/F value of pdmweb3 to the non-secure web engine and the value of pdmweb_d1 to the web director, then create two copies of pdmweb.exe. Rename the first copy to pdmweb3.exe, and then rename the second copy to pdmweb_d1.exe.

Secondary Servers and Web Capabilities

All systems have one primary server. The primary server hosts the core system services, like the communications hub and database. Many enterprise users spread the system to other servers referred to as secondary servers. In this type of setup, you can add web capabilities to the secondary servers.

How to Prepare Web Capabilities

The prerequisite tasks for adding web capabilities to secondary servers (and if you are implementing an SSL Login environment) are as follows:

1. Verify that the primary and secondary servers are installed and configured.
2. Verify that the SSL-enabled web server that is to host the Secure-Login web engine has imported an SSL certificate.
3. On the web server hosting the Secure-Login web engine, create a physical copy (including subdirectories) of the \$NX_ROOT/bopcfg/www/wwwroot directory, and name it \$NX_ROOT/bopcfg/www/wwwrootsec.

4. Define a new virtual directory for the web server named CAisdsec.
5. Point the virtual directory to the physical directory
\$NX_ROOT/bopcfg/www/wwwrootsec.
6. Verify that the virtual directory permissions for CAisdsec match the CAisd virtual directory permissions for script execution.
7. Enforce SSL for the CAisdsec virtual directory.

Note: The name CAisdsec is user-defined and you can rename it.

How to Create a System with Secondary Servers

To create a system with secondary servers, do the following:

1. Using `pdm_edit.pl`, add or edit webdirector entries. One webdirector can distribute load among several web engines. Add webengines and webdirectors to other servers. When specifying a secondary server, use the same name with which the secondary server was configured. This name is the `NX_LOCAL_HOST` value in the `NX.env` file on the secondary server.

Note: Auto-Assignment of webdirector CGI I/F assigns names from within `pdm_edit.pl`: The first webdirector on the primary server is automatically assigned the name `pdmweb_d1.exe`. The second webdirector on the primary server is named `pdmweb_d2.exe`. The third is named `pdmweb_d3.exe`, and so on. A secondary server can also have webdirectors that have CGI I/F values of `pdmweb_d1.exe`, `pdmweb_d2.exe`, and `pdmweb_d3.exe`. No naming conflicts exist because the webdirector names are unique for the local computer on which the webdirectors are running.

2. Each webengine must reference an object manager. An object manager named `domsrvr` always exists. You can add more object managers on the secondary server and then assign webengines to them as follows:
 - **Number of Web Engines to Add**—When implementing any web engine load-balancing scheme, SSL-Login, or both, assign at least two web engines to the same webdirector.
 - **SSL-Login Environment**—When defining the Secure-Login web engine in `pdm_edit.pl`, this web engine must be under the control of a webdirector. Set the Secure-Login web engine protocol parameter value to `https`. This setting causes the webdirector to construct the correct redirecting URL value for the web client.

3. Save changes, exit the `pdm_edit.pl` script, and specify which file to use for a `web.cfg` template when prompted.

Note: Webdirectors do not use a `<Host_Name>-web[#].cfg` file. However, web engines require a unique `<Host_Name>-web[#].cfg` file. Sample `web.cfg` files are automatically generated using the `pdm_edit.pl` script when saving and exiting the script. You are prompted to specify a `web.cfg` file to be used as a template in the creation of the web configuration files. You can import the customizations in the original `web.cfg` to the new web configuration files by specifying the original `web.cfg` as the template file you want to use.

4. On the primary server, save the following backup copies. They are useful if you want to restore the original environment:
`$NX_ROOT/pdmconf/pdm_startup.tpl`
`$NX_ROOT/pdmconf/pdm_startup`
`$NX_ROOT/bopcfg/www/*web*.cfg`
`$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF/web.xml*`
5. On the secondary server, save backup copies of any existing
`$NX_ROOT/bopcfg/www/web.cfg` or `<secondary_name>-web[#].cfg` files.
`$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF/web.xml*`
6. On the primary server, move `$NX_ROOT/samples/pdmconf/pdm_startup.rmt` to `$NX_ROOT/pdmconf` and rename it to `pdm_startup.tpl`.
7. For each web engine that is assigned to a webdirector, verify that the web engine `<Host_Name>-web[#].cfg` webdirector parameters are set correctly by examining the file in a text editor. If necessary, modify the webdirector parameter values to specify the web engine role you want.
8. On the primary server; move all `$NX_ROOT/samples/pdmconf/primary-web[#].cfg` files to the `$NX_ROOT/bopcfg/www` directory.
9. Move all `$NX_ROOT/samples/pdmconf/'secondary_server_name-web[#].cfg'` from the primary server to the secondary server `$NX_ROOT/bopcfg/www` directory.
10. (Servlet servers like Tomcat) Specify a `web.xml.tpl` in each `$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF` directory. When you configure this server, the configuration process creates a `web.xml` file. The `pdm_edit.pl` utility created sample `web.xml.tpl` files. Each file is named `<hostname>-web.xml.tpl`. You can use these files as a sample: copy and rename them to `web.xml.tpl` for each of your servers.
11. (HTTP servers like IIS or Apache) On each CA SDM server (primary, secondary, or both) that you want to host a web engine or webdirector, create copies of `pdmweb.exe` in the `$NX_ROOT/bopcfg/www/wwwroot` directory. A `pdmweb[#].exe` for each web engine and a `pdmweb_d[#].exe` for each webdirector is defined in `pdm_edit.pl` to run on that server. Verify that `pdmweb[#].exe` and `pdmweb_d[#].exe` are named with the CGI I/F values from the `pdm_edit.pl` script definitions (for example, `pdmweb1.exe`, `pdmweb2.exe`, `pdmweb_d1.exe`, and so on).

12. If you are using IIS and want to add server extensions for each CGI interface, copy the file `primary-site.dat` file to your `$NX_ROOT/bopcfg/www` directory as `site.dat`. When the system is reconfigured these sites are added to IIS.
13. Reconfigure the primary server without reinitializing the database, and start services.
14. Reconfigure each secondary server.
15. For Knowledge Management to CA SDM integration, if SSL has been enforced for CA SDM, the CA SDM URL protocol value must be changed.
 - a. From the Knowledge Management, Settings Manager, General Integration, change the CA SDM URL protocol value from `http` to `https`.
 - b. Save and exit.
16. Open a web browser to the CA SDM login page and verify that a user can log in and that the expected redirect or login behavior is observed.

Server Configuration

`pdm_edit.pl` lets you add functionality and increase the performance of the enterprise system. You can configure the CA SDM services. CA SDM has basic daemons on the primary server. You can add more daemons on the primary server, and then expand your system across several servers by adding secondary servers, which increase volume, performance, and enhance security.

If you use `pdm_edit.pl` to modify the daemon configuration, the customizations available using `pdm_configure` are disabled. For example, after you change the `pdm_startup.tpl` file with `pdm_edit.pl`, you cannot enable a web director or event converter through the `pdm_configure` GUI. You can customize these features through the `pdm_edit.pl` utility.

The `pdm_edit.pl` top menu appears after you run `pdm_edit.pl` and enter operating system-specific questions. This menu lists the capabilities of this utility. Entering a number from this menu takes you into submenus for the selected feature. From the submenus, you can return to the top menu. At the end of your session, you select “Quit NO SAVE” or “Save and quit.” The action creates a `pdm_startup.dat` file, which is a summary of all your selections. Keep this file so you can resume later. The utility creates a `pdm_startup.rmt` file, and serves as a replacement for the `$NX_ROOT/pdmconf/pdm_startup.tpl` file on your primary server. After installing the `pdm_startup.rmt` file as `pdm_startup.tpl` on the primary server, reconfigure without reinitializing your database to enable the changes.

Note: Host names of all secondary servers are case sensitive and represent the value of `NX_LOCAL_HOST` in the `NX.env` file on the secondary server

Run pdm_edit.pl on a UNIX Server

You can run `pdm_edit.pl` on a UNIX server. You must have a primary server installation and, optionally, a secondary server installed and configured for all relevant CA SDM documents.

Note: Run `pdm_edit.pl` on the primary server only. For more information about secondary servers running the web interface, see the *Implementation Guide*.

To run `pdm_edit.pl` on a UNIX server

1. From the command line on the primary server, change to the following location:
`$NX_ROOT/samples/pdmconf` directory.
2. Issue the following command:

```
'$NX_ROOT/bin/pdm_perl pdm_edit.pl'
```

CA SDM's `pdm_perl` executable runs the `pdm_edit.pl` script. A series of text-based menu selections and prompts let you add, modify, or delete the object managers, webdirectors, and web engines, and then, save the changes or additions into a `pdm_startup.dat` and `pdm_startup.rmt` files on the primary server.

Run pdm_edit.pl on a Windows Server

You can run `pdm_edit.pl` on a Windows server. You must have a primary server installation and, optionally, a secondary server installed and configured for all relevant CA SDM documents.

Note: Run `pdm_edit.pl` on the primary server only. For more information about secondary servers running the web interface, see the *Implementation Guide*.

To run `pdm_edit.pl` on a Windows server

1. Open a command prompt window and change to the following directory:
`%installation-directory\samples\pdmconf`.
2. Issue the following command:

```
%installation-directory\samples\pdmconf\pdm_perl pdm_edit.pl
```

The `pdm_perl` executable runs the `pdm_edit.pl` script. A series of text-based menu selections and prompts let you add, modify, or delete the object managers, web directors, and web engines, and then, save the changes or additions into a `pdm_startup.dat` and `pdm_startup.rmt` files on the primary server.

Object Managers

Object managers manage all the CA SDM objects. An object manager always exists on the primary server. Each object manager has a name, which is synonymous with the name it uses to communicate with other objects. The default object manager on the primary server is always named `domsrvr`. Enterprise systems with multiprocessor servers or secondary servers can add more object managers to the primary server and secondary servers. You can navigate, add, edit, and delete object managers.

Each object manager is assigned a host. Object managers that are assigned to the primary server always enter `primary` as the host name.

You can group object managers so that the groups can be assigned to provide service to specific groups of web engines and Java clients. Typically, *group* can be left as none and the load balancing feature of the system spreads the load across all object managers. Users that require this feature often have web engines geographically separated from the primary server and would like to co-locate an object manager near the webengine. The users group the object managers that are assigned to the local webengines, and then assign the webengines to this group of object managers.

The Accept mask is an advanced feature that tells the object manager from which clients it accepts connections. Usually, a webengine attempts to connect to an object manager with a name like `web:seattle:1`, `web:seattle:2` or `web:texas:1`. The administrator can specify an Accept mask like `web:seattle.*` to accept all seattle connections and reject others. The administrator can also specify a mask like `web:.*` to accept connections from webengines and reject connections from clients.

The Display value is sometimes useful. Most users let the system default this value to the host name. This value appears on the clients to indicate which object manager it is connected to.

Enter and Exit Object Manager

The CA SDM installation delivers, by default, one object manager on the primary server. The default object manager, `comm. name = 'domsrvr'`, does not display in the Object Manager list; however, be aware that it is present and starts when the CA SDM Daemon Manager starts.

To enter and exit Object Manager

1. From the top menu of `pdm_edit.pl`, select 1, and press Enter.

The Object Manager submenu appears.

2. Press Enter.

You exit the Object Manager submenu and the `pdm_edit.pl` top-level menu appears.

Add Object Managers

You can add object managers to the Object Manager submenu.

To add object managers

1. From the Object Manager submenu, select a, and press Enter.
A prompt requests the host name.
2. Press Enter to select the default value of Primary or enter a value for the host name of the secondary server.
Note: Host names of all secondary servers are case sensitive and are the value of `NX_LOCAL_HOST` in the `NX.env` file on the secondary server.
3. Press Enter.
A prompt requests the Group.
4. Press Enter to accept the default value of None or enter a group name, and press Enter.
A prompt requests the Accept Mask.
Note: An Accept Mask is a filter which restricts object manager communication to only those processes with a name that matches the accept mask.
5. Press Enter to accept the default value of no accept or enter the accept mask string you want to use.
A prompt requests the scoreboard Display Name value.
6. Press Enter to accept the default or enter a Display Name value, and then press Enter.
The Object Manager submenu appears and the newly added object manager is listed.

Edit an Object Manager

You can edit object managers.

To edit an object manager

1. From the Object Manager submenu, select e, and press Enter.

A prompt requests Please Enter Key.

Note: This key value is shown in the object manager list and is constructed automatically for you whenever an object manager is added. The key is in a hostname: # format. # is an integer in sequential order with the maximum number equal to the total of object managers on the specified host. For example, if two object managers exist on the primary server and two object managers exist on a secondary server (host name='godzilla'), the keys would be: 'primary:1', 'primary:2', 'godzilla:1', and 'godzilla:2'.

2. Enter the key (godzilla:1) you want, and press Enter.

A prompt requests Group.

3. Press Enter to accept the displayed value 'as is', or modify the value by entering a group value, and press Enter.

A prompt requests Accept Mask.

4. Press Enter to accept the displayed value 'as is', or modify the value by entering an Access Mask value, and press Enter.

A prompt requests Display.

5. Press Enter to accept the displayed value 'as is' or modify the value by entering a Display value, and press Enter.

The Object Manager submenu appears and the object manager is edited.

Delete an Object Manager

You can delete an object manager.

To delete an object manager

1. From the Object Manager submenu, select d, and press Enter.

A prompt requests Please Enter Key.

2. Enter the key you want, and press Enter.

Note: If you entered the wrong key (which deletes the wrong domsrvr entry), exit from the pdm_edit.pl script *without* saving the changes.

The object manager is deleted and the Object Manager submenu appears.

How to Specify Web Directors

Webdirectors are optional, and are used when systems require two or more webengines. Webdirectors act to balance the load between the webengines. Zero or more webdirectors can exist. Each webdirector balances the load among two or more webengines.

To specify a webdirector, do the following steps:

1. Indicate the hostname.

If the webdirector is on the primary server, enter “primary”; otherwise, enter the name of the secondary server. The secondary server name is the NX_LOCAL_HOST entry in the NX.env file on the secondary server.

Note: Case is important.

2. (Optional) Specify the CGI interface.

In most cases, the default value is appropriate. If you are using an HTTP server, like IIS or Apache, this value is the CGI executable that runs in response to a CGI request (the name of the CGI executable). Create copies of the pdmweb.exe in the \$NX_ROOT/bopcfg/www/wwwroot directory with this name. If you are using a servlet server like Tomcat, a sample web.xml.tpl file is created. You use this file to replace the web.xml.tpl file in the \$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF directory of the specific server. Each sample web.xml.tpl file is created in the current directory and is named <hostname>-web.xml.tpl.

Enter and Exit Web Director

Defining a web director only adds web director information. This web director is used by CA SDM only if one or more web engines are assigned to it. You can enter and exit Web Director.

To enter and exit Web Director

1. From the top menu of pdm_edit.pl, select d, and press Enter.

The Web Director submenu appears.

2. Press Enter.

You exit the Web Director submenu and the pdm_edit.pl top-level menu appears.

Add Web Directors

You can add web directors. Each web director defined in `pdm_edit.pl` must have a unique CGI I/F value. The CGI I/F name is automatically generated for you; however, the suggested CGI I/F name can be manually overridden to reflect any CGI naming scheme you want.

Consider the following before you add web directors:

- If you are using an HTTP server like IIS or Apache, you must have a correctly named `pdmweb_d[#].exe` in the `$NX_ROOT/bopcfg/www/wwwroot` directory.
- If you are using a servlet server like Tomcat, you get a sample `web.xml.tpl` file. This file is named `<hostname>-web.xml.tpl` and is located in the working directory. Use this file to replace the `$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF/web.xml.tpl` file on the appropriate server.

To add web directors

1. From the Web Director submenu, select `a`, and press Enter.
A prompt requests the hostname.
2. Press Enter to select the default value of Primary or enter a value for the hostname of the secondary server.
3. Press Enter.
A prompt requests a name for the CGI Interface.
4. Press Enter to accept the displayed value or manually enter the value you want, and press Enter.
The Web Director submenu appears and the newly added web director is listed.

Edit a Web Director

You can edit a web director.

To edit a web director

1. From the Web Director submenu, select `e`, and press Enter.
A prompt requests Please Enter Key.
2. Enter the key you want and press Enter.
A prompt requests a CGI I/F value.
3. Press Enter to accept the displayed value or modify the value by entering the CGI I/F value you want, and press Enter.
The Web Director submenu appears and the web director is changed.

Delete a Web Director

You can delete a web director entry.

To delete a web director entry

1. From the Web Director submenu, select d, and press Enter.

A prompt requests Please Enter Key.

2. Enter the key you want and press Enter.

Note: If you enter the wrong key (which deletes the wrong web director entry), exit from the pdm_edit.pl script *without* saving the changes.

The Web Director submenu appears and the web director entry is deleted.

Change the CGI I/F Value for a Web Director

You can change the CGI I/F value for a web director.

To change the CGI/IF value

1. From the Web Director submenu, select c, and press Enter.

A prompt requests Please Enter Key.

2. Enter the key you want and press Enter.

A prompt requests a CGI I/F value.

3. Press Enter to accept the displayed value 'as is' or modify the value by entering the CGI I/F value you want, and then press Enter.

The Web Director submenu appears and the CGI I/F value for the web director is changed.

Web Engines

Web engines help prepare web pages for the web client. All systems have one or more web engines. Zero or more web engines can exist on each server. You can navigate, add, edit, and delete web engines. Each web engine connects to an object manager for processing all requests to CA SDM objects. The default is to connect to the default object manager, but if you have more than one object manager, you can set this value to ANY.

When you define a web engine, you specify the host name. Enter **primary** for all web engines that start on the primary server; otherwise, enter the host name of the secondary server. The host name is case sensitive and must match the NX_LOCAL_HOST entry in the NX.env of the secondary server.

Each web engine can run and be accessed directly. In direct access, every web browser enters the specific CGI interface for that web engine and the users determine the system loading (load balancing is not automated). All clients can connect to one web engine and overburden this web engine, while the other web engines go unused. A better approach is to assign two or more web engines to a single web director. All requests that go to one of these web engines is directed to the web director for load balancing and then redirected to the most available web engine in the group.

Other configurations can include web directors. One configuration directs all logins to a single SSL secured web engine that handles all logins. After login, the client is directed to another available web engine for further processing.

CGI Interface

Each web engine has and uses a CGI interface as follows:

- When using an HTTP server like IIS or Apache, the interface is an executable that is stored in the `$NX_ROOT/bopcfg/www/wwwroot` directory. The name of the CGI interface is the name of the executable. Browsers enter the CGI interface on the address line and the request is directed to the appropriate web engine.
- When using a servlet server like Tomcat, the CGI interface is simulated by a servlet. Each servlet server has a `web.xml.tpl` file that describes all the CGI interfaces. The `pdm_edit.pl` utility creates sample `web.xml.tpl` files named `<hostname>-web.xml.tpl`. Use the files to replace the `web.xml.tpl` file on the appropriate server, and reconfigure the server.
- When using a web director, it is often useful to set one CGI interface for one web engine to a known CGI name. All requests go to this web engine and are redirected to a more available web engine. This setup lets the users remember only a single CGI interface name and all requests are balanced across all available web engines.

Enter and Exit the Web Engine

The CA SDM installation, by default, installs and configures one web engine on the primary server (key= primary:1; object manager = domsrvr; CGI I/F= 'pdmcgi', protocol= not specified).

To enter and exit the web engine

1. From the top menu of `pdm_edit.pl`, select `w`, and press Enter.

The web engine submenu appears.

2. Press Enter.

You exit the web engine submenu and the `pdm_edit.pl` top-level menu appears.

Considerations for Adding Web Engines

Consider the following information when you add web engines:

- **Directing clients to web engines**—You can specify how web clients are directed to web engines. For example, consider two web directors that are being used, miko:1 and godzilla:1. miko:1 has four web engines assigned to it, godzilla:1, godzilla:2, godzilla:3, and primary:3. Any web client can request `https:godzilla:8080/Caisd/ComputerAssociates.exe` and be redirected to any of the four web engines in this group. If a web client enters `http://localhost:8080/CAisd/pdmweb6.exe`, the client is always directed to this web engine. Also, notice that some web engines are assigned to the object manager named domsrvr and others are set up to attach to ANY object manager.
- **Using CGI I/F values**—Each web engine defined in `pdm_edit.pl` must have a unique CGI I/F value:
 - If you are using an HTTP server like IIS or Apache, you must have a correctly named `pdmweb[#].exe` executable in the `$NX_ROOT/bopcfg/www/wwwroot` directory.
 - If you are using a servlet server like Tomcat, a servlet replaces the CGI interface. These servers must have a `web.xml.tpl` file in the `$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF` directory that describes each CGI interface. The `pdm_edit.pl` utility creates sample `web.xml.tpl` files named `<hostname>-web.xml.tpl`. Use these files to replace the `web.xml.tpl` file and reconfigure the server.

Add Web Engines

You can add web engines.

To add web engines

1. From the web engine submenu, select a, and press Enter.
A prompt requests the hostname.
2. Press Enter to select the default value of Primary or enter a value for the hostname of the secondary server.
3. Press Enter.
A prompt requests the name of the object manager.
4. Press Enter to accept the default value of `<primary domsrvr>` or manually enter a combination of aliases and domsrvr communication names, and press Enter.
A prompt requests the web directory key.
5. Press Enter to select the default value of Null (no web director), or enter one of the listed web directors by key value, and press Enter.
A prompt requests the name for the CGI Interface.

6. Press Enter to accept the displayed value, or manually enter a value, and press Enter.

If a web director was not specified for this web engine, the web engine submenu appears. If a web director value was specified or previously existed in the web engine definition, a prompt requests that you specify the protocol to use as a prefix for the web engine redirecting URL value.

7. Press Enter if you do not want to pre-append a protocol prefix to the URL redirection value.
8. Enter **http** to attach the protocol prefix http for the redirection value unless SSL is enforced for this web engine. For example, enter **https** to attach the protocol prefix https into the redirection value.

Note: If a web director has been defined for this web engine, the web director controls redirection to this web engine using the redirecting URL value. If the SSL protocol is not used for this web engine, specify **http** or none. If SSL protocol is to be enforced for this web engine, specify **https**. Failure to specify https for a web engine that is to be under SSL can result in a failed redirection attempt by the web director.

9. Press Enter.

The web engine submenu appears and the newly added web engine is listed.

Edit a Web Engine

You can edit a web engine.

To edit a web engine

1. From the web engine submenu, select e, Enter.
A prompt requests Please Enter Key.
2. Enter the key you want, and then press Enter.
A prompt requests that you enter the object manager.
3. Press Enter to accept the displayed value 'as is' or modify the value by entering the object manager value you want, and then press Enter.
A prompt requests a web director.
4. Press Enter to accept the displayed value 'as is' or modify the value by entering the web director value you want, and then press Enter.
A prompt requests a CGI I/F value.

5. Press Enter to accept the displayed value 'as is' or modify the value by entering the CGI I/F value you want, and then press Enter.

A prompt requests the HTTP protocol.

6. Press Enter to accept the displayed value 'as is' or modify the value by entering the HTTP protocol value you want, and then press Enter.

The web engine submenu appears.

Delete a Web Engine

You can delete a web engine.

To delete a web engine

1. From the web engine submenu, select d, and press Enter.

A prompt requests Please Enter Key.

2. Enter the key you want and press Enter.

Note: If you entered the wrong key (which deletes the wrong web engine entry), exit from the pdm_edit.pl script *without* saving the changes.

The web engine is deleted and the web engine submenu appears.

Make Changes Using pdm_edit.pl

You can use the pdm_edit.pl utility to add and move converters, servers, daemons, and aliases in your CA SDM system.

Add a TNG Converter

The pdm_edit.pl utility lets you add a single TNG converter to any one of the servers in your system.

Note: If the Daemon Manager manages the TNG converter; this daemon is started and stopped with the other daemons. If the TNG converter is required to catch events after the CA SDM daemons are shut down, the TNG converter cannot satisfy this requirement, so start and stop the TNG converter as a service.

To add a TNG converter

1. From the pdm_edit.pl top menu select T, and press Enter.

The TNG converter submenu appears.

2. From the TNG converter submenu select A, and press Enter.

A prompt requests a hostname.

3. Enter **primary** for the primary server or the hostname of a secondary server, and press Enter.
A prompt requests an IP address.
4. Enter **HOST_IP_REPLACE** if the TNG converter is on the primary server; otherwise, enter the IP of the secondary server.
The TNG converter is added to the server.

Add a UNIX Converter

The `pdm_edit.pl` utility lets you add a single UNIX converter to any one of the servers in your system.

To add a UNIX Converter

1. From the `pdm_edit.pl` top menu select **N**, and press Enter.
The UNIX converter submenu appears.
2. Select **A**, and press Enter.
A prompt requests a hostname.
3. Enter **primary** for the primary server or the hostname of a secondary server, and press Enter.
A prompt requests an IP address.
4. Enter **HOST_IP_REPLACE** if the UNIX converter is on the primary server; otherwise, enter the IP of the secondary server.
The UNIX converter is added to the system.

Move Login Validation

The login validation server uses the local operating system to validate user requests. The validation daemon starts on the primary server by default, but you can move it to a secondary server where all users have operating system accounts, or to a server that has a different operating system type.

To move login validation

1. From the `pdm_edit.pl` top menu select **U**, and press Enter.
The User Validation server submenu appears.
2. Enter **E** to specify the hostname of the server, and press Enter.
3. Enter **primary** for the primary server or the hostname of the Secondary server, and press Enter.
The login validation is moved.

Move an LDAP Server

The Lightweight Directory Access Protocol (LDAP) Server is installed on the primary server by default, but you can move it to any secondary server.

To move the LDAP server

1. From the `pdm_edit.pl` top menu select **L**, and press Enter.
The LDAP server submenu appears.
2. Enter **E** to specify the hostname of the server or enter **primary** for the primary server or the hostname of the secondary server.
The LDAP server is moved.

Add Repository Daemons

One repository daemon exists on the primary server by default, but you can add more.

To add repository daemons

1. From the `pdm_edit.pl` top menu select **R**, and press Enter.
The repository daemon submenu appears.
2. Enter **A** to specify the hostname of the server.
3. Enter **primary** for the primary server or the hostname of the secondary server.
The repository daemon is added to the server.

Create an Alias Example

Aliases are optional in CA SDM. This example demonstrates how you prepare to create an alias, and then how to create the alias.

Example: Prepare to Create an Alias

Before you create aliases, do the following:

1. Define object managers and add some to groups. The names of the object managers are: `domsrvr:group1:11`, `domsrvr`, `domsrvr:seattle:12`, `domsrvr;seattle:13`, `domsrvr:Tacoma:11`.

2. Define web engines and assign object managers to the web engines. If you accept the default, the object manager is assigned domsrvr.
3. Enter a regular expression that matches a group of object managers. If the web engine connects to any object manager, enter ANY. ANY is a default alias. The web engines and Java clients use the aliases to specify ANY or a subset of the running object managers to which they can connect. For example, you want the Java clients located in Washington to connect to object managers located in Seattle. You connect them to /domsrvr:seattle.*. Or, you can define an alias named SEATTLE and assign it the value /domsrvr:seattle.*.

Example: Create an Alias

This example assumes that you defined an object manager and associated it with a group named seattle. If not, the alias is created, but a warning is displayed. This warning indicates that the defined alias does not match any existing object managers.

1. From the pdm_edit.pl top menu select A, and press Enter.
The Alias submenu appears.
2. Enter a name for your alias. In this example, enter **SEATTLE**.
You are prompted for the regular expression that is to be associated with this alias.
3. Enter **/domsrvr:seattle.***, which matches any object manager with a name that starts with domsrvr:seattle. You can use this alias when you configure web engines and Java clients.
4. Save and exit the pdm_edit.pl utility.
An alias install script (named alias_install.bat or alias_install.sh) is created.
5. Run the alias install script on your primary server. This creates the aliases and transmits them to all secondary servers and java clients.

Move Knowledge Daemons

Knowledge daemons provide the knowledge base for CA SDM. They are installed on the primary server by default, but you can use the pdm_edit.pl utility to move them to secondary servers.

To move knowledge daemons

1. From the pdm_edit.pl top menu select K, and press Enter.
The Knowledge Daemons submenu appears.
2. Enter **E** to change the host.
3. Enter **primary** for the primary server or the name of the secondary server.
The knowledge daemons are moved.

Customization Installation Considerations

Consider the following when you install customizations:

- `pdm_edit.pl` creates a `pdm_startup.dat` file. Keep this file so that you can edit your current selections.
- `pdm_edit.pl` creates `pdm_startup.rmt` file. This file is a replacement for the `pdm_startup.tpl` file found on the primary server in the `$NX_ROOT/pdmconf` directory.
- If you are using a servlet server like Tomcat, a servlet replaces the CGI interface. These servers need a `web.xml.tpl` file. `pdm_edit.pl` creates one `<hostname>-web.xml.tpl` file for each web server. Use the appropriate file to replace the `web.xml.tpl` file in the `$NX_ROOT/bopcfg/www/CATALINA_BASE/webapps/CAisd/WEB-INF` directory on each of the servlet servers.
- If you are using an HTTP server like IIS or Apache, create CGI interface executables in your `NX_ROOT/bopcfg/www/wwwroot` directory by copying and renaming the `pdmweb.exe` file. Do this copy and rename for every CGI interface for web engines and web directors on every server hosting a web client.
- Each web engine needs a `web.cfg` file. `pdm_edit.pl` lists `web.cfg` files and the names to use for saving them on the appropriate server. For example, on the primary server, six of these files exist. Name one `web.cfg` and name another `primary-web2.cfg`. `pdm_edit.pl` creates sample `web.cfg` files for each of the web engines. Copy them to the `$NX_ROOT/bopcfg/www` directory of the appropriate server.
- If you are using IIS and require server extensions added for each CGI interface, `pdm_edit.pl` creates `<hostname>-site.dat` files. Copy the files to the appropriate servers and rename them to `site.dat` in the `$NX_ROOT/bopcfg/www` directory.
- If you defined aliases, run the `alias_install.sh` or `alias_install.bat` file on the primary server.
- Reconfigure all secondary servers.
- Reconfigure the primary server without reinitializing the database.

Save and Exit pdm_edit.pl

You can save pdm_edit.pl changes and exit the utility.

To save and exit pdm_edit.pl

1. From the pdm_edit.pl top menu, select x, and press Enter.
You are prompted to select a template file. This template file serves as the model for the web.cfg files that are created automatically for the web engines.
2. Select one of the values displayed. For example, select 2 to select the web.cfg.tpl file currently in use on the primary server. This selection helps ensure that customizations that are installed on the system are captured in the new sample web.cfg.tpl files.
3. Press Enter.
Changes are saved and the command prompt appears.

Start Web Director

You can start Web Director from the command line, or configure your pdm_startup file to start it automatically.

To start Web Director

1. Open a command prompt.
2. Enter the following command:

`webdirector [-S slump_name] [-c cgi]`
3. Use variables as follows:

slump_name

Specifies the slump name of the web director. The slump name uniquely identifies the web director process. This argument is required only if you want to start more than one web director. Otherwise, we recommend leaving it at its default value of web:director.

Cgi

(Optional) Specifies the name of the CGI process used for the web director. If this argument is not supplied, it defaults to pdmweb_d1.exe.

Web Director starts and sends a message to all running web engines. The web engines associated with that web director (that is, those configured with UseDirector Yes, AfterLogin, or BeforeLogin, and WebDirectorSlumpName equal to the web director's slump name) respond with their configuration settings. When a web engine that is configured to use web director starts, the web engine sends its configuration settings to its associated web director. This action lets web director keep an up-to-date list of all active web engines.

How Web Director Handles User Sessions

You can let users use the URL of any web engine or that of the web director to access CA SDM.

Note: If all web engines associated with a web director are configured with UseDirector Yes (that is, web engines not using AfterLogin or BeforeLogin), web director functions strictly as a load balancer.

Web director handles user sessions as follows:

1. The following web director URL is published to users:

Windows and UNIX: `http://hostname/CAisd/pdmweb_d1.exe`
2. When a web director receives a request, it does the following for the web engine:
 - a. Determines if is most able to handle the request.
 - b. Examines the current session count.
 - c. Calculates the current willingness by using the following formula:

$$\text{current willingness} = (\text{configured willingness}) / (\text{session count} + 1)$$
3. The web director redirects the request to the web engine with the highest current willingness.
4. The selected web engine handles the entire session from login to logout.

If a user in the environment attempts to access a web engine directly by using its own URL, the web engine sends a message to the web director requesting a referral. If the web director responds with the URL of a different web engine, the original web engine redirects the session to the recommended web engine.

How to Set Up Web Director in an SSL Environment

Consider the following information if one or more of the web engines associated with a web director are configured with UseDirector AfterLogin, and the rest of the web engines are configured with UseDirector BeforeLogin:

- **UseDirector AfterLogin**—When a web director receives a request, it determines the web engine most willing to handle the request, but only selects from the web engines with UseDirector AfterLogin. For login, the web director considers web engines with a willingness of zero. When a web engine configured with UseDirector AfterLogin receives a request (either by referral from the web engine or by a user directly accessing its URL), it authenticates the user by whatever method is configured in the user access type. After the user is authenticated, the web engine asks the web director for a referral, and the web director again selects a web engine, this time from all web engines associated with it, and configured with UseDirector BeforeLogin (except the web engines with a willingness of zero). When the authenticating web engine receives the referral, it transfers the session to the recommended web engine.

- **UseDirector BeforeLogin**—If a user in this environment attempts direct access to a web engine configured with UseDirector BeforeLogin, the web engine sends a message to the web director asking for a referral. The web director responds with a referral to a web engine configured with UseDirector AfterLogin, and the original web engine transfers the request to the recommended web engine.

To set up web director in an SSL environment, we recommend that you configure one HTTP server for secure sockets, and associate a single web engine with that server as follows:

1. Read and understand how [web director handles sessions](#) (see page 354).
2. Configure the single web engine with UseDirector AfterLogin and WillingnessValue 0.
3. Configure one or more additional web engines to connect with a second HTTP server that uses standard HTTP protocol.
4. Configure these web engines with UseDirector BeforeLogin and WillingnessValues appropriate to the relative capacity of their host computers.

The web director can direct all logins to the secure server (because servers with willingness zero are still eligible for login), but refers the rest of the session to one of the other web engines.

Improve Performance With Browser Caching

The CA SDM web interface uses many JavaScript, style sheets, and image files, which can be fairly large and affect performance.

To improve performance of the web interface, set up your HTTP server so that the user browser caches these files, causing them to be loaded only once a day.

The web interface performance improves.

Note: The default installation automatically configures caching for Apache and IIS; however, you can configure it manually.

Configure the Microsoft Internet Information Server

You can configure Microsoft Internet Information Server (IIS) to notify the browser that files loaded from the CA SDM directory expire one day after loading. This means that the browser queries the server for these files only once a day, regardless of how many times they are used.

To configure the Microsoft Internet Information Server

1. Launch the Internet Services Manager application (for Windows 2000 and XP, select Programs, Administrative Tools, Internet Services Manager).
2. Navigate to the CA SDM file folder, which is typically CAisd:
 - a. Click the plus sign adjacent to the server running the CA SDM web interface.
 - b. Click the plus sign adjacent to Default Web Site.
 - c. Scroll down to CAisd.
3. Right-click the CAisd folder, and select Properties.

The Properties page appears.
4. Click the HTTP Headers tab.
5. Select the Enable Content Expiration check box.
6. Select the Expire After option, enter 1 into the text field, and select a day from the drop-down list.
7. Click OK.

The properties are saved and the changes take effect immediately.

Configure Apache

You can configure Apache to notify the browser that files loaded from the CA SDM directory expire one day after loading. This configuration means that the browser queries the server about these files only once per day, regardless of how many times they are used.

You configure Apache by updating a text configuration file. The default installation modifies your active configuration file in the apache conf directory (typically httpd.conf) to contain the statement:

```
Include installation-directory/bopcfg/www/CAisd_apache.conf
```

Installation-directory must be replaced with a full path. On Windows, this path is typically c:\Program Files\CA\CA SDM. On UNIX, replace *installation-directory* with the value of \$NX_ROOT.

The file CAisd_apache.conf, which is referenced by the Include statement, contains the following text. Again, *installation-directory* is replaced with the full path as it was in the Include statement.

```
<IfModule mod_alias.c>
    Alias /CAisd installation-directory/bopcfg/www/wwwroot/
    <IfModule mod_expires.c>
        <Directory installation-directory/bopcfg/www/wwwroot>
            ExpiresActive On
            ExpiresDefault "access plus 1 day"
        </Directory>
    </IfModule>
</IfModule>
```

To configure Apache manually for browser caching of CA SDM files, include statements similar to those in `CAisd_apache.conf` in your Apache configuration file. You can either add them directly to the file, or add an Include statement referencing a separate file, like the default installation.

Changes to Apache configuration files take effect only after you recycle Apache.

Clear the Cache

If you change a JavaScript, image, style sheet, HTML, or help file loaded by the HTTP server itself, you must instruct your users to clear their browser cache.

Note: For changes to HTML files to take effect, you must either recycle the web engine or use the `pdm_webcache` utility. In a development environment, you can avoid this task by specifying the configuration file property `SuppressHtmlCache`.

To clear the browser cache for Internet Explorer

1. Select Tools, Internet Options.
A Internet Options dialog appears.
2. Click Delete Files.
A confirmation window appears.
3. Click OK.
The browser cache is cleared.

To clear the browser cache for Firefox

1. Select Tools, Clear Private Area.
2. Click the Clear Private Data Now button.
The browser cache is cleared.

Record Locking Behavior in the Web Interface

When a user edits a database record using the web interface, the user is given an exclusive lock at the record for the default time of two minutes. You can change the default time by using the `ExclLockSeconds` property in the `web.cfg` file.

The following conditions affect if a database record is updated by a user's changes:

- If a user *can* edit and submit changes within an allotted time, the changes are included the database.

During the time that the database record is locked, other users (both web and non-web) can view the record, but cannot edit it. If another user tries to edit the record while it is locked, an error message appears.

- If a user *cannot* edit and submit changes within an allotted time, the record lock automatically drops and other users can edit the record.

When the user submits the updates, time stamps are checked to ensure that nobody else has changed the record and the following occurs:

- If the record has not been changed since the exclusive lock was dropped, the user's updates are saved to the database.
- If another user has edited the record after the lock expired, the user receives an error response from the attempt to save, and the changes are not saved. The user must restart the edit process and re-enter the changes.

More information:

[Configuration File Modification](#) (see page 359)

Print CA SDM Web Pages

CA SDM uses background graphics to format its buttons and notebook tabs. The default setting for many browsers is *not* to print these background images. Therefore, if you select File, Print on either the browser menu or the CA SDM menu, the printed page shows only the corners of the buttons or tabs.

To print CA SDM web pages on Internet Explorer

1. Select Tools, Internet options.

The Internet Options dialog appears.
2. Select the Advanced tab.
3. Scroll down to the Printing heading, and select the Print Background Colors and Images option.

The printed CA SDM web pages include background graphics.

To print CA SDM web pages on Firefox

1. Select File, Page Setup.
2. Select the Format & Options tab.
3. Select the Print Background (color & images) check box.

The printed CA SDM web pages include background graphics.

Configuration File Modification

When you install the CA SDM web interface, a sample web engine configuration file (*web.cfg*) is installed that you can modify to suit your needs. The *web.cfg* file itself contains helpful comments that you can read by viewing the file. You can open the *web.cfg* file from the appropriate directory:

- (Windows) %NX_ROOT%\bopcfg\www\
- (UNIX) \$NX_ROOT/bopcfg/www/

Note: Some additional configuration variables, such as charset, are also available in Options Manager. These are accessible using the Administration tab in the web interface. For more information, see the *Online Help*.

AllowInactiveSrelEntry

Specifies whether a record can be saved when it references inactive records in a reference table.

- When this property is omitted or set to zero, inactive reference table entries (such as request status or change category) are not included in drop-down selects, and cannot be specified for lookup or hierarchical search fields.
- When this property is set to 1, the inactive flag is ignored in reference table entries.

Regardless of the setting of this flag, records already containing a reference to an inactive reference table entry can be saved without changing the reference; the flag affects only new field values.

AnnouncementLength

Specifies the maximum number of announcements to display on the opening screen for both customer and analyst interfaces. CA SDM begins the display with the most recent announcement, continuing for the number of announcements specified by this parameter. Users of the analyst interface can view additional announcements by selecting Announcements from the Search menu.

Default: 10, meaning that the ten most recent announcements display.

AnonymousPrio

Specifies valid priorities for tickets created by guest users. Such users can specify only one of the priorities in the AnonymousPrio list for their tickets. Entries in the list of priorities are separated by spaces. Each entry must be either a number between 1 and 5 or the word “none” (without quotes). The default priority for tickets created by guest users should be specified first, and can be repeated in the list.

The values valid for AnonymousPrio correspond to the symbolic names of the priorities as distributed. You can use the java client to customize these symbolic names; however, this does not affect the specification for AnonymousPrio, which must continue to reference priorities by their default names, where 1 corresponds to the highest priority.

Default: None, meaning that all requests created by a guest user have a priority of none.

Autofill

Specifies that the web interface should auto-fill lookup fields when a user keys data into them and presses Tab to exit the field. When a user does this and the Autofill option is selected, the browser asks the server to confirm that the update is correct. This results either in the full name filling in the field (if the user provided a partial name), or a pop-up search window appearing (if the user’s selection was incorrect or ambiguous).

This property is optional. Autofill is enabled by default so if this property ID is omitted or set to Yes, tabbing out of a lookup field automatically searches the database. If this property is set to No, no auto-fill occurs, and lookup fields are not verified until the record is saved.

CAisd

Specifies the path (including a leading slash) to the alias or virtual directory in your HTTP server that contains the files needed by the CA SDM web server. This property typically has a value of /CAisd in both UNIX and Windows installations. For Apache servers, it should be defined in an Alias statement in a configuration file. For IIS, it should match an Alias field in the Directory Properties Window.

CGI

Specifies the name of the CGI executable program supplied with the web interface (without the .exe suffix).

Default: pdmweb

Note: If you rename this program, you must update this property.

CgiReport

Specifies the name of the CGI executable program for web reports supplied with the web interface (without the .exe suffix).

Default: pdm_cgireport

Note: If you rename this program, you must update this property.

ContactAutoDesc

Specifies whether the contact's name should be inserted into the description of new issues and requests created in the customer and employee interfaces. If this property is omitted or specified as 0, no automatic information is added to the description of new issues and requests. If this property is specified as 1, the contact's name is automatically inserted into the description of issues and requests created in the customer and employee interfaces. This property has no effect on the analyst interface.

ContactAutoDescWithIP

Specifies whether the contact's IP address should be inserted into the description of new issues and requests created in the customer and employee interfaces. If this property is omitted or specified as 0, no IP address information is added to the description of new issues and requests. If this property and the ContactAutoDesc property are both specified as 1, the contact's name and IP address are automatically inserted into the description of issues and requests created in the customer and employee interfaces. This property has no effect on the analyst interface. It is ignored unless ContactAutoDesc is 1.

CstPrio

Valid priorities for issues created with the customer web interface. Users of the customer interface can only specify one of the priorities in the CstPrio list for their issues and cannot update the priority of an issue if an analyst has altered it to a value that is not in the list.

Entries in the list of priorities are separated by spaces. Each entry must be either a number between 1 and 5 or the word "none" (without quotes). The default priority for issues created with the customer interface should be specified first (and can be repeated in the list).

Default: none, 3, 4, 5

The values valid for CstPrio correspond to the symbolic names of the priorities as distributed. You can use the java client to customize these symbolic names; however, this does not affect the specification for CstPrio, which must continue to reference priorities by their default names, where 1 corresponds to the highest priority.

DateFormat

Defines the order of elements in dates.

Default: MM/DD/YYYY hh:mm a(am,pm)

Symbol	Description
M	Print 1 or 2 digits of month
MM	Print 2 digits of month
D	Print 1 or 2 digits of date
DD	Print 2 digits of date
YY	Print 2 digits of year
YYYY	Print 4 digits of year
H	Print 1 or 2 digits of hours on 24 hour clock
HH	Print 2 digits of hours on 24 hour clock
H	Print 1 or 2 digits of hours on 12 hour clock
hh	Print 2 digits of hours on 12 hour clock
M	Print 1 or 2 digits of minutes
mm	Print 2 digits of minutes
S	Print 1 or 2 digits of seconds
ss	Print 2 digits of seconds
a(am, pm)	Print am and pm as a string

DateFormatNoTime

Specifies the same definition as DateFormat, but without specifying the time portion.

DebugSource

Enables the standard browser right-click menu on CA SDM forms. When this property is not set, you can right-click a form to display a CA SDM menu. You should use caution when setting this property, because some of the options on the standard browser right-click menu can cause execution errors (this is the reason why it is usually disabled). On Internet Explorer, you can display the standard browser right-click menu although the DebugSource property is not set by pressing the Ctrl key when you right-click.

DebugTrace

Causes the web engine to write trace information to the stdlog file.

Important! This property should *not* be set for typical use. It should only be used when requested by CA Support.

EmpPrio

Valid priorities for requests created with the employee web interface. Users of the employee interface can only specify one of the priorities in the EmpPrio list for their requests and cannot update the priority of a request if an analyst has altered it to a value not in the list.

Entries in the list of priorities are separated by spaces. Each entry must be either a number between 1 and 5 or the word “none” (without quotes). The default priority for requests created with the employee interface should be specified first (and can be repeated in the list).

Default: none, 3, 4, 5

The values valid for EmpPrio correspond to the symbolic names of the priorities as distributed. You can use the java client to customize these symbolic names; however, this does not affect the specification for EmpPrio, which must continue to reference priorities by their default names, where 1 corresponds to the highest priority.

ExclLockSeconds

Specifies the maximum number of seconds that a user is given an exclusive lock on a record after clicking Edit. After this period elapses, the web engine releases the lock, allowing other users to update the record. The web engine attempts to retake the lock if a user asks to save after ExclLockSeconds has expired. This attempt succeeds only if no other user has updated the record while the lock was available. If the attempt to retake the lock fails, the user must re-enter the updates.

Default: 120 (two minutes)

This argument is optional. If omitted, the default value is assumed.

Note: The ExclLockSeconds setting must be shorter than the Timeout setting. ExclLockSeconds is specified in seconds and Timeout is specified in minutes.

FormCacheMax

Specifies the maximum number of forms to be retained in web engine memory for each user. The web engine always retains the last FormCacheMax forms used by each user. Forms beyond this number are eligible to be timed out. Timed out forms cannot be accessed by the Back or Forward buttons on the main page, and can no longer be submitted on a pop-up form.

Default: 10

Timed out forms save memory in the web engine, but they occasionally require users to manually refresh them. You can set FormCacheMax to -1 to disable the FormTimeout feature.

FormTimeout

Specifies the minimum number of seconds that a form is retained in the web engine before it is eligible for removal from the cache. Users always have at least the number of seconds specified in this parameter to work on a form before submitting it. In addition, the web engine always retains the most recently used FormCacheMax forms for each user.

You can use the StayCacheList property to prevent specified forms from timing out.

Default: 180 (3 minutes)

FormTitle

Specifies a string to be included in the title bar of a web browser displaying a CA SDM web form. The value of FormTitle supplements the title of the specific form displayed.

Default: CA SDM

For example, if the default value is retained, and Microsoft Internet Explorer is used to display the Announcement Detail form, the title bar displays the following:

Announcement Detail—CA SDM—Microsoft Internet Explorer

This property is optional. If omitted, the analyst web interface does not use a constant value in the title. The customer and PDA web interfaces revert to the default value.

HitTrackFile

Specifies the full path to a file that receives a log of all web pages used. One line is written to this file each time a user requests a page. The file can grow indefinitely, so be cautious when specifying this property.

Note: Records containing a time stamp, user ID, database record ID, and HTML form name are appended to this file. The format of the records may change. You must periodically maintain this file so that it does not get too large.

This property is optional. If omitted, no hit tracking file is written.

HtmlCacheSize

Specifies the size of the HTML cache. When this size is exceeded, the least used form is removed from the cache.

Default: 1000.

ListAllMaximum

Specifies the maximum number of records that can be displayed in a list before a request to display the entire list produces a pop-up warning message advising the user that the request adversely impacts performance and is not allowed.

Default: 2500

ListAllWarn

Specifies the maximum number of records that can be displayed in a list before a request to display the entire list produces a pop-up warning message advising the user that the request may adversely impact performance, and asking for confirmation.

Default: 1000

ListBottomMaximum

Specifies the maximum number of records that can be displayed in a list before a request to scroll to the bottom produces a pop-up warning message advising the user that the request adversely impacts performance and is not allowed.

Default: 2500

ListBottomWarn

Specifies the maximum number of records that can be displayed in a list before a request to scroll to the bottom produces a pop-up warning message advising the user that the request may adversely impact performance, and asking for confirmation.

Default: 1000

ListPageLength

Specifies the maximum number of found records to be shown on a list page after performing a search.

Default: 10

LogoutURL

Specifies the complete URL of a web page to be displayed after a user logs out of CA SDM. This property is optional. If it is not specified, logging out returns the login form.

Lr_Refresh

Specifies the log reader refresh interval in seconds. If this property is non-zero, the Notification Log Reader automatically refreshes itself at the specified interval (with a minimum of thirty seconds).

This property is optional. If omitted, the log reader refreshes itself every 5 minutes (a default value of 300 seconds). If this property is specified as zero, the log reader does not automatically refresh at all.

MacroPath

Specifies a list of directory paths that the web engine searches to find files requested by the PDM_MACRO tag. You can specify multiple directories separated by spaces. You can include environment variables in the directory names by prefixing them with a dollar sign (for example, \$NX_ROOT). For both Windows and UNIX, separate path components with a forward slash (/), not a backslash (\). This property is required. It is typically set as follows:

```
$NX_ROOT/site/mods/macro $NX_ROOT/bopcfg/www/macro
```

MatchesFound

Specifies the text of the message to display under a field when a user's key for a lookup field is ambiguous and the edit form must be redisplayed with a drop-down select list. This property is optional; if omitted, it defaults to Multiple Matches.

MaxRecordsAutoSuggest

Specifies the number of records that autosuggest displays when search as you type or autosuggest displays suggestions below a lookup.

Default: 25

MaxSelectList

Specifies the maximum number of matches to display in the drop-down select list shown when a user's key for a lookup field is ambiguous and the edit form must be redisplayed. If more than this many matches are found, the message specified for Too Many Matches is displayed.

MinCharsAutoSuggest

Specifies the minimum number of characters to enter in the lookup fields before search as you type or autosuggest displays suggestions.

Default: 3

MouseoverPreviewDelayTime

Specifies the delay time (in milliseconds) between hovering the mouse pointer over a link and the display of the mouseover preview.

If the mouse moves away from the link before the delay time expires, the preview does not appear.

Default: 1000

NoMatchesFound

Specifies the text of the message to display under a field when a user's key for a lookup field is incorrect and the edit form must be redisplayed. This property is optional; if omitted, it defaults to No Matches Found.

PreLogin Timeout

Specifies the maximum number of minutes the web engine keeps a session active before login. The web engine automatically starts a session when a user requests a login form, in anticipation of the user completing the login. If the user does not login within the time period specified, the web engine destroys the session. If the user subsequently logs in, the web engine creates a session that is transparent to the user.

This property has no end-user impact. Its sole purpose is performance—balancing web engine memory usage versus the overhead of destroying and recreating a session. This property is optional; if omitted, it defaults to one minute.

RedirectingURL

Specifies the URL the WebDirector should use to send requests to this web engine. This property specifies the full URL of the web engine, including http. This property is required if you are using WebDirector. It is ignored if you are not.

SchedExpMaximum

SchedExpMaximum

Specifies the limit of the number of schedule events that can be exported at a time.

Default: 1000

Important! The default is the maximum exports that CA SDM can handle at a time. Increasing this default could cause system instability. If you attempt to export more than the value specified in SchedExpMaximum, a message appears refusing your exporting request.

SellistCacheExclude

SellistCacheExclude

Specifies the names of the factories (objects) to be excluded from caching for <PDM_SELECT> lists. To improve performance, the web engine usually caches in its own memory of the contents of small tables used in <PDM_SELECT> (drop-down) lists and hierarchical search lists. You may want to suppress caching for a table if you are using data partition constraints to specify that different users should receive different views of the table. In addition, including tables in the value of this property eliminates the need for the web engine to query their record count at startup, slightly improving startup performance. This property is optional. If specified, it should contain one or more object names separated by spaces.

SelListCacheMax

SelListCacheMax

Defines the maximum number of records in a table that can be cached in the web engine. The web engine keeps the entire contents of tables at or below its cache size in its own memory, improving its performance in building <PDM_SELECT> lists using these tables. Specifying a higher value for this property improves performance at the expense of memory usage.

Default: 10

SelListCacheMax is ignored for tables used in hierarchical search lists, such as category on requests, issues, and change orders. The web engine always stores the entire contents of tables used in hierarchical search lists in its own memory. If you have a large number of values in any of these tables, you may want to specify the SelListCachePreload property.

SelListCachePreload

Specifies one or more tables to be loaded into the web engine's select cache at startup time. Tables not specified in this property are loaded the first time they are used. If SelListCacheMax is large, or if you have a large number of records in a hierarchical search list (such as category), you may want to specify the table in SelListCachePreload. This avoids a response time delay the first time a user accesses a form using the table.

The specification for the SelListCachePreload property is a blank-separated list of object names. Each object name can be followed by an optional list of attribute names in parentheses. The attributes specified in the list are loaded into the web engine. If no attributes are specified, only the common name and rel attr value of the object are loaded. This is sufficient for drop-down selects, but may not be sufficient for hierarchical searches. If you modify the hierarchical search forms (hiersel_xx.html, where xx is an object name), be sure that the SelListCachePreload property specifies every attribute used in the form. If you omit an attribute, the cache is reloaded when the form is used.

The SelListCachePreload property is optional. If it is omitted, nothing is loaded into the select cache until a user requests a form using a drop-down select or a hierarchical search.

```
chgcatt(description owning_contract) chgstat crs isscat(description owning_contract) issstat pcat(description  
cr_flag in_flag pr_flag owning_contract) pri tskstat urg pcat_cr(description cr_flag in_flag pr_flag  
owning_contract) pcat_pr(description cr_flag in_flag pr_flag owning_contract) pcat_in(description cr_flag  
in_flag pr_flag owning_contract)
```


StayCacheList

Specifies the names of forms that are never removed from the forms cache, regardless of the length of time they have been displayed. This property ensures that the fixed frames on a frame display remain for the lifetime of a session. It can be used with caution to cause other forms to be permanently cached. The default is as follows:

scoreboard.html top_splash.html buttons.html hiersel_admin_tree.html

SuppressHtmlCache

Specifies that the web engine should reread all files defining the contents of a page each time the page is requested. Parsing an HTML file takes a significant amount of web engine processing time, and usually involves reading several physical files (because most pages use PDM_INCLUDE tags). The web engine normally saves parsed files in its own memory so that future requests for the same page can be satisfied immediately. This markedly improves performance, but can be inconvenient for users in the process of developing new or updated pages, as the web engine must be recycled for changes to take effect.

This property is optional and requires no value. If it is specified, the web engine does not cache parsed files, and changes to HTML files take effect immediately. Because of its impact on performance, this property should not be specified in a production environment.

SuppressLoginAndLogoutMsg

Specifies that the web engine should not log a message to the CA SDM log file each time a user logs in or logs out of the web interface.

This property is optional. If it is not specified, the web engine logs a message each time a user logs in or logs out.

SuppressMacroCache

Specifies that the web engine should discard all saved macros each time a new page is requested. The web engine normally saves parsed macros in its own memory so that future requests for the macro can be satisfied immediately. This improves performance, but can be inconvenient for users in the process of developing new or updated macros, as the web engine must be recycled for changes to take effect.

This property is optional. If it is specified, the web engine does not cache parsed macros, and changes to macros take effect immediately. Because of its impact on performance, this property should not be specified in a production environment.

Timeout

Specifies the number of minutes that a user's session can be idle before it is automatically terminated, freeing up all server resources.

Note: The Timeout setting must be longer than the ExclLockSeconds setting. ExclLockSeconds is specified in seconds and Timeout is specified in minutes.

TooManyMatches

Specifies the text of the message to display under a field when a user's key for a lookup field is ambiguous and the number of matches for the key exceeds the value of MaxSelectList. This property is optional; if omitted, it defaults to Too Many Matches.

UpdatedAnnouncementsPopup

The interval that browser checks for a new announcement. When a new announcement is found, it automatically shows the announcement in a popup window. The interval value is in minutes. To reduce the impact to the browser performance, it's recommended to set this variable to the value greater than 5 (minutes).

UseDirector

Specifies when the WebDirector is controlling this web engine. The following table defines possible values:

Value	Description
No	The web engine is independent of the WebDirector. This is the default value.
Yes	The WebDirector must initiate all sessions, including the login form. If a user attempts to make a direct connection to the web engine, the web engine asks the WebDirector for a referral.
AfterLogin	The web engine refers a session to the WebDirector after authenticating a user. A web engine configured with UseDirector AfterLogin is responsible solely for authentication, and is thereby a candidate for the use of secure sockets (SSL) for maximum security.
BeforeLogin	The web engine refers a session to the WebDirector before authenticating a user. A web engine configured with UseDirectory BeforeLogin never displays a login page, and never accepts a login password.

This property is optional. If omitted, the web engine does not use the WebDirector.

UseNestedTabs

Specifies whether to display the nested tab control on detail forms.

Default: Enabled

WebDirectorSlumpName

Specifies the name of the WebDirector servicing this web engine. This property is needed only if you are running more than one WebDirector, or if you have configured your WebDirector to use a slump name other than its default of web:director.

This property is optional if you are using WebDirector. It is ignored if you are not.

WillingnessValue

Specifies the willingness of this web engine to accept sessions, based on a scale of 0 to 10. This property is used only if you are using WebDirector. This value is meaningful only in comparison with the willingness of other web engines associated with the same WebDirector. The WebDirector transfers sessions to web engines in proportion to their willingness values. A web engine with a willingness value of twice that of another web engine's, on average, services twice the number of sessions.

A WillingnessValue of zero means that the web engine does not accept any sessions. This value can be useful when UseDirector is AfterLogin.

This property is optional if you are using WebDirector. It is ignored if you are not. If omitted, the web engine sets its willingness to 5.

WorkFrameTimeout

Specifies the maximum number of seconds the web engine waits for a response to an internal server request before concluding the request has failed. The workframe is used for CA SDM web interface features requiring server data other than normal web pages. This includes such features as autofill, loading category properties, and updating scoreboard counts. Workframe requests to the CA SDM are unlikely to fail. However, workframe requests to other servers (such as integrated products, such as Knowledge Management) may fail if the targeted server is not running, or if a network problem prevents access to it. The WorkFrameTimeout property specifies a length of time before the request is considered to have failed, and the workframe is available for other requests.

Note: WorkFrameTimeout is not checked unless a workframe is needed and all workframes are in use. Therefore, it is quite likely that a remote server is going to have more time than that specified for WorkFrameTimeout to respond. The value of WorkFrameTimeout is a minimum.

This property is optional. If omitted, the web engine uses a 30 second work frame timeout.

Chapter 10: Configuring Auto Assignment

This section contains the following topics:

[Auto Assignment](#) (see page 373)
[Auto Assignment Relationships](#) (see page 374)
[Auto Assignment Methods](#) (see page 374)
[How to Begin Implementing Auto Assignment](#) (see page 375)
[Default Group and Assignee](#) (see page 380)
[Auto Assignment Enablement](#) (see page 380)
[Auto Assignment Override](#) (see page 381)
[Assignment Controls](#) (see page 382)
[Activity Logging](#) (see page 384)
[Enable Activity Logging for Additional Attributes](#) (see page 385)
[Auto Assignment Tracing](#) (see page 385)
[Stored Queries](#) (see page 386)
[How Auto Assignment Assigns Tickets](#) (see page 386)
[How Auto Assignment Assigns Workflow Tasks](#) (see page 392)
[Configuration Item-Based Auto Assignments](#) (see page 395)

Auto Assignment

CA SDM *Auto Assignment* can greatly reduce the time required to manage tickets by automating the process of assigning them to analysts. This automation can make the tasks associated with balancing workloads and coordinating work schedules much more efficient.

You can configure auto assignment to assign tickets based on the following factors:

- Which analyst groups work on which tickets or tasks
- When the work must be done
- Which locations service the affected customers
- The workload and availability of each analyst
- The value of an attribute of a configuration item associated with the ticket

Note: [Configuration item-based auto assignment](#) (see page 395) lets you create group-specific assignments for Request/Incident/Problem tickets only.

You do not have to implement auto assignment all at once. You can develop a strategy for phasing it in gradually. For example, you can begin by enabling it only for selected ticket types, analyst groups, or locations.

Auto Assignment Relationships

The auto assignment process can involve many CA SDM elements. The element relationships are as follows:

- Areas and categories relate to groups, locations, and workshifts.
- Groups relate to locations and workshifts.
- CA Workflow templates relate to contacts.

To make auto assignment easy to administer, all relationships can be maintained from either related element. For example, when relating an analyst group to a change category you can maintain the association from either the Change Category Detail page or the Group Detail page.

Auto Assignment Methods

The following basic methods let you auto assign tickets:

Location-based auto assignment

Creates assignments for all ticket types based on the following:

- Areas and categories related to groups, locations, and workshifts
- Groups related to locations and workshifts
- CA Workflow templates related to contacts

Note: Location-based auto assignment is referred to simply as auto assignment.

Configuration item-based auto assignment

Creates group assignments for request, problem, and incident ticket types based on the following:

- Areas related to tickets associated with configuration items
- Configuration item attributes used to record contact/group information

Location-based auto assignment and configuration item-based auto assignment are exclusive options because you can select only one algorithm for use on a given Request/Incident/Problem Area. Location-based auto assignment and configuration item-based methods both serve to assign tickets when they are created; however, Configuration item-based auto assignment is distinct because it also reevaluates the assignments for a ticket whenever the Area or configuration item of a Request/Incident/Problem ticket is changed. If configuration item-based auto assignment is specified but does not successfully generate a group assignment for a ticket, the Area_Defaults option is consulted to determine if the default Group and Assignee values on the Area should be used to assign the ticket.

How to Begin Implementing Auto Assignment

Follow these guidelines to begin implementing auto assignment for selected analyst groups and analysts:

1. Identify one or more areas or categories for which you want to enable auto assignment.

Note: To review your site's area and category configurations, examine their settings in the CA SDM web interface. For instructions, see the *Online Help*.

2. By default, auto assignment is disabled. Enable it only for the [areas and categories](#) (see page 375) where you want to use it.
3. Build relationships between an identified area or category and the [analyst groups](#) (see page 375) that are eligible to be assigned tickets through that area or category.
4. Mark individual members of the [analyst](#) (see page 376) groups as available.

Areas and Categories

Note: While configuring areas and categories, consider setting [default groups and assignees](#) (see page 380).

To configure auto assignment of request, incident, and problem tickets, use the following controls on the Auto Assignment tab of Request/Incident/Problem Area Detail page:

- Update Groups
- Update Locations
- Update Workshifts

Note: The check box for enabling auto assignment is also located on the Auto Assignment tab of each of these pages. It is visible only while the page is in edit mode.

The following pages on the interface provide the same controls for configuring auto assignment of change orders and issues:

- The Change Category Detail
- The Issue Category Detail

Analyst Groups

To configure auto assignment you must, at a minimum, define the relationships between analyst groups and areas or categories. Assignees are chosen from groups that meet *all* of your specified auto assignment criteria. If no additional constraints are defined, tickets are auto assigned to the group member with the fewest active tickets.

If no groups are associated with an area or category, the default group and assignee are assigned. If these defaults are not defined, the ticket is left for manual assignment.

You can maintain the relationships between groups and areas or categories on the area or category detail page.

Alternatively, you can maintain the same relationships on the Auto Assignment tab of the Group Detail page by using the following controls:

- Update Request Areas
- Update Change Categories
- Update Issue Categories
- Update Locations

Analysts

Fields on the Analyst Detail page help determine whether the analyst is eligible for auto assignment.

Analysts are eligible for auto assignment only if they are marked as available.

The Analyst Detail page provides an Available check box that enables auto assignment.

Consider allowing analysts to control their own auto assignment availability. You can monitor analyst availability by using stored queries.

Note: The Available check box is not considered during auto assignment of workflow tasks.

The Work Schedule field allows analysts to have tickets auto assigned only during their scheduled workshift. Analysts that have no work schedule assigned but are marked as available are eligible for auto assignment at any time provided there are no other constraints that result in an ineligible status.

How to Auto Assign Tickets to a Group and Not to Contacts Within the Group

Auto assignment in CA SDM assigns tickets to contacts that have the Available option selected in the contact record. However, you can auto assign incidents/issues/change orders/requests to a group and not to contacts within the group. The `NX_AUTOASG_GROUP_ONLY` option controls the auto assignment behavior for groups. You install this option to auto assign the tickets to the group instead of individual contacts. `NX_AUTOASG_GROUP_ONLY` is not available on the web interface, you install it from the command prompt.

To auto assign tickets to the group and not to individual contacts in the group

1. Verify that your CA SDM installation is at a minimum level of Release 12.7 Cumulative 2 patch.
2. Open the command prompt on the CA SDM server.
3. Run the following command:

```
pdm_options_mgr -c -s AUTOASG_GROUP_ONLY -v 1 -a pdm_option.inst
```

By default, new options that you add to the nx.env file are not saved after you run pdm_configure. You can save the changes permanently by specifying the -t option as follows:

```
pdm_options_mgr -c -s AUTOASG_GROUP_ONLY -v 1 -a pdm_option.inst -t
```

The commands update the following files in CA SDM with the new option:

- Windows—NX_ROOT/NX.env and NX_ROOT/pdmconf/nx.env.nt.tpl
 - UNIX/Linux—NX_ROOT/pdmconf/NX.env.tpl
4. Open the NX.env file in the CA SDM installation folder and search for the variable @NX_AUTOASG_GROUP_ONLY=1 (it is located the end of the file).
 5. Open the file nx.env.nt.tpl present under NX_ROOT/pdmconf and search for the @NX_AUTOASG_GROUP_ONLY=1 option.
 6. Restart Unicenter ServiceDesk Daemon Server service.
CA SDM auto assigns the tickets to the group only.

Auto Assignment by Location

If your service area is large and consists of many locations that service different customer communities, you can use location as a factor in your auto assignment configuration as follows:

- **Location Assigned to an Area or Category**—If you assign a location to an area or category, tickets in that area or category are auto assigned only if a matching location is found. For example, a request ticket is auto assigned if there is an eligible analyst at the following locations:
 1. The affected asset's location
 2. The affected customer's location

If the affected asset or customer has no specified location, the request is assigned to the default group and assignee. If these defaults are not defined, the request is left for manual assignment. You can use the area or category detail pages to maintain relationships between locations and areas or categories.

- **Location Assigned to a Group**—If you associate a location with a group, only members of that group are eligible for auto assignment of tickets that pertain to their location. You can use group detail pages to maintain relationships between groups and locations.

To maintain location relationships with areas, categories, or groups, use the following controls on the Auto Assignment tab of Location Detail page:

- Update Request Areas
- Update Change Categories
- Update Issue Categories
- Update Groups

Examples: Use Location in Auto Assignment Configuration

The following are examples of how you can use location in your auto assignment configuration:

- **Auto assign tickets only at a specified location**—Tickets from other locations receive the default group and assignee, or are left for manual assignment. For example, you can have many users at your company headquarters, and smaller groups of users at regional offices. An analyst group stationed at headquarters services their local users, while mobile analyst groups visit the regional offices. You can configure auto assignment of tickets to the headquarters analysts only, and manually assign tickets to the mobile analysts.
- **Auto assign tickets by user or asset location**—You can restrict auto assignment eligibility to analyst groups at locations that match the affected user's location, or the affected asset's location. For example, your organization can have many offices, and tickets from each office are handled only by groups located in that office. You can relate each group to appropriate areas or categories *and* to the appropriate location. The auto assignment logic selects eligible analysts only from groups at the correct location.

Auto Assignment by Workshift

You can constrain auto assignment by relating an area or category to a workshift. The workshift determines the timeframe within which tickets are eligible for auto assignment. Tickets opened outside the hours of the workshift are assigned to the default group and analyst or left for manual assignment.

You can also use workshifts to control group and analyst eligibility for assignment:

- If you assign a work schedule to a group, analysts in that group are eligible for auto assignment only to tickets opened during their work schedule.

Note: A group's workshift is specified in the Work Schedule field of the Group Detail page.

- If you assign a work schedule to an individual analyst, the analyst is eligible for auto assignment only to tickets opened during that work schedule.

Note: An analyst's workshift is specified in the Work Schedule field of the Analyst Detail page.

During ticket creation, the auto assignment logic attempts to identify the analyst with the fewest active tickets in a group that meets the assignment eligibility criteria. If an appropriate analyst is not identified, the ticket is assigned to the default group and assignee. If these defaults are not defined, the ticket is left for manual assignment.

For workflow tasks associated with change orders or issues, auto assignment uses a simpler selection strategy. Assignees are selected from groups associated with workflow templates. Workflow task assignees can be of any contact type except group. When a task changes to pending status, auto assignment selects the contact that has the fewest change order or issue workflow tasks assigned. To prevent unwanted results, if the parent change order category or issue category is not enabled for auto assignment, tasks are not auto assigned. Because workflow tasks could potentially encompass individuals external to the service desk organization, relying on them to reflect their availability with the available flag could be problematic.

Note: Auto assignment of workflow tasks does not evaluate the Available flag, and is not available for categories configured to use CA Workflow processing.

You can use the area or category detail pages to relate a workshift to an area or category, or you can use the following controls on the Workshift Detail page:

- Update Request Areas
- Update Change Categories
- Update Issue Categories

Examples: Use Workshifts in Auto Assignment Configuration

The following are examples of how you can use workshifts in your auto assignment configuration:

- **Auto assign tickets to the shift on duty**—If your service desk operates 24 hours per day, you may want to configure auto assignment such that network outage problems are assigned to the shift on duty when the ticket is opened.
- **Allow auto assignment only during a specified shift**—You may want to further restrict auto assignment of tickets in some areas or categories. For example, if a site's application analysts are on duty only during day shift, you may want to auto assign application problems only during the day shift.

Default Group and Assignee

When the auto assignment logic you have configured is unable to identify a suitable group or assignee, the ticket is assigned to the default group and assignee. You can specify these defaults in the Group and Assignee fields of the following web interface pages:

- Request/Incident/Problem Area Detail
- Change Category Detail
- Issue Detail

If auto assignment cannot identify a suitable group or assignee and the defaults are not specified, the ticket is left for manual assignment.

Auto Assignment Enablement

Options and controls let you configure auto assignment. You can control whether auto assignment processing occurs as follows:

- For requests, incidents, and problems assigned to that area, use the Auto Assignment Enabled tab on the Request/Incident/Problem Area Detail page.
- For change orders, issues, and workflow tasks, select the Auto Assignment tab and Auto Assignment Enabled check box on the Change Category Detail page, Issue Category Detail page, and CA Workflow Template Detail pages.

Note: Click Edit to make the Auto Assignment Enabled check box visible.

Example: Enable Auto Assignment for a Request/Incident/Problem Area

This example demonstrates how to enable auto assignment for a Request/Incident/Problem area.

1. On the Administration tab, browse to Service Desk, Requests/Incidents/Problems, Areas.

The Requests/Incidents/Problems Area List page appears.

2. Click the area for which you want to configure auto assignment.

The Requests/Incidents/Problems Area Detail page appears.

3. Click Edit.

The Update Requests/Incidents/Problems Area page appears.

4. Select the Auto Assignment tab, and complete the following fields as follows:

Auto Assignment Mode

Specifies how auto-assignment occurs. You use the Configuration Item Based option to base the auto assignment on the CI Assignable Attribute value.

- **Disabled**—Bases the auto-assignment on the Area Defaults option when it is installed.
- **Configuration Item Based**—Bases the auto-assignment on the CI Assignable Attribute value.
- **Location Based**—Bases the auto-assignment on the location value.

Assignable CI Attribute

Specifies the configuration item attribute to use for the group assignment. You can enter a value directly or click the magnifier to search for an attribute.

5. Click Save.

The Request/Incident/Problem area is enabled to use default values that are entered automatically on request, incident, or problem tickets assigned to the area.

Auto Assignment Override

You can use the Auto Assignment Override (autoasg_override) option in Options Manager to determine whether auto assignment overrides an analyst or group that may have been set during the creation of a ticket.

Other CA SDM features may have set the assignee and/or group before auto assignment obtains control. You can customize your system by setting this option to one of the following values:

0

Honors the existing assignee and/or group. Auto assignment processing will not occur if the assignee and/or group was set during the creation of the ticket.

1

Ignores the existing assignee and/or group. Auto assignment processing will continue and attempt to find an assignee and/or group.

You can set the Assignee and/or Group in various ways:

- Manually by the Analyst
- Area Defaults and the Assignee set options

- Request Templates
- CA NSM integration

Note: Uninstalling the Auto Assignment Override option causes it to operate in its default mode, which is 0.

Assignment Controls

Options and controls let you configure auto assignment. CA SDM features can affect the assignee, the group fields, or both on a ticket before auto assignment processing occurs. We recommend that you review assignment controls before implementing auto assignment.

Manual Assignment

The analyst can manually select the assignee and/or group while creating a ticket.

Assignee_set Option

By default, CA SDM will automatically set the logged in user as the assignee of the request if that user is an analyst. An Options Manager option, Assignee_set, lets you control this behavior. This option is typically installed by default.

Iss assignee_set

The Iss assignee_set option automatically sets the logged in user as the assignee of the Issue if that user is an analyst. It is similar to Assignee_set, except it is for Issues instead of Requests.

Area_Defaults

The Area_Defaults option determines what happens when the request area is specified on a request's detail page. This option lets you set the assignee and group whenever the request area changes. The default assignee and group from the request area are used, and this processing occurs before auto assignment processing.

This option is not installed by default.

Note: The Category_Defaults option provides similar functionality for change orders. The Iss_Category_Defaults option provide similar functionality for issues.

Require Assignee and Group Options

Several options are available in Options Manager for controlling the creation of unassigned tickets. These options are global in scope. They affect the creation of all of the indicated ticket types, regardless of whether auto assignment is in effect. If the indicated condition is not satisfied, the new or updated ticket cannot be saved.

The following options control the creation of tickets without an assignee specified:

require_change_assignee

Application: Change Order Mgr

Description: Makes the Assignee field of a change order ticket required

require_issue_assignee

Application: Issue Mgr

Description: Makes the Assignee field of an issue ticket required

require_incident_assignee

Application: Request Mgr

Description: Makes the Assignee field of an incident ticket required

require_problem_assignee

Application: Request Mgr

Description: Makes the Assignee field of a problem ticket required

require_request_assignee

Application: Request Mgr

Description: Makes the Assignee field of a request ticket required

The following options control the creation of tickets without a group specified:

require_change_group

Application: Change Order Mgr

Description: Makes the Group field of a change order ticket required

require_issue_group

Application: Issue Mgr

Description: Makes the Group field of an issue ticket required

require_incident_group

Application: Request Mgr

Description: Makes the Group field of an incident ticket required

require_request_group

Application: Request Mgr

Description: Makes the Group field of a request ticket required

require_problem_group

Application: Request Mgr

Description: Makes the Group field of a problem ticket required

Templates

You can use templates to set the values on a new request, change order or issue. Assignee and group are among the fields that can be affected by templates.

CA Network and Systems Management Interface

When CA NSM and CA SDM are integrated and you are creating requests from CA NSM events, the user_parms parameter in writer rule definitions is passed to the Text API. The CA SDM writer process (tngwriter) defines its own replacement parameters for changing the text before sending it to the Text API. The keyword LOG_AGENT is added to the end of the input to set the log_agent for the request.

Note: You need to update the Text_API.cfg file for all additional fields that are passed from CA NSM Alert Management Systems to CA SDM. This file is used for integrations with web services, email and AHD.DLL.

More information:

[Keywords](#) (see page 426)

Activity Logging

Auto assignment logs information as events in the ticket's activity log. The success or failure of auto assignment is logged, and any anomalies that may have occurred during processing.

Enable Activity Logging for Additional Attributes

Activity logs display changes attribute values of objects. Objects such as all ticket types, classic workflow tasks, Surveys, Contacts, and Configuration Items feature activity logging. You can enable activity logging for additional attributes in WSP.

Note: You can view activity log simple attributes, including strings, integers, date, duration, and SREL types. You cannot view activity logs for list types, such as QREs and BREs.

Follow these steps:

1. Open WSP and start the Schema Designer.
2. Select the table that you want to modify, and add additional attributes.
3. For each attribute, add +AUDITLOG() to the site-defined UI_INFO field.
Important! When you add the AUDITLOG flag, you *must* also remove the val_fieldupdate_site() function to prevent duplicate activity logs.
4. Save and publish the schema
5. Open WSP again and add the new attributes to the appropriate detail forms.
6. Open CA SDM and define an Activity Association for the additional attributes.
7. Test your changes.

For example, open a previously-saved instance of the object and modify the affected attributes to verify that the appropriate activity logs appear.

Auto Assignment Tracing

In a complex auto assignment implementation, auto assignment may not make the decisions that were expected. Tracing can be enabled to help you understand the processing flow. Tracing is normally turned off but when turned on, numerous messages are written to the stdlog files in \$NX_ROOT\log that describe the various decisions made by auto assignment.

Use caution when implementing this in a high-volume installation, as the number of messages generated could cause the log files to grow and eventually wrap. On very active systems, performance degradation may occur. Tracing is controlled with the pdm_logstat utility. The parameters used by this utility are case-sensitive. Be sure to enter them as shown.

To turn tracing on, run the following command on the server:

```
pdm_logstat -f auto.pm milestone
```

To turn off tracing, run the following command on the server:

```
pdm_logstat -f auto.pm
```

Stored Queries

Two stored queries are provided for monitoring the availability of analysts. CA SDM managers can add these to their Scoreboard:

- **Available Analysts**—Analysts that are marked as available for auto assignment
- **Unavailable Analysts**—Analysts that are marked as unavailable for auto assignment

How Auto Assignment Assigns Tickets


Auto-assignment assigns tickets as follows:


1. The initial save operation of a new ticket invokes auto assignment.
If an area or category is not configured for auto assignment, processing stops.
2. Auto assignment determines whether `Autoasg_override` is installed.
If it is not installed and the ticket has an assignee or group, processing stops.
3. If work shifts are related to the ticket, the open date is evaluated to determine if a work shift includes the ticket.
If a work shift does not include the ticket, processing stops, and auto assignment attempts to assign the default group and assignee.
4. Auto assignment determines whether any groups are related to the ticket.
If no groups are related, processing stops, and auto assignment attempts to assign the default group and assignee.
5. Auto assignment filters out any groups with an associated work shift, where the open date is outside the time frame of the work shift. Groups without a work shift bypass filtering.
6. The following occurs for locations that are related to the ticket:
 - If this ticket is a request and it has a configuration item
The location of the configuration item is matched against the locations related to the request area. If no match occurs, processing stops, and auto assignment attempts to assign the default group and assignee. Otherwise, the customer location is matched against the locations related to the area or category. If no match occurs, processing stops, and auto assignment attempts to assign the default group and assignee.


- If locations are associated with the request area or category


If locations are associated with the request area and the configuration item (during request auto assignment processing) or the customer has no location, auto assignment stops processing.

7. Auto assignment filters out the eligible groups that have related locations that do not match the configuration item location (for requests only) or customer location. If no groups remain, processing stops, and auto assignment attempts to assign the default group and assignee.
8. Auto assignment creates a list of analysts from each of the remaining eligible groups.
9. Unavailable analysts are filtered out.
10. The remaining analysts are checked for Work Schedules. Those analysts that have Work Schedules are filtered out when the open date is not within the work schedules of the analyst.
11. If no analysts remain, processing stops and auto assignment attempts to assign the default group and assignee.
12. All the remaining analysts are ranked according to the number of active tickets assigned to them.
13. The analyst (and associated group) with the least number of active tickets is assigned to the ticket. If a tie occurs, the first analyst that occurs in the group is selected.










How Auto Assignment Assigns Workflow Tasks

Processing logic is used by Auto Assignment to assign workflow tasks as follows:

1. Auto assignment is invoked when the status of a workflow task changes to pending. If the workflow template that the workflow task was created from is not enabled for auto assignment, processing stops. If the parent change order category or Issue category is not enabled for auto assignment, processing stops.
2. Auto assignment checks to determine if Autoasg_override is installed. If not installed and the task has an assignee or group, processing stops.
3. The workflow template that the workflow task was created from is checked to see if any contacts are associated with it. If there are no contacts, processing stops.
4. Auto assignment builds a list of all the contacts that are members of the groups that are currently associated to the workflow template that the workflow task was created from. Any contacts in this list that are groups are filtered out.
5. All of the remaining contacts are ranked according to the number of active change order tasks or issue tasks assigned to them.
6. The contact and associated group with the least number of active tasks is assigned to the task.





Configuration Item-Based Auto Assignments

Configuration item-based auto assignment lets you create group-specific assignments that apply to specific scenarios. You can specify that for Request/Incident/Problem tickets opened for a particular Area, the value of an attribute of the configuration item associated with the ticket controls its assignment.

Configuration item-based and location-based auto assignments are exclusive options because you can select only one algorithm for use on a Request/Incident/Problem Area. Configuration item-based and location-based auto assignment modes both serve to assign tickets when they are created; however, configuration item-based auto assignment is distinct because it also reassigns tickets whenever the Request/Incident/Problem Area of a ticket or configuration item is changed.

Example: Request/Incident/Problem Area Assigns Tickets to a Group

When you configure the Network Area (for Requests/Incidents/Problems) to perform configuration item-based auto assignments using the `network_contact_uuid` (Network Operations) attribute as the Assignable CI Attribute value, any tickets that are opened for the Network Area are assigned automatically to the group specified in the Network Operations field of the CI associated with the ticket. If no CI is associated with the ticket, or if the Network Operations field value of the CI is blank or does not specify a group, the assignment does not occur. In such cases, the system acts in accordance with the Option Manager option, Area Defaults, and assigns the ticket using the Group and Assignee fields of the Category.

How Configuration Item-Based Auto Assignment Works

Request/Incident/Problem tickets must specify the following for configuration item-based auto assignment to occur:

- A configuration item and an Area.
- The Area must have the Auto Assignment option set to Configuration Item Based.

When an analyst creates and assigns a ticket to this Area, or changes the Area on an existing ticket, CA SDM examines the Assignable CI Attribute field of the Area. CA SDM uses the value of Assignable CI Attribute as the name of an attribute and then attempts to find an identically named attribute on the configuration item that is associated with the ticket. If the attribute on the configuration item includes a group, the ticket is assigned to that group.


The following diagram describes the process for configuration item-based auto assignment in more detail:

1. CA SDM checks for the following:
 - a. Is the `autoasg_override` option set?
 - b. Does a ticket specify an Area with its Auto Assignment `auto_assign_mode=2`?
 - c. Does the category specify an assignable CI attribute?
 - d. Does the ticket specify a configuration item?
 - e. Does the configuration item have an assignable CI attribute value?
 - f. Does the value specify a group in the CI attribute?
 - g. Is the assignee of the ticket a member of the ticket group?

2. When the answers are positive for all questions in the previous step, CA SDM sets the group attribute of the ticket to the group attribute of the CI, and logs the assignment in the activity log.

Note: The diagram shows how Configuration Item-Based Auto Assignment uses the `NX_CHECK_ASSIGNEE_IN_AREA_DEFAULTS` variable to determine if the Area option is enabled. `NX_CHECK_ASSIGNEE_IN_AREA_DEFAULTS` is a variable in the `NX.env` file, which is located in the `$NX_ROOT\` directory.

3. CA SDM assigns the ticket to the group.



More information:

[Auto Assignment Override](#) (see page 381)

Enable Configuration Item-Based Auto Assignments

Configuration item-based auto assignment lets you create group-specific assignments that apply to specific scenarios. You can specify that for Request/Incident/Problem tickets opened for a particular Area, the value of an attribute of the configuration item associated with the ticket controls its assignment. Configuration item-based auto assignment reassigns tickets whenever the Request/Incident/Problem Area of a ticket or configuration item changes.

Note: The Options Manager `autoasg_override` option controls the circumstances under which auto assignment takes place. When you set this option to 1, CA SDM ignores any existing assignee and group settings and auto assigns tickets in all cases. If you want CA SDM to auto assign tickets only if they are not already assigned, set the option value to 0.

To enable configuration item-based auto assignment

1. On the Administration tab, browse to Service Desk, Request/Incident/Problems, Areas.

The Request/Incident/Problem Area List appears.

2. Click the area to edit.

The Area Detail page appears.

3. Click Edit.

The Update Area page appears.

4. Select the Auto Assignment tab, and complete the fields as follows:

Auto Assignment Mode

Specifies how auto-assignment occurs. You use the Configuration Item Based option to base the auto assignment on the CI Assignable Attribute value.

Assignable CI Attribute

Specifies the configuration item attribute to use for the group assignment. You can enter a value directly or click the magnifier to search for an attribute.

Click Save.

Auto assignment is enabled. CA SDM performs configuration item-based auto assignments using the attribute that you specified in the Assignable CI Attribute field.

Chapter 11: Managing Your Database

This section contains the following topics:

[Database Management Utilities](#) (see page 401)

[Select and Configure your Database](#) (see page 401)

[Database Loader](#) (see page 402)

[Database Backup](#) (see page 404)

[Database Restore](#) (see page 405)

[Database Table Replacement](#) (see page 405)

[Data Extraction](#) (see page 406)

[Data Dereferencing](#) (see page 407)

[Use the Dbadmin Mode](#) (see page 410)

[Archive and Purge Rules](#) (see page 411)

Database Management Utilities

You can run utilities to manage your database while CA SDM is shut down. If you are running database management utilities and using a database other than the CA SDM default repository, the database environment variables must be set.

Note: For information about setting environment variables, see the documentation for the database you are using.

Select and Configure your Database

You can select the database and configure it for CA SDM use.

To select and configure a database

1. Select the database from the Database Type drop-down list and click Next.
The configuration page for the selected database appears.
2. Enter the database information on the database configuration page and click Next.
Your database configuration is complete.

Note: Use the Help button on this page to get configuration information about the fields for the selected database.

The following buttons appear on the configuration pages.

Help

To bring up the Help for the current page.

Cancel

To cancel the configuration settings you have made so far.

Back

To go to the previous page without saving any of the changes made so far in the current page.

Next

To go to the next page after saving all the changes made in the current page.

Database Loader

The database loader utility, `pdm_userload`, adds, updates, and deletes CA SDM database records. You create a formatted ASCII input file for the `pdm_userload` utility and select the tables to load and the optional fields to add.

Important! The `pdm_userload` utility accesses the CA SDM database and does not directly interact with application software processes. Inventory items added to the database with `pdm_userload` do not update the helper/selection lists until the application has been halted and restarted.

Although the `pdm_userload` utility can run with the application active, system performance is degraded. For best results, ensure that the CA SDM server is running but that no users are using the client interface before you run `pdm_userload`.

Note: If you need to add records with cross-referenced fields, use the `pdm_deref` utility.

More information:

[Data Dereferencing](#) (see page 407)

How to Create and Use an Input File

You can use an input file and the `pdm_userload` utility to populate database tables.

The input file format is as follows:

- Surround field values with double quotes ("value") and separate values with a comma and a space ("value1", "value2").
- Precede double quotes with a backslash (\) to embed them in text strings. To set a date/time field to the current date and time, use `@NOW@` for the input value.
- Surround each record with curly braces separated by spaces, as follows:
({ record field values })
- If properly delimited, input records can span more than one line in the input file as long as individual fields stay on one line. For a multi-line field such as comments, values can include a new line character (\n) to force a new line when the database field is displayed.
- Explicit new line characters are needed only for special formatting. Ordinary running text is automatically displayed with appropriate line breaks, as shown by the following example:

```
"Record status is \"COMPLETE\""
```

To create an input file for the `pdm_userload` utility, do the following:

1. Determine which table you want to load, and the fields you want to populate in that table.

You must populate the Name or Symbol key field for each record that you load.

2. Make a copy of the appropriate *filename.dat* file for the table that you are loading.
3. Edit your newly created copy of the *filename.dat* file, as follows:
 - a. Add an entry for each record to be loaded.
 - b. From under the TABLE line (see the following example), remove fields that you do not want to populate.

```
TABLE table_name
fieldname1 fieldname4 . . . fieldnameN
```

4. Save your file and exit from the editor.
5. Run the `pdm_userload` utility and specify the file, as shown by the following example. In this example, the input file name is `myData.dat`:

```
pdm_userload -f myData.dat
```

The database table is populated.

More information:

[pdm_userload--Add, Update, and Delete Database Records](#) (see page 998)

Drop and Restore Constraints

Some of the mdb tables that start with ca_ (such as ca_contact_) have referential constraints that may impact mass loading of data using pdm_load, pdm_userload, and pdm_restore tools. If mass loading these tables is required, you may need to drop the referential constraints prior to mass loading of the data.

Two SQL scripts are provided for each DBMS type to drop and restore constraints. Prior to mass loading of data affecting ca_* tables, run the Drop version of the script. After the data load has completed, run the Add version of the script.

For SQL Server the scripts are located in *installation-directory*\samples\views\SQLServer directory. Run the following command to drop constraints:

```
osql -E -e <SQLDropConstraints.sql
```

Run the following command to add back constraints:

```
osql -E -e <SQLAddConstraints.sql
```

For Oracle, the scripts are located in *installation-directory*\samples\views\Oracle directory.

Run the following command to drop constraints:

```
sqlplus mdbadmin/ <password> < OracleDropConstraints.sql
```

Run the following command to add back constraints:

```
sqlplus mdbadmin/ <password> < OracleAddConstraints.sql
```

Database Backup

You can back up the contents of a single database table, multiple database tables, or your entire CA SDM database, using the database backup utility, pdm_backup. The output of the backup utility is an ASCII file that the pdm_restore utility can use.

As part of its processing, pdm_backup first shuts down the daemons (UNIX) or services (Windows).

More information:

[pdm_backup--Write Database to ASCII File](#) (see page 958)

Database Restore

The database restore utility, `pdm_restore`, loads an output file from `pdm_backup` to the CA SDM database. The `pdm_restore` utility first shuts down the daemons (UNIX) or services (Windows). Then it restores, clears, and replaces all existing database records. Use the formatted ASCII file created by `pdm_backup` as input for the `pdm_restore` utility.

You can also use the `pdm_restore` and `pdm_userload` utilities to gain access to the CA SDM application in case of a catastrophic database corruption. If your database is damaged to the point where you cannot gain any access to the application and if all other measures have failed, rerun configuration and reinitialize the database to rebuild your database and populate reference data and system tables.

This procedure initializes your database in the same way that you did during the original installation. You can now access CA SDM. The `pdm_restore` utility can be used to restore the last backed up copy of your database.

More information:

[pdm_restore--Restore a Database](#) (see page 990)

Database Table Replacement

The `pdm_replace` utility is a quick and easy way to replace the entire contents of a CA SDM database table with new information. This utility can be useful for massive revisions of look-up tables.

Note: `pdm_replace` takes the same input file format as `pdm_userload`. You can create an input file for `pdm_replace` using `pdm_extract`; however, you cannot use the output of `pdm_backup` as input to `pdm_replace`.

More information:

[How to Create and Use an Input File](#) (see page 403)

[pdm_replace--Replace a Database Table](#) (see page 988)

Data Extraction

The `pdm_extract` utility extracts data from the CA SDM database and produces output in various formats. You can further process this data or enter it into other applications, such as a spreadsheet or another database.

With the `pdm_extract` utility, you can do the following:

- Dump the entire CA SDM database
- Dump one or more database tables
- Extract specific information from the database and produce output in one of the following three formats:
 - Output compatible with `pdm_userload`
 - Comma-separated value (CSV) output
 - Informal report-style output

More information:

[pdm_extract--Extract Data from Database](#) (see page 969)

Use the Data Extractor on UNIX

Before you use the CA SDM data extractor on UNIX, you must do the following:

1. Set the value of the `$NX_ROOT` environment variable to the full path name of the CA SDM installation directory that you defined during installation.
2. Append `$NX_ROOT/bin` to your `PATH` environment variable.
3. Set `umask` to `000`.

Data Selection for Extraction

To select the data for extraction, the data extractor uses an integral SQL subset with the following rules:

- The following SQL functions are specifically supported in the subset:
 - `IS NULL`, `IS NOT NULL`, `LIKE`
 - `SELECT` statements and `WHERE` clauses

- The following SQL functionality is not supported in the subset:
 - Joins
 - Embedded tabs and new-line characters
 - Clauses other than SELECT, FROM, and WHERE
 - Asterisks (*) in SELECT statements
 - Nested SELECTs
 - Aggregate functions
- All SQL reserved words must be in uppercase characters
- Every token in a WHERE clause must be surrounded by white space
- All date/time specifications must be in one of three formats:
 - Elapsed seconds from 12:00:00 a.m. 1/1/1970 Greenwich Mean Time (GMT):
`start_date< or >174182431500`
 - SQL DATE format:
`start_date< or >DATE '2001-11-28'`
 - SQL TIMESTAMP format:
`start_date< or >TIMESTAMP '2001-07-04 15:45:00'`

Note: TIMESTAMP format uses GMT. You can adjust time zones by adding or subtracting the appropriate number of hours, such as: 2001-03-23 14:00:00+02:00 or 2001-06-06 04:45:00-09:00.

Data Dereferencing

The `pdm_deref` utility is a dereferencing tool that converts data from various sources into a format suitable for loading into the CA SDM database. Dereferencing extracts internal IDs for cross-referenced fields. The utility can also be used to calculate down time and SLA down time values.

The `pdm_deref` utility converts data into one of the following formats:

- Output compatible with `pdm_userload`, suitable for loading into the CA SDM database
- Comma-separated value (CSV) output
- Informal report-style output

More information:

[pdm_deref--Dereference ASCII Data](#) (see page 962)

How to Use pdm_deref Example

The following example shows how to use the `pdm_deref` utility in a CA SDM ticket tracking system.

Assume that an existing ticket tracking system that you implemented on a spreadsheet has columns labeled Trouble Description, Technician First and Last Name, and Entry Date. These columns correspond to the `description`, `assignee`, and `open_date` fields in the CA SDM `Change_Request` table. The Trouble Description field contains the same data type as the `description` field. But the `assignee` field is a numeric field, and the Technician field on the spreadsheet is in a “last name, first name” format.

To use `pdm_deref`

1. Load the technicians’ names into the Contact table.
2. Prepare a `pdm_deref` input file with the existing information.
3. Build a specifications file to map the new contact names to assignee values.
4. Prepare a `pdm_userload` input file to be used to update the `Change_Request` table.

This process is described in more detail in the following steps:

1. Prepare a `pdm_userload` input file, `location.dat`, for the Location table as follows:

```
TABLE ca.location
location_name address_2 address_2
{"Boulder NCC - NQ", "716 Main
Street", "Boulder, CO 84302"}
{"Colorado Springs NCC", "2765 Spring Street",
"Colorado Springs, CO 84303"}
{"Denver NCC", "3765 Stoneridge Way", "Denver,
CO 80254"}
```

2. Load the data as follows:

```
pdm_load -f location.dat
```

3. Prepare an input file, `contact.dat`, with the original information as follows:

```
TABLE ca.contact
last_name first_name middle_name location pri_phone_number
{"Harrison", "Frank", "Harold", "NCC - HQ", "303-555-2333"}
{"Hertzog", "William", "L.", "Colorado Springs NCC", "303-966-1987"}
{"Lyman", "Jeanie", "L.", "Denver NCC", "303-966-5301"}
```


4. Prepare a dereferencing tool specifications file, `contact.spec` as follows:

```
Deref
{
input = c_location
output = location_uuid
rule = "SELECT id FROM ca.location WHERE location_name=?"
}
```

Important! Do not place a blank space in front of the SELECT keyword. Deref uses the new contacts' first and last names to obtain the appropriate numeric ID fields for loading the Change_Request table. In addition, the "hooks" represented by question marks (?), correspond to the specified input fields. You must have the same number of hooks as input fields, and they must be in the same order.

5. Run `pdm_deref` as follows:

```
pdm_deref -s contact.spec < contact.dat > contact.out
```

The output file, `contact.out`, looks like the following:

```
TABLE ca.contact
last_name first_name middle_name location.uuid pri_phone_number
{"Harrison", "Frank", "Harold", "69499D5A2424884887E62EC9823F5E47", "303-555-2333"}
{"Hertzog", "William", "I.", "86873FA40BA4234A8CF7A418D7C8B2DB", "303-966-1987"}
{"Lyman", "Jeanie", "L.", "58AA42789957734E8BEE146D07F7AD49", "303-966-5301"}
```

6. Load the `contact.out` file into the CA SDM database as follows:

```
pdm_load -i -f contact.out
```

Note: You must use the `pdm_load` command to use the `-i` option.

7. (UNIX only, optional) Write a script, `Convert_Ticket`, as shown in the following:

```
#!/bin/sh
pdm_load -i $1
cat $2 | pdm_deref -s $3 | pdm_load -i
```

You can run this script, as shown by the following:

```
Convert Ticket location.dat contact.dat contact.spec
```

In this example, `pdm_load` with the `-i` flag is used to speed the process. If you are making these updates on a regular basis, you can drop the `-i` flag so that `pdm_load` checks for duplicate records.

The following are additional examples of dereferencing tool specification files:

```
Deref
{
    input = first_name, last_name, middle_name
    output = assignee
    rule = " SELECT id from ca.contact \

        WHERE first_name=? \
        AND last_name=? \
        AND middle_name=? "
}
```

This rule converts three fields labeled first_name, last_name, and middle_name to the appropriate contact UUID. If all three input fields are not present, the rule is not applied. No match produces an error message and processing continues. For multiple matches, the first value is used; an error message is produced, and processing continues.

Use the Dbadmin Mode

The dbadmin mode is a utility that starts the data manipulation layer of the CA SDM system without starting the object layer, providing the ability to lock the entire database to perform low-level data maintenance without jeopardizing data integrity.

For example, you may want to use pdm_extract, pdm_load, pdm_deref, and pdm_replace for performing batch data updates on the system. By using dbadmin, the administrator is essentially placing a database lock on the entire system. The backup (pdm_backup) and restore (pdm_restore) utilities, both automatically place the system in dbadmin mode to guarantee a consistent backup and restore.

The dbadmin mode is also useful if you want to customize the system without starting the object layer until the data is modified. For example, making an attribute “required” in majic on an existing system can confuse the animator if it needs to update an object that has the required attribute null. You can put the system into dbadmin mode and update the objects using pdm_load, and then start the system as usual.

To place the system into dbadmin mode

1. Halt CA SDM either from Windows Service Manager or by running pdm_halt from the command line.

Note: It is a good practice to send an announcement to alert users and to check for logged in users before halting the system.

2. From the command line, enter the following command using the same case as shown:

```
pdm_d_mgr -s DBADMIN
```

Note: There is no return message, but a pause occurs before the command prompt returns. If there is no pause, check your spelling to verify that you entered data correctly.

3. Run `pdm_status` to verify that the system is in dbadmin mode.

Note: When the system is in dbadmin mode, the system returns the following status, which indicates that it is safe to work on the system:

```
C:\>pdm_status  
The Daemons are not running.
```

4. When all work is complete, run `pdm_halt` to shut down dbadmin mode.
5. Restart the system by following your usual procedures.

Archive and Purge Rules

You can define the rules you want to use for running archive and purge jobs on your CA SDM database. You can modify and activate existing rules instead of creating ones.

Important! The default Archive and Purge rules are set to Inactive. You must set the search filter to look for inactive rules. You can run a rule by selecting it from the Archive and Purge Rules node and clicking Run Now.

Example: Activate the Audit Log Rule

You can activate and run the rule for archiving and purging audit log records.

To activate the audit log rule

1. Navigate to Archive and Purge, Archive and Purge Rules
The Archive and Purge Rule List appears.
 2. Click Show Filter.
Select Inactive from the Status drop-down list, Click Search.
The Inactive rules appear.
 3. Select the Audit Log rule, click Run Now.
The archive and purge process begins for audit log records.
- Note:** You can [edit the rule](#) (see page 414) by clicking the Audit Log link and clicking Edit.

Run Archive and Purge Rules

You can run both active and inactive archive and purge rules manually at any time.

To run an archive and purge rule from the List page

1. From the Administration tab, select Archive and Purge, Archive and Purge Rules from the left pane.

The Archive and Purge Rule List appears.

2. Select the rules you want to run by clicking the box to the left of the rule name, choose the number of hours (1-24) after which you want the rule to stop, and click Run Now.

A message appears indicating that the archive and purge process has started.

3. Click OK to close the message box.

To run an archive and purge rule from the Detail page

1. From the Administration tab, select Archive and Purge, Archive and Purge Rules from the left pane.

The Archive and Purge Rule List appears.

2. Select the desired rule by clicking the rule name.

The Archive and Purge Rule detail page appears.

3. Choose the number of hours (1-24) after which you want the rule to stop and click Run Now.

View Archive and Purge Rules

You can view the summary information for each archive and purge rule on the Archive and Purge Rule List page. To display this page, select Archive and Purge, Archive and Purge Rules from the left pane on the Administration tab.

You can set up a filter to display only the rules that match certain criteria.

Search the Archive and Purge Rule List

You can set up filters to restrict the display of archive and purge rules to only those that match certain criteria.

To filter the Archive and Purge Rule List

1. Click Show Filter on the Archive and Purge Rule List page.

The top portion of the page displays filter fields.

2. Complete the fields as necessary to display only the rules of interest.
3. Click Search.

The Archive and Purge Rule List displays the rules that match your filter criteria.

Start Archive and Purge Using a Third-Party Scheduler

You can use a third-party scheduler to start and stop an archive and purge rule. The scheduler must support issuing commands from a command line. The arcpur.frg file is located in the following directory:

```
$NX_ROOT\bopcfg\interp
```

To start archive and purge with a third party scheduler

1. Open the third-party scheduler.
2. Set the start time.
3. Run the following command:
`bop_cmd -f arcpur.frg "start_arcpur(<rule_name>, 'hours')"`
4. Set the stop time.
5. Run the following command:
`bop_cmd -f arcpur.frg "stop_arcpur(<rule_name>)"`

The archive and purge is scheduled.

Archive and Purge Rule Definitions

You can define the rules you want to use for running archive and purge jobs on your CA SDM database. The product provides a set of predefined Archive and Purge rules. You can choose to modify and activate existing rules instead of creating ones.

Create an Archive and Purge Rule

You can set up rules for automatically archiving outdated records and purging them from the database.

To create an archive and purge rule

1. From the Administration tab, select Archive and Purge Rules.

The Archive and Purge Rule List appears.

2. Click Create.

The Create New Archive and Purge Rule page appears.

3. Fill in the fields as appropriate.

4. Click Save.

The new rule appears on the Archive and Purge Rule List page.

More information:

[Archive and Purge Rule Fields](#) (see page 414)

[Run Archive and Purge Rules](#) (see page 412)

Archive and Purge Rule Fields

You can use the fields on the Archive and Purge Rule pages to define and edit rule definitions.

Rule Name

(Required) Specifies a unique identifier for the rule. Can be up to 30 alphanumeric characters.

Status

(Required) Indicates whether this rule is Active or Inactive.

Schedule

Specifies a workshift in which the rule should be in effect. Workshifts are defined by the Administrator.

Recurrence Interval

Specifies how often this rule is run, in the format hours:minutes:seconds (HH:MM:SS).

Archive File Name

Specifies the name of the file where you want the out-dated records stored. Its location is controlled by the NX_RULE_ARCHIVE_PATH variable stored in the NX.env file. Archive data files can be restored to the database using the PDM_LOAD utility.

Operation Type

Specifies one of the following types of operation that the rule should execute:

Archive/Purge

Purges outdated records from the database that are written to the file specified in the Archive File Name field.

Purge Only

Purges outdated records from the database, but they are not written to the archive file.

Archive Only (Test Run)

Writes outdated records to the archive file without purging them from the database. Use this option to test a newly created or edited archive and purge rule.

Config. Object Name

Specifies the name of the database object this rule can archive and purge.

Object Name

(Read-only) Populates your selection in the Config. Object Name field automatically.

Days Inactive

Specifies the number of days a record is inactive to be eligible for archive and purge from the database.

Additional Query

Archives and purges specific inactive records among the existing inactive records. Use this field when you want to create different rules for archiving and purging the subsets of expired records for the same object. Use the same syntax as you use for stored queries.

Examples: Specify Additional Queries

The following query archives and purges only assigned inactive request records with a priority of 1:

```
priority = 1 AND (assignee IS NOT NULL OR group IS NOT NULL) and active = 0
```

The following query format archives and purges records based on time-span:

```
close_date < EndAtTime('\LAST_YEAR\')
```

Edit an Archive and Purge Rule

You can edit an archive and purge rule that has already been created.

To edit an archive and purge rule

1. From the Administration tab, select Archive and Purge Rules.
The Archive and Purge Rule List appears.
2. Select the rule you want to edit and click the rule name.
The Archive and Purge Rule Detail page appears.
3. Click Edit.
The Update Archive and Purge Rule page appears.
4. Edit in the fields as appropriate.
5. Click Save.
Your edits are saved and the Archive and Purge Rule List page appears.

More information:

[Archive and Purge Rule Fields](#) (see page 414)

[Run Archive and Purge Rules](#) (see page 412)

Archive and Purge History

You can view a summary of the history information for each archive and purge rule.

To view archive and purge history

1. Select Archive and Purge, Archive and Purge History on the Administration tab.
The Archive and Purge History List appears. This list displays an entry for each time the rule has been executed.
2. The following fields are shown:
 - Rule Name**
Displays the name of the rule used in the history list.
 - Start Time**
Displays the time that the rule was started.
 - End Time**
Displays the time that the rule completed.

Main Objects Purged

Displays the main objects that were purged by the rule. For example, a call request.

Child Objects Purged

Displays the related child objects that were purged by the rule. For example, the activity logs of a call request.

3. (Optional) Click Show Filter and specify filter criteria to limit the list to the information of interest.
4. Click a rule name if you want to review the rule's configuration.

The Archive and Purge Rule Detail page appears.

Filter the Archive and Purge History List

You can filter the Archive and Purge History List to display only the entries of interest.

To filter the Archive and Purge History List

1. Select Show Filter on the Archive and Purge History List page.

The top portion of the page displays the filter fields.

2. Complete the fields as necessary to display the entries you want to see:

Rule Name

Select the name of the archive and purge rule whose history you want to see.

Earliest Start Date

Enter the earliest date and time to specify the beginning of a range for filtering the history to show only entries for a specified time frame.

Latest Start Date

Enter the latest date and time to specify the beginning of a range for filtering the history to show only entries for a specified time frame.

Earliest End Date

Enter the earliest date and time to specify the end of a range for filtering the history to show only entries for a specified time frame.

Latest End Date

Enter the latest date and time to specify the end of a range for filtering the history to show only entries for a specified time frame.

Note: To display the Additional Search Arguments field, click the spigot icon. This field is intended only for expert users who understand SQL and Majic and can use it to specify search arguments that are not available in the standard search filter fields. To specify an additional search argument, enter a SQL WHERE clause in this field.

3. Click Search.

The Archive and Purge History List displays the entries that match your filter criteria.

Attachment Handling (arcpur)

Attachment records are saved in the database and attachment files are saved in the repository directory. When you archive and purge attachment records in the database, they are marked as deleted, but cannot be removed. These records are necessary for the repository daemon to locate the attachments.

Set the Archive Type

You can set the Archive Type for attachments in a repository from the Administration tab.

To set the archive type

1. Navigate to Attachments Library, Repositories.

The Repositories page appears.

2. Right-click on a repository (such as Service Desk), select View.

The Repository Detail page appears.

3. Click Edit.

Select one of following from the Archive Type dropdown:

None

No archive and purge process is performed.

Archive and Purge

Outdated records are written to the file specified in the Archive File Name field and purged from the database.

Purge Only

Outdated records are purged from the database, but are not written to the archive file.

4. Click Save.

The Repository Detail page appears, showing your changes.

Note: You cannot use archive and purge with Knowledge Management. This feature is only supported for CA SDM.

Set the Archive Path

You can set the archive path for attachments from the Administration tab.

To set the archive path

1. Navigate to Attachments Library, Repositories.
The Repositories page appears.
2. Right-click a repository (such as Service Desk), select View.
The Repository Detail page appears.
3. Click Edit.
Enter the path in the Archive Path field.
The following is the default directory:
`$NX_ROOT/site/attachments/default/archived_files`
4. Click Save.
The Repository Detail page appears, showing your changes.

How to Restore Archived Data

When restoring archived data, you should start the daemons in dbadmin mode with the `pdm_d_mgr` utility. Dbadmin mode allows limited access, so you can safely run `pdm_load` to restore archived data.

The following describes how archived data is restored:

1. Shutdown CA SDM.
2. Start the daemons in dbadmin mode, using the following command:

```
pdm_d_mgr -s DBADMIN
```
3. Locate the archived data file. By default, the file is located under `$NX_ROOT/site/data/archive`.
4. Run `pdm_load` against the data file. For example:

```
pdm_load -a -f 2004611T1726_Call_Request.dat
```
5. If there is a problem with the load, check the command line and the log for errors. Arcpur.log files can be found in `$NX_ROOT/log`.

Note: Size limitation of arcpur.log files are defined in `$NX_ROOT/NX.env` as

```
# The size limit for the Archive and Purge log file and data file.  
@NX_ARCPUR_FILESIZE=2000000000
```

Archive and purge creates arcpur.log.0, arcpur.log.1 through arcpur.log.9 after reaching the file limit for each log files.

6. Run `pdm_halt`.

This shuts down the daemons.

7. Restart CA SDM.

Note: After restoring the record, it is archived and purged in the next cycle of the archive and purge process.

8. (Optional) To prevent the record to be archived and purged at the next cycle, do the following:
 - a. Update the record to make it active again.
 - b. Inactivate the archive and purge rule.

Archiving and Purging KPI Data

Predefined Archive and Purge Rules are available for archiving and purging KPI data. Navigate to Archive and Purge, Archive and Purge Rules on the Administration tab. Then, filter the list by searching for Inactive rules. You can select one of the following rules and set it to Active:

- `KPI_Data(SQL)`
- `KPI_Data(Stored Query)`
- `KPI_Data(System)`
- `KPI_Ticket_Data`

An optional configuration rule can be created inside the `arcpur_cfg.xml` and `itil_arcpur_cfg.xml` files to archive and purge the `KPI_Ticket_Data` table.

You can find the files in the following directories:

- `arcpur_cfg.xml`
`$NX_ROOT\site\cfg\`
- `itil_arcpur_cfg.xml`
`$NX_ROOT\site\cfg`

Use the following methods to archive and purge the `KPI_Ticket_Data` table:

KPI_Ticket_Data Table

- Use `end_time (last_mod_dt)` in `KPI_Ticket_data` to determine whether a record needs to be archived and purged.
- Link the records in the `KPI_Ticket_Data` table to the records in the `Ticket` table (such as `cr`, `chg`, or `iss`). This ensures that all ticket related records in the `KPI_Ticket_Data` table are archived and purged.

KPI_Ticket_Data table does not have a SREL relationship with any Ticket table and it relies on two fields, obj_name, and obj_id, to link with a ticket. The obj_name value can be cr, chg, or iss and the obj_id value is the ticket id. Define a main_obj for each ticket object.

The following is a sample main_object definition for the ticket object, cr:

```
<!-- KPI Ticket Data -->
<main_obj>
<name>KPI Ticket Data</name>
<internal_name>KPI Ticket Data</internal_name>
<factory>ktd</factory>
<default_query>obj_name='cr'</default_query>
<date_field>end_time</date_field>
<ref_by value="obj_id">cr.id</ref_by>
</main_obj>
```

Note: The configuration rule can only select records for cr. The ref_by tag can match the value of obj_id in KPI Ticket Data to the value of id in cr. If a match is found, it means that a KPI Ticket Data record is referenced by a cr record, so the KPI Ticket Data record are not archived and purged.

After adding the configuration rules for all ticket objects, restart the CA SDM service. These configuration rules become selectable configuration object names in the Archive and Purge Rule detail form.

Archive and Purge Knowledge Management Forums

You can use archive and purge rules to remove deleted forums from the database. You can create an archive and purge rule to archive and purge only forum type documents.

Note: If you do not create an archive and purge rule, the existing Knowledge Document rule can archive and purge all deleted documents including the forum type documents.

To create a rule to archive and purge forums

1. Navigate to Archive and Purge, Archive and Purge Rules on the Administration tab.
The Archive and Purge Rule List appears.
2. Click Create New.
The Create New Archive and Purge Rule page appears.

3. Do the following:
 - Select "knowledge document" in the Config. Object Name drop-down list.
 - Add "KS_TYPE=20" to the Additional Query field.
4. Click Save.

The new archive and purge rule is created.

Chapter 12: Using the Text API

This section contains the following topics:

[Text API](#) (see page 423)

[The Configuration File](#) (see page 435)

Text API

The *Text API* is an interface that lets you use text-based input to create and update objects in the CA SDM database, such as issues, requests, contacts, and assets. Using the Text API, you can assign values to most fields that are accessible to users.

Important! CA SDM requires that all input be in UTF-8 format, or data can be corrupted. The [pdm_unconv utility](#) (see page 996) lets you convert data from a local charset to UTF-8 and from UTF-8 to a local charset.

You can access the Text API by using the following interfaces:

- Command line
- Email
- CA NSM

Note: You can use web services as the alternative to the Text API for cross-application integration.

More information:

[Conversion Methods](#) (see page 433)

Command Line Interface

Use the *pdm_text_cmd* command to activate the Text API command line interface. You can then specify information such as the table to process and the operation to perform by using parameters to the *pdm_text_cmd* command.

The input to the Text API is passed to the *pdm_text_cmd* command in the form of an input file, or directly from STDIN.

Note: When passing the parameters from command prompt, use Ctrl+Z in Windows and Ctrl+D in POSIX.

Important! You cannot use single or double quotation marks as parameters for the *bop_cmd* and *pdm_text_nxd* commands.

More information:

[pdm_text_cmd--Text API Command Line Interface](#) (see page 993)

CA Network and Systems Management Interface

When CA NSM and CA SDM are integrated and you are creating requests from CA NSM events, the *user_parms* parameter in writer rule definitions is passed to the Text API. The CA SDM writer process (*tngwriter*) defines its own replacement parameters for changing the text before sending it to the Text API. The keyword *LOG_AGENT* is added to the end of the input to set the *log_agent* for the request.

Note: You need to update the *Text_API.cfg* file for all additional fields that are passed from CA NSM Alert Management Systems to CA SDM. This file is used for integrations with web services, email and AHD.DLL.

More information:

[Keywords](#) (see page 426)

Input Format

Input to the Text API is specified in the following ways:

- In the command-line interface, input is typically specified in a text file passed to the `pdm_text_cmd` command.
- In the email interface, input is specified in the text of the email. You specify a regular expression to find the target object identifiers.

You format the Text API input in the same way no matter which interface you use.

The basic format for input is as follows:

`%keyword=value`

or

`%PROPERTY={{property_label}}value`

The normal behavior of Text API commands has the following exceptions, where the last-appearing of two or more conflicting commands takes precedence:

- When a message contains multiple valid ticket ID artifacts matching the mailbox rule filter string, or multiple Text API ticket ID commands, the first one encountered is used. Also, a ticket ID artifact, which is identified using the mailbox rule filter string, overrides any Text API ticket ID command, regardless of which appears first.
- When a message contains multiple log comment Text API commands, all comments are posted, although the order in which they appear in the ticket activity log can vary.

All ticket ID artifacts that match the filter, valid or otherwise, and Text API ticket ID commands within the message, applicable or otherwise, applied or otherwise, are commented out before the message is posted. The ticket ID artifacts identified through the mailbox rule filters appear as `-(...)-`. Leading percentage signs (%) in Text API ticket ID commands are converted to two opening parentheses ((, and two closing parentheses)) follow the command. If a Text API ticket ID command appears after another Text API command with a log comment (`%LOG=...`), then the commented-out Text API ticket ID command is made into a separate log comment.

Note: The log comment is the only Text API command that can appear multiple times in one message and still have each occurrence applied. For any other commands, the Text API uses the last occurrence only, because multiple occurrences of other commands conflict with each other. Multiple log comment commands post separate log comment messages to the ticket, and not necessarily in any particular order.

In addition, if a Text API ticket ID command appears in the message either at the beginning of the message or in between two other Text API commands, it is converted into a log comment. If the previous command is a log comment (%LOG=...) or update description (%DESCRIPTION=...), it is appended to that command, rather than becoming a separate log comment.

Incoming messages that are sent HTML-only, without a plain-text version included, lose their message body. If the message matches any mailbox rule filters with an empty message body, a ticket can be created with an empty Description, or with the quoted message subject as the entirety of the ticket description.

Keywords

You can use two types of keywords as input to the Text API.

- Definitions in the [KEYWORDS] section of the text_api.cfg file—This type is a group of keywords that are related directly to the fields for the various tables that you can update. For example, most of the fields on the Issue Detail form are defined the [KEYWORDS] section. Using these keywords, you can set values for fields in the record that you are updating or creating. For example, the following line sets the priority of the issue to 5:

```
%PRIORITY=5
```

The [KEYWORDS] section of the text_api.cfg file lists all keywords. You can define additional keywords (for example, to allow Text API access to fields that you have added when customizing your database schema).

- The following special keywords are always defined as follows, regardless of the contents of the text_api.cfg file:

Keyword	Description
ASSET	Used to attaches an item to a ticket (valid for requests, issues, and change orders). The value specified is the item name, which must already exist. You can specify this keyword multiple times, because a ticket can have multiple items attached to it.
ATTACHMENT	Used internally by the email interface to add email attachments to a ticket.
DESCRIPTION	Specifies the value to use for the ticket's description field. This keyword is assumed if input is sent to the Text API without an explicit keyword. This keyword is applied automatically by the Mail Eater when the message does not begin with a keyword but does contain a ticket ID artifact or keyword. You can change how the DESCRIPTION keyword is handled for updates using the following entry in the [OPTIONS] section of text_api.cfg: <pre>UPDATE_DESC_IS_LOG</pre> If this option is set to YES, the value is used to create a log comment. If the value is set to NO, the value overwrites the existing description field.

Keyword	Description
FROM_EMAIL FROM_EMAIL_OVERRIDE	Used by the email interface to match against the Email Address field in the ca_contact record. It is also used as the log_agent for the ticket. If both are supplied, FROM_EMAIL is ignored. Note: FROM_EMAIL is set automatically by the Mail Eater with the sender address of the message.
FROM_PERSID	Used by the command line interface to define the log_agent for an operation (for example, when a ca_contact record does not have a User ID). This keyword is passed automatically by pdm_text_cmd if the -p parameter is specified. The value is matched to a ca_contact record persistent_id.
FROM_USERID	Used only in the command line interface to define the log_agent for an operation. This keyword is passed automatically by pdm_text_cmd if the -u parameter is specified. The value is matched to a contact's User ID.
LOG	Used to create a log entry (valid for requests, change orders, issues, and contacts). This keyword is applied automatically by the Mail Eater when the message does not begin with a keyword but does contain either a ticket ID artifact or keyword, or a DESCRIPTION keyword.
LOG_AGENT	Used by the CA NSM interface to define the log_agent for an operation. The value is matched to a contact record's ID field.
PROPERTY	Used to set the value of a property (valid only for requests, change orders, and issues). Unlike other keywords, which are followed by an equal sign and a value, the PROPERTY keyword syntax must include the property label, as follows: PROPERTY={{property_label}}value You must specify the <i>property_label</i> exactly as it appears in the database.
SEARCH	Used only in the command-line interface and the CA NSM interface to supply a list of keywords for use in a query to update multiple tickets for an asset. The value is a list of keywords used in the search. The SEARCH keyword is automatically set by the CA NSM interface.
SEARCH_EXPLICIT	Used only in the CA NSM interface to override the SEARCH keyword supplied by the CA NSM interface. The values supplied are the same as the SEARCH keyword.

More information:

[pdm_text_cmd--Text API Command Line Interface](#) (see page 993)

Keyword Input Conventions

The following conventions apply to keyword input formatting:

- Prefix every keyword (including PROPERTY) with a percent (%) sign. The percent sign must be in column position one. If the first nonempty line of the input does not have a percent sign at the start of the line, either %DESCRIPTION= or %LOG= is used as the prefix for the incoming data, depending on whether a ticket ID artifact or keyword was found. If %DESCRIPTION is set, the contents of the message up to the first Keyword is posted as a ticket description. If %LOG= is set, the contents of the message up to the first Keyword is posted as a log comment.
- Do not use any intervening spaces within the keyword between the percent sign and the keyword, or between the keyword and the equal (=) sign.
- Do not quote values; all data after the equal sign is assumed to be the value.
- Keywords are not case sensitive.
- If the input includes duplicate keywords, the last keyword is used; otherwise, the order in which you specify the keyword/value pairs is unimportant.
- Specify keyword values as you would for the corresponding field in the web interface. For example, to specify an Analyst contact type, you use %CONTACT_TYPE=Analyst, even though in the database this value is stored as an integer. The CONTACT_TYPE keyword is defined in text_api.cfg so that it [converts the specified value](#) (see page 433) to match the stored value.

Note: Whether the value is case sensitive depends on your underlying DBMS.

- You can extend string data across multiple lines.

Format an Email Message To Update a Ticket

A user can format an email message to create or update a ticket.

To format an email message to create or update a ticket, use the following fields:

To

Specifies the mailbox name assigned to the CA SDM contact set up for the privileged user.

From

Specifies the person sending the email. The person must be defined in the ca_contact table unless the Allow Anonymous option is specified in the applicable mailbox rule.

Note: The From address is typically part of your email program configuration, and it is not typically set on a per-message basis.

Attachments

Attaches documents and other files to the email to send attachments to the Text API.

Subject

Matches keywords in a mailbox rule filter string, particularly when creating a ticket.

Body

Specifies the message body of the email using the Text API. You can specify the keyword `ISSUE_ID`, `REQUEST_ID`, or `CHANGE_ID`, depending on the type of ticket to create or update a ticket.

Start and End Email Message Delimiters

Some email interfaces add information to the beginning or end of mail messages (for example, MIME encoding) that can cause the email interface to malfunction. If your email interface adds information, you can use the following delimiters: `start-request` and `end-request`. The email interface ignores information that is specified prior to `start-request` and subsequent to `end-request`.

Note: The Mail Eater does not support emails in the RTF or HTML-only formats.

Example: Use start-request and end-request Delimiters

```
"start-request"  
message_body  
"end-request"
```

How the Text API Uses Artifacts

The Text API processes the subject or body of email notifications. Mailbox rules let you identify artifacts and values that the Text API uses. For example, you can define the rule for incidents as `Incident:{{object_id}}`, so finding `Incident:1234` translates to `%INCIDENT_ID=1234` for the Text API. 1234 is the `ref_num` for the Incident. Because the artifact must be unique in the email and easy to find, you can make the artifact more distinctive such as `%Incident:{{object_id}}%`.

Follow the `{{object_id}}` keyword with a character which is not a letter, number, comma, forward-slash (/), plus sign (+), or equals (=) sign, because these characters can appear within an artifact. Otherwise, it is possible that characters which follow the artifact can be misinterpreted as part of the value of the artifact, or that a character within the value of the artifact can be misinterpreted as the character which follows the value.

Mail Eater does the following:

1. Finds the artifact within an email (such as Incident:1234) that maps to the appropriate ticket or other object supported by the Text API.
2. Translates the artifact to a Text API token (such as %INCIDENT_ID=1234).
3. Mail Eater submits the tagged message to the Text API. The Text API processes the email, applies the text, commands, or both which it contains to the appropriate ticket, and generates an automatic response email indicating whether the email message it received was successfully applied. Depending on the actions performed, a notification email message is also sent separately to indicate certain specific events, such as the creation of a ticket.

How to Set Up Notification Replies to Update Tickets

The Text API daemon (pdm_text_nxd) creates and updates tickets with information from external interfaces, such as the command line and email. You can set up mail to use the Text API so that users (contacts) can update tickets by replying to email notifications. The text of the reply is added as a log comment activity to the ticket.

To set up notification replies to update tickets, do the following:

1. Set the notification method that the contact uses to `pdm_mail -T reply_email_address` or `pdm_mail -F reply_email_address`. The `reply_email_address` specifies the incoming address for the mailbox. When the contact clicks reply on an email that address is filled in from the From or Reply-To address of the message to which they are replying.

-T sets the Reply-To address. -F sets the From address, which is used as the reply address if a separate one is not specified.

Note: Some mail programs do not or cannot honor a Reply-To address.

2. Create or update a mailbox rule using a Text API keyword.

The user-defined artifacts in the mailbox rule filters replace the following Text API keywords:

Object	Text API Keyword	Identifier
Incident	%INCIDENT_ID	Ref_num
Problem	%PROBLEM_ID	Ref_num
Request	%REQUEST_ID	Ref_num
Chg_ref_num	%CHANGE_ID	Chg_ref_num
Issue	%ISSUE_ID	Ref_num

3. Create or update a notification phrase that matches the rule.
4. Create or update a message template that uses the notification phrase.
5. Update the mailbox rule that you created in Step 2 to specify the message template that you created or updated in Step 4.

Note: For information about performing each step, see the *Online Help*.

After the user receives the notification and replies to it, the following actions occur:

1. When the filter string is found, the relevant ticket ID keyword and value denoted by the placeholder, if any, are appended to the message.
2. If a matching ticket ID artifact is found, the corresponding ticket is updated, with either a log comment, a new description, or other values in accordance with the text, keywords, and commands in the message.
3. If a matching ticket ID artifact is not found, a ticket is created with a description and other parameters in accordance with the text, keywords, and commands in the message.

How to Set Up a Reply to an Incident Notification Example

This example shows how to set up a reply to an incident notification.

To set up a reply to an incident notification, do the following:

1. Create a mailbox rule using the following fields and values:
 - Filter—Body contains
 - Filter String—%Incident:{{object_id}}%
 - Ignore Case—YES
 - Action—Update Object
 - Action Object—Incident
2. Create a notification phrase that includes the rule as follows:
 - Symbol—Incident Reply
 - Code—IncidentReply

- Active—Active
- Description—Comment that embeds the reply for an Incident/Problem/Request.
- Phrase—In order to add a comment to your @{call_req_id.type.sym}, just reply to this email or include the line below (on a line by itself):

%Incident:{call_req_id.ref_num}%

Note: In auto-reply text of the mailbox rule, omit the call_req_id. prefix. This prefix applies a context which the mailbox rule text is already in, and such a context change is not valid when already acting within that context.

3. Create or update a message template that uses the notification phrase as follows:

- Notification Message Body

This is a simple notification.

@{notification_phrase[IncidentURL1].phrase}

4. Update the mailbox rule that you created in Step 1 to specify the message template that you created in Step 3, as follows:

Message Template—*mailbox rule name*

How an End User Updates a Ticket Example

The following example demonstrates how an end user (John Smith) replies to an email notification to update an incident ticket.

The Body or Subject of the email includes the object identifier. The {{object_id}} placeholder within the filter string denotes the object identifier.

1. A notification is sent to John Smith and includes the following instructions:
In order to add a comment to your incident, just reply to this email or include the line below (on a line by itself).
%Incident:1234%
2. John Smith replies to the notification as follows:
This is my response...
3. The Mail Eater receives the following text version of the John Smith's email:

This is my response...

From: Service Desk

Sent: Wednesday, September 18, 2009 10:22 AM

To: Smith, John

Subject: Simple Notification

This is a simple notification.

In order to add a comment to your incident, just reply to this email or include the line below (on a line by itself).

%Incident:1234%

4. The Mail Eater processes rules in order and finds the %Incident:1234% artifact:

This is my response...

From: Service Desk

Sent: Wednesday, September 18, 2009 10:22 AM

To: Smith, John

Subject: Simple Notification

This is a simple notification.

In order to add a comment to your incident, just reply to this email or include the line below (on a line by itself).

%INCIDENT_ID=1234

5. The Mail Eater adds the Text API keywords and the {{object_id}} value to an %INCIDENT_ID= statement and leaves a marker where the {{object_id}} value was found. The following text shows the data that is sent to the Text API. The bold text shows values added by the Mail Eater.

%LOG=This is my response...

From: Service Desk

Sent: Wednesday, September 18, 2009 10:22 AM

To: Smith, John

Subject: Simple Notification

This is a simple notification.

In order to add a comment to your incident, just reply to this email or include the line below (on a line by itself).

%Incident:-((...))-%

%FROM_EMAIL=john.smith@company.com

%INCIDENT_ID=1234

6. The Text API add a log comment for Incident 1234.

Conversion Methods

Many of the keywords defined in text_api.cfg have an associated method to convert the value specified to a value that is appropriate for storage in the database. This feature lets users specify values just as they would in the web interface without having any knowledge of the underlying implementation.

The configuration file has several examples of this type of keyword definition, including ISSUE.PRIORITY and CONTACT.CONTACT_TYPE. If you need to define additional keywords (for example, to allow Text API access to fields that you have added when customizing your database schema), you can use one of the following predefined methods:

Method	Output Type
lookup_actbool	INTEGER
lookup_asset_by_name	UUID

Method	Output Type
lookup_asset_by_persid	UUID
lookup_chg_category	STRING
lookup_chg_status	STRING
lookup_cnt_by_email	UUID
lookup_cnt_by_last_first_middle	UUID
lookup_cnt_by_logonid	UUID
lookup_cnt_by_persid	UUID
lookup_cnt_meth	INTEGER
lookup_cnt_type	INTEGER
lookup_company	UUID
lookup_cr_status	STRING
lookup_cr_template	STRING
lookup_domain	INTEGER
lookup_grc	INTEGER
lookup_group	UUID
lookup_impact	INTEGER
lookup_iss_category	STRING
lookup_iss_status	STRING
lookup_loc	UUID
lookup_mfr_model	UUID
lookup_nr_family	INTEGER
lookup_org	UUID
lookup_person_contacting	INTEGER
lookup_position	INTEGER
lookup_priority	INTEGER
lookup_prob_category	STRING
lookup_product	INTEGER
lookup_resource_status	INTEGER
lookup_service_lvl	STRING
lookup_severity	INTEGER

Method	Output Type
lookup_state	INTEGER
lookup_timezone	STRING
lookup_type_of_contact	INTEGER
lookup_urgency	INTEGER
lookup_workshift	STRING

If the value you need to convert is not addressed by any of these predefined methods, you need to write a customized method. The method should take a STRING value as its input and return a value (either INTEGER, STRING or UUID) as its output. Return a value of -1 (or “-1”) to denote that the value cannot be determined and is therefore, not set. For UUID, return a “(uuid) NULL”.

For example, you might develop a method to convert a user ID to a ca_contact table reference. The incoming value, such as Administrator, would be passed to the method, and the method would return the ca_contact table id for the user ID of Administrator.

The manner in which you define keywords in the configuration file offers you the advantage of defining multiple keyword mappings to the same field, including different conversion methods, depending on the value being specified. For example, assignee can have several different keyword mappings to define how to set its value based on different input values. One input might be user ID, another might be last name, first name, middle name, and still another might be the actual ca_contact id (for example, 793ED69B4E87A545BD8E911834D829FC). Each keyword maps to a different conversion method, except the last one, which does not need to be converted.

The Configuration File

The text_api.cfg file defines the keywords that are related directly to the fields of the various tables that you can update. You use this file both as a reference to find certain predefined values, such as keywords, and as a mechanism for configuring the Text API, although the default configuration file works for most installations without modification.

The text_api.cfg file is located in the following directory:

- UNIX—\$NX_ROOT/site
- Windows—*installation-directory*\site. For example: C:\Program Files\CA\Service Desk\site

The configuration file is divided into sections, with particular attributes defined within each section. Attribute definitions are of the following form:

keyword=value

None of the keywords are case-sensitive, whereas all values (except in the [OPTIONS] section) are case-sensitive.

Note: You can view and modify the text_api.cfg file using any text editor.

Important! If you are integrating with the CA NSM Alert Management Systems component, you must update text_api.cfg for any additional fields that are passed to CA SDM.

More information:

[Keywords](#) (see page 426)

Options

The [OPTIONS] section of the text_api.cfg file defines processing options that may differ from one site to another. For example, there are options to determine the incoming date format, which fields allow linefeeds to be retained, and whether to allow issues, requests, or change orders to be updated using the email interface. All options in this section are configurable. Be aware that although you can remove table names from the VALID_TABLE_LIST, if you do not want to support Text API access to those tables, you cannot add table names to this list.

Defaults

Use the [XX_DEFAULTS] section provided in the text_api.cfg file for each interface using the Text API (for example, [EMAIL_DEFAULTS] for the email interface and [CMD_DEFAULTS] for the command line interface). The [XX_DEFAULTS] section defines the default values for fields and properties that are required in case the user does not supply them directly. XX refers to the interface type, such as CMD or EMAIL.

To set default values, use one of the following formats:

- `table_name.keyword=value`

The *keyword* must be defined either in the [KEYWORDS] section or as properties in your database. Any method associated with the keyword is automatically applied to the *value*. For example:

```
ISSUE.PRIORITY=1
```

The PRIORITY keyword is defined in `text_api.cfg` so that it performs a lookup to convert the value you specify to match the corresponding value that is stored in the database. Here, value 1 is converted to 5, which is the underlying database value for the priority symbol 1. This feature lets users specify the value just as they would in the web interface.

- `table_name.PROPERTY={{property_label}}value`

The *property_label* must be defined as a property in your database.

In both formats, the *table_name* must be one of the values defined by VALID_TABLE_LIST in the [OPTIONS] section, such as Issue, Request, or Contact.

More information:

[Conversion Methods](#) (see page 433)

Ignore Incoming

There are several [..._IGNORE_INCOMING] sections in the `text_api.cfg` file, one for each interface that uses the Text API (for example [TNG_IGNORE_INCOMING] for the CA NSM interface and [EXT_IGNORE_INCOMING] for the external interface used by other CA products). These sections define fields and properties that are ignored in the input (the format is the same as described in Defaults, except no “=value” is specified). This feature lets you prevent users from setting certain values, which in turn, provides you with more security for such times as letting customers use the email interface.

The IGNORE sections work well when used in conjunction with the corresponding [..._DEFAULTS] sections because you can prevent the user from setting a particular value and supply a default value at the same time. For example, if you want to prevent email interface users from setting the priority of an issue, you could set the following values:

```
[EMAIL_DEFAULTS]
ISSUE.PRIORITY=2
[EMAIL_IGNORE_INCOMING]
ISSUE.PRIORITY
```

In this case, any priority that the user specifies in the email message body is ignored, and all issues created by the email interface are automatically assigned a priority of 2.

Example Input

The following examples show input that you can use in the body of an email message or in a file serving as input to the command-line interface.

Example: First Line Does Not Include a Keyword

In this example, because the first line is missing a %keyword in the first column, the literal %DESCRIPTION= is added to the beginning of the message. This addition sets the description field to “This entire text goes to the description field” (with the line break intact, because ISSUE.DESCRPTION is included in the list of fields for the LINEFEEDS_ALLOWED entry in the [OPTIONS] section of text_api.cfg).

```
This entire text goes
into the description field
%PRIORITY=None
```

Example: First Line Includes a Keyword

In this example, the PRIORITY keyword is defined in text_api.cfg so that it performs a lookup to convert the value you specify to match the corresponding value that is stored in the database. Here, the value None is converted to 0, which is the underlying database value for the priority symbol 1. This feature lets users specify the value as they would in the web interface.

```
%description=This is my description
%priority=None
%CATEGORY=Upgrade.PC
%PROPERTY={{Current CPU}}266 mhz
%PROPERTY={{Current Harddrive}}1 gig
%PROPERTY={{Requested Upgrade}}4 gig harddrive
```

The specified values are used to set the description and priority fields for the ticket, similar to the previous example (notice that keyword case is unimportant).

The value of Upgrade.PC is searched, and the category field for the ticket is set appropriately.

Matching the following labels sets the three property values:

- Current CPU
- Current Hard drive
- Requested Upgrade

Chapter 13: Managing Version Control

Note: In this chapter, most references to “client” apply to either a client or a secondary server.

This section contains the following topics:

[How Version Control Works](#) (see page 439)

[Version Control Files](#) (see page 440)

[Primary Server Control Files](#) (see page 440)

[Secondary Server and Client Control Files](#) (see page 441)

[Version Control for Installation Customizations](#) (see page 441)

[Version Control Server Modes](#) (see page 442)

[Version Control File Syntax](#) (see page 443)

How Version Control Works

CA SDM version control helps you manage system customizations that affect clients and secondary servers. You can use version control on client and secondary server installations.

Version control works as follows:

1. When a client connects to the server, the client sends a list of its controlled components to the server.
2. The server compares the list to its own master list.
3. If the server finds any differences, it can do the following:
 - Notify the client.
 - Fail the connection attempt.
 - Upgrade the client to the correct version of the controlled components.

Note: Version control is not a software delivery system; the upgrade feature simply provides a way to correct minor synchronization problems between client and server.

Version Control Files

CA SDM maintains two sets of version control files. One set is for secondary servers and the other for all clients. A controlled component can represent a file, a directory, or the `client_nx.env` environment variable file. A component can also be generic; that is, not be associated with any external object. You can use a generic component to provide version control for the client as a whole, or for a file or directory too large for an automatic upgrade. A generic component version mismatch is not updated automatically. A mismatch always results in a client fault and must be remedied manually.

Primary Server Control Files

Version control on the CA SDM server consists of one executable, `pdm_ver_nxd`, and two version control files, `server_default.ver` and `server_secondary.ver`. The `ver_ctl` server option determines how version control responds to version discrepancies between installed components on the server and client. The default setting of the `ver_ctl` Options Manager option is `UPGRADE`, causing version control to upgrade when a discrepancy is detected.

`Pdm_ver_nxd` validates version request messages from clients by comparing the data in the request with the following version control files:

- `server_default.ver`
- `server_secondary.ver`
- `server_custom.ver`
- `server_secondary_custom.ver`

If `pdm_ver_nxd` detects a component version discrepancy, it logs a message about the problem and because `UPGRADE` is enabled, the client is upgraded. The contents of the notification message depend on the setting of the `ver_ctl` option.

If you want, you can create and maintain the following two *custom* version control files to manage customizations:

- `server_secondary_custom.ver`
- `server_custom.ver`

Secondary Server and Client Control Files

Version control on a CA SDM secondary server or client consists of one executable (pdm_ver) and a version control file (secondary.ver). If you want, you can create and maintain a second version control file to manage customizations: server_secondary_custom.ver.

The client software includes one executable, pdm_ver, and a default version file, default.ver. Version control automatically creates a second version file (custom.ver) on the client if a server_custom.ver file exists on the server. The two client version files track the versions of standard and customized components installed on the client.

The default setting of the ver_ctl Options Manager option is UPGRADE, causing version control to upgrade the secondary server when a discrepancy is detected. Version checking is performed only for secondary servers and remote clients (there is no need for version control on clients running on the primary server because the server contains the master copies of all components). Version control is always performed against the primary server. When a secondary server or client starts, it reads its version files and sends a message to the primary server listing all installed components and their versions. If a component version discrepancy occurs, the server determines its response by checking the ver_ctl option setting. The server response tells the client whether to continue, terminate, or upgrade components.

Note: Most clients receive a download of files and option variables upon initial startup. This reflects the options chosen by the administrator when the server was configured and options installed.

Version Control for Installation Customizations

If you customize your CA SDM client, you can use version control to monitor and control the installation of the customizations on the clients. To do this, create a server_custom.ver (or server_secondary_custom.ver for secondary servers) file describing your customizations.

Version control uses the two version files (server_default.ver and server_custom.ver) independently to verify the version of components on the clients. Your components should be designed so that a file is never controlled by two or more components. Version control processes each component separately. Version control verifies a file each time it encounters a component that controls the file, which sometimes has unexpected results if the version number or date is not the same.

To use version control to monitor and control the installation of the customizations on the clients, follow these steps:

1. Create the `server_custom.ver` file in one of the following locations:
 - Linux—`$NX_ROOT/site/mods/server_custom.ver`
 - Windows—`installation-directory\site\mods\server_custom.ver`
2. Add components to `server_custom.ver` for each of your changes, using one of the following setups:
 - Create `dir_ctl` components to implement version control at the directory level. When you do this setup, all files in the directory are considered to be at the same version and are upgraded at the same time. This setup minimizes maintenance to `server_custom.ver`, but may sometimes result in upgrading files that are already up-to-date on the client.
 - Create `file_ctl` components to implement version control at the file level. This setup provides better documentation and control, but the `server_custom.ver` file may be larger.
3. Update the version number in `server_custom.ver` whenever you make a change to a customized component. The version number lets version control recognize clients that have a different version.

Version Control Server Modes

The `ver_ctl` server option determines how version control responds to version discrepancies between installed components on the server and client. You set this option with the Options Manager.

The `ver_ctl` option can have one of the following values:

NOTIFY

(Windows client) Prompts the user to continue or quit if a version discrepancy is detected.

(Linux client) Logs a version discrepancy, but the client always completes initialization. No files are changed to correct the discrepancy.

FAIL

Terminates the client if a version discrepancy is detected. The user is prevented from using CA SDM until the client is manually updated.

DISABLE

Ignores version discrepancies. All clients are allowed to connect.

Note: Clients can set their own `ver_ctl` option to `DISABLE` to bypass version control.

UPGRADE

Prompts the client to upgrade the affected components to the correct version if a version discrepancy is detected. If the upgrade is successful, the client connection continues; otherwise, the client terminates. The method of upgrade depends on the version control type of the component.

Valid settings are as follows:

Setting	Description
dir_ctl	The client adds, updates, or deletes files in the directory from copies provided by the server. The directory itself is never deleted.
file_ctl	The client adds, updates, or deletes the file from a copy provided by the primary server.
nxenv_ctl	The client updates its NX.env file.
ver_ctl	A ver_ctl component cannot be upgraded. A discrepancy in this type of component causes the client to terminate without connecting to the server.

Version Control File Syntax

Each version control file describes one or more components. A component can represent a file, directory, or the client_nx.env environment variable file. A component can also be generic; that is, not directly associated with any external object.

Note: For more information about the structure and syntax of version control files, look at the .ver files that are installed on your system in the site subdirectory of the main CA SDM installation directory. These files have useful comments and example settings that may help you.

Use the following syntax to define each component. Items in *italic* print represent data that you supply. The *component-name* and version parameters are always required. Other parameters may be required, depending on the value of *control-type*. All other items represent keywords that you must enter exactly as shown in the following example:

```
[ component-name ]  
version = "x.x, yyyymmdd"  
control-type  
filename = "filename"  
directory = "directory"  
link = "link-directory"  
source = "source-directory"  
effectivity = "effect-spec"  
checksum  
min_release = "release"  
max_release = "release"  
component_type = "file-type"  
o_mode = "owner-mode"  
g_mode = "group-mode"  
w_mode = "world-mode"
```

Version Control Parameters

The following parameters apply to version control:

[**component-name**]

Specifies the name of an item under version control. The name must be unique and enclosed in square brackets. *component-name* is not case-sensitive. This parameter is required to begin a component definition.

version="x.x. *yyyymmdd*"

Specifies a version number (x.x) and a date (*yyyymmdd*) that define the version of the component. This parameter is required, and must be enclosed in double quotes. Version control verifies the version of a component by comparing the version number and date on the server with the version number and date on the client. Both version number and date must match for the component to be considered in sync between client and server. Optionally, if the checksum property is enabled, the file is verified by checksum verification before being updated.

control-type

Specifies the type of version control for this component. The following settings are valid for control-type:

Setting	Description
dir_ctl	Specifies that the component represents a directory. You must provide the directory parameter to specify the path to the directory. You can also provide the filename parameter to specify a mask for files in the directory to upgrade when an upgrade is required. Subdirectories are not upgraded on either UNIX or Windows.
file_ctl	Specifies that the component represents a file. You must provide the directory and filename parameters to specify the path to the file.
Nxenv_ctl	Specifies that this component represents the client_nx.env file, which is used to store internal CA SDM environment variables. CA SDM version control and the Options Manager automatically maintain this file. There is one nxenv_ctl component, and its component name must be CLIENT_NXENV.
ver_ctl	This is the default control type. It specifies that the component is generic; that is, not associated with any specific external object. You can use a generic component to provide version control for the client as a whole, or for a file or directory too large for an automatic upgrade. Components with a control type of ver_ctl cannot be upgraded; a version mismatch on a ver_ctl component when the server is in UPGRADE mode causes the client connection to fail.

filename="filename"

Specifies the name of a file under version control. It does not contain directory specifications. This parameter is required for file_ctl components, but is optional for directory (dir_ctl) control components. When used with directory components, the filename parameter acts as a file mask to restrict the files associated with the dir_ctl component. For example, if the filename for a dir_ctl component is *.README, then an upgrade from that directory includes only files ending with ".README".

directory="directory"

Specifies the path to the directory for dir_ctl components, or to the directory containing the file for file_ctl components. This parameter is ignored for ver_ctl and nxenv_ctl components. The directory path must be enclosed in quotes, and can contain references to environment variables preceded with a \$.

Note: Always use forward slashes (not backslashes) to separate subdirectories in the path name, even on a Windows server.

link="link-directory"

Specifies a link directory on the client in the same format described previously for directory parameter. This parameter is optional for file_ctl and dir_ctl components, and ignored for ver_ctl and nxenv_ctl components. If it is specified, an upgrade to a Linux client causes a symbolic link to be placed in the link directory, pointing to the actual file copied to the location specified by the directory parameter. An upgrade to a Windows client causes the actual file to be copied to both the link and directory locations.

source="source-directory"

(Optional) Specifies a different directory on the server where the server can retrieve files for delivery. This parameter has the same format described previously for the directory parameter. It is useful if the files that are to be delivered to the client are different from the same files in the directory location on the server. This parameter is used to tell the server to retrieve the file from *source-directory* and deliver it to the location on the client specified by the directory parameter. The directory parameter is required if you specify the source parameter.

effectivity="effect-spec"

(Optional) Specifies whether the client should get this component. It lets you exclude download to some clients. If a client is not included in the effectivity specification, it does not get the component. If this parameter is omitted, all clients receive the component. The effectivity specification uses the following symbols:

+ (plus sign)

Indicates to add a client group.

- (minus sign)

Indicates to exclude a client group.

The following client groups are valid:

- SUN4SOL
- AIX
- LINUX
- LINUX390
- HP
- WINDOWS_CLIENTS
- UNIX_CLIENTS

For example, the following specification indicates that only UNIX clients get the files:
`effectivity = "+ UNIX_CLIENTS"`

checksum

Directs the component to upgrade if the checksum of the component on the client does not match the corresponding checksum on the server. If it is applied to a directory, then checksum is applied to each file.

min_release="release" and max_release="release"

Specifies the oldest and latest client to which this component should be distributed. Statements in the `server_default.ver` file define releases. These parameters are in the following form, where `Gaxx` indicates the release, and the values following are genlevels associated with the release.

`! Release GA50 50MVV000900 50W7T000900`

`! Release GA45 45MW000900 50WTT000900`

The order indicates that GA50 is newer than GA45.

[component-name]

Specifies the name of an item under version control. The name must be unique and enclosed in square brackets. *component-name* is not case-sensitive. This parameter is required to begin a component definition.

Setting	Description
file	This is the default component type. It specifies that files copied to the client be obtained directly from the location on the server indicated by the <code>directory</code> parameter.
exe_file	Specifies that files copied to the client be obtained from a location on the server that is dependent on the client's operating system, as shown by the following: ■ windows (Windows) ■ sun4Sol (Solaris) ■ hp (HP-UX) ■ aix—AIX) ■ linux (Linux) ■ linux390 (Linux390) Locations for these subdirectories are dependent on the <code>directory</code> parameter setting. If this parameter is set, then subdirectories are located under the indicated <i>directory</i>. Otherwise, they are located under the <code>bin</code> directory of the main CA SDM installation directory.

o_mode="owner-mode"

Specifies file access permissions for the file's owner.

g_mode="group-mode"

Specifies file access permissions for users in the file owner's group (used for UNIX clients only).

w_mode="world-mode"

Specifies file access permissions for users not in the file owner's group (used for UNIX clients only).

The three mode parameters allow different versions of the same executable to be maintained on the server. They specify access controls on the file when copied to the client. These parameters are used only during an upgrade operation. They consist of one to three characters, with the following significance:

Setting	Description
R	Read
W	Write
X	Execute

PC clients ignore Write and Execute permissions.

You can specify any combination of one or more file access modes. On UNIX clients, the file is given the access mode of specified. On PC clients, the file is made writable or read-only, depending on whether w_mode has been specified.

Uncontrol a Component

You may need to remove control from a component, but if you simply delete the component from the primary server's server_default.ver file, the associated files delete, as well. Instead, you can define an *uncontrol* block in the server_default.ver file, which lets you delete the component without also deleting the associated files on the clients. These statements take the following form:

```
! uncontrol component-name
```

This statement indicates that this component no longer has control. Clients connecting to this server need to remove this component from their version file, but also need to remove the files associated with this component. You can remove the uncontrol block after all the clients have successfully connected at least once.

Note: You should update all clients with the most recent version control client. Older version control clients can ignore uncontrolled components.

Chapter 14: Managing Configuration Items

This section contains the following topics:

[Using the Web Interface](#) (see page 449)
[Contact, Location, and Organization CIs](#) (see page 453)
[CI Relationships](#) (see page 456)
[Versioning](#) (see page 464)
[View CI Attributes in Other CA Products](#) (see page 489)
[Using the CMDBf Viewer](#) (see page 490)
[CMDB Visualizer](#) (see page 490)
[Add a Discovered Asset](#) (see page 493)
[Asset and CI Flags](#) (see page 494)
[CI Reconciliation](#) (see page 495)
[Manage Staged Transactions](#) (see page 514)
[CA CMDB Data Maintenance](#) (see page 538)
[Configuration Audit and Control Facility \(CACF\)](#) (see page 551)

Using the Web Interface

Depending on your role, you can perform user and administration tasks on configuration items (CIs) by using the following features:

- Scoreboards allow you to perform user tasks and manage CIs and CI relationships.
- The Administration tab lets you perform administration tasks with CIs and relationships..
- CMDB Visualizer lets you perform user and administration tasks for visualizing CIs and relationships.

More information:

[View Configuration Items](#) (see page 450)
[Create a Configuration Item](#) (see page 450)
[Update a Configuration Item](#) (see page 451)
[Associate a Maintenance Window with a CI](#) (see page 451)
[View Associated Change Windows](#) (see page 452)
[View Configuration Item History](#) (see page 452)
[Inactivate a Configuration Item](#) (see page 452)
[Reactivate a Configuration Item](#) (see page 453)

View Configuration Items

The Scoreboard lets you open a configuration item and view information about it.

To open and view a configuration item

1. From the Scoreboard, click Search, Configuration Items.
The Configuration Item Search pane appears.
2. Enter search criteria, and click Search.
The Configuration Item List appears and lists the search results.
3. (Optional) Click a configuration item in the Name column to open it and view its detailed information.
Configuration item details appear.

Note: After you create a CI, the CI Detail page does not display some identifying attributes if they are blank (have no value) and do not apply to the CI family. However, identifying attributes with non-blank values always display.

Create a Configuration Item

The Scoreboard lets you create a configuration item.

To create a configuration item

1. From the Scoreboard, click File, New Configuration Item.
The Create New Configuration Item page appears.
2. Enter a unique name for the configuration item in the Name field.
3. Enter the class for the configuration item in the Class field or click the search icon above the field to locate a class.
4. Complete any remaining fields that apply to the new configuration item.

Note: Only Name and Class are required values for creating a configuration item.

5. Click Continue.

The Create New Configuration Item page displays additional tabs and fields.

6. Enter data as required in the appropriate fields on the Attributes tab.

The attributes that appear on the tab are determined by the family of the class that you selected for the configuration item. The information that you enter here is determined by your business processes and the information you want to store and view for a configuration item.

7. Click Save.

A message appears that confirms the configuration item creation.

Update a Configuration Item

The Scoreboard lets you modify a configuration item.

To modify a configuration item

1. From the Scoreboard, locate and open the configuration item.
The Configuration Item Detail page appears.
2. Click Edit.
The Update Configuration Item page appears.
3. Make changes to the configuration item information by selecting the appropriate tabs on the page and entering new data into the fields. Click Save.
A message appears that confirms the configuration item modification.

Note: Whenever a CI is updated by using the –a option, the CI's Last Change Date and user displayed in the Configuration Item List are updated even if no attributes were changed. This update occurs whether a CI was edited in the user interface (and saved without changes) or updated by using GRLoader.

Associate a Maintenance Window with a CI

To control maintenance access to a CI, you can associate that CI to one or more non-global maintenance windows.

To associate a maintenance window with a CI

1. Navigate to the CI Detail page.
2. Click the Maintenance Windows tab.
The Maintenance Windows List appears.
3. Click Update Maintenance Windows.
The Maintenance Windows Search page appears.
4. Specify information for required maintenance windows.
5. Click Search
The Available Windows page appears.
6. Move any required windows in the Available Change Windows list to the Associated Change Windows list.
7. Click OK.
The CI, change windows, and new window associations are saved.

View Associated Change Windows

To determine how a CI can be affected during specific time periods, you can view the maintenance windows that are associated with that CI.

To view change windows associated with a CI

1. Navigate to the CI Detail page.
2. Click the Maintenance Windows tab.

The associated maintenance windows are displayed.

View Configuration Item History

You can view the history of a configuration item.

To view the history of a configuration item

1. Search for and open the configuration item.
The Configuration Item Detail page appears.

2. Click the Versioning tab.
The Versioning page appears.

3. Click Show Log.
The configuration item history is listed.

Note: Show Log ignores the left pane selections.

Inactivate a Configuration Item

If a configuration item is no longer used, you can edit the configuration item details to make it inactive. An inactive status removes the configuration item from the Scoreboard Configuration Item List. You cannot delete a configuration item.

Note: Inactivating a Family or Class does not affect existing CIs in that Family and Class. It only prevents new CIs from being created in that Family and Class.

To inactivate a configuration item

1. Search for and open the configuration item.
The Configuration Item Detail page appears.

2. Click Edit.
The Update Configuration Item page appears.

3. Select Inactive from the Active drop-down list. Click Save.

A message appears that confirms the status change. The configuration item no longer appears in the list of active configuration items. The configuration item continues to appear in any existing relationships until the relationships are inactivated.

Reactivate a Configuration Item

If you need to use a configuration item that is inactive, you can reactivate it.

To reactivate a configuration item

1. Search for and open the configuration item.

The Configuration Item Detail page appears.

2. Click Edit.

The Update Configuration Item page appears.

3. Select Active from the Active drop-down list. Click Save.

A message appears that confirms the status change. The configuration item appears in the list of active configuration items.

Contact, Location, and Organization CIs

A CA CMDB installation can define CIs for *base objects*: Contacts, Locations, and Organizations. These CIs can be managed accordingly, and they can establish relationships with other CIs as desired. For some customers, this approach supports the configuration management process better than employing base objects only as CI attributes.

More information:

[Create a CI from a Base Object](#) (see page 454)

[Select a Base Object for a CI](#) (see page 454)

[Edit CI Details for a Base Object](#) (see page 455)

[Edit CI Attributes for a Base Object](#) (see page 455)

[Create a Base Object CI Using GRLoader](#) (see page 456)

Create a CI from a Base Object

The Scoreboard lets you create a configuration item from a base object.

To create a CI from a base object

1. Click File, New Configuration Item.

The Create New Configuration Item page appears.

2. Complete the CI fields. Name and Class are required. Click Continue.

Note: The common CI attributes Host Name, Serial Number, MAC Address, and DNS Name are not relevant to a CI for a base object.

The Create New Configuration Item page appears.

3. Click the base object link to define the object that this CI represents.
4. Select the object. You can search for an existing object or click Create New to create an object. If you want to create an object, click Save to continue.

The selected object appears at the top of the Configuration Item Detail page.

5. Click Save.

The main attributes of the selected object appear on the Attributes tab of the Configuration Item Detail page.

Select a Base Object for a CI

You can select a base object for a configuration item.

To select a base object for a CI

1. Select a CI in the base object's family (Contact, Location, or Organization).

The Configuration Item Detail page appears.

2. Click Edit.

The Update Configuration Item page appears.

3. In the top section, complete the required fields for the base object. You can enter a value directly or click the magnifying glass to search for the object. For example, to search for a contact, complete one or more of the search fields on the Contact Search window, and click Search; then select a contact from the displayed list.
4. Click Save.

If the selected object is already represented by a different CI in the same family, an error message appears. Select a different object before you can modify the CI.

Edit CI Details for a Base Object

You can edit the attributes of the base object that the CI represents, such as Alt CI ID and Notes, in the same way as in any other Configuration Item Detail window.

To edit the attributes of the base object that the CI represents

1. Select a CI in the base object family.
The Configuration Item Detail page appears.
2. Click the object link at the top of the page.
The object Detail page appears.
3. Click Edit.
The object Update page appears.
4. Edit the fields as desired and click Save.
The Update page closes.

If you have modified one of attributes that appears in the Attributes tab of the Configuration Item Detail page, the change is displayed there after the page is refreshed.

Edit CI Attributes for a Base Object

You can edit the configuration item attributes for a base object.

To edit the attributes for a CI

1. Select a CI in the object's family (Contact, Location, or Organization).
The Configuration Item Detail page appears.
2. Click Edit.
The Update Configuration Item page appears.
3. On the Attributes tab, click Edit.
The Update page appears.
4. Edit the fields as desired, and click Save.
The Update page closes.
5. Click Save.
The Configuration Item Detail page displays the modified CI.

Create a Base Object CI Using GRLoader

You can use GRLoader to create a base object CI from an existing base object.

To use GRLoader to create a base object CI from an existing base object

1. Write XML that identifies the following:

- CI name
- Family
- Class

2. Submit the XML.

The base object is created and the GRLoader displays that one CI was read and inserted.

Example: Create a CI From an Existing Contact

The following example creates a CI from the existing contact **Gibbs, Keith**.

```
<GRLoader>
<ci>
  <name>Gibbs, Keith</name>
  <family>Contact</family>
  <class>Technical</class>
  <base_contact>Gibbs, Keith</base_contact>
</ci>
</GRLoader>
```

CI Relationships

CA SDM functionality depends on the relationships among configuration items. CI relationships have the following characteristics:

- The *relationship type* defines how two configuration items are related.
- Hierarchical relationship types are paired as *provider/dependent*. For example, A (provider) *hosts* B; B (dependent) *is hosted by* A.
- Non-hierarchical relationship types are defined as *peer-to-peer*. For example, *is connected to*, which is the same in either direction.

You can do the following relationship functions by using the CMDB Relationships tab:

- View Active or Inactive relationships by searching on the Active? field.
- From the CI Relationship List, you can click any link for a relationship to launch its CI Relationship Detail page.

More information:

[CI Relationship Types](#) (see page 457)
[Create a Relationship Type](#) (see page 457)
[Manage a CI Relationship](#) (see page 458)
[Create a CI Relationship](#) (see page 459)
[View Relationships for a CI](#) (see page 459)
[Inactivate a CI Relationship](#) (see page 460)
[Reactivate a CI Relationship](#) (see page 460)
[Inactivate CI Relationships \(Edit in List\)](#) (see page 461)
[Inactivate a CI Relationship Using GRLoader](#) (see page 461)
[Reactivate a CI Relationship Using GRLoader](#) (see page 462)
[Delete a CI Relationship from the Database](#) (see page 463)
[CI Relationship History and Comparison](#) (see page 463)

CI Relationship Types

CA SDM provides a list of predefined relationship types that can be used to describe a relationship or association between configuration items. How the relationship is expressed depends on which configuration item is the focus. For example, a provider *supports* a dependent configuration item, but the dependent *is supported by* the provider. These are different expressions of the same relationship, and the relationship type label that you see depends on whether you are viewing a provider-to-dependent or a dependent-to-provider relationship.

Note: For information about the relationship types provided by CA CMDB, see the *CA CMDB Technical Reference Guide*.

Create a Relationship Type

You can create a relationship type, and make it available to create or update relationships.

To create a relationship type

1. On the Administration tab, expand the CA CMDB folder in the left pane.
2. Click CI Relationship Types.
The CI Relationship Type List appears in the right pane.
3. Click Create New.
The Create New CI Relationship Type page appears.
4. Complete the following fields:

Provider to Dependent Label

Describes the relationship from the provider to the dependent, and is the name that appears in the list of relationship types. For example, "services" or "manages" is a Provider to Dependent relationship.

Dependent to Provider Label

Describes the relationship from the dependent to the provider, and is the name that appears in the list of relationship types. For example, "is serviced by" or "is managed by" is a Dependent to Provider relationship.

Peer-to-Peer?

Specifies that the relationship does not have a provider CI and dependent CI but instead consists of two equal CIs. For example, "connects to" and "fails over" are Peer-to-Peer relationships.

Active?

Makes the relationship type available for selection in configuration item relationships and to appear in lists.

5. Click Save.

The relationship type is created and you can use it to create or update relationships.

Manage a CI Relationship

The Configuration Item Relationship page provides all necessary settings in one location to simplify the CI relationship process.

You can use the page to do the following functions:

- Toggle the provider/dependent relationship between two CIs.

Click Reverse to switch the Provider CI and Dependent CI for the relationship. The focal CI changes. If the reversed relationship is not valid, clicking Reverse clears the Relationship field.

- Set Optional fields such as Symbol, Description, and Cost.

Click Optional to set optional fields.

If a Symbol is not supplied, a unique Symbol is automatically generated with prefix **cmdb bmhier:** and a unique number when the relationship is created.

Note: If your installation is integrated with CA Network and Systems Management (CA NSM), you are prompted to specify the CA NSM repository and cost.

Create a CI Relationship

You can add a CI relationship by starting from a focal configuration item.

To add a relationship from a focal CI

1. Select the CI to which you want to add a relationship.
The CI Detail page appears.
2. Click Edit.
The Update CI page appears.
3. Click the CMDB Relationships tab, and click Create New.
The Create CI Relationship page appears.
4. Specify the Provider CI. You can click the magnifying glass link to search for the CI.
Note: Depending on the relationship you want, you can click Reverse to toggle the provider/dependent positions.
5. Specify a Relationship Type. You can click the magnifying glass link to search for the relationship type.
6. Specify the other CI (by default, the Dependent CI). You can click the magnifying glass to search for the CI.
7. Click Save.
The relationship is created.

View Relationships for a CI

You can view relationships for a CI.

To view relationships for a CI

1. Navigate to the configuration item.
The Configuration Item Detail page appears.
2. Select the CMDB Relationships tab.
A list of the relationships for the CI appears.

Inactivate a CI Relationship

You can Inactivate a CI relationship by setting its status to Inactive.

To Inactivate a relationship

1. Navigate to the CI Detail page, click the CMDB Relationships tab, and click Edit.

The CMDB Relationships tab displays in Edit mode.

2. Choose a relationship and click any link on its row.

The CI Relationship Detail page displays.

3. Click Edit.

4. Set the Active? attribute to Inactive. Click Save.

The relationship is made Inactive. On the CI Relationship Detail page, the “Active?” attribute displays the Inactive attribute value.

Reactivate a CI Relationship

You can reactivate a deleted CI relationship.

To reactivate a deleted CI relationship

1. Navigate to the CI Detail page.

The CI Detail page displays in Read-Only mode.

2. Click Edit.

The CI Detail page displays in Edit mode.

3. Click the CMDB Relationships tab.

4. Click any link on a row to open that CI Relationship Detail page.

5. Click Edit.

The Update CI Relationship page appears.

6. Change the Active? attribute to “Active”. Click Save.

The relationship is reactivated.

Inactivate CI Relationships (Edit in List)

You can Inactivate one or more CI relationships by using the Edit in List feature.

To Inactivate one or more CI relationships

1. From the Administration tab, open the CA CMDB folder in the left pane.

The CA CMDB node expands to display subfolders.

2. Click the CI Relationship List folder.
3. (Optional) Specify search filters
4. Click Search.

CI relationships are displayed in the right-hand pane.

Note: By default, all relationships are displayed.

5. Click Edit In List at the top right of the list.

The Relationship Edit in List controls are displayed.

6. Select the row that contains the relationship you want to modify.

The row displays and the controls are populated with the relationship's values.

7. Set the Active? attribute to Inactive.

Click Save if you want to Inactivate only the highlighted relationship or click Change All if you want to Inactivate all relationships in the filtered search.

Inactivate a CI Relationship Using GRLoader

You can Inactivate a CI relationship by using GRLoader to submit XML.

To Inactivate a relationship using GRLoader

1. Write XML that uses a Relation tag to identify the following:

- Provider CI
- Dependent CI
- Relationship type

2. Set the delete_flag to true (or yes or 1).

3. Submit the XML.

The CI Relationship is deleted.

Note: For more information about GRLoader, see the *CA CMDB Technical Reference Guide*.

Example: Delete a CI Relationship Using GRLoader

In the following XML example, the "connects to" relationship between ci_1 and ci_2 is deleted.

```
<GRLoader>
  <relation>
    <dependent>
      <name>ci_2</name>
    </dependent>
    <type>connects to</type>
    <provider>
      <name>ci_1</name>
    </provider>
    <delete_flag>true</delete_flag>
  </relation>
</GRLoader>
```

Reactivate a CI Relationship Using GRLoader

You can reactivate a CI relationship by using GRLoader to submit XML.

To reactivate a CI relationship using GRLoader

1. Write XML that uses a relation tag to identify the following:
 - Relationship type
 - Dependent CI
 - Provider CI
2. Set the delete_flag to false (or no or 0).
3. Submit the XML.

The CI Relationship is reactivated.

Note: For more information about GRLoader, see the *CA CMDB Technical Reference Guide*.

Example: Reactivate a CI Relationship

In the following XML example, the "connects to" relationship between ci_1 and ci_2 is reactivated.

```
<GRLoader>
  <relation>
    <dependent>
      <name>ci_2</name>
    </dependent>
    <type>connects to</type>
    <provider>
      <name>ci_1</name>
    </provider>
    <delete_flag>false</delete_flag>
  </relation>
</GRLoader>
```

Delete a CI Relationship from the Database

You can delete a CI relationship so that it is not just Inactivated but removed from the database permanently.

To remove a relationship from the database

1. Navigate to the CI Relationship List.
The CI Relationship List appears.
2. Right-click on any relationship row and select Delete.
The relationship is removed from the database.

Important!: If the deleted relationship has a Symbol defined, then any relationship with the same Symbol is also removed from the database.

CI Relationship History and Comparison

The Versioning tab displays CI relationships and lets you do the following:

- View the relationship history for any CI in the CMDB. Changes made to a relationship are logged automatically when the relationship is created, updated, or deleted. You can view all current and historical relationships for a CI.
Note: Relationships created with previous versions of CA CMDB do not have audit history information.
- Compare relationships at any snapshot or milestone.

Versioning

Versioning promotes control over the IT infrastructure by tracking and managing the life cycles of all the CIs that make up the authorized state of the CMDB. Versioning also applies to auditing the history of an object. Versioning provides the following functions to control the IT infrastructure:

Snapshots

Recorded automatically for changes that you made to the value of any update to an object, such as when you update a CI. Versioning creates a snapshot that records the complete state of the hardware CI after the memory size of a computer is modified. Versioning can display the snapshot and can indicate which CI attributes have changed. Versioning can also compare it against all other historical snapshots of that CI. You require this critical information when you want to understand the impact of changes on the availability and performance of a CI.

Milestones

Recorded as special-purpose snapshots with a user-defined label, such as First Day of Production or January Baseline. These labels can help find specific snapshots more quickly.

Important! Milestones do *not* apply to change specifications, verification policies, managed change states, and managed attributes.

Comparisons

Compare a CI to another CI that acts as a standard. You can use any CI as a standard for comparison purposes, but we recommend that you dedicate specific CIs for the purpose of acting as *Standard CIs*. A Standard CI lets you define a standard configuration to which other operational CIs in the same family can be compared. For example, an organization can decide to define Server Large as a Standard CI to define the attribute values that define this type of server. The Standard CI for a family can be made an attribute of an operational CI in that family, which lets you compare the current state or any historical state of a CI to its associated standard configuration. Like other comparisons, you can print this information or export it as a comma-delimited file.

Whereas milestones can act as unshared, unchangeable baselines, Standard CIs provide shared, dynamic baselines for your CIs.

Important! A Standard CI must never contain values for any attribute that is used for CI reconciliation.

Change Specifications

View pending scheduled or unscheduled change specifications that a user specified for Change Orders for the CI. You can compare the change specification with the current state of the CI. For example, compare a planned value to the current state of a CI. This comparison can show conflicts with changes, identify overlapping changes in the schedule, and so on.

Versioning lets you complete the following tasks:

- Automatic capture of all CI modifications
- Snapshots of a CI at any time based on date or attribute changes
- Various levels of display including detailed Log, Basic, and Advanced views
- Multiple attribute comparisons with other snapshots, user-defined milestones for CIs, or a Standard CI
- Automatic snapshots whenever CA SDM change tickets change states

Note: This automatic snapshot does not apply to change specifications, verification policies, managed change states, and managed attributes.

- Advanced filtering and comparison with Print and CSV Export support

More information:

[Uses of Versioning](#) (see page 465)

[Shared Asset and CI Audit Trail Records](#) (see page 466)

[Versioning Terminology](#) (see page 466)

[Sources of Versioning Data](#) (see page 469)

[CA SDM Change Management Integration](#) (see page 470)

[CA APM Integration](#) (see page 471)

[CI Versioning Management](#) (see page 472)

[CA SDM Change Management](#) (see page 488)

[Configuration Audit and Control Facility \(CACF\)](#) (see page 551)

[CI Versioning and Future State](#) (see page 573)

Uses of Versioning

Versioning includes the following uses:

Problem Determination

To correct a problem with a CI, you *must* understand what changed in the environment of the CI. Versioning shows its current defective state, and its state at any point in the past, which lets you compare the two states to help identify potential problems.

Performance and Capacity Planning

By reviewing the history of a CI, a performance or capacity planner is better able to determine the causes of performance impasses and to plan for future system growth.

Compliance

You can compare the state of a CI to its family Standard CI. Compare the state at any point in its change management life cycle to help ensure that the CI is in compliance and to help identify attributes that need corrections.

Change Verification

You can view the audit history of the object, such as which user modified a verification policy, the date of the request, the attributes and values that the user modified. You can compare and contrast the changes between specific dates.

Shared Asset and CI Audit Trail Records

The Versioning feature includes audit trail changes made to CA APM. The Versioning feature also supports launch in context from CA CMDB to CA APM directly from an attribute log entry associated with each CA APM change. The CA APM audit trail information is only available for CI/Assets from the CA CMDB families with CA APM auditing logging enabled.

The Versioning tab that includes the CA APM audit trail information supports all families, provided logging for that family is enabled.

By default, CA APM does not log Asset attribute changes. To use CA CMDB Versioning with a CA APM-managed Asset/CI, the Store Audit Trail Data option must be enabled on each asset attribute that requires logging.

Note: For more information about how to enable auditing, see the CA APM documentation.

When asset attributes are modified in CA APM, the Versioning tab data on the CI Detail page can be updated before the Attributes tab data due to caching activity. To resynchronize the values, click the Edit button.

Versioning Terminology

Versioning is the practice of representing and tracking *service transitions*. Versioning typically uses a naming convention to identify the dates, sequences, and meanings of such transitions. These records can be used to identify and compare specific states of a configuration item (CI).

Example: Version Comparison

For a Payroll Application CI, a comparison of Version 3 against Version 2 indicates the enhanced features and other differences of the newer version.

More information:

[States](#) (see page 467)

[Versions](#) (see page 467)

[Snapshots](#) (see page 467)

[Categories](#) (see page 467)

[Milestones](#) (see page 468)

[Standard CIs](#) (see page 468)

States

In the context of versioning, the *state* of a CI represents all the attribute values of that CI at a single point in time. The state of a CI can be the result of attribute changes from multiple MDRs.

Versions

A *version* is an identified instance of an object such as a CI within a product breakdown or configuration structure. Use versions for purposes of tracking and auditing change history.

Snapshots

A *snapshot* is a representation of the complete state of an object at a single point in time. For example, a typical CI has many snapshots associated with it. A snapshot consolidates all modification events that occur to an object during a one-minute time interval. For example, if you update a CI (edited and saved) several times within one minute, CA SDM creates a single snapshot that captures all of the updates during that minute.

Snapshot is a general term that refers to time and date-based snapshots, as well as milestones or Standard CIs. To help you to locate meaningful points in time, CA SDM lets you label snapshots meaningfully. These named snapshots are called milestones.

Every time a change is made to an object, CA SDM creates a snapshot automatically, without the need for manual action. Versioning captures all changes to the object. For example, snapshots created using the user interface, or snapshots originating at an MDR and came to the CMDB through GRLoader or CA CMDB Web Services.

Categories

A *category* identifies a class of attributes. The category is typically the name of the tab that displays the attributes.

Note: Items that are not found on a tab (for example, name, serial number, active flag) are assigned a category that is not a tab name.

Examples: Tabs and Categories

The following examples show how categories correspond to tabs:

- Because *disk space* appears on the Attributes tab, its category is *Attributes*.
- Because *IP address* appears on the Inventory tab, its category is *Inventory*.
- *Milestone* appears in the *General* category.
- A *Standard CI* appears in the Classification category.
- *Relationships* appear in the Relationship category.

Milestones

A *milestone* is an on-demand labeled snapshot of a CI that is created to mark an event, a logical breakpoint, or an accumulation of changes. The milestone contains the actual state of the CI at the time that the snapshot was created. Milestones let you quickly identify and navigate to significant points in the history of a CI. Creating a milestone, creates an equivalent “date” snapshot.

Milestones are CI-specific, not shared. When you create a milestone for a high-level object such as a service, the subcomponents of that high-level object do not automatically create milestones. To take a milestone for an object which is composed of several subcomponents, you can generate several independent milestones.

Milestones are static and can never be changed or be deleted. When you create a milestone, the snapshot is of the current state of the CI. Later, when displayed or used in a comparison, milestones always reference a point in the past. A good naming convention for your milestones helps you easily identify critical points in the life of a CI.

Important! Milestones do *not* apply to change specifications, verification policies, managed change states, and managed attributes.

Standard CIs

A *Standard CI* is an abstracted configuration for a CI family that you can use for baseline comparisons to “real” CI instances in the same family. A Standard CI can be a real CI in the sense that it represents a physical object or it can exist only for comparison purposes. Standard CIs can have the following associations:

- A Standard CI can be shared among several CIs.
- A CI can only have a single Standard CI associated with it at any given time. When the Standard CI is changed, that change is reflected in any comparison between any of the associated CIs and the Standard CI.
- A family can have multiple Standard CIs. For example, one family might have standard CIs named “Standard Test Server,” “Standard Production Server,” “Standard Acceptance Server.” We recommend that you name Standard CIs in such a way that they can be easily searched for and identified as Standard CIs.

Because a Standard CI is itself a CI, it can be managed. It has an audit trail, security, a history of changes, and so on, just like any other CI. After a Standard CI has been defined for a particular CI, you can use it to verify compliance with corporate standards.

The concept of “date” or “time” does not apply to when comparing a Standard CI with a snapshot or milestone. Only the Standard CI attribute values are used for comparison purposes, snapshots and milestones specific to the Standard CI do not apply.

Example: Standard CI Use

A company defines an **Employee Workstation** configuration as a Standard CI to compare with its actual desktop computers in the Hardware.Workstation family. A comparison reveals that a specific computer has only 1 GB of RAM, instead of the Standard CI memory value of 2 GB.

Sources of Versioning Data

Versioning data generates whenever you create or modify an object by any component that supports the Audit Log facility. Snapshots generated from multiple sources are indistinguishable from one another. When the environment is properly configured, snapshots appear automatically whenever changes are made to CIs and their relationships.

All CA CMDB families are enabled for versioning. CIs in CA Service Desk base families must be converted to CA CMDB families to take advantage of versioning.

Versioning data includes the following sources:

- **CA CMDB CI updates**—Updates to common CI attributes and family-specific CI attributes by using the user interface.
- **CA CMDB relationship updates**—Updates to relationships on the Relationships tab and the CI Relationship List by using the Edit and Edit In List feature.
- **GRLoader**—CI insert and attribute updates, relationship updates, Standard CI assignments, and Milestones imported by using GRLoader. When updates from an MDR are sent using GRLoader, the MDR source is also recorded, enabling you to track attribute changes directly to their source.

Note: For more information, see the *Technical Reference Guide*.

- **CA SDM Change History**—A snapshot is automatically created for all CIs associated with a change ticket. Up to four snapshots are taken when the CI is opened, closed, active or resolved.

- **CA Asset Portfolio Management (CA APM) CI changes**—Changes made by CA APM to CIs with logging enabled.
- **Change Verification**—Updates to change specifications, verification policies, managed change states, and managed attributes.

Note: CIs created in CA SDM or previous CA CMDB releases only contain modification changes. Initial attribute values such as name, family and class were not logged and do not appear in snapshots. In addition, Relationships created with CA Service Desk or previous versions of CA CMDB do not have audit history information.

CA SDM Change Management Integration

Versioning data is generated for each CI that is associated with a CA SDM change ticket. As the change ticket moves from open to active to resolved to closed, a snapshot is taken for each of the associated CIs. Thus, snapshots are created for opened, closed, active, or resolved statuses. If no CIs are associated with a change request, no snapshots are taken.

Note: The versioning feature requires no special configuration and is supported automatically in integrated installations.

The versioning feature handles the change ticket as follows:

1. Compares the state of the CI at the time the ticket was opened with its current state.
2. Verifies that all the required changes have been properly executed and that no additional changes were introduced.
3. Closes the ticket after the changes have been validated.

For auditing purposes, you can readily compare the state of the CI before and after each change was made. The comparison information helps answer questions including the state of a CI before and after a change request, for example:

- Did the appropriate change occur as requested?
- What other changes occurred to the CI not related to the change ticket?
- What time did the changes occur?
- How does this CI compare against the standard configuration, both before and after the change?

CA APM Integration

Versioning displays changes made by CA APM to its CIs. Versioning also supports launch in context from CA CMDB to CA APM directly from any attribute log entry that is associated with a CA APM change. CA APM audit information is only available for CIs from CA CMDB families when CA APM auditing logging is enabled.

Note: When CI attributes are modified in CA APM, the Versioning tab data on the CI Detail page can be updated before the Attributes tab data, due to caching activity. To resynchronize the values, click the Edit button.

By default, CA APM does not log CI attribute changes. To use Versioning with a CA APM-managed CI, the Store Audit Trail Data option must be enabled on each asset attribute that requires logging.

By default, a ten-minute delay exists from the time that a CI is updated in CA APM and the time it becomes available to versioning in the CMDB. You can modify this delay by changing the @NX_DBMONITOR_TIMER_MINUTES variable in an integrated installation.

Note: For information about how to enable attribute logging for assets, see the CA APM documentation.

More information:

[CA APM Integration Capabilities](#) (see page 471)

CA APM Integration Capabilities

CA SDM provides the following integration capabilities with CA APM:

- Shared records and audit log that differentiate between CA APM CIs and assets
- CA CMDB and CA APM launch in context to CI/asset information
- Shared extension table fields
- A CA SDM contact updates when CA APM changes a primary contact
- CA APM asset type modification to use CI families

CI Versioning Management

You can display and manage the history of the CI, associated snapshots, milestones, unscheduled changes, and other views. You use the left pane to navigate and the right pane to display CI details and the audit log.

Note: Versioning works similarly for managed attribute history and managed change state history. For example, the Managed Attribute History tab displays versioning information about a managed attribute.

Use versioning information to identify change specifications and the associated Change Orders. Versioning identifies unscheduled changes in the change specifications in the snapshot view that correspond to Change Orders with no scheduled start date. You can view the snapshot and log of a change specification in CACF.

Example: Unscheduled Change

In this example, a user tries to change the value of a managed attribute. You view the Versioning tab on the CI detail page that displays Unscheduled Change in the snapshot. This snapshot displays information about the change, such as the date, time, username, and attribute value.

When you select a snapshot or milestone in the left pane, the right pane displays the attribute values of that CI state, event, or comparison. The left pane can display CI snapshots or milestones by date and time (Basic mode) or by CI characteristics (Advanced mode), and includes the following links to help you manage a CI:

Create Milestone

Labels the state of a CI.

Show Log

Displays unfiltered CI history.

Advanced/Basic

Toggles between snapshot views.

Hide Empty Values

Permits filtering of blank data fields. When not selected, all CI attributes are displayed.

Print and Export

Print or cut-and-paste to save the versioning data on display.

Show Log

The log lets you view configuration item history. The Filter option permits filtering of log rows. Print and Export allow you to print or cut-and-paste to save the versioning data on display.

The Versioning information pane displays the following fields:

- Category
- Date
- Log
- Attribute
- New Value
- Old Value
- Changed By
- MDR

Note: The log only displays attribute *updates* for CIs that were created in previous releases of CA SDM or CA CMDB; it does not display their initial CI attribute values. Relationships that were created in previous releases of CA SDM or CA CMDB do not include audit history information and do not appear in the log.

Log Filtering

The log can be filtered by entering a simple string in the Filter field. The Filter is case insensitive. If the filter string appears anyplace in a log row, the row is displayed. No special characters or wildcards are used in the filter. Log rows matching the filter criteria are displayed hiding rows that do not match.

You do not need to press Enter or refresh to update the display. The log view is filtered as each keystroke is entered and applies to all fields; Date, Log, Attribute, Old value, New Value, MDR, and Name.

View Configuration Item History

You can view the history of a configuration item.

To view the history of a configuration item

1. Search for and open the configuration item.
The Configuration Item Detail page appears.
2. Click the Versioning tab.
The Versioning page appears.
3. Click Show Log.
The configuration item history is listed.

Note: Show Log ignores the left pane selections.

Print a CI Log

You can print the log for a CI.

To print the log for a CI

1. Select a CI and click its Versioning tab.
2. Click the Show Log link.
3. (Optional) Enter the desired filter criteria to hide/show log rows.
4. Click Print.
A printer friendly window appears. The Versioning report for the CI displays.
5. Use the web browser window to click File, Print to print a copy of the report.
The formatted text is sent to the defined printer.

CSV Export Support

You can export the CA CMDB log content to a Comma Separated Value (CSV) format file that third-party reporting tools or applications can import. The export is based on what is currently displayed, which can be filtered or unfiltered.

More information:

[Export Data](#) (see page 485)

Integrated CA SDM and CA APM Logs

In addition to changes that were made using CA CMDB facilities, the CI log also includes update events from CA SDM change ticket activities and CA APM updates.

More information:

[CA APM Integration](#) (see page 471)

MDR Launch-in-Context and Source Identification

CA CMDB provides a way to launch in context, directly from the log to the MDR data provider for a particular CI log entry. You can also launch in context to the change specification. This launching feature does the following:

- Tracks attribute changes back to the source MDR, when the MDR used the <mdr_name> <mdr_class> and <federated_asset_id> tags in the GRLoader input.

Note: For more information, see the *CA CMDB Technical Reference Guide*.

- Identifies when a CI attribute is being updated by more than one MDR. This situation occurs when multiple MDRs contribute data independently to a CI definition.
- Identifies which MDR is acting as the authoritative source.

Note: Log entries created in CA SDM or previous CA CMDB releases do not contain MDR launch in context information. In addition, CA Cohesion ACM provides MDR launch information for most but not all attributes or relationships.

Attribute Names

Attribute names displayed in the log match the attribute name shown in the user interface; they do not reflect the internal object name. For example, CA CMDB displays “Maintenance Type” instead of “mtce_type”.

Logging Custom Families and Attributes

To enable logging for customized families and attributes, specify new MDR attributes and audit triggers. If you create new families or attributes, you also must create metadata files to specify human readable attribute names to display in the log; otherwise, the internal object name appears.

Note: For more information, see Extending CA CMDB in the *Administration Guide*.

Create a Milestone

You can create a milestone from the user interface or by using GRLoader.

To create a milestone from the user interface

1. Select a CI and click its Versioning tab.
2. Click Create Milestone.
3. Enter descriptive text for the milestone. Click Save.

The Create Milestone page closes.

4. Select View, Refresh.

The Basic view displays a new snapshot with the current date and time that represents the milestone that you created. Switching to the Milestone view, displays only those snapshots with assigned names. Milestones are displayed in descending date/time sequence, with the most recent at the top.

To create a milestone using GRLoader

1. Add the <milestone> tag to the XML for that CI.
2. Apply the same reconciliation rules apply as when associating a milestone with a CI.

Note: The value of the milestone tag is the label associated with that milestone.

3. Save and close the XML for the CI.

The milestone is created.

Note: For more information, see the *CA CMDB Technical Reference Guide*.

Example: Use GRLoader to Create a Milestone

The following GRLoader sample creates the milestone **Fiscal year end 2008**.

```
<ci>
<name>server1 </name>
<milestone>Fiscal year end 2008</milestone>
</ci>
```

Create a Standard CI

You create a Standard CI like you create any other CI.

To create a Standard CI in the user interface

1. Select Create New Configuration Item from the File menu.
2. Select the family you want to use for the new Standard CI.
3. Define the attributes, as with any regular CI.

The Standard CI is created.

Note: Any CI can act as a Standard CI, but the following cautions apply:

- A Standard CI should follow a naming convention so that it is not confused with regular CIs.
- A Standard CI must only act as a baseline for CIs in the same family as the Standard CI.
- A Standard CI should *not* be assigned values for any attributes used in reconciliation; for example, MAC Address, Serial Number, Alt CI ID, DNS, and so on.
- Because standard CIs are indistinguishable from regular CIs, they appear in your Scoreboard and count in the total number of CIs.
- It is best if a Standard CI does not represent any actual instance of a business object.

Assign a Standard CI to a CI

You can assign a Standard configuration item to a configuration item by using the CI Detail page. You can assign a Standard CI to a list of CIs or using Edit in List.

To assign a family's Standard CI to a CI using the CI Detail page

1. Select a CI and click Edit.
The Update Configuration Item page appears.
2. Click the icon next to the Standard CI field.
The CI List displays.
3. Select a CI to serve as the Standard CI. Click Save.
The Standard CI is assigned.

To assign a Standard CIs to a list of CIs, using edit in list

1. From the Scoreboard or Administration tab specify the search filter to show a single CI or multiple CI's in the detail list and click Search.
The CI's matching the search criteria are listed.
2. Click the Edit In List button at the top right of the list results.
The CI Edit in List controls is displayed.
3. Select a row containing CI to assign the Standard CI too.
The row is shown in highlighted color.
4. Click the icon next to the Standard CI field.
The CI List displays.
5. Select a CI to serve as the Standard CI. Click Save to update the single CI or Change All to update all CIs in the list.
The Standard CI is assigned to the single CI (Save) or all CIs listed (Change All)

Note: Any CI can act as a Standard CI, but the following cautions apply:

- A Standard CI should follow a naming convention so that it is not confused with regular CIs.
- A Standard CI must only act as a baseline for CIs in the same family as the Standard CI.
- A Standard CI should *not* be assigned values for any attributes used in reconciliation; for example, MAC Address, Serial Number, Alt CI ID, DNS, and so on.
- Because standard CIs are indistinguishable from regular CIs, they appear in your Scoreboard and count in the total number of CIs.

It is best if a Standard CI does not represent any actual instance of a business object.

Assign a Standard CI to a CI Using GRLoader

To assign a Standard CI to a CI, include the <standard_ci> tag in the description of a CI.

Note: For more information, see the *CA CMDB Technical Reference Guide*.

Example: Assign a Standard CI to a CI

If you have a standard workstation configuration stored in a CI named **standard workstation configuration**, you can assign that Standard CI to Joe's Workstation by using GRLoader to load the following XML:

```
<ci>
  <name>Joe's Workstation</name>
  <class>Server</class>
  <standard_ci>standard workstation configuration</standard_ci>
</ci>
```

Use the Basic View

You can select snapshots in either Basic or Advanced view; the same kinds of versioning data can be displayed in either mode. The Basic view lets you easily view the state of a CI.

To use the Basic view to view the state of a CI

1. Open the CI in the user interface and click the Versioning tab.
Existing snapshots are listed on the left side of the page.
2. Select a snapshot type from the Snapshot type drop-down list:

Note: If you select a Standard CI for the focal CI, it displays at the top of the snapshot or milestone list and is indicated by a "Standard CI" indicator at the right of the CI name. If you click a Standard CI, the attributes for the Standard CI appear and not the focal CI.

Snapshot

Lists all CI snapshots identified based on date/time and the Standard CI. This is the default option.

Milestone

Lists all user-defined Milestones and the Standard CI.

The state of the CI displays in the right pane.

The informational text area at the bottom displays information about the current item the mouse pointer is hovering over. Hovering over a snapshot, milestone, or standard CI shows information about the date and time it was created and the type of snapshot.

Note: You can select two or more snapshots, milestones, or Standard CIs to compare them.

Use the Advanced View

You can select snapshots in either Basic or Advanced view; the same kinds of versioning data can be displayed in either mode. Advanced view lets you view snapshots based on attribute type, value, and time stamp. Advanced view also lets you make any-to-any comparisons with attributes, milestones, or a Standard CI. A tree hierarchy shows a folder for each attribute of the CI, and each attribute folder contains a history of the values of the attribute. The hierarchy is organized as follows:

Root

Attribute name

Attribute value1

Attribute value2

This hierarchy lets you view history of unique values for any particular attribute at a glance.

To use Advanced view

1. Open the CI in the user interface and click the Versioning tab.
Existing snapshots are listed on the left side of the page.
2. Click Advanced.
Advanced Selection shows a folder hierarchy.
3. Navigate the folder hierarchy, and click the folder that includes the information you want to view:

Date

Lists Snapshots based on date/time, which is identical to the Basic view. The snapshots can be subdivided further based on year/month if there are 30 or more snapshots for the CI. The Standard CI is also listed in this folder if one has been assigned to the focal CI.

Milestone

Lists all user defined Milestones, which is identical to the Basic view. The Standard CI is also listed in this folder if one has been assigned to the focal CI.

Standard CI attributes are displayed as special tree leaf nodes with a "Standard CI". Standard CI values only appear for attributes for which there is a standard CI value; they do not appear for attributes where the value has not been set.

Relationship

Contains all the relationships (past and present) for the CI. The folder hierarchy conveys the following information:

```
Relationship
  Relationship Type
    Partner CI
      Status and Date
```

Relationship Type specifies the kind of relationship, such as “is contained by”, “hosts” or “communicates with”.

Partner CI is name of the CI associated with the relationship.

Status and Date specify the Status of the relationship at the specified Date and Time. Status values include the following ones:

- Relationship Created—The state of the CI when the Relationship was created.
- Relationship Terminated—This CI is no longer involved in the relationship. The relationship still exists at the partner end of the relationship, but the focal CI is not involved.
- Relationship Deleted—The relationship was marked as deleted.
- Relationship Changed—The relationship was reactivated from deleted state.
- New Partner and Type—The partner end of the relationship was assigned to a new CI, and the relationship type was changed at the same time.
- New Relationship Type—The relationship type between two CIs has changed.
- Partner CI Assigned—The partner end of the relationship has changed.

4. Click an attribute value.

The complete state of the CI at the time that the attribute value was set displays.

Example: Use Advanced View to Show Disk Space Attributes

In the following example, Advanced Selection shows that disk space increased from 10 to 20 to 100 GB.

Root
Disk space
10 GB
20 GB
100 GB

Click any value to see when the change occurred, the state of the other attributes when the change occurred, and the change request number that was open when the disk space was changed.

View Snapshot Details

You can view snapshot details in the Basic or Advanced view. For example, the basic view displays a snapshot of an unscheduled change to a CI.

To view snapshot details

1. Open the object in the user interface and click the Versioning tab.

A list of existing snapshots appears on the left side of the page.

2. Click a snapshot.

Details are displayed in the right pane of the Basic and Advanced views. When you select only single item, the right pane shows information about the selected snapshot, milestone (only for CIs), or standard.

The displayed data includes the following details and indicators:

Hide Empty Values

Permits filtering of blank data fields. When unchecked, all CI attributes are displayed.

Bolded Value field text

Indicates that an attribute or relationship has change since the last snapshot was taken. When viewing details for a Standard CI, all values are bold.

Value

Shows the attribute's previous value and the time of the last change when you hover the cursor over the value. The information appears in text area at the bottom of the Versioning tab.

“(blank)” Value

Indicates whether a previous value was cleared.

Relationship Category

Shows information about the relationship, including the type and partner CI information.

MDR Launch in Context and Source Identification

Provides launch-in-context to a provider MDR from the CI detail entry.

Note: CIs created in CA SDM or previous versions of CA CMDB can lack MDR and Changed-By information. In addition, CA Cohesion ACM provides MDR launch-in-context for most but not all attributes or relationships.

Tracks attribute changes back to the source MDR.

Detects when a CI attribute is updated by more than one MDR. This situation occurs when multiple MDRs contribute data independently to a CI definition.

Identifies which MDR acts as the authoritative source.

View the State of a CI on a Specific Date (Basic View)

You can use the Basic view to see the state of a CI on a specific date.

To view the state of a CI on a specific date

1. Select a CI and navigate to its CI Detail page.
2. Click the Versioning tab.

The Basic view of the CI displays.

3. Select the date you want.

The right side of the display shows the state of the CI on the selected date.

View the State of a CI on a Specific Date (Advanced View)

You can use Advanced Selection to view the state of a CI on a specific date.

To view the state of a CI on a specific date

1. Select a CI and navigate to its CI Detail page.
2. Click the Versioning tab, and click Advanced.

The Advanced Selection tree displays.

3. Expand the Date folder.
4. Select the date on which you want to view the CI.

The state of the CI on that date displays.

Launch the MDR That Set an Attribute

If the history for a CI references an MDR, you can launch an MDR to view CI details.

To launch an MDR to view CI details

1. Select a CI and navigate to its CI Detail page.
2. Click the Versioning tab.
3. Click Show Log, or select the CI snapshot or milestone you want to view.
4. Select the attribute row you want to investigate.
5. If the CI history references an MDR, click the MDR link.

The provider MDR shows additional CI details.

Print a Snapshot

You can print a snapshot as of a selected date.

Follow these steps:

1. Select the object and click the Versioning tab.
2. Select Basic view.
3. Select Snapshot from the Snapshot type drop-down list.
4. Select the date in the date list on the left.
5. Click Print.

A printer friendly window appears and displays the Versioning report for the object.

6. Use the web browser window to click File, Print to print the report.

The formatted text is sent to the defined printer.

Print a CI Milestone

You can print a milestone of a CI as of a selected date.

To print a milestone of a CI as of a selected date

1. Select a CI and click its Versioning tab.
2. Select the milestone you want from either the Basic or Advanced view.
3. Click Print.

A printer friendly window appears and displays the Versioning report for the CI.

4. Use the web browser window to click File, Print to print the report.

The formatted text is sent to the defined printer.

Export Data

The Versioning Export to CSV page lets you export the snapshot and log information in a Comma Separated Value (CSV) format. Data displays in an Export page. The formatted data corresponds to the view you obtained, with rows shown on separate lines with comma-separated column values enclosed in double quotation marks. The filtering and data comparison you select on the Versioning tab is what is formatted for export.

To export CI data

1. Select an object and click the Versioning tab.
2. Click the view you want of the CI.

For example, you click Show Log, a snapshot comparison, milestone (CIs only), and so on.
3. Click Export.

The Export Log to CSV for CI page appears.
4. Use the Select All action (from either context menu or keyboard shortcut).
5. Use cut and paste to transfer the data from the formatted page to a CSV file or third-party application. For example, you can Export to an Excel spreadsheet.

The data is exported.

CI Family Changes and Snapshots

Family changes can affect a snapshot of a CI.

A snapshot includes the following different kinds of attributes:

- Common attributes
- Family-specific attributes

When you change the family of a CI, the following snapshot changes occurs:

- Family-specific attributes associated with the CI change.
- Snapshot details reflect the family at the time of the snapshot.

The family of a CI determines the attributes of the CI. If you change a *family-specific* attribute and then change the family of the CI, the result is a CI that no longer possesses the *family-specific* attribute you changed. Changing the family of a CI has no impact on its *common* attributes.

Example: Change the Family Associated with a CI

This example shows how CI family changes affect a snapshot of a CI.

Time	Status/Changes
0	The CI is in Family1.
1	One Family1-specific attribute value is changed.
2	Several common attribute values are changed.
3	The CI's family is changed from Family1 to Family2. When this occurs, all Family1-specific attributes become unavailable to this CI.
4	One Family2-specific attribute value is changed.
5	Several common attributes are changed.
6	The CI family is changed back to Family1. When this occurs, all Family2-specific attribute values become unavailable to this CI, but previous Family1-specific attribute values are restored.
7	Several common attributes are changed.
8	One Family1-specific attribute is changed.

In the example, a snapshot at Time=7 does not contain information from the changes made at Times 3 or 4. Those changes represent changes made to Family2-specific attributes.

Compare Snapshots (Basic View)

You can compare two or more snapshots, milestones, or Standard CIs.

To compare snapshots, milestones, or Standard CIs

1. Select a CI and navigate to its CI Detail page.
2. Click the Versioning tab.
3. Select a Milestone or Snapshot from the Snapshot type drop-down list.
Milestones or snapshots are listed.

4. Select two or more snapshots, milestones, or Standard CIs.

The comparison view displays the following results:

- A column is displays how the snapshots or milestones differ.
- Each row displays the attribute values for the selected snapshots.
- Only attributes with differences are displayed. Date is always displayed unless snapshots are identical.
- Standard CI comparisons are not time-specific, so their Date fields are set to say **Standard CI**.
- Comparison results can be printed or exported to a CSV file.

Compare Snapshots (Advanced View)

You can compare two or more snapshots, milestones, or Standard CIs.

To compare snapshots, milestones, or Standard CIs in Advanced view

1. Select a CI and navigate to its CI Detail page.
2. Click the Versioning tab, and click Advanced.

Advanced Selection displays the CI.

3. Navigate the folder hierarchy to the value you want to use for the snapshot comparison.
4. Select a value for each snapshot comparison.

The comparison for the CI displays; for example, snapshot comparison based on the folder values selected.

The comparison view displays the following results:

- A column is displays how the snapshots or milestones differ.
- Each row displays the attribute values for the selected snapshots.
- Only attributes with differences are displayed. Date is always displayed unless snapshots are identical.
- Standard CI comparisons are not time-specific, so their Date fields are set to say **Standard CI**.
- Comparison results can be printed or exported to a CSV file.

Note: In Advanced view, two CIs can have a relationship with no relationship type assigned (for example, CIs created by CA NSM). Such relationship nodes are identified as (blank).

CA SDM Change Management

Versioning data is automatically generated for each CI that is associated with a CA Service Desk change ticket. As the change ticket moves from open to active to resolved to closed, a snapshot is taken for each of the associated CIs. Versioning lets you do the following:

- Compare the state of the CI to its Standard CI at any point in a change management life cycle.
- Help ensure that the CI is in compliance and help identify those attributes which need remediation.
- Perform comparisons between states of the CI at any point in life cycle of the change order, but also at any date or time in-between.
- View progress at each stage of the change, including any milestones.

The integration between CA Service Desk and CA CMDB is illustrated by demonstrating how you can audit changes performed during the service management life cycle of a CI. In this example, a change request is used to upgrade components of a hardware server.

To create a change order and associate it with a CI

1. Create a CA Service Desk change ticket by clicking File, New Change Order.
2. Enter the change request information and order summary, for example: Upgrade hard drive to 500 GB.
3. Click the Configuration Items tab, and then click Update CIs.
4. Specify CI search criteria (for example, the name of the hardware server) and click Search.
5. In the Affected Configuration Items Update form select the CIs associated with the change ticket (for example, the hardware server from Step 4), and add them it to the Affected Configuration Items list. Click OK.
6. Save the change request.

To perform the changes

1. Make the changes for the CI. For example, perform the installation of the hard drive on the physical computer, and then update the Disk Capacity attribute for the hardware server CI.
2. Close the change request.
3. Ensure that the change was completed: View the CI, click the Versioning tab, and compare snapshots of the CI before and after the change.

A comparison of the state of the CI between the times when the change order was opened and closed shows the hardware server hard drive size was 100 GB before the change, 500 GB after the change was completed, and the date and times that the changes occurred. Any other attributes that were modified between the open and close times also display.

View CI Attributes in Other CA Products

You can use the Common Asset Viewer page to view configuration item attributes in other CA products.

To view configuration item attributes in other CA products

1. Locate and open the configuration item.
2. Click Asset Viewer.
The Common Asset Viewer page appears.
3. Click the links on the Owned Resources tab to locate attribute information from other CA products.
4. Click Close Window when you are done.

You have viewed configuration item attributes in other CA products.

Using the CMDBf Viewer

CA SDM provides the CMDBf Viewer to display the results of CI federation across MDRs. From a CI Detail page (or the CI right-click menu on the CI List), click CMDBf Viewer to see CI attributes of federated CMDBs and MDRs in parallel. On the Federated View page, you can click Retrieve to update the information from any of the federated MDRs. For better readability, CA CMDB metadata files can reconcile MDR attribute names and CA CMDB attribute names.

Note: This feature requires MDRs that support Query. You configure the MDR CMDBf Endpoints to display their results on Federated View. For more information, see the *Implementation Guide*.

If the CI does not have any federated data, the viewer displays only CA CMDB attributes.

CMDB Visualizer

CA SDM lets you align your IT components (*configuration items*, or CIs) with your business services. CA CMDB defines *relationships* among CIs, as when a group of CIs work to provide a business service. CMDB Visualizer lets you see the entire picture of your CI relationships, and provides you with functions to manage the relationships. Working from a *focus CI* (any CI of interest), you can use the Visualizer to display up to nine levels of related CIs.

CA uses a provider/dependent model to define relationships among CIs. All CIs that contribute to a business service are *providers* to that business service (the *dependent*). In much the same way, providers can also be *dependents* that rely on other CIs. You can use Visualizer to perform the following provider/dependent analyses:

Browse

Displays unfiltered view of all CIs.

Impact Analysis

Starts with a focal CI (provider) and searches for its dependents.

Root Cause

Starts with a business service (dependent) and view all the CIs that are providers to that service.

Cause and Effect CIs

Combines impact analysis and root cause in one search.

Trace Relation

Displays all possible relationships based on levels. If you select only one CI, this filter displays the Browse view.

Shortest Path

Displays the shortest chain of relationships based on levels.

CMDB Visualizer lets you do the following actions:

- Visualize multiple levels of CIs from a configurable graphical view
- Monitor or cancel rendering progress
- Search using flexible criteria
- Filter based on CI families, relationship types, and other attributes
- Display CI relationships
- Trace a relationship between two CIs
- Visualize a dependency chain
- Invoke CA CMDB directly from Visualizer
- Display CI attributes and properties
- Save graph metadata
- Print the graph layout
- Find a specific CI on a displayed graph
- Create a CI (depending on role)
- Create new CI relationships (depending on role)
- Use the Scratchpad to store key CIs
- Display CI status
- Hide or reveal a CI in the Visualizer layout
- Role-based data security
- Launch external MDRs
- Obtain online help for Visualizer features

More information:

[Perform Root Cause Analysis](#) (see page 492)

[Visualizer Administration](#) (see page 492)

Perform Root Cause Analysis

Using Visualizer with the CMDBf Viewer, you can perform a root cause analysis.

To conduct a root cause analysis

1. In Visualizer, locate a Service CI in a particular problem condition (for example, slow or unavailable).
2. Right-click the CI and select Make Focal CI.
3. Select a Root Cause filter using the following criteria:
 - Class Type: n/a
 - Dependent to Provider Relationship: All relationships to be displayed for the root-cause analysis.
4. Click View.

In the resulting graph, all CIs that are related to the focal CI are displayed as specified in the filters. All paths between CIs include intermediary CIs to the default level.

5. Navigate to the CIs and inspect them for incidents, problems, or recent changes as candidates for the root cause of the focal CI condition.
6. Launch CA CMDB in context for a particular CI.
7. Click CMDBf Viewer.

The Federated View is displayed with the list of MDR providers for the CI.

8. Click Retrieve to obtain the latest MDR attributes.

The MDR attributes are updated.

Visualizer Administration

The Visualizer administration interface can be used to edit CMDB Visualizer settings. These functions are only available to roles with administrator privileges.

The Visualizer administration interface provides the following tabs:

Visualizer Configuration Tab

Permits setting of server and CI display information.

Relationship Style Tab

Defines relationship graphical characteristics.

Note: This page is disabled on a Visualizer secondary server.

CI Status Tab

Defines CI graphical characteristics.

Note: This page is disabled on a Visualizer secondary server.

Filters Tab

Creates, edits, and deletes filters for CI analysis.

Note: This page is disabled on a Visualizer secondary server.

Icon Configuration Tab

Maps a CI family to its respective icon image when CMDB Visualizer has been upgraded from r11.2 to Release 12.7.

Note: For more information about Visualizer definitions, see the CMDB Visualizer online help.

Add a Discovered Asset

You can change non-owned assets in the Management Database (MDB) into owned configuration items.

To add a discovered asset

1. Click the Discovered Assets button on the Configuration Item List page.
The Discovered Asset Search page appears.
2. Click Search to display the Discovered Asset List.
3. Select the asset from the list that you want to add as a configuration item, and then right-click the asset and select Create New Configuration Item from the pop-up menu that appears.

The Create New Configuration Item page appears with information about the item populating some of the fields.

4. Complete any remaining fields that apply to the new configuration item and click Continue.

Note: Only name and class are required values for creating a configuration item.

5. Enter data as required in the appropriate fields on the Attributes tab.

The family of the class that you selected for the configuration item determines the attributes that appear on the tab. The information you enter here is determined by your business processes and the information you want to store and view for a configuration item.

6. Click Save.

The discovered asset is added.

Asset and CI Flags

Two common attribute flags can categorize further the type of CI/asset for filtering purposes and control what entities are visible in CA SDM or other products such as CA APM.

The flags are as follows:

CI? (YES/NO)

Identifies configuration items.

Asset? (YES/NO)

Identifies assets.

By default, a CI created by CA CMDB is flagged as a CI and not as an Asset. You can override this if necessary. A CI can also be an Asset. CA CMDB does not allow the Asset flag to be changed once it has been set to Yes. The Asset flag is typically set to Yes when an asset has been created by CA APM.

These flags appear on all CI detail forms and in the CI search facility. The values can be updated in the CI detail form, GRLoader, and CA CMDB web services. The Edit in List feature also supports updating the new flags in multiple records.

The object attribute name for the CI flag is **is_ci**. The object attribute name for the Asset flag is **is_asset**. These names are SREL attributes that reference bool. The attribute names can be used as GRLoader XML tags. For example, to change the default value of the Asset flag when creating a CI with GRLoader, add the following XML to define the new CI:

```
<is_asset>YES</is_asset>
```

CI Reconciliation

Reconciliation helps ensure that updates from multiple data sources that refer to the same business object only update a single CI, even if they possess different identifying information. *Ambiguity* represents the possibility that a CI is not unique. CIs are *ambiguous* when they represent the same real business object. CI transactions are ambiguous when they have more than one possible target CI. Ambiguous CIs can lead to incorrect data in the CMDB, which negates the value of the CMDB and can lead to incorrect business actions.

Automatic CI reconciliation is a key CA CMDB advantage that saves significant time compared with manual data maintenance. The process of reconciling CIs uses several specific identifying attributes. However, automatic reconciliation can result in the following:

- False positive matches
Existing CIs are updated instead of creating a CI.
- False negative matches
New CIs are created instead of updating an existing CI. The set of CIs with similar identifying attributes are ambiguous because they resemble the same real business object with similar identifying attributes.

CA Service Desk Manager supports the following reconciliation approaches:

MDR-Based Reconciliation (Passive)

Allow the CMDB to reconcile any ambiguous data based on the MDR-Based Reconciliation process.

Identify and Resolve ambiguous CIs (Reactive)

Identify and resolve ambiguous CIs through identification and management of existing CIs in the CMDB.

Review and modify inbound data using a transaction work area (Proactive)

Review and modify inbound data before loading into the CMDB using a transaction work area (TWA).

More information:

[MDR-Based Reconciliation](#) (see page 496)

[How to Identify and Resolve Ambiguous CIs](#) (see page 498)

[Review and Modify Inbound Data Using Transaction Work Area \(TWA\)](#) (see page 510)

MDR-Based Reconciliation

MDR-based reconciliation is performed at the management data repository, to reduce further the occurrence of multiple CIs in the MDB that refer to the same object in the physical world.

MDR-based reconciliation treats the MDR as a trusted source that always uses the same federated asset ID when it communicates information about a single CI. All updates from a given MDR to a given federated asset ID always update the same CI, even when identifying attributes are changed.

MDR-based reconciliation, reconciliation management, and the Transaction Work Area (TWA) described in these sections enable you to take control of the reconciliation process. However, to use reconciliation management and the TWA successfully, first understand how CA SDM uses reconciliation attributes.

Important: If you reinstall or reinitialize any external data provider (for example, CA Cohesion ACM), inactivate and reactivate its MDR definition in the CMDb. If the MDR reuses its federated asset IDs, inadvertent CI data overlay can occur.

More information:

[How MDR Reconciliation Matches CIs](#) (see page 496)

[How MDR Reconciliation Identifies CIs](#) (see page 497)

How MDR Reconciliation Matches CIs

MDR-based reconciliation uses the following process to identify the appropriate matching CI:

1. If the transaction specifies an ID, the CI is identified and reconciliation is complete.
2. If the transaction does not specify an ID, CA SDM checks whether federated identification attributes are specified and match a CI. If there is a match, the transaction reconciles to the matching CI.
3. If the transaction does not specify an ID or federated identifying attributes, the CI transaction uses the [identification attributes](#) (see page 497), as listed in the following section.

How MDR Reconciliation Identifies CIs

MDR-based reconciliation uses the following precedence to identify a CI:

1. ID (if a transaction specifies an ID, reconciliation is successful)
2. Federated identification attributes
 - Federated asset ID
 - MDR name
 - MDR class
3. CI identifying attributes
 - Tenant (if multi-tenancy is installed)
 - Name
 - Serial number
 - MAC address
 - System name
 - Alternate asset ID
 - DNS name

How to Identify and Resolve Ambiguous CIs

We recommend managing ambiguity on a CI by CI basis. In each case, the Configuration Administrator must detect when ambiguity exists and determine the optimal approach. CA SDM takes a broad approach to identify and manage ambiguous CIs that are already in the CMDB.

The *ambiguity index* is an operational measure of the potential nonuniqueness of a configuration item (CI) or a CI transaction item, based on its identifying attributes. The ambiguity index measures the probability of one or more CIs representing the same real business object, or the probability that a CI transaction has more than one possible target CI.

Important! If you manually delete CIs with database queries, errors can occur for CIs ambiguity indexes. To avoid these errors, execute the `cmdb_update_ambiguity` utility and set the `-full` parameter to 1. This parameter helps ensure that you receive the accurate ambiguity indexes when you execute `cmdb_update_ambiguity.bat` or the shell script.

Use the following process to identify and resolve ambiguous CIs:

1. Identify ambiguous CIs.
 - Calculate the ambiguity index for all CIs
 - Review the list of ambiguous CIs from the Scoreboard
2. Determine if the identifying attributes for each CI are valid.
3. Resolve the ambiguous CI with one of the following actions:
 - Modify the identifying attributes
 - Exclude a CI from ambiguity management
 - Reject a CI Update by inactivating the CI
 - Supersede a CI

More information:

[CI Ambiguity Example](#) (see page 499)

[How to Identify Ambiguous CIs and Determine if their Identifying Attributes are Valid](#) (see page 500)

[How to Resolve Ambiguous CIs](#) (see page 507)

CI Ambiguity Example

Data from four different data sources is loaded into the CMDB. Each data source uses its own subset of identifying characteristics. Because of this inconsistency, more CIs exist in the CMDB than are desired.

The following are examples of CI ambiguity:

Example: Ambiguous CIs

The following four CIs reside in the CMDB:

- Name(Server1) DNS Name(dns1) Serial Number(serial1)
- Name(Server1) DNS Name(dns1)
- Name(Server2) DNS Name(dns1) Serial Number(serial1)
- Name(Server3) MAC Address(mac1)

Due to shared identifying characteristics, the two instances of Server1 and Server2 are ambiguous with each other. Server3 is not ambiguous.

Every CI has an ambiguity index associated with it. The ambiguity index is approximately the number of existing CIs that match on any of the identifying attributes. The greater the index, the more CIs that match on the identifiers, and therefore the greater the probability that CI data was entered inconsistently and that additional CIs are incorrectly created. CIs with an ambiguity index of zero are unique across all identifiers and are therefore unambiguous.

The ambiguity index of each of the previous CIs is

- First Instance for Server1: Count of other CIs matching name + dnsname + serial number = $1+2+1=4$
- Second Instance of Server1: Count of other CIs matching name + dnsname = $1+2=3$
- Server2: Count of other CIs matching name + dnsname + serial number = $0+2+1=3$
- Server3; Count of other CIs matching name + mac address = $0+0=0$

More Information:

[How To Calculate the Ambiguity Index](#) (see page 501)

[View Ambiguous CIs from the Administration Tab](#) (see page 505)

[View Ambiguous CIs from the Scoreboard](#) (see page 506)

How to Identify Ambiguous CIs and Determine if their Identifying Attributes are Valid

Investigate any CI with a nonzero ambiguity index to determine if it is unique or was inadvertently created because of incorrect reconciliation. CIs are not ambiguous by themselves, they are ambiguous with other CIs. A system can have many sets of ambiguous CIs, each set containing CIs with common values for the identifying attributes.

When you research the ambiguity of a CI, you also research the ambiguity of other CIs in the same set. When you resolve the ambiguity of a single CI, you reduce or eliminate the ambiguity of other CIs in that set.

CA SDM has reconciliation management tools which enable you to find ambiguous CIs and the set of CIs which the ambiguous CIs is a member. You can research the underlying cause of the ambiguity and resolve it.

When managing ambiguity, do the following:

1. Calculate the ambiguity index for all CIs.
2. Use the scoreboard to view the list of ambiguous CIs. The scoreboard lists all CIs which are ambiguous, in descending order of ambiguity.
3. Starting with the CI that has the highest level of ambiguity in the scoreboard,
 - a. Inspect all CIs in its ambiguity set. The reconciliation tab on CI detail form of a single CI lists all other CIs in the ambiguity set.
 - b. Determine if all CIs in the set are legitimate or if there was an error made in reconciliation. Review the identifying attributes and determine if the CIs in the set are created correctly or are created due to a false negative reconciliation match.
4. Determine which CIs in the set, if any, are false negatives.

When you have determined that a false negative reconciliation match has occurred, and additional CIs have been created, determine which CI is the valid one and which CIs were incorrectly created. Consider factors such as identifying attributes and other attributes such as last update date, related problems, issues, and other aspects of the CI.

How To Calculate the Ambiguity Index

Before you can begin managing ambiguity, update the ambiguity index for the existing CIs and CI Transactions. You update the ambiguity index for CIs and CI transactions by running the `cmdb_update_ambiguity` command.

Important! Run `cmdb_update_ambiguity` at least once to measure CI or CI transaction ambiguity. If you do not run the command, all ambiguity indexes are 0 (zero), which implies no ambiguity.

You can run the utility before and during ambiguity management. Schedule the utility to run periodically, so the ambiguity indexes reflect the current state of your system.

To use the ambiguity index

1. Determine the necessary parameters to run the utility.
2. Run the utility.
3. Start the CA SDM Web Client and navigate to the Ambiguous CI or Ambiguous CI Transaction Lists to manage the ambiguity.

More information:

[How to Identify and Resolve Ambiguous CIs](#) (see page 498)

cmdb_update_ambiguity Command

You can calculate CI and CI transaction ambiguity indexes by entering command syntax similar to the following:

```
cmdb_update_ambiguity [parameters] -m { ci | twa | all }
```

For command parameters, see [cmdb_update_ambiguity Parameters](#) (see page 503).

By default, the CI scan is done from the last scan date. If "-full 1" is specified, a complete scan is performed.

Example: Calculate CI ambiguity on a Microsoft SQL Server database

The following command calculates the ambiguity index for all existing CIs and Transactions in the TWA:

```
cmdb_update_ambiguity -m all -d MSSQL -u servicedesk -p dbpassword -s dbserver1
```

Example: Calculate CI ambiguity on an Oracle database

The following command calculates the ambiguity index for CIs and specifies database information.

```
cmdb_update_ambiguity -m ci -d Oracle -u mdbadmin -p dbpassword -s server1 -port 1521 -SID orcl
```

How To Use a Configuration File

You can specify many cmdb_update_ambiguity parameters in a configuration file. You can use the configuration file to secure parameter settings in an encrypted form using operating system tools. The valid keywords and the corresponding command-line options are listed in the parameter table.

Note: If you specify a parameter both on the command line and in a configuration file, the command-line value overrides the value in the configuration file.

Example: Use a configuration file to specify the parameters for a Microsoft SQL Server database

The following command runs the configuration file ambiguity_mssql.cfg.

```
cmdb_update_ambiguity -m all -c ambiguity_mssql.cfg
```

Configuration File Format

Configuration file options are specified as *keyword=value*. On Windows, the directory separator can be a double backslash (\\) or a single forward slash (/). The path name must not be enclosed in double quotation marks ("").

Valid keywords and the corresponding command-line options are listed in the parameter table.

Note: A hash mark (#) in column 1 starts a comment line.

Example: Microsoft SQL Server configuration settings

```
#Sample configuration file for Microsoft SQL Server
DBType=MSSQL
DBUser=servicedesk
DBPassword=dbpassword
DBHost=dbserver1
LogLocation=C:\Program Files\CA\Service Desk Manager\log
LogLevel=ERROR
SchemaName=dbo
```

Example: Oracle configuration settings

```
#Sample configuration file for Oracle
DBType=Oracle
DBUser=mdbadmin
DBPassword=dbpassword
DBHost=dbserver1
LogLocation=/tmp/ambiguity/log
LogLevel=INFO
DBPort=1521
DBSID=orcl
SchemaName=mdbadmin
```

cmdb_update_ambiguity Parameters

You can specify parameters on the command line or in a configuration file (some parameters are command-line only). On the command line, use quotation marks (") to enclose any path name with spaces; in the configuration file, do not use quotation marks. If any parameter is specified on the command line and also in the configuration file, the command-line value overrides the configuration file value.

The `cmdb_update_ambiguity` command takes the following parameters:

Option	Configuration File Keyword	Values	Notes
-m	(none)	tw, ci, all	(Required) Command line only tw = compute ambiguity on TWA only. ci = compute ambiguity for CIs only. all = compute ambiguity for CIs and TWA.
-d	DBType	MSSQL or Oracle	(Required. Windows only.) Database type. On Linux/UNIX, only Oracle is supported and this option is not needed.

-u	DBUser	<db user name>	(Required) Database user name. A user name with spaces must be enclosed in double quotation marks (for example: -u "sys as sysdba"). The quotation marks are not required in the configuration file.
-p	DBPassword	<db password>	(Required) Password for database user. A password with spaces must be enclosed in double quotation marks (for example: -p "secret code"). The quotation marks are not required in the configuration file.
-c	(none)	<configuration file>	(Optional) Command line only Full pathname of configuration file. Must be enclosed in double quotes if there is a space in the pathname.
-log	LogLocation	<directory to place log file>	(Optional) Log file directory. The Default is the NX_ROOT\log directory.
-level	LogLevel	INFO, ERROR, DEBUG	(Optional) Level of detail to write to the log file. Default value is ERROR.
-s	DBHost	<server name>	(Required) Database server host name. To use a Microsoft SQL Server named instance, specify <i>server\instance</i> on the command line, or <i>server\\instance</i> in the configuration file.
-CI	CI	<CI uuid>	(Optional) Compute ambiguity for the specified CI and all CIs that are ambiguous with it.
-full	(none)	0,1	(Optional) Command line only Optimizes the performance of the scan by only considering those CIs updated since the last time the utility was run. If set to 1, forces a full scan of all CIs in the computation of the ambiguity index. The default is 0. This parameter does not apply to the calculation of transaction ambiguity. The utility always evaluates all transactions in the TWA.

-port	DBPort	<port number>	(Required. Oracle only) Oracle port number
-SID	DBSID	<SID name>	(Required. Oracle only) Oracle SID name
-h	(none)		(Optional) Prints help usage message.
-schema	SchemaName	<db schema name>	(Optional) Default is mdbadmin for Oracle; or dbo for Microsoft SQL Server.

More information:

[cmdb_update_ambiguity Command](#) (see page 501)

View Ambiguous CIs from the Administration Tab

After you run the cmdb_update_ambiguity utility, you can view and manage Ambiguous CIs from the Administration tab.

To view ambiguous CIs, on the Administration tab, browse to CA CMDB, Reconciliation Management, Ambiguous CIs.

The Ambiguous CI List page displays the list of ambiguous CIs that have an index greater than 0 (zero).

Note: When you click Clear Filter in the Ambiguous CI List page or Ambiguous CI Transaction List page, it does not flush the Additional Search Arguments field (ambiguity > 0).

More information:

[View Ambiguous CIs from the Administration Tab](#) (see page 505)

View Ambiguous CIs from the Scoreboard

After you run the `cmdb_update_ambiguity` utility, you can view and manage Ambiguous CIs from the Scoreboard.

To view ambiguous CIs from the Scoreboard

1. If you are in the CA SDM Scoreboard, expand the CMDB folder.
2. Browse to Reconciliation Management, Ambiguous CIs and click one of the following:
 - All
 - Updated In The Last Day
 - Updated in The Last Week
 - Updated in The Last Month

The list of ambiguous CIs that have an index greater than 0 (zero) displays.

Note: When you click Clear Filter in the Ambiguous CI List page or Ambiguous CI Transaction List page, it does not flush the Additional Search Arguments field (ambiguity > 0).

More information:

[How to Identify and Resolve Ambiguous CIs](#) (see page 498)

How to Resolve Ambiguous CIs

After you identify ambiguous CIs and determine if their identifying attributes are valid, resolve the ambiguity among the CIs in the ambiguity set using one or more of the following:

Modify the identifying attributes

If you determine the identifying attributes are not complete or invalid, set the CI identifying attributes so that the CI is unique using either the Web interface, GRLoader, or CMDBf.

Note: When the MDR updates the CI, the MDR can undo the reconciliation changes that were manually performed.

Exclude a CI from ambiguity management

If you determine that the CIs identifying characteristics are correct and represents a known, valid ambiguity, you can remove the CI from the ambiguity management lists and the ambiguity calculation of other CIs, and the CI can be marked as not ambiguous (exclude ambiguity).

Reject a CI Update by inactivating the CI

When you determine that the identifying attributes of a CI are incorrect and that updates using those attributes cause data corruption, the CI can be inactivated, preventing further updates. The user or MDR generating the information receives an error and the entire transaction is rejected.

Supersede a CI

Sometimes the updates to the CMDB are beyond the control of the administrator and invalid identifying data needs to be in the system with valid nonidentifying attribute data. The incoming transaction attribute data can be transparently redirected to a superseding CI.

Note: The transparent redirection of attribute data from one CI to another can cause confusion because transaction data may not be stored in the same CI as the transactions identifying attributes would lead you to believe. Use this method with discretion, when the previous methods cannot be used.

How to Exclude CIs from Ambiguity Management

Sometimes you recognize that although a particular CI receives an ambiguity index greater than zero, it should be left as is. Any CI with its Exclude Ambiguity check box selected is not considered part of the ambiguity index or ambiguity management features.

Exclude Ambiguity for a CI

You can update the exclude ambiguity option for a CI from the web interface.

To exclude ambiguity for a CI

1. Select the CI that you want to remove from ambiguity management from the Ambiguous CI List page.
The Configuration Item Detail page appears.
2. Click Edit.
The Update Configuration Item page appears.
3. Select the Exclude Ambiguity check box on the Reconciliation tab, and click Save.
The CI is excluded from the ambiguity index calculation and other ambiguity management features.

How To Exclude Ambiguity Using GRLoader

You can update the exclude ambiguity option for a CI by using GRLoader to set the not_ambiguous flag for the CI.

not_ambiguous values are as follows:

YES (1)

Removes the CI from ambiguity management. The CI is identified as unique regardless of the identifying attributes of other CIs. The ambiguity index of the CI remains zero.

NO (0) (default)

This CI is eligible for ambiguity management. The uniqueness of the CIs identifying attributes determines the ambiguity index of this CI. The identifying attributes of this CI are considered when evaluating the ambiguity index of other CIs.

Note: For more information about GR Loader, see the *CA CMDB Technical Reference*.

How To Reject Updates

When a CI is marked as Inactive, no updates can be performed on it. Set the CI to Inactive if you want CA CMDB to reject data for a given CI (whether by name, serial number, MAC address, and so on) and you want the MDR to know of the rejection.

In this way, you can reject problematic data immediately and reflect it back to the source, where the MDR can correct the rejected input.

To set a CI to Inactive in the web interface, edit the CI, set it to Inactive, and click Save. You also can set a CI to Inactive by using GRLoader or CMDBf web services. Inactivated CIs can also be reactivated.

Supersede Ambiguous CIs

You can supersede an ambiguous CI to redirect data to a specific target CI. You can display superseded CIs, which are Inactive and all web interface updates to them are ignored.

Note: Updates sent to superseded CIs by GRLoader or CMDBf web services are redirected to the superseding CI. However, updates to identifying attributes are ignored.

To supersede an ambiguous CI

1. Select the CI that you want to be the focal (superseding) CI from the Ambiguous CI List page.
The Configuration Item Detail page appears.
2. Click Edit.
The Update Configuration Item page appears.
3. Click the Reconciliation tab.
All CIs that are ambiguous with the focal CI display.
Determine which CIs you want to supersede by the focal CI. Inspect each CI in the list by clicking each CI to launch its Configuration Item Detail page.
4. Select one or more ambiguous CIs that you want to supersede with the focal CI and click Supersede.
The focal CI supersedes the selected CIs.

View Superseded CIs from the Administration Tab

You can view superseded CIs from the Administration tab.

To view superseded CIs from the Administration tab, on the Administration tab, browse to CA CMDB, Reconciliation Management, Superseded CIs.

The Superseded CIs list appears.

View Superseded CIs from the Scoreboard

You can view superseded CIs from the Scoreboard.

To view superseded CIs from the Scoreboard

1. If you are in the CA SDM Scoreboard, expand the CMDB folder.
2. Expand the Reconciliation Management folder
3. Click the Superseded CIs node.
The Superseded CIs list appears.

Review and Modify Inbound Data Using Transaction Work Area (TWA)

You can stage CI and relationship transactions before execution, by copying data into the TWA staging area. Once in the staging area, you can manipulate CIs and relationships by using the web interface or native SQL.

You also can validate the CI transactions to prevent the creation of new CIs when you should update existing CIs. In this approach, you view each transaction and the potential CIs that it can update so that you can reconcile the transaction manually to the target CI. Likewise, relationship transactions can be validated to reference the correct CIs.

The *ambiguity index* is an operational measure of the potential nonuniqueness of a configuration item (CI) or a CI transaction item, based on its identifying attributes. The ambiguity index measures the probability of one or more CIs representing the same real business object, or the probability that a CI transaction has more than one possible target CI.

Review and modify ambiguous CI transactions using the following process:

1. Identify ambiguous CI transactions.
 - Calculate the ambiguity index for all CI transactions
 - Review the list of ambiguous CI transactions from the Scoreboard
2. Determine if the identifying attributes for each CI transaction reconcile to the intended CI.
3. Resolve the ambiguous CI transactions with one of the following actions:
 - Modify the identifying attributes
 - Specify the target CI

More information:

[CI Transaction Ambiguity Example](#) (see page 511)

[How to Identify an Ambiguous CI Transaction](#) (see page 512)

[How to Resolve Ambiguous CI Transactions](#) (see page 514)

[Manage Staged Transactions](#) (see page 514)

[Transaction Work Area](#) (see page 515)

[Populating the TWA](#) (see page 517)

[How to Use the Web Interface to Update Data in the TWA](#) (see page 528)

[Manual Reconciliation](#) (see page 530)

[How To Load Transactions into the CMDB](#) (see page 532)

[TWA Administration](#) (see page 534)

CI Transaction Ambiguity Example

Transaction data from different data sources is loaded into CA CMDB. Each data source uses its own subset of identifying characteristics and may not fully identify the target CIs for the transactions. Because of this inconsistency, more CI Transactions may exist in CA CMDB than are valid.

The following is an example of CI transaction ambiguity:

Example: Ambiguous CI Transactions

The following CI transaction resides in the TWA:

- Name(Server1) DNS Name(dns1), Serial Number(serial1)

The following CIs reside in CA CMDB:

- Name(Server2) DNS Name(dns1)
- Name(Server3) DNS Name(dns1), Serial Number(serial1)
- Name(Server4) MAC Address(mac1), Serial Number(serial1)

Due to shared identifying characteristics, Server1 transaction is ambiguous with Server2, Server3, and Server4 in CA CMDB.

Every CI Transaction has an ambiguity index associated with it. The ambiguity index is approximately the number of existing CIs that match on any of the identifying attributes, minus one, specified in the CI transaction. The greater the index, the greater the number of other CIs that match on the transaction identifiers, and therefore the greater the probability that CI data was entered inconsistently and the possibility that additional CIs are incorrectly created. CI Transactions with an ambiguity index of zero have identifying attributes that are unique across all CIs or have a target CI specified and are therefore unambiguous.

Example: Calculate the Ambiguity Index

The following CI transactions reside in the TWA:

- Name(Server1) DNS Name(dns1), Serial Number(serial1)
- Name(Server2) DNS Name(dns1), Serial Number(serial1)

The following CIs reside in CA CMDB:

- Name(Server1) DNS Name(dns1), Serial Number(serial1)
- Name(Server2) DNS Name(dns1), Serial Number(serial1), Mac Address(mac1)

The first transaction (Server1) ambiguity is 0 because there is an exact match with the Server1 CIs identifying attributes. The only possible target CI to this transaction is Server1 CI.

The second transaction (Server2) is ambiguous with Server1 CI and Server2 CI.

The ambiguity index for Server2 transaction consists of the following components:

- Number of matching CIs with Name (Server2) = 1
- Number of matching CIs with DNS Name(dns1) = 2
- Number of matching CIs with Serial Number (serial1) = 2

Based on shared CI identifying characteristics, the ambiguity index for server2 transaction is $(1-1) + (2-1) + (2-1) = 2$

How to Identify an Ambiguous CI Transaction

Review and modify transactions that have ambiguous CI targets before loading the transaction data. You can also resolve CI transactions that do not have a target CI but have identifying attributes for one or more existing CIs. This resolution step verifies that the appropriate CI is updated and prevents incorrect or inconsistent data from being created.

The Reconciliation Tab on the CI Transaction Detail page lists all CIs that are ambiguous with this focal CI Transaction. Inspect the CI attributes such as last update date, related problems, issues, and other aspects of the CI to identify the target CI that corresponds to the transaction.

Note: To view ambiguity either for CIs or CI transactions, you must calculate the ambiguity index by using the [cmdb update ambiguity utility](#) (see page 501).

For more information, see [Review and Modify Inbound Data Using Transaction Work Area \(TWA\)](#) (see page 510).

View Ambiguous Transactions from the Administration Tab

You can view and manage ambiguous transactions from the Administration tab.

To view ambiguous transactions, on the Administration tab, browse to CA CMDB, Reconciliation Management, Ambiguous CI Transactions.

The list of CI transactions that have an index greater than 0 (zero) display.

Note: When you click Clear Filter in the Ambiguous CI List page or Ambiguous CI Transaction List page, it does not flush the Additional Search Arguments field (ambiguity > 0).

View Ambiguous Transactions from the Scoreboard

You can view and manage ambiguous transactions from the Scoreboard.

To view ambiguous transactions from the Scoreboard

1. If you are in the CA SDM Scoreboard, expand the CMDB folder.
2. Browse to Reconciliation Management, Ambiguous CI Transactions and click one of the following:
 - All
 - Updated In The Last Day
 - Updated in The Last Week
 - Updated in The Last Month

The list of CI transactions that have an ambiguity index greater than 0 (zero) displays.

Note: When you click Clear Filter in the Ambiguous CI List page or Ambiguous CI Transaction List page, it does not flush the Additional Search Arguments field (ambiguity > 0).

How to Resolve Ambiguous CI Transactions

After you identify ambiguous CI transactions and determine if their identifying attributes or the target CI are valid, resolve the ambiguity among the CI transaction using one of the following:

Modify the identifying attributes

If you determine the identifying attributes are not complete or invalid, you can set the CI identifying attributes so that the CI is unique using the web interface.

Specify the Target CI

If you determine the target CI is not set or not valid, you can specify a CI for an ambiguous transaction so that its data is directed to a specific target CI regardless of the values of the identifying attributes specified in the transaction.

Specify a Target CI for an Ambiguous CI Transaction

You can specify a CI for an ambiguous transaction so that its data is directed to a specific target CI regardless of the values of the identifying attributes specified in the transaction.

To specify the target CI for an ambiguous CI transaction

1. Select a CI transaction from the Ambiguous Transaction List page.
The Configuration Item Transaction Detail page appears.
2. Click Edit.
The Update Configuration Item Transaction page appears.
3. On the Reconciliation tab, select a CI from the list to designate the CI as the Target CI for the transaction. Click Set Target and then click Save.
The selected CI becomes the target CI for the transaction.

Manage Staged Transactions

CA CMDB provides a facility where you can store data before you load it into the CA CMDB. This "staged" data is stored as transactions in the transaction work area (TWA). A staged transaction is a unit of work that creates or updates a CI or Relationship. The TWA can contain many transactions for a given CI or relationship.

You can capture data being loaded into the CMDB before that data is committed so that you can modify:

- Nonstandard data can be cleaned up and standardized.
- Incomplete data can be supplemented. For example, CI names starting with "NY" can have their location set to "New York".

- Data that does not match existing lookup tables (SRELs) can be modified.
- Transactions can be scheduled for implementation at a future time.
- Reconcile transactions to existing CIs in the CA CMDB before you load the data.
- Validate the CI and relationship transactions to prevent the creation of invalid data or new CIs when a single or existing CIs should be updated. View each transaction and the potential CIs that it can update so that you can reconcile the transaction manually to the target CI.

You can use the TWA to help you proactively manage the reconciliation process. You can configure GRLoader to load the data into the TWA, where CA CMDB lets you modify the transaction data to handle transactions that can create potentially update or reference the wrong CI.

For more information, see [Review and Modify Inbound Data Using Transaction Work Area \(TWA\)](#) (see page 510).

Transaction Work Area

CA SDM provides the transaction work area (TWA) for inspecting and modifying CI and relationship data before you load the data into the CMDB.

The TWA process works as follows:

1. [Load the data](#) (see page 517) into the TWA. The content can include CI transactions or relationship transactions. The input process does not attempt to reconcile records; it simply populates the work area with CI and relationship transaction records in one of the following ways:
 - GRLoader – Reads XML data and stores it in the TWA using the -littwa or -littwar options.
 - Native SQL – Places data into the TWA using standard SQL processing.
 - Create a CI transaction or relationship transaction using the web interface.

2. (Optional) Manually reconcile transactions to existing CIs in the CMDB before you load the data. You can also [simulate loading the data](#) (see page 522) to predetermine whether a set of transactions can create new CIs (and therefore create new ambiguities for other CIs) or relationships.

For more information, see [Review and Modify Inbound Data Using Transaction Work Area \(TWA\)](#) (see page 510).

3. Modify the data. You can modify the TWA from the following sources:

CA SDM user interface

The web-based user interface lets you view and modify transactions in the work area.

Native SQL

When performing many changes to multiple transactions, native SQL can modify the data in the TWA.

Note: Changes made using native SQL can bypass all CA SDM security.

4. GRLoader loads the data from the TWA to the CA CMDB using the -lftwa or -lftwai options. Each TWA row is treated as a separate transaction.
 1. If there is an error after a GRLoader run, the error code is populated into the transaction to indicate that it is incomplete (to facilitate future retries).
 2. All other records are identified as completed transactions.
5. Review the transaction results and correct any errors using the web interface. Repeat steps 3 and 4 as needed.

Note the following:

- If you created any custom families or attributes and want their columns to appear in the work area, modify the work area to include the customized columns. For more information, see [Extend the TWA Object](#) (see page 534).
- Use the latest version of GRLoader to take advantage of the TWA. If you are using a product that includes a previous release of GRLoader, upgrade GRLoader to the latest version.
- Changes made to transactions in the TWA are not audited.

Populating the TWA

Input data can come from a number of sources, including:

- CA products like CA Cohesion ACM, CA Wily CEM, CA Wily Introscope, and CA Spectrum that import data by using GRLoader
- Any other MDR that imports data by using GRLoader
- Another CMDB
- Microsoft Excel spreadsheets
- Database tables
- An ETL tool chosen by a vendor

See [Database Limitations](#) (see page 537) for more information.

How to Load Data into the TWA Using GRLoader

The following GRLoader options support the uses of the TWA:

-littwa

Loads XML data to TWA in initial state.

-littwar

Loads XML data to TWA in ready state.

-lftwa

Loads transactions to the CMDB from TWA.

-lftwai

Loads transactions to the CMDB from TWA and inactivates successful transactions.

GRLoader input XML documents can be loaded into the CMDB or the TWA without modification.

Loading data into the transaction work area can be accomplished by using GRLoader with the `-littwa` (load to TWA) option. In the GRLoader configuration file:

```
grloader.loadtotwa=yes
```

In TWA mode, instead of creating CIs and relationships directly, GRLoader inserts the information into the transaction work area tables. Data that has been loaded into the TWA can be edited, changed and verified. After the data modification process is complete, individual transactions can be loaded into the CMDB by using `-lftwa` or `-lftwai`.

When loading CI data into the CMDB, CIs undergo automatic reconciliation, which, if successful results in data from a single business object updating a single CI in the CMDB. When loading *CI transactions* into the TWA, no automatic reconciliation occurs. Multiple transactions with identifying attributes targeting a single CI can appear in the TWA. Existing rows in the TWA are not updated when new rows are added in load to TWA mode, even if they are an identical match to an existing CI. Even identical CI definitions can appear many times in the TWA.

GRLoader validates data values when running in load from TWA mode (-lftwa or -lftwai) to create objects in the CMDB or when running using the simulation mode (-simci or -simrel), see [How to Simulate TWA Operations](#) (see page 522) for more details. The data values are not validated when GRLoader is run in load to TWA (-lftwa or -lftwar) mode.

The -lftwai parameter (or configuration option grloader.loadtotwa.inactivate) inactivates successfully loaded TWA transactions. After a successfully loaded transaction has been marked inactive, it can be permanently purged from the TWA by using Archive/Purge.

Example: Load to CMDB

The following command loads data directly into the CMDB:

```
grloader ... -i mydata.xml -a -n
```

The XML file mydata.xml contains:

```
<GRLoader>
  <ci>
    <name>server1</name>
    <class>Server</name>
  </ci>
</GRLoader>
```

Example: Load to TWA

The following command loads the data into the TWA, so that it can be manipulated before loading into the CMDB:

```
grloader ... -i mydata.xml -lftwa
```

More information:

[XML Input](#) (see page 519)

[lookup](#) (see page 519)

[Date Format](#) (see page 520)

[EMPTY](#) (see page 521)

XML Input

GRLoader uses XML input data. When loading to the TWA, the following XML columns are mapped to prevent conflicts when the column names overlap with the standard column names in the database.

FROM XML Name	TO Database Column
tenant	tgt_tenant
id	tgt_id
delete_flag	tgt_delete_flag

For relationships, the following mapping applies:

FROM XML Name	TO Database Column
name	provider_name or dependent_name
dns_name	provider_dns_name or dependent_dns_name
delete_flag	tgt_delete_flag

Mappings are reversed when loading from the TWA database tables.

lookup

When you specify data for lookup (SREL) attributes in the TWA, they must have the same format as if specified in native XML for GRLoader. Lookup attributes accept only a specific set of values that must be defined in related tables in CA SDM. These attributes also can have additional restrictions and exceptions that must be met for the assignment to occur. The specified values are the same whether specified using XML, SQL, or the TWA web interface.

To determine whether an attribute is an SREL, refer to the following *CA CMDB Technical Reference* sections:

- Families and Classes (SREL attributes are identified)
- Contact and Other Lookup Fields
- Fields Validated Against Data in Existing Tables (SREL)

Example: lookup XML

In the following example, the alternate SREL column is set by specifying the lookup parameter:

```
<ci>
<name>server1</name>
<owner lookup="userid">mckpe99</owner>
</ci>
```

In the TWA, the same thing can be accomplished by suffixing the alternate SREL column to the data value, surrounded by delimiters (see the following `grloader.workarea.delimiters`). The equivalent transaction is represented in the transaction work area as:

ID	Name	Owner
100	server1	mckpe99 {userid}

You also can set the delimiter character used above in the `grloader` configuration file:

```
grloader.workarea.delimiters=xy
```

x and y are different characters that typically do not appear in the work area.

Date Format

GRLoader supports the following date format options in XML:

- `datefmt=UTC`
- `datefmt=localtime` (the default)

Examples: datefmt XML

```
<ci>
<name>server1</name>
<purchase_date datefmt="UTC">1241197235</purchase_date>
</ci>
```

```
<ci>
<name>server3</name>
<purchase_date>2009/05/01</purchase_date>
</ci>
```


The equivalent transaction is represented in the transaction work area as:

ID	Name	purchase_date
101	server1	1241197235
102	server2	2009/05/01

Note: If a date contains a special character such as a slash (/), then the format is assumed to be "localtime". If a date contains no special characters, the date is assumed to be UTC.

EMPTY

GRLoader supports the `update_if_null` option in the XML which clears a field in the CMDB. The following example clears the owner field for server1. Without that attribute, the owner field is not affected. When using the TWA, you can use the keyword `EMPTY` instead.

Example: `update_if_null` XML

```
<ci>
<name>server1</name>
<owner update_if_null="yes"></owner>
</ci>
```

In the TWA, the database value is cleared by specifying the keyword `EMPTY` as the string value. The equivalent transaction in the work area is:

ID	Name	Owner
102	server1	EMPTY

The keyword value can be set by using the `grloader.emptyvalue` configuration option:

```
grloader.emptyvalue=xxxx
```

`xxxx` represents any string that typically does not appear in the work area data.

How To Simulate TWA Operations

You can predetermine whether a set of transactions can create new CIs or relationships (and therefore create new ambiguities for other CIs) by using the following options:

–simci

Simulates processing CI transactions only. Can be used to determine whether transactions create or update CIs. When the -simci option is used, GRLoader performs data validation.

–simrel

Simulates processing relationship transactions only. Can be used to determine whether relationship transactions create or update relationships. The –simrel option checks relationships for the existence of the provider and dependent CIs, validates relationship types, and so on.

The output from simulation mode is directed to the TWA or to the _err.xml file. In normal load mode, the _err.xml file contains the CI input and a comment indicating whether the CI was inserted or updated. When simulation is used, the GRLoader message on the CI Transaction List indicates whether the CI or relationship was inserted or updated, with other relevant error messages. The transaction state remains unchanged.

Simulation can also be enabled in a configuration file by using the **grloader.simulateloadci** and **grloader.simulateloadrelation** options.

Note: If the GRLoader input creates CIs and relationships at the same time, the –simrel option can only process actual CIs, not CIs that are scheduled to be created. Because of this limitation, -simci and –simrel are mutually exclusive.

How to Use SQL to Update Data into the TWA

The TWA is stored in the MDB and you can update it directly with SQL queries.

More information:

[TWA Tables](#) (see page 523)

[TWA SQL Column Names](#) (see page 523)

[Insert New Records in the TWA](#) (see page 526)

[How to Use the ODBC Driver](#) (see page 526)

[CI Identifiers in the TWA](#) (see page 527)

[How to Set Transaction Status](#) (see page 527)

[How to Share Tables with CA SDM](#) (see page 528)

TWA Tables

Transaction Work Area (TWA) tables include:

ci_twa_ci

A single table that contains all attributes across all CI families. The table stores data in a de-normalized form to enable customers and services to understand and manipulate the content.

ci_twa_relation

Complements the ci_twa_ci table. This table contains CI relationship information.

TWA column names correspond to the attributes that are used with GRLoader, with some exceptions.

Note: For CI attribute information, see the *CA CMDB Technical Reference Guide*.

TWA SQL Column Names

Column names in the ci_twa_ci and ci_twa_relation tables match the CI and Relationship attribute names with some exceptions.

TWA fields common to ci_twa_ci and ci_twa_relation

apply_after_dt (Apply After Date)

GRLoader executes the transaction only if the current date is after the Apply After Date. Measured in number of seconds since the epoch. If the value is zero (0), this field is ignored.

tran_chg_ref_num (Change Order)

Specifies the change order identifier associated with this transaction. The GRLoader -chg option uses this attribute to select only transactions with the specific change order number.

Note: For more information about using the -chg option option, see [Filter By Change Order Number](#) (see page 534).

tran_status (Transaction Status)

Specifies the status of the transaction must be one of the following values (only transactions with tran_status=1 are executed).

Value	Short Description	Long Description
0	Initial	Default value. Requires intervention to move to "Ready" for GRLoader to act upon this transaction
1	Ready	The transaction is ready to be loaded into either the appropriate CI or relationship tables

Value	Short Description	Long Description
2	Successful	Transaction processed successfully
3	Warning	Warning detected during load into ca_owned_resource/bmhier
4	Error	Error detected during load into ca_owned_resource/bmhier
40000+	For customer use	

tran_message (Message)

GRLoader message showing results of the load or simulation.

tran_dt (Transaction Date)

Prevents accidental overlaying of current information by old transaction data. The transaction date is compared with the last modification date of the CI and GRLoader rejects the transaction if the CI is more recent.

If GRLoader rejects a transaction as older than the target CI, verify that the transaction is still applicable and set the transaction date manually to be more current and run GRLoader again.

tgt_delete_flag (Active?)

Inactivates the target CI or relationship. Set this value to 1 (one). Not the same as setting delete_flag to 1, which indicates that the transaction should be deleted.

ci_twa_ci columns

In addition to the column names below, you can also specify CI attribute names as column names in your SQL statements.

tgt_id (Target CI)

Specifies the UUID of the target CI. Used when performing manual reconciliation.

superseded_by (Superseded By)

Specifies the UUID of the superseded CI.

tgt_tenant (Tenant)

Specifies the tenant name assigned to the target CI. Tenant only can be assigned at CI creation time if the GRLoader user has proper authorization. Applies only when multi-tenancy is enabled.

ci_twa_relation columns

In addition to the column names below, you can also specify relationship attribute names as column names in your SQL statements.

Use the following database column names when using SQL to define relationship transactions:

provider_name (Provider Name)

provider_mac_address (Provider MAC Address)

provider_serial_number (Provider Serial Number)

provider_system_name (Provider Host Name)

provider_asset_num (Provider Asset Number)

provider_dns_name (Provider DNS Name)

provider_tenant (Provider Tenant)

AND

dependent_name (Dependent Name)

dependent_mac_address (Dependent MAC Address)

dependent_serial_number (Dependent Serial Number)

dependent_system_name (Dependent Host Name)

dependent_asset_num (Dependent Asset Number)

dependent_dns_name (Dependent DNS Name)

dependent_tenant (Dependent Tenant)

If provider or dependent CI UUIDs are known, you can use the following fields:

provider_id specifies the UUID of the provider CI

dependent_id specifies the UUID of the dependent CI

Note: provider_tenant and dependent_tenant only apply when multi-tenancy is enabled.

Insert New Records in the TWA

You can insert a new record into the ci_twa_ci and ci_twa_relation tables from outside of CA SDM.

To insert records into the ci_twa_ci and ci_twa_relation tables, specify 0 in the ID column as:

```
insert into ci_twa_ci values(id, name) values(0,'test-ci');
insert into ci_twa_relation values(id, type) values(0,'test-relation-type');
```

Caution: If multi-tenancy is enabled, provide tenant, tgt_tenant, provider_tenant or dependent_tenant attributes as appropriate. If you create TWA transactions without tenant specified, the Public tenant is assumed.

How to Use the ODBC Driver

You can use SQL instead of GRLoader to load the TWA. We recommend that you use the ODBC driver that is provided with CA SDM, instead of native SQL services. Using the ODBC driver provides the following benefits:

- Performance
- Security
- Database independence
- Data integrity in an environment where there can be database updates coming from the CA SDM server at the same time that the bulk editing tool is running.

Cautions

- Only CA SDM ODBC drivers honor security features such as multi-tenancy, data partitioning, and role access. SQL that is executed using the native DBMS drivers does not enforce security.
- Make frequent data backups when performing updates using native SQL.

CI Identifiers in the TWA

In the `ci_twa_ci` table, three identifiers are associated with a CI transaction: `ID`, `tgt_id`, and `superseded_by`. These attributes are:

- `ID` is the transaction sequence number. **Do not modify this field.** IDs 1-2,000,000,000 are reserved for use by systems that modify the TWA outside the CA SDM server. IDs from 2,000,000,001 to approximately 4 billion are reserved for use by CA SDM. When the system runs out of IDs, reinitialize the table.
- `tgt_id`, `provider_id`, or `dependent_id` are the UUID of the target CI. Set this field whenever you perform manual reconciliation.
- `superseded_by` is the UUID of the superseding CI. Set this field whenever you manually designate the superseding CI for the transaction.

How to Set Transaction Status

You can set the transaction status (`tran_status`) of each transaction in the TWA. For GRLoader to process a transaction, it must be set to Ready (`tran_status=1`).

Value	Short Description	Long Description
0	Initial	Default value. Requires intervention to move to "Ready" in order for GRLoader to act upon this transaction
1	Ready	The transaction is ready to be loaded into either the appropriate CI or relationship tables
2	Successful	Transaction processed successfully
3	Warning	Warning detected during load into <code>ca_owned_resource/bmhier</code>
4	Error	Error detected during load into <code>ca_owned_resource/bmhier</code>
40000+	For customer use	

If GRLoader -lftwa is run two or more times, the status is updated after the first run and all subsequent runs do not find any transactions. Reset the status of the ready transactions before every GRLoader -lftwa execution.

For information about how to represent XML data in the TWA tables, see the XML examples in previous sections.

How to Share Tables with CA SDM

SQL Server Only

When using SQL to update the ci_twa_ci and ci_twa_relation tables, set the last_mod_dt column manually, as follows:

```
SET last_mod_dt = DATEDIFF(s,'19700101', GETUTCDATE())
```

Example: Set last_mod_dt

The following example sets the last_mod_dt:

```
UPDATE ci_twa_ci
    SET tran_status=1,
        last_mod_dt = DATEDIFF(s,'19700101', GETUTCDATE())
    WHERE tran_status=1
```

Note: TWA table updates can take several minutes to appear in the web interface or GRLoader. For more information about update timing, see the NX_DBMONITOR_TIMER_MINUTES option.

How to Use the Web Interface to Update Data in the TWA

You can stage CI and relationship transactions before execution, by copying data into TWA staging area. Once in the staging area, you can manipulate CIs and relationships by using the web interfaces.

You also can validate the CI and relationship transactions to prevent the creation of invalid data or new CIs when a single or existing CIs should have been updated. In this approach, you view each transaction and the potential CIs that it can update so that you can reconcile the transaction manually to the target CI.

To manage the transaction work area, select the Administration tab and navigate to the CA CMDB, Transaction Work Area node. From this node, you can manage the Transaction Work Area pages.

Note: For more information, see the *Online Help*.

More information:

[Manage CI Transactions](#) (see page 529)

[Manage Ambiguous CI Transactions](#) (see page 530)

[Manual Reconciliation](#) (see page 530)

Manage CI Transactions

The CI Transaction Detail page displays all nonblank attributes in the transaction. If you clear the Hide Empty Values check box, all CI attributes are displayed. For additional attribute information, point to the attribute name.

You can do the following:

- View a transaction
- Create a transaction
- Edit the transaction
- Save or Cancel transaction changes (Edit mode)
- Reset the original transaction data (Edit mode)
- Show or hide data (Attributes tab)
- Set the target CI (Reconciliation tab)

After you save a single transaction, click the Reconciliation tab for that transaction to verify that there is no ambiguity in its intended target CI. If there is ambiguity, you can [specify a target CI](#) (see page 514). When all data entry is complete, you can [simulate loading the data](#) (see page 522) to validate the data and check reconciliation.

Important! Data that you enter into the TWA using this web interface must follow the rules described in the section "[How to load data into the TWA Using GRLoader](#) (see page 517)".

The web UI does not validate case-sensitive data in the TWA, verify that the values you enter exactly match the lookup values. Alternately, when you run GRLoader, you can specify the `-I` option to enable case insensitive lookups.

Note: Even though the `tgt_id` or `superseded_by` are stored as UUIDs in the transaction, they are displayed as CI names in the web interface.

Manage Ambiguous CI Transactions

Review and modify transactions that have ambiguous CI targets before loading the transaction data. You can also resolve CI transactions that do not have a target CI but have identifying attributes for one or more existing CIs. This resolution helps ensure that the appropriate CI is updated and prevents incorrect or inconsistent data from being created.

After you save a single transaction, click the Reconciliation tab for that transaction to verify that there is no ambiguity in its intended target CI. If there is ambiguity, you can [specify a target CI](#) (see page 514). When all data entry is complete, you can [simulate loading the data](#) (see page 522) to validate the data and check reconciliation.

For more information about reconciling the data in the TWA, see [Review and Modify Inbound Data Using Transaction Work Area \(TWA\)](#) (see page 510).

Manual Reconciliation

To reconcile a CI transaction manually, determine the target CI that matches the identifying attributes of the transaction and set the `tgt_id` of that transaction to the UUID of the target CI.

When GRLoader processes a transaction that specifies a `tgt_id`, it updates the target CI with the transaction information and registers that CI again when its identifying attributes have changed.

You can specify the target CI explicitly in the transaction. You also can use the Reconciliation tab to select a target CI.

For more information, see [Review and Modify Inbound Data Using Transaction Work Area \(TWA\)](#) (see page 510).

Specify a Target CI for an Ambiguous CI Transaction

You can specify a CI for an ambiguous transaction so that its data is directed to a specific target CI regardless of the values of the identifying attributes specified in the transaction.

To specify the target CI for an ambiguous CI transaction

1. Select a CI transaction from the Ambiguous Transaction List page.
The Configuration Item Transaction Detail page appears.
2. Click Edit.
The Update Configuration Item Transaction page appears.
3. On the Reconciliation tab, select a CI from the list to designate the CI as the Target CI for the transaction. Click Set Target and then click Save.
The selected CI becomes the target CI for the transaction.

Manage Relationship Transactions

The Relationship Transaction List displays relationship transaction data input.

Using this page, you can do the following:

- Search for relationship transactions

On the list, you can do the following:

- Create a relationship transaction
- Edit in list (except for identifying attributes)

Update a Relationship Transaction

You can view and modify data for a single relationship transaction. On the Relationship Transaction Detail page, you can do the following:

- Create a relationship transaction
- View the transaction
- Edit the transaction
- Save or Cancel transaction changes (Edit mode)
- Reset the original transaction data (Edit mode)

To update a relationship transaction

1. Click Edit.

The Update Relationship Transaction page appears.

2. Modify and complete fields.
3. Select Active.
4. Click Save.

The relationship transaction is saved.

When all data entry is complete, you can simulate loading the data to validate the data and check reconciliation.

Important! Data that you enter into the TWA using this web interface must follow the rules described in the section "[How to load data into the TWA Using GRLoader](#) (see page 517)".

The web UI does not validate case-sensitive data in the TWA, verify that the values you enter exactly match the lookup values. Alternately, when you run GRLoader, you can specify the `-I` option to enable case insensitive lookups.

Note: Even though the `provider_id` or `dependent_id` are stored as UUIDs in the transaction, they are displayed as CI names in the web interface.

How To Load Transactions into the CMDB

Instead of using XML input to create CIs and relationships, you can specify the following options to select the TWA as its input source:

`-lftwa` (load from TWA)

`-lftwai` (load from TWA and inactivate)

Note: Unlike other GRLoader modes that use an `err.xml` file, loading from the TWA returns error messages to the transaction Message field for viewing.

As transactions are processed, GRLoader sets the transaction status to indicate success or failure of each TWA transaction. If GRLoader `-lftwa` is run multiple times in succession, set the status of the transaction to Ready between runs of GRLoader, because it does not attempt to load data that is already loaded.

The transaction is executed from the TWA as if GRLoader were running using XML input. Transaction security is based on the user that performs the load from the TWA, not the user that loaded the TWA transaction originally.

Auditing occurs when the transactions are processed successfully.

Note: Before you perform the load, you can predetermine whether a set of transactions will create new CIs or relationships by using the simulation options (-simci and -simrel), see [How to Simulate TWA Operations](#) (see page 522) for more information.

More information:

[How To Prevent Data Regression](#) (see page 533)

[Filter By Change Order Number](#) (see page 534)

How To Prevent Data Regression

While transaction data for a CI is in the TWA, some other action can update that CI in the CMDB and cause the TWA transaction data to become invalid. To prevent unwanted data regression, GRLoader compares the TWA tran_dt (Transaction Date) field with the CI last_update_dt (Last Modified Date) or relationship last_mod_dt (Last Modified Date). If the last_update_dt is more recent, GRLoader rejects the TWA transaction and posts an error message to the tran_message column.

You can resolve this situation by doing one of the following:

- Use the web interface to edit the CI or relationship tran_dt (Transaction Date) field to a more appropriate value.
- Set the GRLoader option grloader.workarea.ignore_transaction_dates="yes" to ignore the transaction dates and insert the data.

Note: If GRLoader rejects a transaction as older than the target CI, and then you verify that the transaction should run, manually set the transaction date to be more current and re-run GRLoader.

Warning! The ignore transaction dates option can cause data backleveling or regression. When a single CI or relationship is updated multiple times in the same GRLoader -lftwa run, the last_update date of the CI or relationship is set to the current time during the first update. To allow for multiple updates to the same CI or relationship, an additional check is made to see if the CI or relationship was updated after GRLoader started. If it was, then the update is allowed.

Filter By Change Order Number

The GRLoader -chg option can be used to select only transactions that are associated with a change ticket. To filter by change number, you must set the Change Order attribute (tran_chg_ref_num) to the appropriate change request number.

Note: The Change Order string is not validated when loaded into the CMDB.

Example: Filter By Change Ticket

The following option loads only transactions related to Change Order 23434.

```
grloader -lftwa -chg 23434
```

TWA Administration

The TWA may require the following administration:

- Extend the TWA – required when you extend any OOB families or define your own custom families.
- Archive and purge of the TWA
- How to use the TWA with CA Cohesion ACM
- Database limitations

More information:

[Extend the TWA Object](#) (see page 534)

[Archive and Purge of the TWA](#) (see page 535)

[TWA GRLoader Command Options](#) (see page 536)

[TWA GRLoader Configuration File Options](#) (see page 537)

[How to Use the TWA with CA Cohesion ACM](#) (see page 537)

[Database Limitations](#) (see page 537)

Extend the TWA Object

Any modifications to the CA CMDB family schema require parallel modifications to the TWA schema. If no custom families or attribute have been defined, the TWA requires no modifications.

To extend the TWA object

1. Customize the user-defined families and attributes using the procedures in Extending CA CMDB in the *Administration Guide*.
2. Add the new family attributes to the ci_twa_ci schema using Web Screen Painter Schema Designer.

To add a new attribute to ci_twa_ci schema:

1. Using Web Screen Painter Schema Designer, open the schema for the ci_twa_ci table.
2. Add the desired family attribute to the extension table.
3. Publish the modified ci_twa_ci schema.

Note: The new attributes must be type STRING and hold the longest possible text value.

3. Add the metadata for the custom family to the TWA as follows:
 1. Using Web Screen Painter, open the cmdb_metadata_site_families.html file.
 2. Add an PDM_INCLUDE for the appropriate cmdb_metadata_extension.html file that was created for the custom family.

Archive and Purge of the TWA

To maintain the TWA, CA SDM provides the following archive/purge rules:

CI Inactive Transactions

Archives and purges CI transactions that are marked for deletion.

CI Successful Transactions

Archives and purges CI transactions that have completed successfully.

Relationship Inactive Transactions

Archives and purges relationship transactions that are marked for deletion.

Relationship Successful Transactions

Archives and purges relationship transactions that have completed successfully.

You can select how often the rules run or whether they are enabled.

You can customize and enable the rules for enabling archive and purge of the TWA. If you are a heavy user of the TWA, be aware of the limitations on the number of records in the TWA imposed by the DBMS.

If you have to reinitialize the TWA to clear out all data in the TWA, do the following:

- Set all transactions to inactive
- Run archive purge

Important! Do not use SQL to delete all records in the TWA because it deletes required header records.

TWA GRLoader Command Options

CA SDM provides the following command options for use in GRLoader TWA processing:

-littwa

Loads XML into the initial state.

-littwar

Loads XML into the ready state.

-lftwa

Loads transactions from the TWA.

-lftwai

Loads transactions from the TWA and inactivates successful transactions.

-chg *nnnn*

Used with **-lftwa** and **-lftwar**. Loads only those transactions associated with change order *nnnn*.

Note: The Change Order string is not validated when loaded into the CMDB.

-l

Ignore case of lookup attributes. By default free-form text for lookup attributes is processed as case sensitive unless you use this option.

-simci

Determines whether CI transactions can create or update CIs. When this option is used, GRLoader perform error checking against CIs only.

-simrel

Determines whether relationship transactions can create or update relationships. This option verifies relationships for the existence of the provider and dependent CIs, validates relationship types, and so on.

TWA GRLoader Configuration File Options

CA CMDB provides the following TWA options for the GRLoader configuration file:

```
grloader.emptyvalue=EMPTY
grloader.loadfromtwa=yes
grloader.loadfromtwa.inactivatesuccessful=yes/no
grloader.loadtotwa=yes
grloader.loadtotwa.ready=yes/no
grloader.workarea.delimiters={ }
grloader.workarea.ignore_transaction_dates=yes
grloader.simulateloadci=yes/no
grloader.simulateloadrelation=yes/no
```

How to Use the TWA with CA Cohesion ACM

To use the TWA with CA Cohesion ACM processing, apply patch RO08739 to CA Cohesion ACM r5. This patch allows you to specify the `-littwa` option in the Other Options field on the Export Options tab of the CA CMDB Export Report definition. Upgrade the GRLoader component. Contact CA support for questions.

Database Limitations

The TWA is subject to the data limitations of the underlying database, as follows:

- When using TWA and Microsoft SQL Server, the total length of one data transaction must not exceed 8060 bytes.
- When using the `dbload` and `pdm_load` utilities, a load operation is restricted to 512 columns.
- Oracle has a 1000-column limitation. If you customize CA CMDB, verify that the total number of attributes does not exceed 1000 across all families (both CA-supplied and user-defined).
- SQL Server 2005 limitations as described in the following table.

Limitation Type	Limit (32-bit and 64-bit)
Columns per wide table	30000
Columns per base table	1024

Limitation Type	Limit (32-bit and 64-bit)
Columns per SELECT statement	4096
Bytes per row	8060
Rows per table	Limited by available storage

CA CMDB Data Maintenance

The tasks you perform to maintain CMDB data require administrator privileges. These tasks include how to set up the families, classes, and relationship types used to manage configuration items and relationship information.

More information:

[CA CMDB Family/Class Structure](#) (see page 538)

[Change Family/Class of a Single CI](#) (see page 539)

[Change the Family/Class of a List of CIs](#) (see page 540)

[Change CI Family/Class Using GRLoader](#) (see page 540)

[Extending CA CMDB](#) (see page 541)

CA CMDB Family/Class Structure

CA CMDB provides default configuration item families and the CI classes they contain.

Note: For information about the default CA CMDB family and class structure, see the *CA CMDB Technical Reference Guide*.

Note: The following Service Desk and CA APM base families do not have their own CA CMDB extension tables:

- Computer
- Contact (as base object)
- Hardware
- Location (as base object)
- Organization (as base object)
- Other
- Project
- Software

In CA CMDB, CIs in these base families receive CI Detail pages with some extraneous fields and lacking an Attributes tab. To take advantage of CA CMDB advanced features such as the ability to track family-specific attributes, versioning, snapshots and baselines, we recommend that you use the Change Family and Class capability to convert these CIs to CA CMDB families.

Change Family/Class of a Single CI

You can change the Family and Class for a single CI.

To change the Family and Class for a single CI

1. Select the CI that you want to change.
2. Click Edit.

The Update Configuration Item page displays.

3. Change the Class value for the CI. Selecting the class also determines the family for the CI. You can enter a value directly or click the magnifying glass to select from a list of classes.
4. Click Save.

The Family and Class of the single CI is changed.

Change the Family/Class of a List of CIs

You can change the Family and Class of a list of CIs.

To Change the Family/Class of a List of CIs

1. In the Configuration Item List, click Show Filter.
2. Complete one or more of the Search fields to search for the CIs you want to edit.
3. Click Search.

The Configuration Item List populates with all the CIs that match your search criteria.

4. Click Edit in List.

The Configuration Item List displays the fields that can be changed for all selected CIs.

5. Edit the Family and Class fields as desired.
6. Click Change All.
7. Click Save.

When you refresh the Configuration Item List, it displays the updated CIs.

Change CI Family/Class Using GRLoader

You can change the Family and Class for a CI by using GRLoader.

To change the Family or Class of a CI using GRLoader

1. Create XML code to change the CI attribute.
2. Using GRLoader, open the CI you want to modify.
3. Run the XML through GRLoader to change the CI attribute.

The Family and Class for a CI are changed.

Example: Set the Class of an Open CI to Document

The following example shows an XML fragment that sets the class of an open CI to *Document*.

```
<ci>
  <name>document number 1
  <class>Document
</ci>
```

Important! Do not attempt to update both the Family and Class attributes of a CI at the same time. To change both, you must use two separate CI updates in two separate invocations of GRLoader).

Extending CA CMDB

CA CMDB is a highly flexible system that can be extended to include additional families, classes, and attributes according to your business needs. New attributes can be family-specific or *common* (applicable across all families). While CA CMDB provides predefined families with many classes and attributes based on industry standards, some business cases require one or more of the following activities:

- Extend one or more of the CI families by adding new attributes. For example, to add a GPS coordinate for every device on your office campus, you can define a `gps_coordinate` attribute to add to any desired family. If you only want to extend one family, use Web Screen Painter Schema Designer to define the new attributes in the existing extension table. In addition, whenever you add an attribute, you also must modify the Detail page, Attribute tab, and metadata forms that use the attribute. For more information, see [Add Family Attributes](#) (see page 542).
- Extend all CI families by adding a common attribute. For more information, see [Add Common Attributes](#) (see page 543).
- Add new classes to an existing family to support more classification detail in your system. For example, instead of the generic Server class, you can create a separate class for every model of server device. For more information, see [Define a New CI Class](#) (see page 544).
- Add a new family by using an existing extension table and its attributes. A new family provides an alternative way of organizing or classifying CIs. For more information, see [Define a New CI Family](#) (see page 545).

- If the existing class or family structure does not match your requirements, you can start over with a minimum set of attributes. If you want to add a new family using a new extension table, define the new extension table and its attributes using Web Screen Painter Schema Designer and also create the forms that are required for display and update. For more information, see [Constructing a New Attribute Framework](#) (see page 545).

Important: Extending CA CMDB requires specialized knowledge of CA SDM data structures and tables, and familiarity with Web Screen Painter (WSP). We recommend that you contact CA Services to assist in this activity and also read and understand thoroughly the following sections before attempting to extend CA CMDB families and attributes.

Adding New CA CMDB Attributes

Adding a new CA CMDB attribute requires careful planning. CA CMDB already provides many attributes in its predefined families and classes, so determine whether they are sufficient for your needs before considering customization.

If you determine that a new attribute is necessary, we recommend proceeding conservatively by asking whether you need to:

- adding one or more new attributes to an existing family?
- adding the new attribute(s) to more than one family?
- adding a new common attribute, which applies across all families?

You can extend an existing family by adding new attributes. For example, if some devices on your office campus are assigned a GPS coordinate, you can add a `gps_coordinate` attribute to any applicable CI family. For more information, continue to [Add Family Attributes](#) (see page 542).

After your attribute requirements are identified, if you determine that they require the use of new families and classes, then see [Adding a New CA CMDB Family or Class](#) (see page 544).

Add Family Attributes

You can add a new attribute to one or more existing families by using the Web Screen Painter Schema Designer, which updates the database to include the new attribute and also updates the associated CA SDM tables and files. This method is preferred over updating the tables and file changes manually.

To add a new attribute to a family

1. Using Web Screen Painter Schema Designer, open the schema that corresponds to the family extension table.
2. Add the desired attribute to the extension table.
3. Publish the modified extension table.

Note: GRLoader and Versioning automatically pick up new attributes without further action. However, we also recommended that you enable logging. Logging is required for auditing purposes and Versioning enabled to record all snapshots.

To enable logging, verify that UI_INFO for the attribute is set to **AUDITLOG**.

After a new family attribute is created, it must be added to each display form so that users can see and update that attribute and to the metadata form specific to that extension table. For more information, go to [Add Attributes to Forms](#) (see page 547).

In addition, Versioning requires metadata, including information about column headings and the related Standard CI attributes. You define new metadata for all new attributes that you create; for instructions, continue to [Create Metadata](#) (see page 548).

Add Common Attributes

Common attributes are stored in the **ca_owned_resource** table in the **nr** object. This table provides the following customizable fields:

- smag_1
- smag_2
- smag_3
- smag_4
- smag_5
- smag_6

If you require less than seven customized attributes among all CI families, these fields provide a convenient solution.

As supplied, these custom fields do not display on any form. To allow CA CMDB users to view or update a smag_*n* field, use the Web Screen Painter (WSP) to add it to any display form as desired. All logging and GRLoader functions are already enabled.

Note: For Web Screen Painter Schema Designer procedures, consult the "Customizing" chapter of the *CA Service Desk Implementation Guide* and the Customize Database Schema section of Web Screen Painter online help.

Adding a New CA CMDB Family or Class

Adding a new CA CMDB family requires careful planning. CA CMDB provides many pre-defined families and classes, so determine whether these are sufficient for your needs before considering customization. If not, some new requirements can be met by one of the following:

- defining a new class in an existing family
- defining a new family that uses an existing extension table

If you determine that existing extension tables are insufficient for your needs, skip this section and continue to the section [Constructing a New Attribute Framework](#) (see page 545) to create an extension table and the forms that it requires.

Define a New CI Class

You can add new classes to support higher levels of classification granularity. For example, instead of using any of the CA CMDB-provided classes in the Hardware.Server family, you can define additional classes for different server devices.

Before you create a new configuration item class, determine whether a suitable class already exists in the Configuration Item Class List.

Note: If you also require a new family to govern a new set of classes, skip this topic and proceed to [Define a New CI Family](#) (see page 545). Then you can return and populate the new family with classes.

To define a CI class using the Administration interface

1. Navigate to the Configuration Item Class List.
2. Click Create New.
The Create New Configuration Item Class page appears.
3. Enter a unique name for the new class.
4. Enter the appropriate family name in the Family field, or click the icon over the field to search for a family.
5. Verify that the Record Status field is set to Active.
6. Click Save.

The new CI class is defined and ready for you to create new CIs.

Define a New CI Family

Before you create a configuration item family, determine whether a suitable family already exists in the Configuration Item Family List. If not, you can create new families as an alternative method of organizing or classifying CIs using existing attributes.

If the new family that you need cannot use an existing extension table (or the default table), complete all the additional preparations outlined in [Constructing a New Attribute Framework](#) (see page 545). The new extension table can be used to define new families.

To define a CI family using the Administration interface

1. Navigate to the Configuration Item Family List.
2. Click Create New.
The Create New Configuration Item Family page appears.
3. Enter a unique name for the new family.
4. Verify that the Record Status field is set to Active.
5. In the Extension Name field, select the extension table that identifies the type of family that you want to create. This may be a new extension table that you created recently.

For example, if you are adding a family for an unspecified type of hardware, select `ci_hardware_other`. This type ensures that when you create configuration items whose classes use the new family, the appropriate attributes display on the Attributes tab. If you do not select a table name in the Extension Name field, the default table is used and only default attributes appear when you create a configuration item within the new family.

6. Type a description in the Description field.
The description you enter displays in the Configuration Item Family List for informational purposes.
7. Click Save.
The configuration item family is defined.

Now that your new family is defined, you can add classes to it as described in [Define a New CI Class](#) (see page 544).

Constructing a New Attribute Framework

If none of the existing class or family structures match your requirements, you can start over with a minimum set of new attributes. This requires construction of a new extension table and other supporting structures.

Before using the CA CMDB administration interface to define a new CI family based on a new extension table, use the Web Screen Painter Schema Designer to create the new extension table.

To use the new extension table, you must also create new HTML forms for:

- new CI Detail page
- new Attributes tab with its associated attributes
- new metadata form and metadata

CA CMDB supplies templates that you can use to build these HTML forms. The following sections provide more detailed information about what is required.

Note: For Web Screen Painter Schema Designer procedures, consult the "Customizing" chapter of the *CA Service Desk Implementation Guide* and the Customize Database Schema section of Web Screen Painter online help.

To use the new framework in the CA CMDB user interface, define the following:

- one or more CI families that use the new extension table
- CI classes to classify your CIs by type

More information:

[Adding a New CA CMDB Family or Class](#) (see page 544)

Create a New Extension Table

Before you can define a family based on a new extension table, you must update the database with the new table and also update the CA CMDB schema with information about that table.

To create an extension table

1. Using the Web Screen Painter Schema Designer, define the new extension table and extension name.
2. Save and publish the new extension table.

Note: WSP Schema Designer automatically creates the logging trigger in CA CMDB.

Continue to the next section to create the CI Detail page.

Create a CI Detail Page

A CI Detail page is required to support attribute display for CIs that are associated with a new extension table.

To create the CI detail page

1. Using the Web Screen Painter Schema Designer, click File, New and create a form based on `detail_extension.template`.
2. Save this new form as **`detail_extension.htmlpl`**, where *extension* is the name of the extension table.
3. Follow the instructions listed in the file, replacing the **`***EXTENSION***`** string with the name of the new extension table (defined previously).
4. Save the file with all changes.

The CI Detail page includes two attribute sections:

- common attribute section named `cmdb_detail.htmlpl`
- family-specific section (the Attributes tab) named `nr_cmdb_extension_tab.htmlpl`, where *extension* is the name of the new extension table.

Continue to the next section to [Create the CI Attributes Tab](#) (see page 547).

Create a CI Attributes Tab

The Attributes tab displays the family-specific attributes for a CI.

To create the Attributes tab

1. Using the Web Screen Painter Visual Editor, click File, New and create a form based on `nr_cmdb_extension_tab.template`.
2. Save this file as **`nr_cmdb_extension_tab.htmlpl`**, where *extension* is the name of the new extension table.
3. Follow the instructions listed in the file, replacing the **`***EXTENSION***`** string with the name of the new extension table (defined previously).
4. Save and publish the file with all changes.

Continue to the next section to populate the Attributes tab.

Add Attributes to Forms

After Web Screen Painter Schema Designed has been used to create a new attribute in an extension table, that attribute must be added to any form that is used for display or update. For new family-specific attributes, the only form that must be changed is the Attributes tab named **`nr_cmdb_extension_tab.htmlpl`**, where *extension* is the name of the extension table. This form must be edited to include any new attributes.

To edit an attribute form

1. Using Web Screen Painter Visual Editor, click File, Open to access the appropriate form.
2. Drag and drop the new attribute(s) on the form.

Note: CA CMDB-provided forms do not lend itself to editing using the Web Screen Painter Visual Editor, so use the Web Screen Painter text editor on the Source tab.

3. Save and publish the form.

If you have not yet created a metadata form, continue to the section [Create a Metadata Form](#) (see page 548). To define metadata for a new attribute on the form, continue to the section [Create Metadata](#) (see page 548).

Create a Metadata Form

A new extension table requires its own metadata form to define column headings and Standard CI information for Versioning.

To create the metadata form

1. Using the Web Screen Painter Visual Editor, click File, Open to access `cmdb_metadata_extension.template`.
2. Save the file as **`cmdb_metadata_extension.html`**, where *extension* is the name of the new extension table.
3. Follow the instructions listed in the file, replacing the **`***EXTENSION***`** string with the name of the new extension table (defined previously).
4. Save and publish the form with all changes.

Continue to the next section to populate the metadata form.

Create Metadata

Metadata includes information about attribute column headings and Standard CI information that the Versioning feature requires.

Important: Metadata requires careful planning to ensure correct data in Snapshots, correct titles in Versioning, and successful Standard CI comparisons.

To create metadata

1. Using the Web Screen Painter Visual Editor, click File, Open to access **cmdb_metadata_extension.html**, where *extension* is the name of the extension table.
2. Following the instructions listed in the form, copy and modify the indicated row for each attribute in the new extension table.

Note: The following attributes, although required, do not need metadata:

- ID
- Last_modified_by
- Etc (to be provided)

3. Save and publish all changes.

If you are adding metadata to an existing CA CMDB family, audit changes will be displayed correctly on the Versioning tab. However, if you are defining metadata for a new extension table, you must have a new family and class for your attributes; for more information, see [Adding a New CA CMDB Family or Class](#) (see page 544).

The structures for your new extension table are now in place. To define a new family for your attributes, continue to [Define a New CI Family](#) (see page 545).

Example

In this sample scenario, you create an **Automobile** family and a **Sedan** class within it for tracking inventory for manufacturing, shipping, rental transportation or other purposes. Of course, you could also create many other automobile classes; this example is for demonstration purposes only.

Step 1: Create the New Extension Table

1. In Web Screen Painter Schema Designer, click Add Extension Table.
2. Enter a name for the new extension table. In this example, **vehicle**.

The **zvehicle** extension table is created and its properties displayed

Note: “z” is appended to the beginning of all new table names to distinguish them from application-provided tables.

3. In the Table Info tab, set the Function Group field to **inventory**. Other fields are populated with default values.
4. Add new columns and attribute information as desired.
5. Save and publish the new extension table.

Step 2: Create a New Family

To create a family

1. Navigate to the CI Family List.
2. Click Create New.
The Create New Configuration Item Family page displays.
3. Enter a unique name for the new family. In this example, **Automobile**.
4. Verify that Record Status is set to Active.
5. Select the extension table name. In this example, **zvehicle**.
6. Click Save.

The new CI family is created.

Step 3: Create a New Class

1. Navigate to the CI Class List.
2. Click Create New.
The Create New Configuration Item Class page displays.
3. Enter a unique name for the new class. In this example, **Sedan**
4. Verify that Record Status is set to Active.
5. Select the Family. In this example, **Automobile**.
6. Click Save.

The new CI class is created.

Step 4: Create the New CI Detail Form

To create the CI Detail form

1. In the Web Screen Painter Visual Editor, open the form detail_extension.template.
2. Save the template with the name **detail_zvehicle.htmlpl**.
The new form is saved under NX_ROOT\site\mods\wsp\project\web\analyst
3. In the new form, replace all instances of *****EXTENSION***** with **zvehicle**
4. Save and publish the new CI Detail form.

Step 5: Create the Attributes Tab

To create the Attributes tab form

1. In the Web Screen Painter Visual Editor, open the form `nr_cmdb_extension_tab.template`
2. Save the template with the name **nr_cmdb_zvehicle_tab.htmlpl**
The new form is saved under `NX_ROOT\site\mods\wsp\project\web\analyst`
3. In the new form, replace all instances of *****EXTENSION***** with **zvehicle**
4. Add attributes to the form as desired.
5. Save and publish the new Attributes tab form.

Step 6: Create the Metadata Form

To create the metadata form

1. In the Web Screen Painter Visual Editor, open the form `cmdb_metadata_extension.template`
2. Save the template with the name **cmdb_metadata_zvehicle.htmlpl**
The new form is saved under `NX_ROOT\site\mods\wsp\project\web\analyst`
3. Replace all occurrences of *****EXTENSION***** with **zvehicle**
4. Using the heading and attribute placeholders, add metadata for all family-specific attributes.
5. Save and publish the new metadata form.

Configuration Audit and Control Facility (CACF)

Configuration Audit and Control Facility (CACF) unifies three disciplines, Change Management, Configuration Management (CMDB), and Discovery Management to verify that changes are executed accurately and no unauthorized changes occur.

Change verification helps ensure that the CMDB reflects any changes accurately, and the Discovery Management tools verify the changes.

Change verification provides the following benefits:

Change Management

- Assurance that authorized changes execute correctly
- Detection of incorrectly executed changes
- Detection of unauthorized changes

- Management reporting
- Full auditing of all changes down the attribute level
- Detection of overlapping or conflicting changes

Configuration Management

- CMDB contains an accurate and current representation of all managed CIs
- Ability to view the future state of the CI with the proposed change specifications
- Full auditing of CIs

CACF has two major sections, the CMDB administrative interface and the change management interface that the following sections cover in more detail.

More information:

[CACF Administration and Policy Definition](#) (see page 553)

[Managed Attributes](#) (see page 566)

[Managed Change States](#) (see page 566)

[Change Specifications](#) (see page 569)

[How Change Verification Occurs](#) (see page 574)

[How to Archive and Purge Audit Data](#) (see page 578)

[Implement a Change Verification Strategy](#) (see page 579)

[Planning and Implementing Change Verification](#) (see page 583)

[Change Verification Best Practices](#) (see page 589)

[Verify a CI Attribute Value Update Manually](#) (see page 593)

[Example: Allow Rogue Updates Only From a Specific Location](#) (see page 597)

[Example: Upgrade Laptops in Your Organization](#) (see page 598)

[Example: Lock Down Nonverified Change Orders](#) (see page 599)

[Example: Allow a CI Update If No Matching Change Order Exists](#) (see page 600)

[Example: Defer All Updates from CA Configuration Automation to the TWA](#) (see page 601)

[Example: Only Log the Policy Results as a Test](#) (see page 601)

[Example: Reject a CI Update](#) (see page 602)

[Example: Allow Change Orders Created Without Specifications](#) (see page 602)

[Example: Do Not Allow Change Orders Created Without Specifications](#) (see page 603)

[Example: Allow Rogue Inserts from Selected Sources](#) (see page 603)

[Example: Allow a Rogue Update for a Nonproduction CI](#) (see page 604)

More information:

[How Change Analysts Define Change Specifications](#) (see page 555)

[Change Verification](#) (see page 555)

[How Change Managers Use Change Specifications](#) (see page 556)

[How Configuration Audit and Control Facility Works](#) (see page 557)

CACF Administration and Policy Definition

The CMDB Administrator defines CIs and attributes which are managed as well as the policy for updating those CIs and attributes. You administer the CACF components in the Configuration Control node in the CA CMDB section of the Administration tab and define and view CACF change specifications from the Change Order, CI, Incident forms, or the Configuration Audit node in the Administration tab.

The administrator must consider allowing standard changes, defining authoritative and trusted sources of data, as well as defining which CIs and CI attributes are under CACF management. If an incorrectly executed change or rogue change occurs (also referred to as a variance), CA SDM can accept the new value, create an Incident, or copy the data to the TWA for later processing. CA SDM can use any combination of these actions.

Important! The Configuration Administrator *must* [establish a change verification strategy](#) (see page 579) for your environment. The default policy for CACF allows all changes to all CIs, even if the change is rogue, or if the change does not match a change ticket.

You can implement change verification policies dynamically or scheduled in advance. You can define these policies as generic or highly specific. The change verification policy describes how CA SDM responds to the following events:

Updates from Unauthorized MDRs

Indicates that specific MDRs are authoritative for specific attributes. CI attribute updates from unauthorized MDRs can be selectively accepted or rejected.

For example, define policies to prevent [assign the value for acm in your book] from updating the IP address of a CI, even if a matching Change Order exists. Allow updates to the IP address to occur only if the source MDR is Spectrum.

Rogue Changes

Detects and manages updates to CIs when no corresponding Change Order exists. Specify a policy that manages rogue changes requesting inserts or updates of CI data.

For example, define a policy whenever a CI named like server* in New York changes, but does not have a matching Change Order, do not update the CI but instead load the data to the TWA.

Incorrectly executed changes

Detects whenever a Change Order is not [implemented](#) (see page 555) correctly.

Auditing facilities provide the Configuration Manager with capabilities to view changes to policies and CACF object definitions. CACF logs each attempted update to the CI, whether or not the update was allowed. This log helps you determine which policy allowed or disallowed a change to a CI.

The Configuration Administrator defines which Change Order statuses represent editable changes, and which Change Order change states represent verification states. By default, you can edit change specifications when the Change Order is in the *RFC* change state, and change verification occurs when a Change Order is in the *Verification in Progress* state. Optionally, change tickets can be automatically promoted or closed when all associated change specifications execute successfully.

Important! Possible performance degradation can occur if the Configuration Administrator enables the Create Incident option with more than one active verification policy.

Note: By default, CACF considers updates to attribute values to Change Orders with change verification active as non-rogue changes which is in effect in the Verification in Progress status. The Configuration Administrator can [modify this behavior](#) (see page 566).

How Change Analysts Define Change Specifications

The Change Management interface portion of change verification integrates into the change ticket order form. Change Analysts typically create change specifications when they create a Change Order. The Change Analyst defines the change in terms of the specific CI and CI attribute that they want to change.

The Change Analyst can describe the change specification using any of the following templates:

- Provide the exact CI, CI attribute, and value.
For example, after the change has been implemented, the IP Address for server1 sets to 10.10.10.10.
- Provide the CI attribute and value, but omit the CI name.
For example, after the change, all CIs attached to the Change Order upgrade to 8GB of memory.
- Provide an attribute value to set when the attribute is not discoverable.
For example: after the change, the CMDB reflects that all CIs are located in NY.
- Use the discovered value.
For example, you do not know the new IP Address in advance of the change, but you know that the IP Address will definitely change. Set the IP Address of the CI to whatever value the authoritative MDR discovers.

When the change order is approved, the change specifications are approved as part of the same approval process. Subsequent modifications to the specification can be prohibited and are fully audited.

Change Verification

Change verification helps ensure that Change Orders execute according to change specifications. Verification occurs after a Change Order enters a change management state such as Verification In Progress and when any kind of CI update occurs. At this time, when data from any MDR or web client user imports into the CMDB, CACF compares the incoming attribute data with active change specifications. If the CI attribute data matches, CACF verifies the change. If the inbound data does not match the Change Order specifications, actions to remediate the variance can occur based on policy.

Executed policies can any of the following actions:

- Rejects a single attribute update from the transaction.
- Rejects the entire transaction, including all attribute and value pairs.

- Loads the entire transaction to the TWA for deferred processing after approval and verification of the transaction.
- Creates an Incident that describes each variance.
Note: This action can trigger existing notification and automation. To prevent multiple incident creation, rogue change detection only creates a single incident for all rogue changes to a single CI.
- Accepts the transaction unconditionally when accepting data from authorized source MDRs.

For example, the CMDB Administrator can define a policy that creates an Incident when a change to a production server executes incorrectly. When imported data from an authorized MDR does not match a pending change specification, CACF creates an Incident. The policy can allow or reject a CI update by the nonmatching import data. In addition, the Change Manager can have the opportunity to accept the newly discovered value, even if the value does not match the planned change specification exactly.

Note: If multiple discoveries occur during change specification verification, discovery data can invalidate a previously verified change. At the end of verification, the CI is in the state that all the change specifications want.

How Change Managers Use Change Specifications

The Change Manager can perform the following actions:


- View and modify the verification status of each change specification.
For example, the Change Manager wants to know which change specifications are still pending for Change Order 12345, which is in *Pending Verification* status.
- Manage Incidents for rogue or improperly executed changes.
For example, the Change Manager wants to see which changes have failed verification and which changes require additional investigation.
- Manage changes which require manual intervention.
For example, a Change Order specifies that manual verification is required for a nondiscoverable change, and the Change Analyst must complete the task.
For example, the Change Manager investigates and determines that Change Order 12345 specifies ten different change specifications. Due to a typo, change specification number 9 was incorrect and can be canceled.

Auditing facilities provide the change capabilities to view the status of each update to the change specification and any overrides to the change specification. The scoreboard shows the status of the running system and any exceptional conditions that require intervention or investigation. CA Business Intelligence reports provide a longer term view of the efficiency of Configuration Manager policies to monitor the health of your environment.

How Configuration Audit and Control Facility Works

CACF governs inbound CI data before loading the data into the CMDB. This verification helps ensure that each requested Change Order executes correctly, and detects and manages rogue changes automatically. The Configuration Administrator (CMDB Administrator) defines verification policies to determine how CA SDM responds when any deviation or variance from the requested change occurs.

The following diagram shows how CACF works:



1. Inbound CI data loads from Web Services, GRLoader, MDR, or the Web Interface to request an attribute update for a CI in the CMDB.
2. Based on the verification policies established by your Configuration Administrator, CACF completes one or more the following actions:
 - CACF accepts the update and modifies the CI data in the CMDB.
 - CACF defers the update to the TWA for further processing.
For example, your Change Manager handles transactions in the TWA.
 - CACF creates an Incident and lets you manage the ticket.
For example, your Incident Management process requires managing the Incident for reporting purposes.
 - CACF rejects the update and the CI data is not changed in the database.

Verification Policy

After the Configuration Manager and the Change Manager agree on a change verification policy, those decisions transcribe into a CA SDM change verification policy. A system can have many verification policies. When a CI update occurs, CACF selects a single verification policy for each updated attribute, and that policy can decide to update the CI. It can also decide to create a TWA entry, or to create an Incident on the existence of a matching change specification.

Note: If you want to monitor updates through the web client, specify the MDR Class Pattern as **Web Client** (case sensitive) in the verification policy. This pattern usage applies to English and localized versions of CA SDM when you want to monitor updates through the web client.

Important! Determining what policies you need in your environment requires significant thought into the goals of the organization and how you want to introduce those policies.

Verification Policy Selection

When an update occurs to a CI, CACF selects a policy for each updated attribute. This selection depends on the source of the update, the target CI, and whether matching change specification exists. When multiple policies match, CACF selects the single policy with the lowest sequence number to manage the attribute update.

Note: CACF can select multiple policies to manage a single transaction that updates multiple attributes. For example, updating a CI with a new description and new IP Address can trigger two different policies.

Each policy has four sections that CACF uses in policy selection:

Identification and Priority

Identifies a policy and specifies the order to evaluate the policy.

Change Order Alignment

Selects transactions based on how closely the attribute update matches an existing change specification.

Transaction Filter

Specifies the data sources to which this policy applies.

CI Filter

Specifies the CIs to which this policy applies.

A policy matches an attribute update when it passes each of these filters. If a policy does not match a transaction, CACF bypasses the policy, and CACF evaluates other policies in the system. If CACF bypasses all policies, the default action lets the CI attribute update occur.

Verification Policy Filter Syntax

To help with establishing generic policies, the Change Order Alignment, Transaction Filter, and CI Filter accept an asterisk as a final pattern character as a wildcard. Use an exclamation point at the beginning of these filters to indicate negation.

Important! Policy pattern rules differ from change specification planned value rules. For policy filters, you cannot use embedded asterisks (*) or the greater than (>) and less than (<) special characters.

Policy filters are case-sensitive for verification policy selection patterns, such as CI Name or CI Class, even if you define those attributes as case insensitive.

Note: A blank "" (no value specified in the field) policy filter matches only the blank value, but a policy filter that uses an asterisk matches all values.

Consider the following information with the following example verification policies:

- Policy1 selects CIs name(prod*) attribute(IP Address) where a matching change specification exists.
- Policy2 selects CIs name(prod*) attribute(IP Address) where there is a rogue update.
- Policy3 selects CIs name(test*) attribute(Memory Installed) from mdr(Cohesion).
- Policy4 selects CIs name("") attribute(Memory Installed).

The following list describes the possible outcomes of these policies:

- When an update to CI(test2) attribute(Memory Installed) occurs, Policy3 manages the update. CACF ignores Policy1 and Policy2.
- When an update to CI(prod3) attribute(Memory Installed) occurs, and a matching change specification exists, Policy1 manages the update.
- When an update to CI(development4) attribute(IP Address) occurs, CACF updates the CI with the new ip address because no matching policy exists.
- When an update to CI(prod1) attribute(Disk Capacity) occurs, no policies for attribute(Disk Capacity) exist, and CACF updates the CI with the new value. In this example, CACF does *not* manage the Disk Capacity attribute.
- Policy4 only matches CIs with blank names. Because blank CI names are invalid, the change never executes.

Verification Policy Prioritization by Sequence Number

When multiple policies match an attribute update, CACF chooses the policy with the lowest sequence number. Consider the following policies and their sequence numbers.

- Policy name(Policy1) sequence(1000) CI(prod*)
- Policy name(Policy2) sequence(2000) CI(prod-NY*)

When an update to CI(prod1) occurs, Policy1 manages the transaction. When an update to CI(prod-NY-1) occurs, Policy1 still manages the transaction because it has a lower sequence number. In this example, Policy2 never takes effect.

We recommend that you change the sequence numbers so that the most specific policy has the lower sequence number, as shown in the following example policies:

- Policy name(Policy1) sequence(2000) CI(prod*)
- Policy name(Policy2) sequence(1000) CI(prod-NY*)

When an update to CI(prod1) occurs, Policy1 manages the transaction. When an update to CI(prod-NY-1) occurs, Policy2 manages the transaction because it has a lower sequence number.

Change Order Alignment

Policies filter on how closely a CI attribute update matches a change specification. After a CI update, CACF searches for a change specification that matches both the CI and the attribute name. This search can verify the change specification. This Change Order Alignment can occur as follows:

- Change specifications exist for this attribute and CI.
- Change Orders exist for this CI, and those Change Orders do *not* contain change specifications.
- A rogue insert transaction occurs which indicates a new CI. A Change Order does *not* exist with this CI attached to it.
- A rogue update transaction occurs which indicates an existing CI. A Change Order does not exist with this CI attached to it.

Only Change Orders in a change state where *change verification is active* are considered when looking for a matching Change Order. Closed, unapproved, unscheduled, or otherwise nonexecuted Change Orders are not considered when searching for a matching Change Order.

After CACF determines the Change Order alignment of a transaction, it performs a search for a corresponding policy. A policy can manage one or more Change Order alignment types:

Change Order with Specifications

Indicates that a Change Order with a change specification which specifies the updating CI and attribute name exists.

Change Orders Without Specifications

Indicates Change Orders that specify this updating CI, and those Change Orders do not have any change specifications.

Rogue Insert

Indicates an inserted CI, by definition, the CI does not have any matching Change Order in a verification state.

Important! Policies which prevent rogue inserts must specify a managed attribute of *Name* or *All Managed Attributes*, where *Name* specifies an active managed attribute).

CACF always allows rogue inserts and updates when a transaction only has unmanaged attributes.

Rogue Update

Indicates an updated CI that has no matching Change Orders in a verification state.

More information:

[Example: Allow Change Orders Created Without Specifications](#) (see page 602)
[Prevent Rogue Changes to the Attribute](#) (see page 587)

Example: Change Order Alignment Policies

The following example policies specify Change Order Alignments:

- Policy name(legacy) filters CI(*) attribute(Any Managed Attribute) alignment(Change Orders Without Specifications) action(Allow Attribute Update)
- Policy name(new) filters CI(*) attribute(ny Managed Attribute) alignment(Change Orders with Specifications) action(Allow Update Only if Matches Change Specification)
- Policy name(no-rogue) filters CI(*) attribute(Any Managed Attribute) alignment(rogue insert or rogue update) action(Always Cancel Entire Transaction, create Incident)

These examples provide the following functionality:

- Policy(legacy) allows CA SDM r12.6 change orders without change specifications to update the CI as they did before a CACF implementation.
- Policy(new) enforces change verification for all new change orders which have matching change specifications.
- Policy(no-rogue) prevents updates to CIs when there are no Change Orders in a Verification Active state.

Transaction Filter

Policies filter transactions based on the source of the transaction, including the [attribute name, MDR name, MDR class, and role](#) (see page 559) of the user that performs the update.

Note: If you want to filter users logged on through the web interface only, specify the keyword **Web Client** for the MDR class. The userid of the contact specifies the MDR name. This method of identifying users applies only to verification policies, and does not appear in any other part of the product.

Example: Transaction Filters

The following example policies specify Transaction Filters:

- Policy name(root_access) filters ci(*) attribute(All Managed Attributes) role(Administrator) action(Allow Attribute Update)
- Policy name (cohesion_not_authorized) filters ci(*) attribute(IP Address) MDRclass(Cohesion) action(Keep Old Attribute Value)
- Policy name(john) filters ci(user1*) attribute(All Managed Attributes) MDRName(user1) MDRClass(Web Client) action(Allow Attribute Update)

These examples provide the following functionality:

- Policy(root_access) lets any user with Administrator access update any value. We recommend this type of policy to have a low sequence number.
- Policy(cohesion_not_authorized) prevents any MDR of MDR class(Cohesion) from updating the IP Address of the filtered CI. This example shows how to prevent an unauthorized MDR from updating data.
- Policy(user1) lets user1 update the CIs that the contact owns, but only when using the web client. This example shows how to provide specific users full control over their data.

CI Filter

Policy filters transactions on the characteristics of the updated CI. This selection criteria includes CI Name, class, priority, service type, and location.

Important! The CI filter is based on the attribute value in the CI before the update occurs, not the value in the inbound transaction data.

Example: CI Filter Policies

The following example policies specify CI Filters:

- Policy name(priority1) filters Service Type(priority1 resolution) action(Allow Update Only if Matches Change Specification, Create Incident)
- Policy name(prod-NY) filters ci(prod*) location(NY) attribute(All Managed Attributes) action(Allow Update Only if Matches Change Specification)
- Policy name(prod-not-NY) filters ci(prod*) location(!NY) attribute(All Managed Attributes) action(Allow Attribute Update)

These examples provide the following functionality:

- Policy(priority1) requires CIs with a service type of *priority1 resolution* to have a matching Change Order. This policy is a best practice because it helps control the most important CIs in the CMDB to be highly controlled. This policy also requires verification for all changes under change management to have change specifications and must be verified to be completed.
- Policy(prod-NY) requires CIs in the NY location to have matching Change Orders. Using location to filter policies can help gradually implement change verification on a site by site basis.
- Policy(prod-not-NY) illustrates using the exclamation point in the location pattern to indicate CIs not located in NY.

Policy Actions

After CACF selects a policy, the action section of the policy determines the outcome of the attribute update. This section identifies the most important part of the verification policy because it affects the integrity of the CMDB. This section also affects the change management workflow, related notifications, and created Incidents.

A policy can have one of the following Update Behaviors:

Allow Update Only if Matches Change Specification

Conditionally applies the inbound attribute update to the CI if it matches a change specification. This update occurs when the Change Order is in a Verification Active state. Selecting this option enables change verification.

Allow Attribute Update

Unconditionally applies the inbound attribute update to the CI. This option effectively disables all change verification. Use this behavior for standard changes, when you have a trusted data source, authoritative, and you do not require a Change Order.

Always Cancel Entire Transaction

Cancels the update and any other attribute update in this transaction, even if a matching change specification exists. Use this behavior to prevent MDRs from updating CIs where they are not authorized to update. If any policy cancels a transaction, the entire transaction is canceled, even if other policies would have allowed the change to occur.

For example, specify this behavior to stop an unauthorized MDR from inserting a CI, and also specify a common attribute name such as *Name*.

Keep Old Attribute Value

Cancels the single attribute update, but allows other attributes in the transaction to update if their policy allows it. Use this behavior when an MDR is unauthoritative at the attribute level.

A verification policy can specify that Incidents close automatically when failed verifications are remediated. The policy can specify that the inbound transaction data copies to the TWA. This type of policy is useful when data from an unauthorized MDR requires review before updating the CMDB.

To let the Configuration Manager identify that CACF policies that created the TWA record, the Change Order field in the TWA is set to the name of the policy that created it. Writing data to the TWA is independent of updating the CI in the CMDB, so a policy can update the CI, write to the TWA, or both.

Policy Scheduling

You can specify activation and deactivation dates on verification policies, which lets the Configuration Administrator schedule unattended policy changes.

Multiple Policies

As you introduce more verification policies into your environment and more attributes are managed, organizing the policies becomes increasingly complex as more overlap exists between policies. For example, you have one policy for Server CIs and another policy for CIs in NY. You must consider the appropriate policy for Server CIs in NY. Use the policy sequence number to force one policy to take precedence when multiple policies potentially could take control over a single attribute. As you manage more attributes, the possibility that more policies manage a single transaction can increase.

Consider the following information when you have multiple active policies for a single transaction:

- If any selected policy requests writing data to the TWA, data also writes to the TWA.
- If any selected policy requests to cancel the transaction, CACF cancels the entire transaction.

Important! Facilities that use the createAsset web service (including GRLoader and CMDbF) break the transaction into three separate transactions, two inserts, and an update. Updates to ca_owned_resource may be allowed, while updates to the CI extension table are canceled. Updates to CI family-specific attributes in the extension table may not be detected as inserts, but rather as updates.

Example: Multiple Policies

The following examples specify multiple policies:

- Policy name(IP Address) attribute(IP Address) action(Allow Attribute Update)
- Policy name(Memory Installed) attribute(Memory Installed) action(cancel transaction)
- Policy name(Disk Capacity) attribute(Disk Capacity) action(Keep Old Attribute Value)

These examples provide the following functionality:

- If a transaction only updates IP Address, the update is performed.
- If a transaction only updates Memory Installed, the update is not performed.
- If a transaction updates both IP Address and Memory Installed, the transaction is canceled in the web interface.
- If a transaction updates both IP Address and Disk Capacity in the web interface, IP Address is updated, but Disk Capacity is not updated.
- If a transaction updates both IP Address and Memory Installed in the web interface, neither attribute is updated.

Incident Consolidation

When a verification policy requests to create an Incident, CACF can reduce the number of open Incidents as follows:

- CACF creates only one open Incident for a single CI for rogue changes. Additional rogue changes update the single open Incident for rogue changes.
- CACF creates only one open Incident for each change specification verification failure. Additional verification failures for the change specification update the open Incident.

Managed Attributes

Managed attributes indicate those eligible CI attributes for change verification by CACF. By default, this list contains *CI Name* and *Any Managed Attribute*. You add CI attributes that you want managed as part of your change verification strategy. Define managed attributes as part of your change verification strategy in the Configuration Control, Managed Attributes node in the CA CMDB section of the Administration tab.

CACF does not consider unmanaged attributes (attributes not listed) for change verification. These unmanaged attributes update as usual.

Important! Change verification ignores unmanaged attributes and lets them update the CI, unless a verification policy specifies the *Always Cancel Entire Transaction* behavior.

Note: Case sensitivity in the managed attribute definition only applies when CACF compares the change specification planned value with the inbound CI transaction data. Case sensitivity does not apply to the selection patterns in the policy, which are always case-sensitive.

For a list of CI attribute names, see the *CA CMDB Technical Reference Guide*.

Managed Change States

CACF uses managed change states to indicate which Change Order statuses CACF manages. CACF uses managed change states to control how or when to apply change verification for CI updates made to the system. You can customize these change states to suit the needs of your organization.

You configure managed change states on the Administration tab in the CA CMDB, Configuration Control, Managed Change States node.

Note: CACF ignores Change Orders in a Change Order state that is listed in the Managed Change States.

The following list describes the Managed Change States options:

Change Specifications Editable

Specifies whether you can edit the change specifications for a Change Order. Typically, after a change request receives approval, you cannot change the request, and updates are restricted to performing a small set of override options in the *Verification in Progress* state.

Change Verification Active

Specifies whether changes discovered for a CI while the Change Order is in this status are considered for change verification. Change orders and their related change specifications are compared with any inbound transactions to verify that they were executed successfully.

CACF monitors all CIs for any changes to their managed attribute values. As CACF verifies each attribute level change to the CI, CACF compares it against a list of change specifications which are in a *change verification active* state.

After a change specification enters this state, any change specification without a specific CI undergoes expansion. This expansion occurs where new change specifications are created using the list of CIs attached to the Change Order.

After a change specification exits this state, the list of change specifications with a verify status of *Set After Change Executed* complete. This action updates the CI with the planned values, as specified in each change specification.

Implementation State

Specifies whether the state represents a state when the changes are being executed or implemented on the CI. When a Change Order enters this transitional state, it is understood that the attribute values in the CI are volatile and may be updated as requested by the pending change specifications. CACF cannot consider these changes rogue changes, but it also cannot consider the changes for final verification. Typically, the change verification process should only compare the inbound attribute data after the Change Order executes completely.

Show Change Specification Override Buttons

Specifies whether the Change Analyst can control change specifications, and what level of control is given. In some implementations, the Change Analyst can edit change specifications as necessary, while in other implementations, the Change Analyst can [override](#) (see page 577) or [cancel](#) (see page 577) the change specification.

Promote Change Order after Verification

Specifies whether a Change Order promotes to the next default state automatically after CACF verifies all the change specifications.

Default Managed Change States

You can define Change Order states during which change specifications are created, overwritten, verified, or ignored.

The following list describes the default managed change state definitions provided by CA SDM as an example:

RFC

Define the Change Order specifications in this state, but are not considered during change verification. Changes to the CI detected while the Change Order is in this state are typically considered as rogue.

Approval in Progress

Change specifications in this state have the same characteristics as RFC and are editable and not considered for verification.

After Change Order approval, the change specifications are approved as part of the same approval process. CACF prohibits additional modifications to the specification and provides full auditing.

Implementation in Progress

Changes to CIs may occur, but are not considered rogue or used for verification. You may want to change the definition of this change state to let verification occur for Change Orders in this state.

Verification in Progress

Change specifications cannot be edited, but an analyst can override them. [Change verification is active](#) (see page 574) in this state.

More information:

[Special Characters](#) (see page 570)

[Defining Bulk Load](#) (see page 572)

[Defining Bulk Changes](#) (see page 572)

[CI Versioning and Future State](#) (see page 573)

Change Specifications

A Change Order contains change specifications that define the specific CI changes that are requested for a CI. CACF uses these change specifications when the Change Order enters a verification state to validate and confirm that the actual changes to the CI completed correctly. Create change specifications from a Change Order, CI, or the Configuration Control, Change Specifications node the CA CMDB section of the Administration tab.

The change specification contains the following main sections:

Change Order Number

Specifies the Change Order that requests the change.

CI Name

Contains the name of the CI that you want to update.

You can leave the CI Name field blank to imply that the change specification applies to all CIs defined for the Change Order. Use this option when all the CIs for a Change Order use the same managed attribute and planned value. The change specification applies to all CIs that are attached to the Change Order when the change order status moves to a managed change state with change verification active, referred to as expansion.

Attribute Name

Specifies the name of the managed attribute that you want to update.

Selecting Any Managed Attribute indicates that any managed attribute can change during the verification of the Change Order. Because you can update multiple attributes in this case, you cannot specify the planned value.

Planned Value

Indicates the expected value of the attribute after executing the change. After the Change Order moves into a verification state, and the CI updates through the web interface, GRLoader, or Web Services, CACF compares the planned value with the inbound data to determine a match.

[Special characters](#) (see page 570) can be embedded in the planned value when the exact value is not known.

Status

Specifies the [status of the change specification](#) (see page 575), such as Verification Pending.

If you do not know the planned value in advance, but the value may change, you can set the status of the change specification to Use Discovered Value. This behavior requires discovery to update the CI before CACF considers the change specification as validated. You require this behavior for numeric fields and SRELs where an asterisk cannot be accepted as a planned value.

To help you with setting the planned value, the Discovered Attribute History tab lists all values that were recently discovered by an MDR, and whether were actually authorized, or loaded into the CMDB. This tab displays the format, case sensitivity, and other information about values so that you can determine an appropriate planned value pattern.

Special Characters

You can embed special characters when you do not know the exact value. Use the asterisk wildcard character to match any number of characters. The pattern *must* match the discovered data in the same sequence, and spaces are significant.

For example, enter `10.*.*` as the Planned Value to match any IP addresses that begins with `10.` and has two periods that follow any value between the periods.

For example, the inbound server_type value contains Windows 2003 (WIN32) 5.2.Service Pack 2 (Build 3790) Intel x86. To verify this value, specify a planned value of `*Service Pack 2*` in the change specification.

A word at the beginning of the planned value indicates that the discovered value must start with that value. Similarly, an asterisk at the beginning of the planned value indicates that the discovered value can begin with any value and end with the value specified following the asterisk.

The following table provides examples about how to use the asterisk:

Planned Value	Discovered Value	Match or No Match
a	b	No Match
a	a	Match
a	aba	Match
a	bab	Match
a*	a	Match
a*	ab	Match
a*	ba	No Match
*a	a	Match
*a	ab	No Match
*a	ba	Match

A pattern that begins with an exclamation point results in the negation of the value. You can only use the exclamation point as the first character in the pattern. For example, you cannot use a pattern of 10.!*.*..*.

To compare numeric values contained within string values, use greater than (>) or less than (<) as the first character in the planned value. If there is a leading exclamation point, it must be the second character.

Important! CACF ignores leading or trailing nonnumeric values in the discovered value and planned value patterns.

The following table provides examples about how to use the exclamation point, and the greater than and less than symbols:

Planned Value	Discovered Value	Match or No Match
>200	aaa 201 bbb	Match
>200GB	aaa 200 bbb	No Match
>200GB	300 GB	Match
!<200GB	200 Bytes	Match
!<200	200 Bits	Match

If you do not know the planned value in advance, but the value may change, you can set the status of the change specification to Use Discovered Value. This behavior requires discovery to update the CI before CACF considers the change specification as validated. You require this behavior for numeric fields and SRELs where an asterisk cannot be accepted as a planned value.

To help you with setting the planned value, the Discovered Attribute History tab lists all values that were recently discovered by an MDR, and whether were actually authorized, or loaded into the CMDB. This tab displays the format, case sensitivity, and other information about values so that you can determine an appropriate planned value pattern.

Defining Bulk Load

Managing new CIs without first creating a Change Order that requires defining those CIs presents an issue within CMDB environments. Consider the following information when you define a large number of CIs

- Define a special bulk loading policy that allows rogue insert to complete successfully, restricted by userid, MDR, or role.
- Define a special bulk loading policy that reroutes all new CIs to the TWA for verification and later processing. This action also requires the previous special bulk loading policy.

Defining Bulk Changes

Creating identical changes to a large number of CIs, for example, when changing the location for a collection of CIs, with the following steps:

1. Create a Change Order.
2. Define a single change specification which describes the change and leave the CI name blank.
3. Attach all the CIs to the Change Order.

If you do not know the exact details of a change in advance of the implementation, for example, you acquire a new Server and you only know the CI Name, complete the following steps:

1. Create a Change Order.
2. Create a change specification that specifies Any Managed Attribute as the attribute name.
3. Leave the planned value empty, since it is unknown.
4. Move the ticket through the change management process.

5. When the Change Order is in a verification active state, and discovery runs, the inbound CI data loads into the CI (assuming that the policy allows it).
6. When all discoveries for the CI complete, the Change Manager must mark the change specification as verified manually.

Follow these steps when there are a large number of dissimilar changes to a large number of CIs. For example, when moving a collection of servers from one location to another, and also assigning each one a unique IP address.

1. Create a Change Order
2. Create a spreadsheet and list on each row the Change Order number, CI and, the new attribute values. Each row results in a distinct change specification.
3. Load the spreadsheet with GRLoader to create the required change specifications.
4. Promote the Change Order and its change specifications as usual.

Note: For more information about using GRLoader, see the *CA CMDB Technical Reference Guide*.

CI Versioning and Future State

View change Specifications in the CI Detail form Versioning tab to provide more information and identify Change Order and corresponding change specifications. This tab also shows the future state of the CI as if the changes were applied. View snapshots of the CI, as of the scheduled date of the Change Order, or if the Change Order has not been scheduled yet, it shows as an unscheduled change.

Versioning enables the following functionality:

- Quickly identify overlapping or conflicting change specification.
- Launch directly to change specification and view the corresponding change details by clicking the Change Order link under the MDR column in the detail view.
- Snapshots showing an *Unscheduled Change* indicated a change that is scheduled for some unspecified time in the future.
- The Change Order number displays with each change specifications shown on the right side of snapshot label.
- Informational text displays in the bottom pane providing the following information about the change specification as hover text:
 - Change Order number
 - Date of the scheduled change, if the Change Order specifies a scheduled date
 - Date need by, only shown if Change Order specifies a need by date

More information:

[Managing Change Specifications](#) (see page 575)

[Change Specification Statuses](#) (see page 575)

[Managing Failed Verifications](#) (see page 577)

[Managing Undiscoverable Attributes](#) (see page 577)

How Change Verification Occurs

When a user requests a save to a CI, CA SDM searches for applicable change specifications as it searches for all matches. The following information describes an applicable change specification:

- The CI in the change specification is the same as the one that is being saved.
- The Change Order is in a managed change state defined with Change Verification Active.
- The attribute in the change specification is the same as the attribute being updated.
- The change specification is active
- If there are no managed attributes being updated or policy in effect, the save proceeds uninhibited.

If the inbound data values match a change specification exactly, CACF considers the change specification as validated. Depending on the policy, verification closes any Incident that change verification created.

Verification take place when information from web clients, web services, or GRLoader update the CI while the Change Order and all underlying change specifications are undergoing verification. After all verifications for a Change Order complete, the Change Order can optionally promote to the next Change Order state automatically.

If the inbound data values do not match the values in the change specification, it is considered an incorrectly implemented or failed change and Incidents are created or appended to, depending on policy.

For example, CA SDM does not find applicable change specifications, but Change Orders exist in a *change verification active* state. These Change Orders specify the target CI and have no change specifications. CACF manages this change by policies which handle *Change Orders without Specifications*.

- If there are no applicable change specifications or Change Orders without specifications, the save operation is considered rogue.
- After all attributes have been processed, any policy that requested the inbound data to write to the TWA triggers.

Managing Change Specifications

After a change has been executed or implemented, the Change Order enters a change state that is managed by CACF and that change state has *change verification active* set. When the Change Order is in this state, CACF reviews all related CI activity and reports the status of each change specification to the scoreboard.

The scoreboard is used by the Change Manager to verify that the required changes have been correctly executed as well as to perform any manual verifications that were required by the change orders.

Once all verifications for a Change Order are complete, the Change Order may optionally be automatically promoted to the next Change Order state (typically closed) where all change specifications are automatically marked inactive.

Change Specification Statuses

Change specifications use status to indicate the type of verification to perform and the current state of verification. The status values are also used in the verification log when recording change verification operations during the verification process.

Initial Statuses

Used when creating a Change Specification to specify the type of verification for CACF to perform. These statuses are non-final and the Change Order is not considered verified when change specifications are in either of these statuses.

Verification Pending—The change specification has not been verified, or the change is waiting for the CI update.

Manual Verification will be required—When the changes for a Change Order are verified, manual verification will be required.

Use Discovered Value—Copy the discovered value to the CI at verification time, and used when the planned value is not known prior to verification.

Set After Change Executed—Sets a CI attribute to the planned value after verification has completed. Use this status for setting CI attributes when the attribute is not discoverable. For example, set the CI Primary Contact to User1 after a change completes.

Final Statuses

Indicates a completed change specification and considered as final. When all change specifications enter either of these final states, the Change Orders are eligible for promotion to the next default state.

Verified—The change specification has been successfully verified.

Was Manually Verified—The change specification has been manually verified.

Used Discovered Value—The CI has been updated with the discovered value.

Was Set To Planned Value—The CI has been updated with the planned value after Change Order verification.

Accepted Planned Value—The CI has been updated with the planned value and overwritten during verification.

Accepted Discovered Value—The CI has been updated with the last discovered value and overwritten during verification.

No Change—The planned value matched the CI at verification time and no verification was required.

Cancel—The change specification was [canceled](#) (see page 577) by the change manager.

Intervention Statuses

Indicates a change specification requires manual intervention for verification. CA SDM highlights these statuses in a red color on list forms. These statuses are not final and the Change Order is not considered verified when change specifications are in either of these statuses.

Failed Verification—Discovery found the value to be other than what was specified in the Change Order. The Change Analyst must determine if the change was correctly executed and the Change Order was incorrectly specified, or if the Change Order was correct and the change requires verification again. Depending on the definition of the Managed Change State and the current change order status, the Change Analyst can override, change, or cancel the failed change specification.

Manual Verification Active—Manual verification is required before the change specification can be marked final.

Action Override Statuses

Indicates an action to initiate during change verification. These statuses are not final and the Change Order is not considered verified when change specifications are in either of these statuses.

Accept Planned Value—Request to override the CI attribute with the planned value.

Accept Discovered Value—Request to override the CI attribute with the last discovered value.

Reporting Status

Indicates the result of policy operations that the verification log displays for logging purposes and not used by change specifications.

Update Was Allowed—A policy allowed a request to update a CI and did not match any change specifications.

Managing Failed Verifications

When a verification fails, the CMDB, Configuration Audit, Failed Verifications node in the scoreboard lists the number of failed verifications and change specifications that require manual intervention in a red color. The Change Specifications can also be viewed under the Change Specifications tab in the Change Order, CI or Administration tab with the verify status of "Failed Verification". When a verification fails, the Change Analyst can intervene or wait until the next discovery occurs.

If the managed change state allows the Show Change Specification Override Buttons action, the Change Analyst can open the change specification, review the data, and click one of the following options to close the verification by moving it to a final state:

Accept Discovered Value

The Analyst determines that the change specification is incorrect and that the discovery tool has discovered the correct value.

Accept Planned Value

The Analyst determines that the discovery tool is wrong or has not performed discovery, and to accept planned value, as if it was discovered correctly.

Cancel

This portion of the Change Order was not executed and this specification was canceled.

Managing Undiscoverable Attributes

When a Change Specification requires manual verification, the Configuration Analyst must manually verify the change. The CMDB, Configuration Audit, Manual Intervention node in the scoreboard shows the number of change specifications that require manual intervention in a red color. The Change Specifications can also be viewed under the Change Specifications tab in the Change Order, CI or Administration tab with the verify status of "Manual Verification Active."

If the managed change state allows the Show Change Specification Override Buttons action, the Change Analyst can open the change specification, review the data, and click one of the following options to close the verification by moving it to a final state:

Mark as Verified

Used when the attribute is not discoverable and manually verified.

Cancel

This portion of the change order was not executed and this specification was canceled.

How to Archive and Purge Audit Data

We recommend that you archive and purge obsolete Verification Log entries, CACF audit history, and CI audit history as part of your periodic database maintenance.

Important! The CACF CA Business Intelligence reports typically present monthly reports with a yearly summary. The default archive purge time value for the archive purge rules for the log, CACF audit history, and CI audit history tables is 30 days. Use the same value in the archive and purge rules to help ensure data consistency between verification log entries, Incidents, and Change Orders. You can probably change the 30 day default to something more in keeping with your reporting requirements.

Complete the following actions to archive and purge audit data:

1. Use the current Incident rule to archive and purge inactive Incidents older than *nn* days.

CA SDM archives and purges verification log entries only after archiving and purging associated Incidents. If you associate a Change order with an Incident, CA SDM does not check if the Change Order is active.
2. Use the current Change Order rule to archive and purge inactive Change Orders older than *nn* days.

CA SDM archives and purges verification log entries and change specifications only after archiving and purging associated Change Orders. If you associated an Incident with a Change Order, CA SDM does not archive and purge the Change Order.
3. Use the Rogue Change Verification Log rule for rogue changes that did not create an Incident.


Use this rule to archive and purge verification log entries older than *nn* days. By definition, rogue changes are not associated with Change Orders.
4. Use the CMDB Audit rule to archive information shown in the Change Specification History, Verification Policy History, Managed Change State History, and Managed Attribute History tabs show in the respective Change Specifications, Verification Policy, Managed Change State and Managed Attributes detail forms. CA SDM stores this information in the `ci_audit` table.

Implement a Change Verification Strategy

The Configuration Administrator determines how aggressively to implement the change verification strategy for your CMDB environment. You identify areas of your Change Management process that require a change verification strategy. These areas include attributes under change control, the Change Order states that indicate change verification is active, when you can modify change specification, and authoritative MDRs.

For example, you only want to allow updates to the IP Address attribute from MDR1. You also determine the appropriate actions when CACF detects a variance and rogue change.

The following diagram explains how a Configuration Administration implements a change verification strategy:



1. [Identify the Change Order States](#) (see page 580).
2. [Identify the Managed Attributes](#) (see page 580).
3. [Identify the CIs that You Want to Manage](#) (see page 581).
4. [Identify the Authoritative MDRs](#) (see page 582).
5. [Determine the Action for a Variance](#) (see page 582).
6. [Determine the Action for a Rogue Change](#) (see page 583).

Identify the Change Order States

The Configuration Administrator identifies the Change Order states for when change verification is in effect for the CI after a change executes. Change states help you determine specific conditions, such as if you can edit change specifications in a particular state. For example, you want to review the RFC change state for Change Orders that request updates to CIs.

Follow these steps:

1. On the Administration tab, click CA CMDB, Configuration Control, Managed Change States.
2. Click RFC to view the change state details or create a Change Order state that has not been already defined.

Note: By default, the RFC Change Order state lets you only edit change specifications. The Implementation in progress state does not let you edit change specifications. The Verification in progress state enables change verification and displays override buttons for the Change Manager or other authorized user.

3. Specify the CACF options and behavior for when a Change Order enters this state.
4. Click Save.

Identify the Managed Attributes

Identify which CI attributes you want to manage for the change verification strategy. For example, you want to manage the IP Address (alarm_id) attribute with change verification.

Follow these steps:

1. On the Administration tab, click CA CMDB, Configuration Control, Managed Attributes.
2. Click Create New.
3. Complete the following steps:
 - a. Enter **alarm_id** as the Attribute Name
 - b. Enter **IP Address** as the Attribute Label.
 - c. Select an Initial Verify Status from the drop-down list.

Default: Verification Pending

- d. (Optional) Select the Case Sensitive option if you want to enforce case sensitivity for change specification planned value comparisons.

Default: Disabled

- e. Click Save.

Identify the CIs that You Want to Manage

Identify the CIs that you want to manage with a verification policy. For example, you want to manage the IP Address (alarm_id) for all high priority CIs with names beginning with NY_Server.

Follow these steps:

1. On the Administration tab, click CA CMDB, Configuration Control, Verification Policies.
2. Click Create New.
3. Complete the following steps:
 - a. Enter a sequence, such as **100**.
 - b. Enter **NY Server IP Addresses** as the Policy Name.
 - c. Select the appropriate Change Order Alignment options.

For example, you want the verification policy for all options except Change Orders Without Specifications.
4. Complete the following steps to specify the transaction and CI filters:
 - a. Select IP Address from the Managed Attribute drop-down list.
 - b. Enter a Role Pattern or leave the asterisk to apply to all roles.
 - c. Enter **NY** as the Location Pattern.
 - d. Enter **NY_Server*** as the CI Name Pattern.

For example, this filter verifies CIs named NY_Server1, NY_Server2, and so on.
5. Select *Allow Update Only if Matches Change Specification* from the Update Behavior drop-down list.
6. (Optional) Use Log Only Mode if you want to experiment with the policy and to view the results only in the standard log, and not affect the active CMDB environment.
7. Click Save.

Identify the Authoritative MDRs

Identify the authoritative MDRs in your CMDB environment. For example, you want to allow updates to CIs from CA Configuration Automation that you identified as MDR1. You consider updates from Web Services that you identified as MDR2 as unauthoritative, and you want to send those requests to the TWA.

Follow these steps:

1. Open the NY Server IP Addresses policy, click Edit.
2. Enter **MDR1** as the MDR Name Pattern.
Note: If you want to exclude an MDR named MDR2, but you want to allow MDR1, MDR3, and so on, enter **!MDR2** as the pattern.
3. Create a separate verification policy, such as MDR2 NY Server.
4. Enter the same information as in the previous policy except the following fields:
 - a. Enter **MDR2** as the MDR Name Pattern.
 - b. Select *Always Cancel the Entire Transaction* as the Update Behavior.
 - c. Select Always from the Write Data to TWA drop-down list.
5. Click Save.

Determine the Action for a Variance

Determine the action for a variance. For example, a change to a CI does not match the values specified in the Change Order. You want the verification policy to create an Incident for the variance.

Follow these steps:

1. Open the MDR2 NY Server policy that you created and click Edit.
2. Select Yes from the Create Incident drop-down list and assign a template.
3. Click Save.
4. An end user creates a Change Order with a change specification to update a CI from MDR2.
5. CACF creates an Incident if MDR2 requests a CI update with a value that does not match the planned value in the change specification defined in the Change Order.

Determine the Action for a Rogue Change

Determine the action for a rogue change. For example, CACF detects a change to a CI that does not have any change specifications from any active Change Orders. The Configuration Administrator wants to reject these types of requests.

Follow these steps:


1. Create a verification policy and enter **Rogue NY Server** as the name.
2. Complete the appropriate pattern fields.
3. Verify that only Rogue Insert and Rogue Update are selected as the Change Order Alignment options.
4. Select *Allow Update Only If Matches Change Specification* from the Update Behavior drop-down list.
5. Select Yes from the Create Incident drop-down list and assign a template.
6. Click Save.

Planning and Implementing Change Verification

The Configuration Manager wants to implement change verification so that the CMDB contains correct data. The Change Manager wants to help ensure that changes execute correctly, and is concerned with rogue changes and understand the number of rogue changes that occur and from which data sources. Both managers agree implementing change verification provides significant value to the organization. The Change Manager wants to help ensure that any changes to the Change Management process do not harm the production environment.

The Configuration Manager and the Change Manager agree to a conservative-phased implementation, so that a Configuration Administrator implements policies for specific locations and specific CIs in those locations. This scenario describes example phases about completing the verification policy implementation to a broader audience for your organization.

The following diagram shows how a Configuration Administrator completes the example phases to implement change verification:



1. [Set Up the Environment](#) (see page 585).
2. [Gather Information About Changes to a Managed Attribute](#) (see page 585).
3. [Determine the Scope of Rogue Changes to the Attribute](#) (see page 586).
4. [Prevent Rogue Changes to the Attribute](#) (see page 587).
5. [Require Change Specifications when Updating the Attribute](#) (see page 588).

Set Up the Environment

You set up the environment so that Configuration Administrators can always update CIs. The Configuration Administrator can create and update CIs in the CMDB as needed without having to specify a change order and is trusted to providing correct CI data.

Follow these steps:

1. Create a verification policy named **Policy0.1** with the following information:
 - Enter **1** as the sequence.
 - Enter **Allow Administrator to Always Update CIs** as the description.
 - Select all Change Order Alignment options.
 - Enter **Configuration Administrator** as the Role Pattern.
 - Enter **Web Client** as the MDR Class Pattern.
 - Enter an asterisk (*) as the CI Name and Class, and all other patterns.
 - Select Any Managed Attribute for the Managed Attribute.
2. Select Allow Attribute Update as the Update Behavior action.
3. Click Save.

The policy is enabled to only allow Configuration Administrators to update CIs always.

Gather Information about Changes to a Managed Attribute

You want to gather information about changes to a specific CI managed attribute to identify, what CIs are being updated, what are the source of the changes, who is making the changes, and the values being specified to help when defining a verification strategy for the attribute and corresponding CIs. For example, you want to understand the scope of all changes to IP Address (alarm_id) and log this to the CA SDM standard log (stdlog). Gathering this information can take several weeks, depending on the number of changes to attributes that occur in your organization.

1. Define IP Address (alarm_id) as a Managed Attribute.
2. Define Implementation in Progress and Verification in Progress as Managed Change States in your environment and enable Change Verification Active in both states.
3. Create a verification policy named Policy1.1 with the following information:
 - Enter **3000** as the Sequence.
 - Enter **Log All Changes to IP Address** as the description.
 - Select all Change Order Alignment options.

- Enter an asterisk as the CI Name and Class and all other patterns.
 - Enable Log Only Mode.
4. Click Save.
 5. After a few weeks, review the stdlogs to view the sources of the updates.
For example, the log displays rogue CI updates and updates with matching Change Orders.
 6. You determine that users located in NY have been updating the IP Address of CIs without creating Change Orders.
 7. After you complete your analysis, inactivate Policy1.1 by editing the policy and setting Active? to Inactive.

Determine the Scope of Rogue Changes to the Attribute

You want to determine the scope of rogue changes to the IP Address attribute. Knowing the scope of rogue changes helps you understand the impact of rogue changes on your organization. For example, based on your previous analysis using the Log Only option, you want to create Incidents whenever a change occurs to IP Address to CIs that begin with "test" located in NY without a Change Order. Communicate with Change Managers in your organization to ignore new Incidents that this process creates.

Follow these steps:

1. Complete the following initial implementation actions:
 - Define Implementation in Progress as a managed change state and enable the Implementation State option.
 - Define Verification in Progress as a managed changes state and disable the Promote Change Order After Verification option.
 - Define IP Address (alarm_id) as a Managed Attribute if not already defined.
2. Create a verification policy named Policy2.1 with the following information:
 - Enter 3001 as the Sequence.
 - Enter a description about this policy. For example, you enter **Rogue inserts and rogue updates will cause Incidents to be created. All other changes will not be impeded.**
 - Select Rogue Inserts and Rogue Updates as the alignments.
 - Select IP Address from the Managed Attribute drop-down list.
 - Enter **test*** as the CI Name and **NY** as the Location and an asterisk for all other Patterns.
 - Select Allow Attribute Update, Yes to Create Incident, and an Incident Template from the actions.
3. Click Save.

After you complete this phase, CACF creates Incidents for all computers in NY. You discover the following information after reviewing the Incidents:

- Several MDRs report different IP addresses for the same CI.
 - Some MDRs are authoritative, other MDRs are not authoritative.
 - Some CIs should not be managed based on Location, Family, Name, Service Type, and so on.
4. The Change Manager meets with other IT Managers to review the CMDB, CI detail forms, verification logs, and filtering by attribute name to see the rogue data.
 5. Your organization decides to update the CI filter in Policy 2.1 to manage Priority 1 Servers only.

Prevent Rogue Changes to the Attribute

You want to prevent rogue changes to the IP Address attribute to help ensure CMDB data integrity and require a Change Order for the update. The Change Order must specify the CI and a change specification is not required for the IP Address in order for the change to occur. Change verification occurs at the CI/Change Order level but not at the attribute level. For example, prevent any rogue updates to and Use Incident creation to track the users that request these changes.

Follow these steps:

1. Set Policy1.1 and Policy2.1 to Inactive to disable these policies so that they are not in effect.
2. Create Policy 3.1 with the following information:
 - Enter **3100** as the Sequence.
 - Enter **Prevent rogue changes without any Change Order** as the description.
 - Select Rogue Update as the alignment.
 - Enter **test*** as the CI Name and **NY** as the Location and Any Managed Attribute and an asterisk for all other patterns.
 - Select the Always Cancel Entire Transaction action.
3. Click Save.
4. Create Policy3.2 with the following information:
 - Enter **3200** as the Sequence.
 - Enter **Require a Change Order for IP address changes** as the description.
 - Select Change Orders Without Specifications as the alignment.
 - Enter **test*** as the CI Name and **NY** as the Location and an asterisk for all other patterns.

- Select IP Address from the Managed Attribute drop-down list.
 - Select Allow Attribute Update as the action.
5. Click Save.

This policy eliminates rogue changes because Change Orders must accompany the change.

The Configuration Manager sees that data not defined in the Change Order is also getting updated since change specifications are not required for changes at the attribute level. For example, requesting a change to increase Memory Installed while changing the IP Address at the same time, would not be considered a rogue change because there is a Change Order. They also would like the Change Orders to be automatically be verified and promoted and determine they want verification at the attribute level.

Require Change Specifications when Updating the Attribute

The Configuration Manager shows concern that data in the CMDB receives multiple updates for changes not specified in the Change Order. For example, a user changes the value of Memory Installed (phys_mem) and IP Address (ip_address) at the same time. In the previous phase, CACF does not consider this request as a rogue change because a Change Order exists.

The Configuration Manager wants to enforce change verification at the attribute level. The Configuration Administrator creates a policy that requires the use of change specifications when updating IP Address and to create an Incident if no rogue update is attempted..

Follow these steps:

1. The Configuration Manager contacts Change Analysts that changes to IP Address must contain a change specification.
2. Set Policy3.1 and Policy3.2 to Inactive to disable these policies so they are not in effect.
3. Create Policy4.1 with the following information:
 - Enter **4000** as the Sequence.
 - Enter **Require Change Specifications for IP Address changes** as the description.
 - Select Rogue Insert, Rogue Update, and Change Orders Without Specifications as the alignment.
 - Enter **test*** as the CI name and **NY** as the Location, and an asterisk for all other patterns.
 - Select IP Address from the Managed Attribute drop-down list.
 - Select Always Cancel Entire Transaction as the action.

4. Click Save.
5. Create Policy4.2 with the following information:
 - Enter **4100** as the Sequence.
 - Enter **Require Change Specifications for IP Address changes** as the description.
 - Select Change Orders With Specifications as the alignment.
 - Enter **test*** as the CI Name and **NY** as the Location pattern.
 - Select IP Address from the Managed Attribute drop-down list.
 - Select Allow Update only if Matches Change Specification as the action.
6. Click Save.

You have successfully completed the appropriate example phases to implement change verification in your environment. You can expand the change verification strategy to include more attributes, CIs, MDRs, and tenants.

Change Verification Best Practices

Consider the following best practices when implementing change verification:

- Define a small number of policies.
- Avoid the use of negative logic (exclamation point in a policy pattern).
- A single CI update save should match a small number of policies.
- Organize policies in a numerically tiered structure.
- Use Log Only Mode policies before you implement a change verification policy.
- Minimize overlapping policies, where multiple policies manage a single attribute being updated.
- Limit the number of Incidents that CACF creates.
- Implement archive and purge of CACF data.

More information:

[Organizing Policies](#) (see page 590)

[Multi-Tenancy Considerations](#) (see page 591)

[CACF Roles and Functional Access](#) (see page 591)

Organizing Policies

You can use several strategies to keep track of policies. We recommend that you use a multi-tier approach. Use increments of 100 between policy numbers to allow for future inserts.

Consider the following example tiers for this strategy:

Fundamental Policies

100,000-199,999

- Allow Change Administrators to perform inserts
- Disallow inserts by MDR Class(xxxx)
- Exclude all test CIs from incident creation

Temporary or Transient Policies

200,000-299,999: Bulk loads this week

Application-Specific Policies

- 301,000-301,999: Exceptions to the following general application-specific policies:

For example, server1 in NY policies

- 311,000-311,999: Policies that relate to Service Type priority1 resolution
- 320,000-320,999: Policies that relate to NY
- 331,000-331,999: Policies that relate to Lisle

Technology-Specific Policies

- 410,000-410,999: Policies related to Servers.
- 411,000-411,999: Policies related to Network

Default Policies

You can use these policies if the previous policies do not apply to your environment.

900,000-999,999: Default policies, such as all updates to managed attributes (Any Managed Attribute) must have a change specification.

Multi-Tenancy Considerations

Consider the following information when using CACF in a multi-tenancy environment:

- Managed attributes, verification policies, Change Orders, and change specifications are tenanted.
- Tenants up their hierarchy can view all CACF objects.
For example, if a tenant creates a managed attribute, this action may block a tenant in the hierarchy from creating a duplicate managed attribute. The hierarchy requires unique policy sequence numbers.
- To determine the policy that manages a change to an attribute, consider only policies from the tenant hierarchy of the object (CI).
- A subtenant can create a policy that overrides the supertenant by assigning a lower sequence number to their local policy.
- CACF only considers policies at the tenant level and its parents in the hierarchy.
- CACF does not consider the tenant of that contact that performs the change.
Note: Verification policies are specific to your system and do not vary by user or role.

CACF Roles and Functional Access

The following table describes the default roles that CACF uses:

Role/Functional Access	Administration	CI	Incident/Problem/Request	Change Order
Administrator	Modify	Modify	Modify	Modify
Configuration Administrator	Modify	Modify	Modify	Modify
Configuration Analyst	View	Modify	Modify	Modify
Configuration Viewer	None	View	Modify	View
Change Manager	View	Modify	Modify	Modify
Service Desk Administrator	Modify	Modify	Modify	Modify
Service Desk Manager	View	Modify	Modify	Modify
System Administrator	Modify	Modify	View	View
Level 1 Analyst	View	View	Modify	View
Level 2 Analyst	View	Modify	Modify	Modify
Incident Manager	View	View	Modify	Modify

Role/Functional Access	Administration	CI	Incident/Problem/Request	Change Order
Problem Manager	View	View	Modify	Modify

Administration

Indicates creating, updating, and viewing CACF administration, policies, and attribute management.

CI

Indicates creating, updating, and viewing CI and relationship attribute data.

Incident/Problem/Request

Indicates modifying and viewing outstanding CACF issues, including rogue changes and incorrect change execution variances.

Change Orders

Indicates creating, updating, and viewing change specifications.


For example, the Change Manager role can view CACF policies and can manage attributes, but cannot update them.

Important! Update access to the Change Order and its status determine who can edit change specifications. For example, the Change Administrator provides this access to the Change Manager.

Verify a CI Attribute Value Update Manually

The Configuration Administrator determines that your CA SDM environment requires a verification policy for the Memory Installed (phys_mem) attribute. The Configuration Administrator creates a Managed Attribute definition with the *Manual Verification will be required* status because CA SDM does not have an MDR to discover the value. The Configuration Administrator reviews the managed states and creates a verification policy. The Change Manager views a Change Order that requests an update to the Memory Installed value of a CI, and this change request requires manual verification.

The following diagram explains how a Configuration Administrator establishes a verification policy and how the Change Manager verifies a CI attribute value update manually:



1. [Create a Managed Attribute definition](#) (see page 594).
2. [Review the Managed Change states in your environment](#) (see page 594).
3. [Create a Verification Policy for the Managed Attribute](#) (see page 595).
4. [Review the Change Specifications list](#) (see page 596).
5. [Accept the Planned Value](#) (see page 597).

Create a Managed Attribute Definition

The Configuration Administrator creates a Managed Attribute definition for the Memory Installed (phys_mem) attribute.

Follow these steps:

1. On the Administration tab, click CA CMDB, Configuration Control, Managed Attributes.
2. Click Create New.
3. Enter **phys_mem** as the attribute name.
4. Enter **Memory Installed** as the Attribute Label.
5. Select *Manual Verification will be required* from the Initial Verify Status drop-down list.
6. (Optional) Select the Case Sensitive option if you want to enforce case sensitivity for change specification planned value comparisons.

Default: Disabled

7. Click Save.

The Managed Attribute is saved.

Review the Managed Change States in Your Environment

The Configuration Administrator reviews the Managed Change States in the CA SDM environment. Change verification initiates when the status of a Change Order changes to a Change State that CACF manages, such as Verification in Progress.

Note: The Configuration Administrator can customize which state in the Change Order lifecycle initiates change verification. The administrator can also customize which states allow you to modify change specifications when implementation starts.

Follow these steps:

1. On the Administration tab, click CA CMDB, Configuration Control, Managed Change States.
2. Click Implementation in Progress to open the managed status detail page.
3. Review the details about the managed status.

For example, you decide to make the Implementation in Progress status allow editing change specifications, and also enable Change Verification Active. This example indicates an active change verification process, and that you can edit values when the Change Order sets to the Implementation in Progress status.

4. Close the page.

Create a Verification Policy for the Managed Attribute

The Verification Policy specifies the CACF action when an MDR discovers an attribute value, and attempts to update the CMDB with that data.

Follow these steps:

1. On the Administration tab, click CA CMDB, Configuration Control, Verification Policies.
2. Click Create New to open the detail page.
3. Enter **1000** as the policy Sequence.
Note: You can enter a higher or lower number, based on other verification policy priorities in your environment.
4. Enter a Policy Name, such as **RAM Management** and a brief description of the policy.
5. Select Rogue Insert and Rogue Updates under Change Order Alignment.
6. If you want the policy to prevent a specific MDR from updating the Memory Installed (phys_mem) attribute, complete the following tasks:
 - Select Memory Installed as the Managed Attribute.
 - Enter a Role Pattern or use an asterisk to apply to all roles.
 - Enter a CI Name Pattern or use an asterisk to apply to all CIs.
 - Enter a Class Pattern or use an asterisk to apply to all classes.
 - Select Keep Old Attribute Value as the Update Behavior.**Note:** If you only want to allow updates to this attribute from a user logged in through the web interface, enter **Web Client** as the MDR Name Pattern.
7. Click Save.
8. Close the page.

Review the Change Specifications List

A Change Order with an RFC status wants to change the value of the Memory Installed (phys_mem) attribute in a CI named server1. The Change Manager sets the verification status of the Change Specification to *Manual Verification will be required*. The Change Manager can view the Change Specification in CA SDM.

Follow these steps:

1. The Change Manager moves the Change Order from RFC to Verification in Progress.
All the change specifications with Manual Verification become Manual Verification Active.
2. From the Change Order, select the Change Specifications tab from the Configuration Management tab.

The tab highlights all change specifications in red to indicate that manual verification is required.

Note: You can also view change specifications from the Administration tab when you click CA CMDB, Configuration Audit, Change Specifications. Additionally, view change specifications from a CI after you select the Change Specifications tab from the Related Tickets tab.

3. Search for change specifications with *requires manual verification* as the Verify Status.
For example, change specifications of Change Order 21 contain this status.
4. Research the appropriate value for the Memory Installed attribute for the server1 CI.

Note: Change specifications that failed verification also display in red to help show change specifications that require further attention.

Accept the Planned Value

The Change Manager researches the correct attribute value for the CI and makes the appropriate decision on the change request.

Follow these steps:

1. Open the Change Specification to view the detail page.
2. You research the specific Memory Installed (phys_mem) attribute for CI and confirm the planned value.
3. Verify that the Change Order of the Change Specification is in Verification in Progress State.
4. Click Accept Planned Value.

The Verify Status changes to Accepted Planned Value.

Note: After you click Accept Planned Value, the CI updates with the Planned Value automatically.

You have now successfully verified a managed attribute value manually.

Example: Allow Rogue Updates Only From a Specific Location

In this example, server CIs in your New York office require repair. The vendor that repairs your servers also resides in New York. The Asset Manager requires that your organization ships all defective servers to New York. The Configuration Administrator wants to allow rogue updates to the CIs when the hardware arrives in New York. The Configuration Administrator creates a verification policy for the Maintenance Vendor attribute so that a Service Desk Analyst in New York can verify that the vendor receives the servers.

Follow these steps:

1. Create the following managed attribute:
 - Enter **vendor_repair** as the Attribute Name.
 - Select verification pending from the Initial Status drop-down list.
 - Enter **Maintenance** as the label and description for details about the managed attribute.
2. Create the following managed change status:
 - Enter **Vendor Hold** as the Change Order Status.
 - Enable the Change Verification Active option.

3. Create the following verification policy:
 - Select Rogue Insert and Rogue Update as the change order alignments.
 - Enter a sequence number for the policy.
For example, you enter **201**.
 - Enter **server*** as the CI Name Pattern.
 - Enter **Server** as the Class Pattern.
 - Enter **New York** as the Location Pattern.
 - Select Allow Attribute Update from the Update Behavior drop-down list.
4. The Service Desk Analyst creates a Change Order and change specifications for the server CIs and specifies New York for the location.
5. The Service Desk Analyst sets the Change Order Status to Vendor-Hold.
6. The Service Desk Analyst updates the location for the server CIs as New York.
For example, defective servers from the Chicago office ship to New York, and the Service Desk Analyst verifies that the servers arrived in New York.
7. The Service Desk Analyst enters the vendor information in the CI.
The Change Order closes after all the CI locations are set to New York and all the change specifications are verified.

Example: Upgrade Laptops in Your Organization

In this example, an Asset Manager wants to upgrade all laptops from Windows XP to Windows 7 in your organization. The Configuration Administrator creates a verification policy for the Product Version attribute to filter all laptops with the Windows XP operating system. The Configuration Administrator creates a Managed Attribute definition with the *Use Discovered Value* status, reviews Managed States, and creates a verification policy. CA SDM receives updates to the CMDB from CA Configuration Automation from laptops that require an operating system update and the Asset Management team completes the upgrades.

Follow these steps:

1. Create the following managed attribute:
 - Enter **product_version** as the Attribute Name.
 - Enter **Product Version** as the Attribute Label.
 - Select Use Discovered Value from the Initial Status drop-down list.
 - Enter a label and description for details about the managed attribute.

2. Create the following verification policy:
 - Enter a sequence that you base on other policies in your organization.
For example, you enter **101**.
 - Enter **Windows* XP*** as the CI Name Pattern.
 - Enter **Workstation** as the Class Pattern.
 - (Optional) Enter a Location Pattern if you want to associate the policy with a specific office in your organization.
 - Select Allow Update Only if Matches Change Specification from the Update Behavior drop-down list.
 - Select Yes from the Create Incident drop-down list.
3. Create a Change Order and change specifications that specify Windows 7 for the Product Version managed attribute for each of the laptop CIs.
4. Move the Change Order to Implementation in Progress status and wait for your Asset Management team to start the upgrades.
5. When the implementations complete, move the Change Order to Verification in Progress.

CA Configuration Automation discovers laptop information and imports the data in the CMDB.
6. View the open CACF Incidents list for any variances that CACF detected.
7. Select one of the Incidents and review the details about the laptop that CA Configuration Automation discovered.
8. Click Create Change Order and associate the CI with the ticket, and specify change specifications for Product Version for the outstanding changes.

Wait until discovery verifies that all remaining changes completed.
9. Repeat Steps 6-8 as necessary for any new Incidents that CACF creates.

Example: Lock Down Nonverified Change Orders

In this example, the Configuration Administrator wants to only allow a CI update for a matching Change Order. Only CIs in the Server class located in NY update for verified Change Orders. Any variance creates an Incident.

Follow these steps:

1. Create a verification policy.
2. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select Any Managed Attribute as the Managed Attribute.

- c. Enter **NY** as the Location Pattern.
 - d. Enter **Server** as the Class Pattern.
 - e. Select Allow Update Only if Matches Change Specification as the Update Behavior.
 - f. Select Yes from the Create Incident drop-down list.
3. Save the policy.

Example: Allow a CI Update If No Matching Change Order Exists

In this example, the Configuration Administrator wants to allow updates for all CIs named test*, even if no matching Change Order exists. This policy accepts updates from all users in an administrative role.

Follow these steps:

1. Create a verification policy.
2. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select Any Managed Attribute as the Managed Attribute.
 - c. Enter **test*** as the CI Name Pattern.
 - d. Enter **Administrator*** as the Role Pattern.
 - e. Select Allow Attribute Update as the Update Behavior.
3. Save the policy.
4. A user with an administrative role creates a change specification for a CI named test5.

The CI updates successfully.

Example: Defer All Updates from CA Configuration Automation to the TWA

In this example, the Configuration Administrator wants to defer all CI updates from CA Configuration Automation to the TWA. This policy does not update the CI in CMDB, and writes the data for all rogue changes to the TWA for further evaluation instead.

Follow these steps:

1. Create a verification policy.
2. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select Any Managed Attribute as the Managed Attribute.
 - c. Select Rogue Insert and Rogue Update as the Change Order Alignment.
 - d. Enter **CCA** as the MDR Class Pattern.
 - e. Select Always Cancel Entire Transaction as the Update Behavior.
 - f. Select Always for the Write Data to TWA option.
3. Save the policy.

Example: Only Log the Policy Results as a Test

In this example, the Configuration Administrator wants to test a new policy before implementing it in your CMDB environment. The Log Only option lets CACF write the potential CI impacts of the policy to the standard log file.

Follow these steps:

1. Create a verification policy.
2. Complete the information for the alignment, filters, and action.
3. Select the Log Only Mode option.
4. Save the policy.
5. View the standard log file after you perform CI updates that match the policy specifications and filter criteria to simulate executing the policy.

Example: Reject a CI Update

In this example, the Configuration Administrator wants to reject updates from an MDR named Cohesion for the IP Address (alarm_id) attribute only. This policy does not update the IP address of the CI in CMDB, while it may update other attributes. CACF writes all attributes to the TWA for further evaluation.

Follow these steps:

1. Add IP Address (alarm_id) to the managed attributes list.
2. Create a verification policy.
3. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select IP Address as the Managed Attribute.
 - c. Select Rogue Insert and Rogue Update as the Change Order Alignment.
 - d. Enter **Cohesion** as the MDR Pattern.
 - e. Select Keep Old Attribute Value as the Update Behavior.
 - f. Select Always for the Write Data to TWA option.
4. Save the policy.

Example: Allow Change Orders Created Without Specifications

In this example, the Configuration Administrator wants to trust Change Orders that users created without specifications. This policy assumes that the CI text describes all changes accurately and lets the CI update.

Follow these steps:

1. Create a verification policy.
2. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select Any Managed Attribute as the Managed Attribute.
 - c. Select Change Orders Without Specifications as the alignment.
 - d. Select Allow Attribute Update as the Update Behavior.
3. Save the policy.

Example: Do Not Allow Change Orders Created Without Specifications

In this example, the Configuration Administrator does *not* want to trust Change Orders that users created without specifications. This policy ignores Change Orders without specifications and creates Incidents.

Follow these steps:

1. Create a verification policy.
2. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select Any Managed Attribute as the Managed Attribute.
 - c. Select Change Orders Without Specifications as the alignment.
 - d. Select Always Cancel Entire Transaction as the Update Behavior.
 - e. Select Yes from the Create Incident drop-down list and select an Incident template.
3. Save the policy.

Example: Allow Rogue Inserts from Selected Sources

In this example, the Configuration Administrator wants to allow new CIs from selected sources. The z/OS MDRs can create CIs without requiring a Change Order.

Follow these steps:

1. Create a verification policy.
2. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select Any Managed Attribute as the Managed Attribute.
 - c. Select Rogue Insert as the alignment.
 - d. Enter **z/OS** as the MDR Class Pattern.
 - e. Select Allow Attribute Update as the Update Behavior.
3. Save the policy.

Example: Allow a Rogue Update for a Nonproduction CI

In this example, the Configuration Administrator wants to allow updates to the IP Address attribute from the Spectrum MDR, but the name cannot begin with *PROD*.

Follow these steps:

1. Create a verification policy.
2. Complete the following steps:
 - a. Enter a sequence number.
 - b. Select IP Address (alarm_id) as the Managed Attribute.
 - c. Select Rogue Update as the alignment.
 - d. Enter **!PROD*** as the CI Name Pattern.
 - e. Enter **Spectrum** as the MDR Class Pattern.
 - f. Select Allow Attribute Update as the Update Behavior.
3. Save the policy.

Chapter 15: Administering MDRs

This chapter describes how to define MDRs, import data, map CIs back to their source, define launch parameters, launch back to the source MDR, and use MDRs to map and view federated CI information.

This section contains the following topics:

[What is an MDR?](#) (see page 605)

[MDR Launcher](#) (see page 607)

[Define a URL to Launch an MDR](#) (see page 607)

[Set Up a CA APM MDR Provider](#) (see page 609)

[Launch in Context from CA CMDB to CA APM](#) (see page 610)

[CI Properties that Support MDR Federation](#) (see page 610)

[CA Cohesion ACM MDRs](#) (see page 613)

[Using GRLoader](#) (see page 617)

[CI Naming Conventions and Restrictions](#) (see page 617)

[system_name Naming Convention](#) (see page 618)

[Using the CMDBf Viewer](#) (see page 620)

[How to Update Metadata Files for CMDBf Mapping](#) (see page 620)

What is an MDR?

The Configuration Management Database Federation (CMDBf), a working group composed of representatives from CA, IBM, HP, Microsoft, and other companies, defines a Management Data Repository (MDR) as anything that collects information about configuration items (CIs).

To create the relationship between an MDR and its CIs when implementing MDR Launcher, do the following:

1. Define the MDR.
2. Define the CIs that reference that MDR.

You cannot have a CI that references a non-existent MDR, but you can define a CI without defining a MDR association. You can add the MDR information during an update or edit to take advantage of the MDR Launcher capability.

The same CI can be discovered by multiple MDRs. After the CI is discovered, each MDR attempts to manage that CI, and an MDR can do the following actions:

- Discover detailed attributes about the CI.
- Attempt to modify the CI state.

Example: A CI Discovered by Multiple MDRs

A CI is discovered by both Network Management software and an Asset Management software package.

- The network management software maintains information about network configuration and network topology.
- The asset management software contains information regarding cost, depreciation, licensing and maintenance contracts for that particular CI.

MDR Classes and MDR Names

An IT enterprise can include many MDRs. Each MDR has an identifier called an *MDR name* (*mdr_name*). Because it is common for an MDR to use the host server name as the *mdr_name* (to allow multiple MDRs to reside on the same host server), an *MDR class* (*mdr_class*) is appended to the *mdr_name* to identify each MDR uniquely.

CA Cohesion ACM is an enterprise discovery tool that integrates tightly with CA CMDB. Each CA Cohesion ACM MDR defined to the CA CMDB must have an *mdr_class* of **Cohesion**.

Note: For additional information about CA Cohesion ACM, see the CA Cohesion ACM online help. For CA Cohesion ACM-CA CMDB integration, see the *CA Cohesion ACM Implementation Guide*.

How does an MDR Complement CA SDM?

An MDR typically contains more detailed CI information than CA CMDB. However, a single MDR is not generally aware of the existence of other MDRs, nor does it focus on the relationships that a particular CI can have with other CIs, especially if they are contained in other MDRs. CA CMDB is well-suited to managing this type of environment because it focuses on managing CIs regardless of their MDR source.

CA CMDB is not intended to store all attributes for all CIs; instead it is used to consolidate the most important attributes that must be managed centrally. Attributes that are under the control of change management are excellent candidates for inclusion in the CMDB. Attributes that CA CMDB does not manage can be accessed by using the MDR Launch capability. In addition, CA CMDB provides the CMDBf Viewer, which allows side by side comparison of that CI's attributes across multiple CMDBs and MDRs.

MDR to CA SDM Definition

The Administration tab allows you to define an MDR to the CA CMDB.

Before you can import any CI to an MDR, the MDR must be defined. Attempts to import a federated CI to a non-existent MDR fails the import.

Note: For instructions on defining an MDR, see the *Implementation Guide*.

MDR Launcher

The MDR Launcher is an open integration tool which permits you to view data from almost any MDR by using a web page, without any coding. MDR Launcher permits anyone viewing a CI to obtain additional details about the CI, and to gain control over it (if the MDR supports such control).

Some uses of the MDR Launcher include the following ones:

- From the hardware.server detail page, launch CA Cohesion ACM to verify a change.
- From an Air Conditioning CI detail, launch to a vendor web page for diagnostic information and incident reporting.
- From a Contract CI, launch a contracts management system to learn contract details.
- From an SLA CI, launch CA Service to review Service Level Agreements before making a change.
- From a Server CI, launch CA Remote Control to take over the server to diagnose and correct a problem.

Define a URL to Launch an MDR

CA CMDB uses a URL to start a web page session with the source MDR to operate the MDR Launcher. You define the URL that CA CMDB uses.

To define an URL for an MDR

1. Click the Administration tab.
2. In the left pane, open the CA CMDB, MDR Management tree.
3. Click MDR List.

The Management Data Repository (MDR) List page appears.

4. Click an existing MDR (or create and save a new one).

The MDR Provider Definition page appears.

5. Click Edit.

The Update MDR Definition page appears.

6. Complete the following parameters:

Hostname

Specifies the Network address or DNS name of the web server that serves up web pages for the MDR.

Port

Specifies the port number used by the Hostname web server.

Path

Specifies the path to the web page, including the page itself.

Parameters

Specifies the any parameters used to identify the desired CI to the MDR. CA CMDB posts this information to the MDR.

Userid

Specifies the userid, if a common userid is allowed access to the MDR.

Shared Secret

Specifies information shared between CA CMDB and the MDR. For CA Cohesion MDRs, the value specified here must match the value specified in the CA Cohesion properties file, for the com.cendura.security.oneclickauth.secret value.

CMDBf Namespace

Specifies the federated_asset_id that is passed to the query as a local ID. For CA CMDB, the value is http://cmdb.ca.com/r1.

CMDBf Timeout

(Optional) Specifies time limit for CMDBf endpoint query. Default is ten (10) seconds.

CMDBf Endpoint

Specifies the Query Service endpoint for the MDR. Required for CMDBf Viewer and retrieving updated MDR data. If you use CA CMDB as an MDR provider, the value is http://cmdb_hostname:cmdb_port/axis/services/QueryPort.

Save the definition.

The URL is defined.

Note: In addition, the URL can contain the substitution variables to further qualify the CI to the MDR. For more information, see the *Implementation Guide*.

Set Up a CA APM MDR Provider

You can set up an MDR to be the CA APM provider.

Follow these steps:

1. On the Administration tab, click CA CMDB, MDR Management, MDR List.

2. Click Create New to specify the CA APM MDR.

The MDR Provider definition appears.

3. Enter the following required MDR provider information:

- Button Name—Specify APM or any other valid button name. We recommend using a Button Name of APM.
- MDR Name—Specify APM for CA Asset Portfolio Management r11.3.4 or ITAM for CA APM r12.6
- MDR Class—Specify GLOBAL.
- Hostname—Specifies the CA APM server name by using the network address or the DNS name of the CA APM web server.
- URL for Launch in Context—Specifies `http://{hostname}:{port}/{path}?{parameters}` and must not be changed.

The MDR provider form automatically populates the path and parameter values with the required CA APM launch in context information.

4. Click Save.

The CA APM MDR provider is set up.

Note: For more information about the MDR launcher, see the *Implementation Guide*.

Launch in Context from CA CMDB to CA APM

The CA CMDB MDR Launcher facility supports launch in context to CA APM when CA APM and CA CMDB are sharing the same MDB. The CA CMDB UI provides a launch in context button on the Attributes tab in the CI detail form when the user creates a special CA APM MDR provider definition.

The CA APM MDR definition has all the capabilities of a traditional MDR. The CA CMDB Versioning feature also supports launch in context directly from an attribute log entry associated with each CA APM change.

Important! Unlike other MDRs, the CA APM MDR is automatically associated with each CI or Asset. The MDR class of GLOBAL and MDR name of APM are used to identify the CA Asset Portfolio Management r11.3.4 MDR. The MDR class of GLOBAL and MDR name of ITAM are used to identify the CA APM r12.6 MDR. Use of the CA APM MDR is fully compatible with other MDRs, even for the same CI.

CI Properties that Support MDR Federation

Configuration item properties (attributes) identify assets for MDR federation purposes.

Federated Asset ID

People are known to different organizations by different identifiers. You can be known by the following identifiers:

- A unique nickname to your close friends
- Driver's license (unique ID associated with you)
- A national service ID (for example, a selective service card)
- A health insurance number
- A national tax identification number

Each of these unique identifiers refers to you; however, the identifier is only valid when used to describe you to the appropriate repository.

Similarly, a CI can have multiple identifiers to associate it with its source MDRs. Each CI is known to an MDR by only a single identifier. We call this identifier the *federated asset ID*. The process of associating a CI with one or more MDRs is *mapping the CI*.

Mapping occurs when CIs are loaded into the MDB in one of two ways:

- Defining the CI mapping using the Administration user interface
- Loading CIs using the GRLoader utility

MDR Name

The MDR name identifies the MDR to CA CMDB when exporting data using XML and GRLoader. The MDR typically has its own naming convention for how it identifies itself: a combination of host server name plus an identifying instance name or number. Because only one MDR exists on a particular host, the MDR name is often set to the host server name. **Required for CMDBf Viewer.**

Important! The MDR Name for CA Asset Portfolio Management release 11.3.4 is APM and the MDR Name for CA APM r12.6 is ITAM. Both products are supported; however, we recommend that you verify product availability at supportconnect.ca.com before implementing the products.

MDR Class

The MDR class is defined by the customer to group MDRs.

Note: An MDR Class of CMDBf is required for CMDBf viewing.

Important! The MDR Name plus the MDR Class must be unique within your enterprise.

MDR Definition with CA Cohesion ACM Installation

CA Cohesion ACM MDRs have the following requirements:

- The `mdr_name` specified in the MDR definition on the CA CMDB server must match exactly the value of the attribute `com.cendura.installation.name` in the `cendura.properties` file on the target CA Cohesion ACM server.
- CA Cohesion ACM MDRs must have an MDR class of Cohesion.
- The MDR must specify the hostname and port number of the CA Cohesion ACM server.

To run MDR Launcher, edit the following portion of the `cendura.properties` file:

```
# -- Configure One-Click Authentication --  
com.cendura.security.oneclickauth.secret=shared_secret  
com.cendura.security.oneclickauth.scheme=  
com.cendura.security.oneclickauth.user=userid
```

Important! The secret that is specified in the `cendura.properties` file must match the *shared secret* in the MDR definition.

The MDR Launcher logs on as the *userid* specified in the properties file and inherits its security attributes. To use functionality such as Refresh for CI attributes, verify that this user has sufficient privileges.

Note: For more information about creating a user and setting security options for that user, and for customizing the properties file, see the *CA Cohesion ACM Implementation Guide*.

CA Cohesion ACM MDRs

CA Cohesion ACM MDRs must be defined before importing data from a CA Cohesion ACM server. A CA Cohesion ACM MDR should specify an MDR Class of Cohesion.

Example: Cohesion1 MDR Definition

In the following Cohesion1 MDR definition, the XML specifies an MDR Name of cohesion_server and an MDR Class of Cohesion. These values are required for the CIs to be imported.

Button Name

Cohesion1

MDR Name

cohesion_server

MDR Class

Cohesion

Active?

Active

Owner

CMDBAdmin

Description

CA Cohesion ACM server in Chicago

Hostname

cohesion_server

Port

8090

Path

CAisd/html/cmdb_cohesion.html

Parameters

hostname={hostname}+port={port}+family={family}+name={name}+secret={password}+federated_asset_id={federated_asset_id}

userid

cohesion_userid

Shared Secret

Chicago01

URL to launch in Context

http://cmdb_hostname:8080/{path}?{parameters}

In addition, because this is a Cendura server, the values above should match the values in the cendura.properties file on that server, as shown in the following example:

```
com.cendura.security.oneclickauth.secret=Chicago01
```

```
com.cendura.installation.name=cohesion_server
```

You can modify the URL syntax to handle special requirements.

How to Associate an MDR to a CI Manually

You can associate MDRs with a CI manually by using the Federated CI Mapping facility found in the CA CMDB Administration tab, in the MDR Management branch of the tree.

Before you associate a CI to an MDR, do the following:

1. Create the MDR definition (if it does not exist).
2. Create the CI definition (if it does not exist).
3. Identify the unique Federated Asset ID of the CI that you want to connect to the MDR. This ID is MDR-specific, so it is beyond the scope of this document.

Note: To identify Federated Asset ID, consult your MDR documentation.

CA Cohesion Automatic Import

You can import CIs directly from CA Cohesion ACM by running a CA Cohesion ACM report specifying the host name, port, user ID, and password to a CA CMDB server. If the MDR is defined in CA CMDB, CA Cohesion automatically generates the XML to create CIs and relationships, with the information necessary to perform the MDR Launch on a CI. The report automatically imports the CIs into CA CMDB.

Note: For more information on exporting CIs from CA Cohesion ACM, see the online help that is available from the Reports, Report Templates tab.

CI to MDR Mapping

Because each MDR uses a `federated_asset_id` to identify a CI, one CI can be related to multiple MDRs. A `federated_asset_id` does not have to be unique across MDRs, but a `federated_asset_id` must be unique within an MDR. Each MDR must have a unique MDR class and MDR name.

Important: Whenever a CI or an MDR provider is made Inactive, all Federated CI Mappings associated with the CI or the MDR provider are also made Inactive.

After you create an MDR Data Provider definition, do the following:

1. Create a CI in the CMDB that references an MDR.
2. Verify that the MDR definition works.

Because you can only launch from within the context of a CI, it is not possible to test directly from the MDR definition, which has no CI context.

You can view the Federated CI Mapping List in the Administration tab, in the Federated CI Mapping node.

To view the CI in a particular MDR context, click an MDR Launch button.

The target MDR is launched in the context of the CI that was opened.

Example: CI Mapping

1. Click the Administration tab.
2. Navigate to MDR Management, Federated CI Mapping.

The Federated CI Mapping List appears.

3. Enter `server1` in CI Name field.

The following columns show the values:

Federated Asset ID

1000234

1000235

CI Name

server1

server1

MDR Provider

Cohesion1

Cohesion2

Active

Active

Active

The Cohesion1 MDR Provider knows of the existence of server1, and the Cohesion2 MDR Provider also knows of the existence of server1. The example shows that they each have independently assigned the server a unique ID.

4. Click the CI Name server1.

The server1 Configuration Item Detail page appears that includes launch buttons would be named Cohesion1 and Cohesion2.

5. Click either MDR Provider launch button.

Additional details about the CI appear.

MDR Definition Administration

Administering MDR definitions is a flexible process. You can modify the parameters in an MDR definition even when CIs reference it. For example, you can modify the button name, host name, user ID, shared secret, and others after the MDR has been defined and the CIs are loaded.

CA Cohesion ACM Report

CA Cohesion ACM provides a facility for scheduling recurring reports. You can use this facility to simplify the process of keeping the CMDB synchronized with the data in the CA Cohesion ACM MDR. Common errors such as invalid password (due to password change or expiration) can prevent successful importation of data. To be notified of any errors that can occur during background data execution of the data import, enable the Notification option in the CA CMDB Export report in CA Cohesion ACM. An email will be sent to you that reports when an import error occurs. If you run the Cohesion report in the background as a scheduled task, we recommend that you enable the Notification option.

Note: For information on scheduling execution of the CA CMDB Export report on a regular basis, see the *CA Cohesion ACM Product Guide*.

Using GRLoader

When you use GRLoader to load a CI, you must populate the following fields in the XML for the MDR Launch to operate:

- <mdr_class>
- <mdr_name>
- <federated_asset_id>

The values provided for <mdr_name> and <mdr_class> in the XML must match the values provided in the MDR definition exactly.

Important! The MDR name and class must be defined using the Administration interface before CIs that reference the MDR can be imported. If the MDR specified in the XML is not defined, the CI is not imported.

The GRLoader supports importing CIs and CI relationships from spreadsheets in the XLS and XLSX formats. To load CI data into CA SDM, you must format the source data into XML or Microsoft Excel spreadsheets.

Note: For more information about using GRLoader to load spreadsheet data, see the *CA CMDB Technical Reference Guide*.

CI Naming Conventions and Restrictions

CIs have the following naming conventions and restrictions:

- CI Name

This is the common or display name used in all CI lists. The total length of the name must not exceed 255 characters. The CI name does not need to be unique, but we recommend that it is globally unique. In addition, for situations where the name is determined by an MDR, we recommend that MDR administrators emphasize human readability when populating this field.

- Software CIs

For third-party software, follow the same naming convention for any CIs as the names that you manually create using the Administration tab. Matching naming conventions lets CA Cohesion ACM-discovered CIs reconcile to the manually created ones. If this convention is not followed, multiple CIs can be created even when there is only one software instance.

- Use of systemname Attribute

The systemname attribute uniquely associates a *single* CI with a particular host name.

If multiple CIs are imported that specify the same systemname as an existing CI, reconciliation helps ensure that only one CI results.

Example

The following output lines show the creation of a Server CI (provider), a Software CI (dependent), and a *runs* relationship between them. The resulting relationship represents a server that runs Apache software.

CI: Name: Server1 Class: Server Systemname: Server1

CI: Name: Apache1 ON Server1 Class: Software

Relationship: Server1 runs Apache1 ON Server1

system_name Naming Convention

We recommend the following naming standards for software CIs and all MDRs so that other MDRs can integrate well with CA Cohesion ACM and each other, and for reconciling CIs properly. CA Cohesion ACM follows the same standards.

System_name

Identifies software CIs uniquely. When you define a relationship that involves a software CI, specify the same system_name as the one in that CI definition. If multiple instances of the same version of the same software are installed in the directory on the same hostname, modify the system_name to enforce uniqueness. The total length of the system_name must not exceed 255 characters; data corruption can occur if this restriction is violated.

System_name must be a unique identifier for this instance of the software on a single host.

Use the following syntax:

hostname | softwarename | version | business-application

pipe (|) character

Separates the various fields in the concatenation of the syntax to let the user use the search facility.

hostname

Specifies the hostname that contains the software.

softwarename

Specifies a common name for the software.

version

Specifies the version number of the software if available.

business-application

Specifies a unique identifier for this instance of the software on *hostname*. If the instance of the software is associated with a business application or service, the name of that service is the qualifier. When you cannot determine *business-application*, you can use the installation directory to identify the software. If the total length of this field exceeds 255 characters, use ellipses (...) to shorten the total length of this field to no more than 255 characters.

Examples: Use the UI Search Facility

You can use the UI search facility to search software CIs, as in the following examples:

Use case	Name	System_name
Find all software CIs on host	xxx	Xxx%
Find all instances of software	yyy	yyy%
Find all instances of software	yyy version 123.0	yyy% % % 123.0%

In the results list returned by the search above, the user sees only the Name field.

```
mac_address    Null. - Inappropriate for software
asset_num      Null - Inappropriate for software
serial_number  Null - Inappropriate for software
dns_name       Null - Inappropriate for software
```

Using the CMDBf Viewer

CA SDM provides the CMDBf Viewer to display the results of CI federation across MDRs. From a CI Detail page (or the CI right-click menu on the CI List), click CMDBf Viewer to see CI attributes of federated CMDBs and MDRs in parallel. On the Federated View page, you can click Retrieve to update the information from any of the federated MDRs. For better readability, CA CMDB metadata files can reconcile MDR attribute names and CA CMDB attribute names.

Note: This feature requires MDRs that support Query. You configure the MDR CMDBf Endpoints to display their results on Federated View. For more information, see the *Implementation Guide*.

If the CI does not have any federated data, the viewer displays only CA CMDB attributes.

How to Update Metadata Files for CMDBf Mapping

To improve readability of attribute comparisons, you can use CA CMDB metadata files to translate between MDR attribute names and CA CMDB attribute names. For any MDR attribute that does not have CA CMDB mapping, the Federated View displays the attribute name sent by the MDR. Metadata can be defined for CMDBf data providers to do the following:

- Display MDR attribute values using CA CMDB attribute names
- Prevent MDR provider attributes from being displayed in Federated View
- Define MDR provider attributes that do not have CA CMDB equivalents

You define metadata using the `cmdb_metadata_federation_viewer_site_attr.html` file. This file contains instructions on how to update the file. Metadata can apply to all CMDB families (common attributes) or family-specific attributes.

To map external MDR attribute names to CA CMDB labels, you update the respective `cmdb_metadata_extensionable.html` form, using the following fields in the `cmdbmetadata` macro:

- `mdr_attr` - the name of the MDR attribute to be translated.
- `mdr_name` - the name of the MDR being translated. Regular expressions are supported.

Example: Attribute mapping

The following three metadata statements defines metadata that equates the CA CMDB "phys_mem" attribute with the provider attribute "mdr_memory" for all providers named "myMdr" or starting with "MDR". In addition, "physical_memory" is equated to with "phys_mem" for all other providers.

```
<macro name=cmdbMetadata attr="phys_mem" provider_attr="mdr_memory" provider_name="myMdr">
<macro name=cmdbMetadata attr="phys_mem" provider_attr="mdr_memory" provider_name_regexp="MDR.*">
<macro name=cmdbMetadata attr="phys_mem" provider_attr="physical_memory" provider_name_regexp=".*">
```

Example: Attribute hiding

The following metadata statement hides the MDR provider "widget_cost" attribute for all providers named "myMdr".

```
<macro name=cmdbMetadata hide_provider_attr="YES" provider_attr="widget_cost"
provider_name="myMdr">
```

Example: Setting an attribute label

The following metadata statement defines an attribute name "ext_mem_capacity" using the label "External Memory Capacity" under the attributes category in the CMDBf Viewer.

```
<macro name=cmdbMetadata attr="ext_mem_capacity" category="Attributes" heading="External Memory
Capacity" help="Total external memory">
```

More information:

[How To Display MDR Attribute Values With CA CMDB Attribute Names](#) (see page 622)

[How To Hide MDR Provider Attributes](#) (see page 623)

[How To Define MDR Attributes Without CA CMDB Equivalents](#) (see page 623)

[Define CMDBf Data Provider Metadata](#) (see page 624)

How To Display MDR Attribute Values With CA CMDB Attribute Names

Metadata can create an association between MDR provider attributes and CA CMDB attributes so they are displayed together to see differences and share labels. By default, MDR attributes that do not have a mapping, are displayed as **Not in CMDB** in the viewer.

cmdbMetadata macro arguments to equate a CA CMDB attribute with an MDR provider attribute include:

- attr – CA CMDB attribute name
- provider_attr – MDR provider attribute name
- provider_name – MDR provider name
- provider_name_regexp – MDR provider name regular expression

provider_name or provider_name_regexp are required.

Example: Associate MDR Attributes with CA CMDB Attribute Names

The following three metadata statements do these respective actions:

- Equate the CA CMDB "phys_mem" attribute with the provider attribute "mdr_memory" for all providers named "myMdr".
- Equate the CA CMDB "phys_mem" attribute with the provider attribute "mdr_memory" for all provider names starting with "MDR".
- Equate "physical_memory" with "phys_mem" for all other providers.

```
<macro name=cmdbMetadata attr="phys_mem" provider_attr="mdr_memory" provider_name="myMdr">
```

```
<macro name=cmdbMetadata attr="phys_mem" provider_attr="mdr_memory" provider_name_regexp="MDR.*">
```

```
<macro name=cmdbMetadata attr="phys_mem" provider_attr="physical_memory" provider_name_regexp=".*">
```

How To Hide MDR Provider Attributes

Some MDR provider attributes do not need to be displayed in the Federated View. Metadata from an MDR provider can be hidden for a particular MDR provider. This option only applies for MDR provider attributes and does not apply for CA CMDB attributes.

cmdbMetadata macro arguments to hide an provider attribute include:

- `hide_provider_attr` – "YES" – hides the MDR provider attribute
- `provider_attr` – MDR provider attribute name
- `provider_name` – MDR provider name
- `provider_name_regexp` – MDR provider name regular expression

`provider_name` or `provider_name_regexp` are required.

Example: Hide an MDR-Only Attribute

The following metadata statement can be used to hide the MDR provider "widget_cost" attribute for all providers named "myMdr".

```
<macro name=cmdbMetadata hide_provider_attr="YES" provider_attr="widget_cost" provider_name="myMdr">
```

How To Define MDR Attributes Without CA CMDB Equivalents

You can define labels and help text for MDR provider attributes that do not correspond with any CA CMDB attributes. Attributes are labeled Not in Family in the Federated View.

cmdbMetadata macro arguments to define an MDR provider attribute include:

- `attr` – CA CMDB attribute name
- `category` - Category name where attribute is displayed
- `heading` - Heading label for attribute
- `help` - Brief description of the attribute

Example: Define an MDR-Only Attribute

The following metadata statement defines an attribute name "ext_mem_capacity" using the label "External Memory Capacity" under the Attributes category in the Federated View .

```
<macro name=cmdbMetadata attr="ext_mem_capacity" category="Attributes" heading="External Memory Capacity" help="Total external memory">
```

Define CMDBf Data Provider Metadata

You can control how data is displayed in the Federated View.

To define metadata for CMDBf Viewer

1. Use Web Screen Painter to open the `cmbd_metadata_federation_viewer_site_attr.html` file.
2. Determine which metadata changes are needed.
Note: Sample templates are provided in the file.
3. Copy the appropriate template and substitute the required arguments according to the instructions in the file.
4. Save and publish the changes.

Chapter 16: Managing Change

This section contains the following topics:

[Change Management in CA SDM](#) (see page 625)
[Change Management Components](#) (see page 626)
[View the Change Calendar](#) (see page 627)
[CAB Responsibilities](#) (see page 627)
[Change Manager Responsibilities](#) (see page 629)
[Define Tasks for the Change Manager Role](#) (see page 631)
[Change Categories, Status, and Risk Levels](#) (see page 632)
[View the Change Order Scoreboard](#) (see page 633)
[Define a Change Order Stored Query](#) (see page 633)
[Configure Change Manager Options](#) (see page 634)
[Change Calendar](#) (see page 635)
[How to Schedule Change Orders](#) (see page 647)
[How to Schedule Change Windows](#) (see page 652)
[Conflict Analysis and Collision Detection](#) (see page 656)
[CA Workflow Visualization](#) (see page 656)
[Change Management Process Definition for CA Workflow](#) (see page 658)
[CAB Console and Reporting](#) (see page 676)
[Risk Assessment](#) (see page 681)
[Impact Explorer](#) (see page 685)

Change Management in CA SDM

CA SDM Change Management is a set of features for Change Managers, Change Coordinators, and Change Advisory Board (CAB) members to coordinate the review and approval of change requests for CI components and services. For example, Change Managers can review and approve all changes to CI components and services to ensure that no new security vulnerabilities are introduced into the production environment. The Change Manager leads the CAB and is responsible for the ultimate approval of change requests.

Change Management includes the following functions:

- Follow ITIL processes relating IT changes to Incident/Problem management.
- View Service Management information for a CI; for example, the number of changes logged on a CI, change implementation dates, and more.
- Detect collisions when multiple change orders call for changes to the same CIs at the same time.
- Create and store change windows in the Change Calendar.
- Create a change order in calendar view.

- Assess the risks associated with a change.

Note: For more detailed information about change management procedures, see the Change Management information in the *Online Help*.

Change Management Components

Change Management includes the following components:

Change Calendar

Displays a [graphical view of change events](#) (see page 635), including all scheduled, failed, and in-progress change requests for CIs and services in a configurable [calendar view](#) (see page 647). The calendar also displays black-out dates, which are freeze periods. Users can create a change order from the context menu on the daily, weekly, and monthly calendar views. The Change Manager, Level 2 Analysts, and CAB members use this feature.

Change Scheduler

Displays [scheduled time periods](#) (see page 653) during which CI changes can or cannot occur in the Change Calendar. The Change Manager uses this feature to view and create change schedules and CI associations during the time period.

Conflict Analysis and Collision Detection

Analyze change orders to help identify potential implementation conflicts that could increase risk of failure. The Change Manager uses this feature.

Workflow Visualization

Displays the approval process for change requests. The Change Manager uses this feature.

CAB Console and Reporting

Displays a [dashboard](#) (see page 676) that facilitates quick online approvals of change orders that require CAB approval.

In CA SDM, the Change Manager can generate a report with details about the proposed requests for changes ready for CAB review and electronically distribute the reports to all CAB members.

Note: Users with appropriate privileges can display Compliance, Forecast, Trend, and Volume predefined reports for change management in BusinessObjects InfoView with CA Business Intelligence.

Risk Assessment

Provides the ability to attach risk assessments for each change request submitted. The Risk Survey Option lets you create surveys to evaluate risks, and associate them with change categories. The risk survey lists a series of single or multiple choice questions.

Impact Explorer

Displays CI-to-CI relationships for CIs that are attached to a change order. Impact Explorer is launched from the Change Order Detail page to facilitate on-demand [impact analysis](#) (see page 685). The Change Manager uses this feature.

Change Verification

Verifies that changes execute accurately and no unauthorized changes occur, so that CMDB contains an accurate and current representation of all managed CIs.

View the Change Calendar

CA SDM can launch various change management features (in compliance with ITIL and COBIT standards) to view or create change orders after selecting a day on the *Change Calendar*. New change orders are automatically scheduled for the selected day.

The Change Calendar reports conflicts, mismatches, and potential risks among change schedules, change orders, and CIs. You can use the calendar to understand upcoming changes, their potential impacts on the IT environment, and much more.

To view the Change Calendar

1. Click the Change Calendar tab.
The Change Calendar filter fields appear.
2. Complete the search fields to display a calendar view showing the schedule for the change orders of interest.
3. Click Search.
The change order calendar displays information that matches your search criteria.

CAB Responsibilities

The Change Advisory Board (CAB) coordinates the review and approval or rejection of change requests for CI components and services. CAB members are responsible for the following:

- Reviewing all major changes to production systems.

- Attending all relevant CAB meetings as required by the Change Manager.
- Reviewing all submitted requests for changes to determine their impact, resources required to implement them and any ongoing costs.
- Participating in scheduling and coordination of the Change Calendar.
- Helping to ensure that all changes are adequately assessed and prioritized.
- Participating in reviews after the installation is complete.

More information:

[How the CAB Process Works](#) (see page 628)

[Assign Members to the CAB Group](#) (see page 629)

How the CAB Process Works

The CAB is responsible for reviewing all major changes to production systems. CAB members are notified that a change request requires their approval, and they do the following:

1. List change requests that the CAB must review.
2. Open the CAB Console and view the information contained in the request.
3. View the business justification, implementation plan, configuration items, and supporting documentation associated with the request and decide to approve it or reject it.
4. Approve or reject the change order, and the next change order in the list appears automatically.

The person requesting that the change is notified automatically that the CAB status is updated for the change request.

More information:

[CAB Responsibilities](#) (see page 627)

Assign Members to the CAB Group

The Members Update page lets you assign members to a CAB group.

To assign members to the CAB group

1. On the Group Detail page, select the Members tab.
2. Click Update Members.
The Contact Search page displays.
3. Enter the search criteria to display the desired contacts and click Search.
The Members Update page displays, listing the contacts that matched the search criteria.
4. From the list on the left, select the contacts you want to assign to this group. To select multiple items, hold down the CTRL key while clicking the left mouse button.
5. When you have selected all the contacts you want, click the selection button (>>).
The selected contacts move to the Members list on the right.
6. Click OK.
The Group Detail page displays, with the selected contacts listed on the Members tab.

Change Manager Responsibilities

The Change Manager is responsible for the overall enterprise change management process and for the ultimate approval of change orders. They also generate the change management metrics analysis reports and they do the following:

- Review change requests and add appropriate stakeholders and approvers as needed.
- Facilitate resolution of issues, such as detecting collisions and scheduling conflicts in the calendar.
- Review installation, back-out, and fallback plans for accessibility and soundness.
- Understand the risk of each change and ensure that the appropriate risk level is assigned to the change.
- Monitor changes for their respective areas to ensure that they comply with Technology Change Management requirements.
- Represent their respective areas and communicate the impact of high-level risk changes at weekly CAB meeting.

- Facilitate in reviews after the installation is complete for problem installations and failed changes.
- Serve as the escalation point for change requesters, stakeholders, approvers, implementers, and support groups.

More information:

[How the Change Manager Role Works](#) (see page 630)

How the Change Manager Role Works

Change Managers are responsible for monitoring change orders to the operational environment to ensure that the operation team members comply with business policies and processes. The Change Calendar can facilitate resolution of issues such as change order conflicts by scheduling black-out dates and freeze periods during which a CI or a set of CIs can be changed. Change Managers also do the following:

1. Oversee the CAB Console from which online and quick approvals of change orders and requests for changes are managed.
2. Organize CABs with members appropriate for the change orders under consideration and conduct regularly scheduled CAB meetings to review incoming change orders.
3. Create reports with details about the proposed requests for changes ready for CAB review and electronically distribute the reports to all CAB members.
4. Perform a real-time review of each request for change and update the record with the CAB decision during the CAB meeting.
5. Use BusinessObjects InfoView to manage Compliance, Forecast, Trend, and Volume reports and create on-demand reports.
6. Represent their respective areas and communicate the impact of high-level risk changes at CAB meetings.
7. Evaluate the risk of each change and ensure that the appropriate risk level is assigned to the change.

Define Tasks for the Change Manager Role

You can define tasks for the Change Manager role.

To define tasks for the Change Manager role

1. On the Administration tab, select Security and Role Management, Role Management, Role List.
The Role List page appears.
2. Select the Change Manager role.
The Change Manager Role Detail page appears.
3. Click Edit.
The Change Manager Update Role page appears.
4. Use the following tabs and fields to configure tasks and access permissions for the Change Manager role:
 - Authorization
 - Function Access
 - Web Interface
 - Knowledge Management
 - KT Document Visibility
 - Tabs
 - Report Web Forms
 - Go Resources
 - Tenant Read Access
 - Tenant Write Access
5. Click Save, Close Window.
The Change Manager role record is updated.

Note: For more information about the tabs that appear on the Change Manager Role Detail page, see the Role Management information in the *Online Help*.

Change Categories, Status, and Risk Levels

You can define how change orders operate within your service environment. You can edit the default values that are installed with CA SDM, or define your own.

To manage Change Order default values

1. Select Service Desk, Change Orders on the Administration tab.
2. Expand the Change Order node and select *one* of the following:
 - Categories
 - Change Types
 - Closure Codes
 - Conflict Status
 - Risk Level
 - Risk Survey
 - Status
 - Workflow Task Status Code
 - Workflow Task Types

The List page for the selected item appears.

3. Select the item to edit.

The Update Details page appears.

4. Use the controls available on the tabs at the bottom of the page to define how change orders operate within your environment.
5. Click Save, Close Window.

The updated item appears in the list.

View the Change Order Scoreboard

The Change Order scoreboard shows the change orders, conflicts, and scheduled tasks that are assigned to Level 2 Analysts, Change Managers, Change Coordinators, or CAB members. Users can view their assigned and unassigned records by priority.

To view the Change Order scoreboard

1. Navigate to Change Orders on the CA SDM scoreboard.
2. Expand the folders to reveal nested folders for the following:
 - Assigned or unassigned Open and Closed items
 - Resolved or unresolved conflicts
 - Scheduled tasks that start today or next week.
3. Select the folder for the items you want to see.
The List page appears.
4. (Optional) Click Show Filter and complete one or more of the fields to specify search criteria that restrict the list to comments of interest.
5. Click Search.
The List page displays summaries of the items that match your search criteria.
6. (Optional) Click the Edit in List button to display some additional fields that can be associated with an item.

Define a Change Order Stored Query

Defining the stored queries that are available to users on the Change Order scoreboard is an administrative task. You can modify the predefined stored queries that are installed with CA SDM, or define your own.

To define a Change Order stored query

1. Select Service Desk, Application Data, Stored Queries on the Administration tab.
The Stored Query List appears.
2. Select the stored query you want to edit.
The Stored Query Detail page appears.
3. Click Edit.
The Update Stored Query page appears.

4. Edit the field values as appropriate.
5. Click Save, Close Window.

The updated stored query appears in the Stored Query List.

Example: Define a Stored Query to List Change Orders Assigned to a CAB to Which the Logged In User Belongs

This example demonstrates how you can create a stored query that lists only change orders that are assigned to a CAB to which the logged in user belongs.

To create the stored query

1. Navigate the Scoreboard to the Update Stored Query page.
2. Edit the field values as follows:
 - a. Select Scoreboard Usage.
 - b. Set Type to Change Order.
 - c. Enter the Where clause:

`cab.[group]group_list.member IN (@cnt.id) AND active = 1`

3. Click Save, Close Window.

The stored query appears in the Stored Query List.

More information:

[CAB Console and Reporting](#) (see page 676)

Configure Change Manager Options

You can configure options for the Change Manager role.

To configure Change Manager options

1. Select Options Manager on the Administration tab.
2. Expand the Change Order Mgr node.

The Option List appears.
3. Right-click the option you want and select Edit from the context menu.

The Update Option page appears.
4. Edit the option as appropriate.
5. Click Save, Close Window.

The updated option appears on the Options list.

Change Calendar

The *Change Calendar* provides a graphical view of change events with implementation start and end times. This calendar view of the change schedule provides analysts and managers a quick way to identify when events occur and how they affect the environment, organization, and resources.

The calendar lets you create change orders from the daily, weekly, and monthly views and also lets you view global change windows for scheduled change orders. If multiple change windows exist in a month, they are grouped together, such as in a single Blackout Window group. Drill down to a weekly or daily view, or view hover information to see the individual window details.

Note: You can only [export](#) (see page 636) Change Order schedules, not Change Windows.

More information:

[Add Schedule Information to a Change Order](#) (see page 635)

[iCalendar Event Templates](#) (see page 636)

[Export Schedules to iCalendar](#) (see page 636)

[Schedule Views](#) (see page 637)

[Scheduling View Configuration](#) (see page 639)

Add Schedule Information to a Change Order

You can add schedule information when you create or edit a change order.

To add schedule information

1. Do one of the following actions:
 - Click File, New Change Order.
 - Open a change order and click Edit.
2. Complete the following schedule fields:

Type

Specifies the ITIL change type as Standard, Normal, or Emergency. A default value may be defined for a change category that CA SDM uses to initialize new change orders of the category.

Schedule Start Date

Specifies the start date of the change order.

Duration

Specifies the duration of the change order, in 00:00:00 format.

The schedule contains only those change orders with a non-empty schedule start date and duration. Type is useful to group change orders on the schedule, but does not directly affect scheduling.

3. Continue creating the change order.
4. Click Save when you finish.

The schedule information is added to the change order.

iCalendar Event Templates

iCalendar event templates control the information that is exported to iCalendar format.

The following predefined templates are installed with CA SDM:

- Change Schedule
- KnowledgeScheduleCreation
- KnowledgeScheduleExpired
- KnowledgeScheduleReview
- KnowledgeScheduleStart

Note: You can edit the predefined iCalendar event template codes, but you cannot delete them or create new ones.

Important! The SchedExpMaximum variable in web.cfg controls the maximum events allowed for an export. Increasing the default (1000) could cause system instability. If you attempt to export more than the value specified in SchedExpMaximum, a message appears refusing your exporting request.

Export Schedules to iCalendar

CA SDM lets you export Change Orders in the standard iCalendar format. This data exchange lets you import Change Order schedules into many widely used calendaring applications, including Microsoft Outlook and Lotus Notes.

Note: When exporting schedules on some calendaring programs, choosing the Open option instead of Save causes the file to import incorrectly. To avoid this issue on Knowledge Management and Change Order schedules, select the Save option instead of Open. After you save the exported file, import it through the interface of the calendar program interface. You cannot export Change (blackout and maintenance) Windows, only Change Orders.

Important! The data you export is based on the current view. If you want to export a custom range of dates, such as 32 days, the export should be done from the list view. Otherwise, the view is truncated to a month or week, and it only exports that amount.

To display the Change Calendar

1. Click the Change Calendar tab.

The Change Calendar filter fields appear.

2. Complete the search fields to display a calendar or list view that includes the change orders of interest.
3. Click Export.

The Schedule Export page appears.

Important! The SchedExpMaximum variable in web.cfg controls the maximum events allowed for an export. Increasing the default (1000) could cause system instability. If you attempt to export more than the value specified in SchedExpMaximum, a message appears refusing your exporting request.

4. Enter the location where you want to save an iCalendar file.

An iCalendar file containing all events in the current view is saved at the specified location.

Schedule Views

CA SDM provides the following scheduling views:

List

Displays a list page sorted by schedule start and end date.

Month

Displays a calendar for a full month.

The view shows change orders in groups, with each entry collecting one or more change orders. You can see detailed information about the change orders in a group by hovering over the group with the mouse; by pressing Alt+Right Arrow when the focus is on the group; or by clicking on the group to display its contents in an n-day view.

Week

Displays a full week in a single column, starting with the day configured as the first weekday.

The view shows changes individually and includes detailed information about each change order scheduled during the week.

Day

Displays change orders for a single day.

The view shows changes individually and includes detailed information about each change order scheduled during the day.

***n* Days**

Displays change orders for the number of days specified by the dropdown selector.

The view shows changes individually and includes detailed information about each change order scheduled during the days selected.

Navigating the Schedule Views

You can use the arrow and tab keys to navigate the scheduling views. The following keyboard shortcuts are available:

Tab

Navigates to a later date. Use this from the last cell in the schedule to navigate to the Search button.

Shift+Tab

Navigates to a previous date. Use this from the first cell in the schedule to navigate to the Next Month button.

Shift+Arrow

Navigates around the calendar from date to date in the direction of the arrow.

Right Arrow

Displays the context menu for the currently focused date or event. If there is no context menu, it navigates to the next higher date (similar to Shift+Right Arrow).

Alt+Right Arrow

Displays a hover information popup for the currently focused date or event. If there is no hover information, it navigates to the next higher date (similar to Shift+Right Arrow).

Down Arrow

Navigates to the next event in the current cell. If there are no events in the current cell, or if the focus is already on the last event, it navigates to the date in the next cell down (similar to Shift+Down Arrow).

Up Arrow

Navigates to the previous event in the current cell. If focus is not on an event, it navigates to the date in the next cell up (similar to Shift+Up Arrow).

Calendar View Hotkeys

Each schedule view supports hotkey access to its buttons. The following are the supported hotkeys:

Alt+0

Switch to list view.

Alt+1

Switch to daily view.

Alt+7

Switch to weekly view.

Alt+3

Switch to monthly view.

Alt+9

Switch to *n*-day view.

Alt+<

Move to previous time period in current view.

Alt+>

Move to next time period in current view.

Scheduling View Configuration

You configure the monthly and weekly scheduling views by specifying `pdm_macro` statements in the `<head>` section of the HTML forms defining the schedule. We recommend using the Source View of Web Screen Painter to edit these forms.

Any form that displays a schedule must contain the following:

- A `schedConfig` macro
- At least one `schedAttr` macro
- At least one `schedGroup` macro

The configuration macros are in a separate source file referenced by a `pdm_include` statement in the main source file. This file lets you configure your schedule without modifying the main source file.

For example, the configuration macros for the Change Calendar form `list_chgsched.html` are in a file named `list_chgsched_config.html`. For the Knowledge Lifecycle Schedule, you can modify the `list_kdsched_config.html` using the same macros.

You can find `list_chgsched_config.html` and `list_kdsched_config.html` in the following directory:

`$NX_ROOT\bopcfg\www\html\web\analyst\`

schedConfig Macro—Configure Schedule

The `schedConfig` macro specifies that a form contains a schedule and provides basic configuration information. The following values are valid macro arguments:

autosearch=1|0

Specifies whether the schedule form reloads data from the server when the user selects a view outside the currently selected date range. Setting the value to 1 (default) causes the form to search automatically when the user selects a view with one or more days outside the date selection range of the search filter. Setting the value to 0 requires the user to press the Search button to initiate a search.

defaultView=0|1|7|30|99

Specifies the default view for the search filter as 0 (list), 1 (day), 7 (week), 30 (month), or 99 (n-day).

The specification for `defaultView` affects only the initial display of the search filter. After the schedule displays, CA SDM automatically keeps the filter view selection aligned with the current view.

Default: 30

firstday=0|1|2|3|4|5|6|7

Specifies the first weekday on the monthly view as a number between 0 (Sunday) and six (Saturday).

Default: 0

export=xxx|0

Specifies the code name of the template used for exporting in iCalendar format. Setting the value to 0 indicates the export feature and button are disabled.

Default: ChangeSchedule

legend=1|2|0

Specifies the location of the schedule legend showing the name and formatting of the groups on the schedule. You can set the value to 1 to position the legend above the schedule, or 2 to position the legend below the schedule. Set the value to 0 to disable the legend.

Default: 2

maxGroups=0/n

Specifies the maximum number of groups to be displayed in a single cell of the calendar month view.

If there are more than maxGroups scheduled for a single day, CA SDM displays only the first maxGroups-1, and replaces the last with a "..nn more changes" hyperlink that the user can mouseover or click to see the full list. Set the value to 0 to disable this feature and allow an unlimited number of events in a calendar cell.

Default: 4

nday=(n,n,...)

Specifies selections for the drop-down list for the n-day view.

The specification is a list of day counts that are to be included in the drop-down list, or 0 to indicate that the n-day drop-down list is omitted from the schedule. The first value specified is the default for the drop-down list.

Default: (3,7,14,28)

round=(hr,min)|0

Specifies whether schedule start and end dates are rounded when collecting change orders or knowledge documents into groups. Specify round=0 to disable rounding.

By default, schedule start and end date groups objects. All CA SDM dates include a time, and without rounding, objects scheduled as little as a minute apart would be in separate groups. Rounding determines the group after adjusting the start date to an earlier hour or minute and the end date to a later hour or minute.

The value of round specifies either an hour or a minute (but not both). Times are rounded to the nearest multiple of the value specified, for example:

round=(0,15) rounds to the nearest quarter hour

round=(0,30) rounds to the nearest half hour

round=1 rounds to the nearest hour

round=12 rounds to the nearest half day (12:00 AM or PM)

round=24 rounds to the nearest day

Default: (0,15)

timefmt=24hr|([am],[pm])

Specifies the format of times on the calendar views of the schedule.

The default value of 24hr specifies that times are displays in 24 hour format (0:01 - 23:59). The alternative value of (am,pm) specifies a suffix for either morning or afternoon times, or both.

Note: All schedConfig arguments are optional.

schedAttr Macro—Specify a Stored Attribute

The schedAttr macro specifies an attribute stored for each item selected for the list. Stored attributes are available for hover information on the monthly view, for the detailed or summary information in views other than the monthly view, and in the setSchedEvents() JavaScript function. The following values are valid macro arguments:

attr=xxxx

(Required) Specifies an attribute from the object on the schedule, such as a change order or Knowledge Document. Dotted attributes are permitted. The keyword attribute name *Clnn* can be used on the Change Calendar to specify that first *nn* Cls associated with the change order are included with the information stored.

Note: This argument is the only required argument for the schedAttr macro.

attrRef=.COMMON_NAME|xxxx

Stores the attribute of the referenced table stored for an SREL attribute (ignored for non-SREL attributes). The attribute name specified must be prefixed with a dot.

Default: .COMMON_NAME

label=

Displays a label for the attribute on the n-day view.

Default: the Majic DISPLAY_NAME of the attribute

ident=1|0

Specifies whether the attribute is an identifier for the object (such as a reference number of a change order). Identifier attributes are displayed without a label in front of the group name in hover information and the n-day view.

Default: 0

detail=1|0

Specifies whether the attribute is included in the detail information shown on views other than the monthly view. Detail information is the information shown when the Summary Only check box on the view is not selected.

Default: 1

hoverInfo=1|0

Specifies whether the attribute is included in the hover information pop-up displayed on the monthly view when the mouse pointer hovers over a group, or the user presses Alt+Right Arrow when focus is on the group.

Default: 0

summary=1|0

Specifies whether the attribute is included in the detail information shown on views other than the monthly view. Detail information is the information shown when the Summary Only check box on the view is not selected.

Default: 0

Note: CA SDM displays attributes in summary, detail, or hover information in the same order as their schedAttr macros.

schedGroup Macro—Specify an Event Group

The schedGroup macro specifies the name and color coding of a group of items. The monthly view aggregates all items in a group into a single event. Views other than the monthly show individual items in the format for the group to which they belong. The following optional values are valid macro arguments:

grpname=xxx

(Required) Specifies the name of the group. The macro automatically assigns a number to the group and assigns the number to a JavaScript variable with a name of the form schedGroup_xxx, where xxx is the name of the group. This variable can be used in the setSchedEvents() JavaScript function to create an event belonging to the group.

Note: This argument is the only required argument for the schedGroup macro.

label=xxx

Specifies a label for the group. If specified, the label is displayed in all views.

legend=xxx|0

Displays a description of the group for the legend that appears at the bottom of the schedule. Groups appear in the legend if at least one example of the group exists in the current view. Specifying 0 causes the group always to be excluded from the legend.

Default: 0

color=black|color

Specifies the color of text in items of this group. You can specify color in CSS format, either a valid web color or a hexadecimal value preceded by a pound sign.

Example: Enter either "#FF0000" or "red" for red.

Default: black

bgcolor=white|color

Specifies the background color of items of this group. You can specify bgcolor in CSS format, either a valid web color or a hexadecimal value preceded by a pound sign.

Example: Enter either "#FF000" or "red" for red.

Default: white.

style=normal|bold|italic

Specifies the style of text of this group in the normal, bold, or italic style.

Default: normal

Configure Blackout and Maintenance Windows on the Change Calendar

You edit PDM_MACRO statements in `list_chgsched_config.html` to modify colors, labels, legends, and icons for change windows.

Note: If you use `schedGroup` macros to change the formatting of windows on the Change Calendar, you must also update `schedule.css` if you want the formatting to be consistent with the Change Scheduler.

To configure change windows

1. Open the `list_chgsched_config.html` form in a text editor or WSP.

You can locate this file in the following directory:

```
$NX_ROOT\bopcfg\www\html\web\analyst\
```

2. On the `schedGroup` macro, modify the following PDM_MACRO statements:

For maintenance windows:

```
<PDM_MACRO NAME=schedGroup grpname=maintWindow  
  bgcolor=lightgreen  
  label="Maintenance"  
  legend="Maintenance Window"  
  icon="confirmation_12.png">
```

For blackout windows:

```
<PDM_MACRO NAME=schedGroup grpname=blackoutWindow style=italic color=white  
  bgcolor=black  
  label="Blackout"  
  legend="Blackout Window"  
  icon="warning_12.png">
```

bgcolor

Specifies the background color of the window.

label

Specifies the window label as Blackout or Maintenance.

legend

Specifies the text of the legend as it appears on the Change Calendar.

icon

Specifies an optional URL to a 12x12 pixel web graphic.

This icon displays with a change order or group on the Change Calendar if the change order lies completely within a maintenance window or overlaps a blackout window.

3. Save the form.

The change windows are configured.

The setSchedEvents() JavaScript Function

The setSchedEvents() JavaScript function creates events in the schedule. Modify this function when you want to view any new group objects. The predefined group objects appear by default.

CA SDM calls setSchedEvents() once for each object (change order or knowledge document) selected by the schedule search filter. The function creates events for the object by calling a second function, schedEvent(), and passing the group ID, start date, and end date of the event.

The function can create any number of events (including zero) for an object. The default setSchedEvents() function for the Change Calendar (list_chgsched.html) creates one event for each change order and groups change orders by change type. This function is coded as follows:

```
1.      function setSchedEvents( chg )
2.      {
3.      var grpnum;
4.      switch( chg["chgtype"] - 0 ) {
5.          case 100: grpnum = schedGroup_std; break;
6.          case 300: grpnum = schedGroup_emer; break;
7.          default: grpnum = schedGroup_norm; break;
8.      }
9.      chg.schedEvent( grpnum, chg["sched_start_date"], chg["sched_end_date"] );
10.     }
```

The case parameter specifies the change type ID. To list the case IDs, see [Create a Change Type](#).

The function has a single argument of a JavaScript object containing the attributes specified by schedAttr macros. The switch statement in lines 4-8 examines the chgtype attribute of the change order, and assigns the appropriate group number from one of the schedGroup_xxxx variables defined by previous schedGroup macros. On line 9, it calls the schedEvent() function to create an event in the schedule, passing the group number previously assigned and the schedule start and end dates. The dates are available in the argument object because they were specified in earlier schedAttr macros.

How to Schedule Change Orders

You can view the schedule of all configuration items associated with a change order. Consider scheduling information when you create, edit, or view a change order, or when you update the Change Calendar. Viewing the schedule can help you avoid scheduling collisions.

To view and update a schedule, do the following:

1. Create or open a change order.
2. Associate CIs with the change order.

Note: If you do not associate a CI with the change order, an error appears when you click the Scheduler button.

3. Click View Scheduler.

The Change Scheduler appears and displays the following views:

Daily

(Default) Displays the time period of the change order for the selected date and time. If no date is associated with the change order, the view defaults on the current day.

Weekly

Displays the time period of the change order in the current week and includes the implementation start date, which can be configured to the Change Calendar start date.

Note: The schedule duration determines the length of shading on the schedule. For example, if you create a change order with a two hour duration, the light-yellow shading on the scheduler displays for the two hour defined duration.

4. (Optional) Click a CI to view its details. Hover over the CI name to view its full name.
5. (Optional) Modify the Implementation Start Date or Duration.
 - a. Click View Schedule to preview the updated schedule.
 - b. Click Update Schedule to update the schedule.

Note: You can only modify the schedule in edit mode.

6. Save the change order.
7. Refresh the Change Calendar to view global and CI-associated change windows.

Use the Change Scheduler Example

The following example demonstrates how to use the change scheduler when creating a change order.

1. On the Service Desk tab, select File, New Change Order.
The Create New Change Order page appears.
2. Select Update CIs on the Config. Items tab.
The Configuration Item Search page appears.
3. Create or search for CIs.
4. Using the Affected Configuration Items Update page, add CIs to the change order.
5. Click OK.
The Create New Change Order page updates.
6. Click the Scheduler button.
The Schedule for Change Order page appears.
7. Select a view and do *any* of the following:
 - Click within the table cells to modify the Schedule Start Date.
 - Click Schedule Start Date to use the Date Helper.
8. Modify the Duration and click View Schedule.
The schedule previews your changes.
9. Click Update Schedule.
The schedule updates with your changes.
10. Save the change order.

Set the Maximum Number of CIs

You can set the maximum number of CIs that display in the change scheduler.

To set the maximum number of CIs

1. Open *NX.env*.

You can locate this file in the following directory:

`$NX_ROOT\`

2. Modify the value of `@NX_CHANGE_SCHEDULER_MAX_CI_CNT` to the maximum number you want.

Note: By default, the variable is set to 100. If you set the maximum number of CIs greater than 100, only the first 100 display and you get a warning.

3. Save *NX.env* and cycle CA SDM.

The maximum number of CIs is set.

Example: Set the Maximum Number of CIs to 25

With the following setting, only the first 25 CIs will display on the change scheduler.

`@NX_CHANGE_SCHEDULER_MAX_CI_CNT=25`

Modify Change Order Duration Background Color

You can change the background color of the change order duration. Modify the following lines in *schedule.css* to change the highlighting of the green change order duration bar.

To change the background color

1. Open *schedule.css*.

You can locate this file in the following directory:

`$NX_ROOT\bopcfg\www\wwwroot\css`

Important! `$NX_ROOT\sdk` also contains a file called *schedule.css* that contains helpful comments about css controls.

2. Modify the following lines with the color code you want:

```
td.noBorderBackColor{border-left:none;border-right:none;background-color:#F6E3CE;}
td.withBorderBackColor{border-right:none;border-left:1px solid;background-color:#F6E3CE;}
```

3. Save *schedule.css*.

In the following example, the change duration background color is `#FF0000`.

Example: Modify the Change Duration Background Color to `#FF0000`

```
td.noBorderBackColor{border-left:none;border-right:none;background-color:#FF0000;}
td.withBorderBackColor{border-right:none;border-left:1px solid;background-color:#FF0000;}
```

Modify Change Window Colors in the Change Scheduler

You can change colors that indicate change windows by modifying lines in *schedule.css*.

Note: If you modify *schedule.css* to change the formatting of windows on the Change Scheduler, you must also modify the schedGroup macros used by the Change Calendar if you want its formatting to be consistent with the Scheduler.

To change the background color

1. Open *schedule.css* in a text editor.

You can locate this file in the following directory:

```
$NX_ROOT\bopcfg\www\wwwroot\css
```

2. Modify the following lines with the color code you want:

```
.schedConflict { background-color: #ff0000; font-size: 4px }  
.schedBusy { background-color: #0176ff; font-size: 4px}  
.schedCurrent { background-color: #008000; font-size: 4px; }  
.schedMW { background-color: #40ff40; font-size: 4px; }  
.schedBW { background-color: black; font-size: 4px; }  
.schedNone { font-size: 4px; }  
.schedDialog { width: 100%; height: 4px; margin-left: 0px; margin-right: 0px; margin-top: 4px; margin-bottom:  
4px; }
```

3. Save *schedule.css*.

The background color is modified.

More information:

[schedConfig Macro—Configure Schedule](#) (see page 640)

How to Schedule Change Windows

You can schedule change windows and view them in the Change Calendar. *Maintenance windows* establish time periods when CI changes should occur, and *blackout windows* establish time periods when CI changes should *not* occur. Change windows help you schedule changes to minimize their effect on critical business processes. CA SDM can implement change windows at the global or system level.

To schedule change windows, do the following:

1. View the Change Calendar to see where you want change windows.
2. Create the appropriate change windows for your organization.

Maintenance Window

Indicates a scheduled time period during which a CI or a set of CIs can be changed. Because changes can involve downtime, these windows can be used to minimize disruptions to critical business processes. In general, scheduled changes should occur inside a maintenance window.

Blackout Window

Indicates a scheduled time period during which CI changes should not occur. This blackout can indicate a reduced support period (for example, a holiday), a corporate event, or a critical business time, such as fiscal year-end. In general, scheduled changes should only occur outside blackout windows.

Global Change Windows

Indicates a blackout or maintenance window that occur for your entire organization. For example, you want to create a [global](#) (see page 656) blackout window named Holiday that starts in November and ends in January in the American - East timezone. This blackout window does *not* allow nonemergency changes between these dates. Global change windows do not associate with specific CIs because they apply to all CIs.

3. (Optional) Specify the recurrence patterns of the change windows.
4. Open the Change Calendar.

The legend displays icons and colors to identify change windows in the schedule.

5. Create change orders associated with CIs.
6. Update the Change Calendar by using the Change Scheduler.

Note: For more information about creating change windows, see the *Online Help*.

View Change Windows

You can view change windows in the Change Calendar to see when their associated CIs can or cannot be changed.

To view change windows

1. On the Administration tab, click Service Desk, Change Orders, Change Windows.
The Change Windows List appears.
2. Select the change window that you want to view.
The change window appears.
3. (Optional) Click Edit.
You can modify the change window.
4. Save or close the window.
The Window Detail page appears.

Associate a CI with a Maintenance Window

To control when a CI undergoes maintenance, you can associate that CI to a maintenance window.

To associate a CI with a maintenance window

1. Navigate to the Change Window Detail page.
2. Click the Associated CIs tab.
The Associated CIs list appears.
3. Click Update Configuration Items.
The CI Search page appears.
4. Specify information for the required CIs.
5. Click Search
The Available CIs page appears.
6. Move any required CIs in the Available CIs list to the Associated CIs list.
7. Click OK.
The Change Window Detail page appears.
8. Click Save.
The CI is associated with a maintenance window.

The change window, the CIs and the selected window associations are saved.

View Associated CIs

To determine which CIs could be impacted during specific periods of time, you can view the CIs that are associated with non-global maintenance windows.

To view CIs associated with a maintenance window

1. Navigate to the Change Window Detail page.
2. Click the Associated CIs tab.

The associated CIs are displayed.

Create a Blackout Window Example

The following example creates a blackout window to indicate that your organization does not schedule change orders for a specific time period. You want this blackout window to last 48 hours and recur every Friday at 6:00 PM in your time zone.

Follow these steps:

1. On the Administration tab, click Service Desk, Change Orders, Change Windows.
The Change Windows List appears.
2. Click Create New.
The Create New Change Window page appears.
3. Complete the appropriate fields as follows:
 - a. Enter **Friday_Blackout** as the window name.
 - b. Select the type Blackout.
 - c. Select the Active status.
 - d. Select the upcoming Friday and 6:00 PM from the Date Helper as the Start Date.
 - e. Select the upcoming Sunday and 6:00 PM from the Date Helper as the End Date.
 - f. Select your time zone.
 - g. (Optional) Enter a description for the blackout window.
 - h. On the Recurrence Pattern tab, select Weekly.
 - Set the recurrence to every **1** week.
 - Select the date to end the recurrence.
4. Save and close the blackout window.
5. Refresh the Change Windows List.
6. Open and refresh the Change Calendar.
The blackout window appears in the schedule.

Create a Global Maintenance Window

Create a global maintenance window for your organization. For example, you do *not* want to define maintenance windows by server or CI and want to complete global changes during weekends.

Follow these steps:

1. Navigate to the Change Windows List page.
2. Click Create New.
3. The Create New Change Window page appears.
4. Enter or modify the necessary fields for the maintenance window.
5. Set Type to Maintenance.
6. Verify that the Global checkbox is checked.

Note: A global maintenance window cannot be associated with CIs.

7. Enter Start Date, End Date, and Timezone.
8. Click Save.

The global maintenance window is saved.

Conflict Analysis and Collision Detection

Conflict analysis detects and shows collisions that occur when two or more changes to the same configuration item are scheduled for implementation at the same time. The change management team handles scheduling collisions in the following ways:

- Uses the Change Calendar to see change order-related CI collisions and makes adjustments to them to reduce potential impact.
- Searches for and detects collisions in a log.

Note: For more information about handling scheduling collisions, see the *Online Help*.

CA Workflow Visualization

CA Workflow visualization lets you monitor the progress of change order workflow tasks. The process viewer graphically displays every step of the CA Workflow process, including complete and incomplete steps and the current location of a workflow progression.

The process viewer displays the status and path of a CA Workflow process when it is associated with a Change/Issue category or a Request/Incident/Problem area.

How to Visualize the Workflow

Visualize the workflow as follows:

1. Associate a CA Workflow process with a change category.
2. Create a change order using that change category.
3. On the Workflow Tasks tab, click the View Process button to launch the process viewer.

Note: For more information about associating CA Workflow processes with categories or areas, see the *Online Help*.

Using the Process Viewer Example

This example demonstrates how to use the process viewer on your system using a default change category, after installing CA Workflow and enabling the options.

1. On the Administration tab, navigate to Service Desk, Change Orders, Categories.
The Change Category List page appears.
 2. Select Add.IT.Other from the Symbol column.
The Add.IT.Other Change Category Detail page appears.
 3. Click Edit.
The Add.IT.Other Update Change Category page appears.
 4. Click the Workflow tab.
Click CA Workflow check box.
The Workflow Definition List appears.
 5. Select Order PC - Service Desk Release 12.7.
The process name populates the CA Workflow Definition Name field.
 6. Click Save.
The Change Category Detail page updates.
 7. Close the window.
 8. Create and save a change order using the Add.IT.Other category.
The Workflow Tasks tab displays the Order PC - Service Desk Release 12.7 process and the Workitem List.
 9. Click the View Process button on the Workflow Tasks tab.
The Process Viewer appears.
- Note:** You can view updated progress of the workflow tasks by refreshing the Process Viewer.

Change Management Process Definition for CA Workflow

The Change Management Process Definition manages Standard, Normal, or Emergency change orders. As part of the CA Workflow, the Change Management Process Definition manages all change order tasks from the initial request for change, through the Post Implementation Review. As each task completes, the Change Management Process Definition updates the change order status and activity log. The Change Management Process Definition also emails the contact or group about tasks that require completion.

The Change Management Process Definition includes the following functionality:

- Analyzes the change order type and associated risk level to determine the necessary level of approvals for change implementation.
- Provides a sample foundation as a basis for change management.
- Aligns with ITIL v3 standards for risk assessment, impact and conflict analysis, and approvals by both the Change Manager and the CAB.
- Incorporates the following CA SDM Change Management functionality:
 - Risk Assessment Survey
 - Business case
 - Conflict analysis
 - Impact analysis
 - Closure codes
 - Status codes
 - Change Type
 - CAB and CAB approval flags

More information:

[CA Workflow](#) (see page 245)

Change Management Process Definition Components

The Change Management Process Definition has the following components:

- **Change Management Process Definition**—The process definition that manages the entire change order process. All activities and subactivities follow the current CA Workflow best practices for exception handling. The process definition file name is *r12_Change_Mgmt_[en_language].xml*. For example, the English process definition name is *r12_Change_Mgmt_en_US.xml*.

Note: For information about CA Workflow best practices, see

https://support.ca.com/iri/portal/anonymous/phpdocs?filePath=0/7956/7956_tecd oc.html.

- **Get Group Mem and Notify Process Definition**—A building block that manages email notifications for groups. The Get Mem and Notify process definition accepts the group UUID as input. Depending on the notification flag, the Get Mem and Notify process definition can email every group member or only the manager of the group. The process definition file name is *r12_Get_Group_Mem_Notify_[en_language].xml*. For example, the English process definition file name is *r12_Change_Mgmt_en_US.xml*.
- **Get Group Mem and Notify Actor**—An actor that invokes the Get Mem and Notify process definition. The actor file name is *r12_Get_Group_Mem_Notify_Actor_[en_language].xml*. For example, the English actor file name is *r12_Get_Group_Mem_Notify_Actor_en_US.xml*.
- **Date and Time Convert**—An actor that accepts the date and time as inputs and supplies a legible date and time string for Scheduled Start Date and Scheduled End Date fields on approval forms. The Date and Time Convert actor is a CA Workflow JavaScript Actor. The actor file name is *r12_Date_Time_Convert_Actor_[en_language].xml*. For example, the English actor file name is *r12_Date_Time_Convert_Actor_en_US.xml*.

How to Set Up the Change Management Process Definition

To set up the Change Management Process Definition, do the following:

1. Use CA Workflow to [review the documentation](#) (see page 660) about activity nodes, Sub Activities, and notifications.
2. Create the [CAB and Implementation groups](#) (see page 661) who are responsible for managing the change process.

3. [Create contacts](#) (see page 662) and assign them to the CAB or Implementation groups.
4. [Configure email notifications](#) (see page 661) for the Change Management Process Definition.
5. If you are not using the default Tomcat server port 8090, [configure the Tomcat server](#) (see page 663) for the appropriate communications between CA Workflow and CA SDM.
6. [Configure one or more change categories](#) (see page 664) to use the Change Management Process Definition.
7. [Install the Category Defaults](#) (see page 664) option in Options Manager.

View Documentation about Activity Nodes and Attributes

Because all activity nodes and attributes within the Change Management Process Definition are documented, you can use CA Workflow to view information about activity nodes and attributes. If necessary, you can change the flow of the process definition to meet the needs of your organization.

Note: For information about viewing activity nodes and attributes, see the CA Workflow documentation.

To view documentation about activity nodes and attributes

1. Open the Change Mgmt – Service Desk r12.1 process definition in the CA Workflow application.
The Change Management Process Definition appears in CA Workflow.
2. Double-click an activity node.
The activity node opens.
3. On the Properties tab, review the Description field for additional details about the node.
4. Exit the activity node properties.
5. (Optional) Navigate to the Attributes tab in the lower window of the process definition.
6. Double-click or mouse-over the attribute.
The Attribute page opens with information about the attribute.

View Sub Activities

The Change Management Process Definition uses several CA Workflow Sub Activities to simplify the work flow. While working with the Change Management Process Definition, you can review and customize the Sub Activities if necessary.

To view Sub Activities

1. Open the Change Mgmt – Service Desk r12.1 process definition in the CA Workflow application.

The Change Management Process Definition appears in CA Workflow.

2. Left-click any node with a plus (+) symbol and select Open Tab or Open Inline.

The details about the Sub Activity appear in a new tab or next to the process definition activities.

How to Set up Email Notifications for the Change Management Contacts

The system sends email notifications about change order task assignments and status changes. Email notifications contain links to the work list or change order for additional review. Because the Change Management Process Definition sends email notifications, email configuration is required.

To set up email notifications for Change Management contacts, do the following:

1. On the Administration Tab, select Options Manager, Email options.

The Option List appears.

2. Configure the Email options as appropriate for your organization.

Note: For information about configuring email, see the *Online Help*.

How to Create Change Management Process Groups

When you create change management process groups who are responsible for approvals and additional tasks, consider the following:

- You can use any group name as long as the CA SDM and CA EEM names match.
- The group name in CA SDM must also match the group folder name in CA EEM.
- Each group requires more than one member.
- The group members match in CA SDM and CA EEM. For example, if the CA SDM CAB group has four members, the CA EEM CAB group has the same members.

Note: For information about creating CA SDM and CA EEM groups, see the *Online Help* and the CA EEM documentation.

To create change management groups, do the following:

1. On the Administration tab, select Security and Role Management, Groups.
The Group Search page appears.
2. Click Create New.
The Create New Group page appears.
3. Create the following groups in CA SDM:
 - **CAB Group**—A CAB group that is associated to the Change Category or the change order.
 - **Implementation Group**—An implementation group that assigned to the Change Category or the change order.
4. Click Save.
The Group Detail page appears.
5. Create the same groups in CA EEM.

More information:

[Assign CAB Groups](#) (see page 678)

[Manage CAB Groups](#) (see page 677)

How to Create Change Management Process Contacts

When you work with the Change Management Process Definition, you create contacts who are responsible for approvals and additional tasks throughout the change process. You also assign the contacts to the [Change Management Process groups](#) (see page 661).

To create the change management process contacts, do the following:

1. On the Administration tab, select Security and Role Management, Contacts.
The Contact Search page appears.
2. Click Create New.
The Create New Contact page appears.
3. Create the following contacts with valid email addresses and assign them to the appropriate groups:
 - **Requester**—One or more contacts who submit the Request for Change.
 - **Change Manager**—The contact who is the manager of the Implementation group. The Implementation group can only have one change manager.
 - **CAB Manager**—The contact who is the manager of the CAB group. The CAB group can only have one CAB manager.

The Contact Detail Groups tab lists the current group assignments.

4. On the Members tab of the Implementation Group detail page, set the Manager flag for the Change Manager and the CAB Manager.
5. Click Save.

The Members tab shows Manager-Yes for the following contacts:

- Change Manager
- CAB Manager

6. Create the same contacts in CA EEM and assign them to the CA EEM groups.

Note: The contact names can use any valid user name or system login. For information about creating CA SDM and CA EEM contacts, see the *Online Help* and the CA EEM documentation.

Specify an Alternate Port for the CA Workflow Tomcat Server

If the CA Workflow Tomcat server is running on a port other than the default port 8090, you configure the port for CA Workflow to communicate with CA SDM.

To specify an alternate port for the CA Workflow Tomcat server

1. Open the CA Workflow application.
The CA Workflow page appears.
2. Click the Actors tab.
The Actor Types/Actors list appears.
3. Select JavaScript, USD_Initializer.
The Activity Options CA Workflow page appears.
4. Double-click Get Global Attributes.
The JavaScript Operation page appears.
5. Change the WFTomcatPort value to the new port used by CA Workflow Tomcat.
6. Click OK.
The Tomcat Server is configured for CA Workflow to communicate with CA SDM.

Configure the Change Category

When you specify the change category for the Change Management Process Definition, every new change order that uses the category automatically launches the Change Management Process Definition. All preconfigured change category values automatically populate the change order.

To configure the change category

1. On the Administration tab, select Service Desk Manager, Change Orders, Categories.

The Change Category List appears.

2. Create or open an existing Change Category.

The Change Category Detail page appears.

3. On the Workflow tab, click Use CA Workflow.

4. Click the CA Workflow Definition link.

The CA Workflow Definition List appears.

5. Click Change Mgmt - Service Desk r12.1.

6. Click OK.

The Change Category Detail page appears.

7. (Optional) Assign the CAB, Group, and Risk Survey.

8. Click Save.

The next new change order that uses this category automatically launches the Change Management Process Definition with the values in the change category.

Note: For information about how to create or edit Change Categories, see the *Online Help*.

Install the Category_Defaults Option in Options Manager

The Category_Defaults option uses category fields to populate the CAB, Group, and Risk Survey on the change order. For example, the CAB group you specify for a category appears on all new change orders that use the category.

To install the Category_Defaults option in Options Manager

1. On the Administration tab, select Options Manager, Change Order Mgr, Category_Defaults.

The Options List appears.

2. Right-click Edit Category_Defaults.

The Category_Defaults Update Options page appears.

3. Click Install.
4. Restart the CA SDM services.

New change orders automatically include values from the category.

How the Change Management Process Definition Works

When a requester saves a change order with a category that uses the Change Management Process Definition, the Change Management Process Definition automatically invokes.

The following steps describe how the Change Management Process Definition operates in CA SDM:

1. The requester opens and saves a change order that includes the following information:

Requester

Identifies the person who opened the change order.

Change Category

Identifies the change order category.

Type

Identifies the change order Type such as Standard, Normal, or Emergency Changes. The Type can also update as a part of the Workflow process.

Summary/Description

Describes the reason for the change order.

Scheduled Start Date

Identifies the start date.

Duration

Specifies the estimated length of time for the change.

CI

Identifies at least one affected CI.

The Change Management Process Definition launches. The requester receives an email notification to complete the Risk Assessment Survey.

2. The requester reviews the change order and completes the [Risk Assessment Survey](#) (see page 669) that describes how the change order affects servers, applications, and users.

The Change Management Process Definition analyzes the answers to determine the risk and assure that the change order follows the appropriate process path. The requester receives an email notification to perform conflict and impact analysis.

3. During [conflict and impact analysis](#) (see page 670), the requester checks for scheduling conflicts and analyzes the impact of the change order on CIs.

The requester receives an email notification to perform change analysis.

4. During [change analysis](#) (see page 671), the requester identifies and submits the type of change, reason for the change, business impact, and overall impact of the change.

The change manager receives an email notification to approve the change order. If the change order is a high risk, the CAB group also receives a notification.

5. The change manager reviews the change analysis and enters comments about the review. The change manager [approves, rejects, or marks the change order as incomplete](#) (see page 670).

The system updates the Change Order Detail page accordingly and responds with one of the following actions:

- If the change order is approved and no CAB approval is required, the implementation group receives a notification. The change order shows Status-Approved.
- If the change order is approved and additional CAB approvals are required, the CAB manager and members receive notification for approval. The change order shows Status-Approval in progress.
- If the change order is rejected, the change order shows Status-Rejected.
- If the change order is marked as incomplete, change order shows Status-Incomplete.

6. For approved change orders, a member of the implementation group confirms the [start of the implementation](#) (see page 673) and implements the change.

The change order shows Status-Implementation in progress.

7. When the work is complete, the implementation manager, completes the [Post Implementation Review](#) (see page 675) (PIR) to describe details about the outcome of the change.

CA Workflow closes the change order. The change order shows a Closure Code and Status-Implemented.

More information:


[How to Set Up the Change Management Process Definition](#) (see page 659)

[How to Implement the Risk Survey](#) (see page 682)

[View Default Risk Surveys](#) (see page 682)


How the CM Process Definition Manages CO Approvals

The following diagram shows how the [Change Management Process Definition](#) (see page 665) manages approvals for a new change order.



How the CM Process Definition Manages CO Implementations

The following diagram shows how the [Change Management Process Definition](#) (see page 665) manages change order implementation.



Complete the Risk Assessment Survey

When you create a change order, the system notifies you to complete the Risk Assessment Survey associated with the change category. As requester, you identify, evaluate, and quantify the risks of change orders that belong to change categories before modifying a system or service.

To complete the Risk Assessment Survey

1. Open the change order by clicking the Risk Assessment Survey email notification link.

The Change Order Detail page appears.

2. On the Workflow Tasks tab, click the Requester Risk Assessment Survey link and follow the directions for completing survey.

The CA Workflow work list appears.

3. On the Tasks Tab, click the Perform link.

The Perform Tasks page appears.

4. Make a note of the instructions and click the link.

The Risk Survey appears.

5. Carefully complete the survey questions so the Change Management Process Workflow can accurately assess the appropriate risk path for the change.

6. Click Submit.

The Perform Task page appears.

7. Click Confirm.

The Change Order Detail page shows Status-RFC. The Requester receives an email notification to perform conflict and impact analysis.

More information:

[How to Access a Risk Survey Directly from a URL](#) (see page 684)

Perform Conflict and Impact Analysis

During conflict and impact analysis, the requester checks for scheduling conflicts and analyzes the impact of the change order on CIs.

To perform conflict and impact analysis

1. Open the change order by clicking the email notification link.
The Change Order Detail page appears.
2. On the Workflow Tasks tab, click the Requester Impact and Conflict Analysis link.
The Tasks page appears.
3. On the Tasks Tab, click the Perform link.
The Perform Tasks page appears.
4. Make a note of the Impact and Conflict Analysis instructions and click the link.
The Change Order Detail page appears.
5. On the Change Order Conflicts tab, click Conflict Analysis to resolve all scheduling and outstanding conflicts. Resolve all conflicts before implementing the change.
The Conflicts List appears.
6. On the Config Items tab, click Impact Analysis to review details about the CIs.
The Configuration Item List appears.
7. On the Config. Items tab, click Impact Explorer to review information about each CI.
The Impact Explorer opens.
8. Click Visualizer for a graphical view of the CI relationships.
Note: The Visualizer button is only available when CA CMDB is installed.
9. Return to the Perform Task page and click Confirm.
The Workflow Tasks tab shows a link for launching change analysis. The requester receives an email notification to perform change analysis.

More information:

[Impact Explorer](#) (see page 685)

[Launch CMDB Visualizer from Impact Explorer](#) (see page 688)

Perform Change Analysis

During change analysis, the requester identifies and records the following information:

- Type of change
- Reason for the change
- Business impact
- Overall impact of the change

To perform change analysis

1. Open the change order by clicking the email notification link.
The Change Order Detail page appears.
2. On the Workflow Tasks tab, click the Requester Change Analysis Form link.
The Tasks page appears.
3. On the Tasks Tab, click the Perform link.
The Change Analysis page appears.
4. On the Chg Analysis Tab, answer the questions to confirm the change order and click Submit. For example, summarize your analysis and explain the purpose of the change order. If necessary, specify whether the change order is a repeat or a rework. Summarize the business impact and overall impact of the change.

Based on the change type and risk, the system performs the following actions:

- If the Change Type is Normal with a high to extreme risk, the Change Order Detail page shows the following:
 - CAB Approval-Yes
 - Status-Approval In progressThe system notifies the Change Manager to approve the change order. Upon Change Manager approval, the CAB group is notified for additional approval.
- If the Change Type is Normal with a medium to low risk, the Change Order Detail page shows the following:
 - CAB Approval-No
 - Status-Approval In progressThe system notifies the Change Manager to approve the change order.
- If the Change Type is Standard, the system notifies the implementation group to implement the change.

- If the Change Type is Emergency, the Change Order Detail page shows CAB Approval-No.

The system notifies the ECAB. After the ECAB processes the change order, the Change Order Detail page shows Status-Approval In progress and the Change Manager receives a notification to approve the change order.

Approve, Reject, or Mark a Change Order as Incomplete

The change order type and risk level drive the approval process for change orders that use the Change Management Process Definition. The following change order values determine the change order approvers:

- If the Type is Normal and the Risk is Medium to Low, the change order requires only the change manager approval.
- If the Type is Normal and the Risk is High to Extreme, the change order requires CAB member approval.
- If the Type is Standard, no approvals are required. The system notifies the Implementation group to work on the change order.
- If the Type is Emergency, the system notifies the ECAB group. The change order shows Status-Approval In progress. When the change order is ready for approval, the system notifies the change manager.

As an approver, you review the change analysis information and determine whether to approve, reject, or mark the change order as incomplete.

To approve, reject, or mark a change order as incomplete

1. Open the change order by clicking the email notification link.
The Change Order Detail page appears.
2. On the Workflow Tasks tab, click the Change Manager Approval Form link.
The Tasks page appears.
3. On the Tasks tab, click the Perform link.
The CAB or Change Manager Approval page appears.
4. Review the Chg Analysis tab for additional information about the change order.
5. Based on the required approval level, select one of the following approval tabs and answer the approval questions:

Chg Mgr Approval

Describes the change manager decision for the change order.

CAB Approval

Describes the CAB approver decision for the change order.

6. Select one of the following approval levels:

Approve

Accepts the change order.

Reject

Rejects the change order.

Incomplete

Marks the change order as incomplete.

The Result page appears. Based on the approval level, the following actions occur:

- If you approve the change order and no additional CAB approval is required, the Change Order Detail page shows Status-Approved and the Implementation group receives a notification to begin work.
- If you approve the change order and CAB approval is required, the Change Order Detail page shows Status-Approval in progress. The CAB group also receives a notification.
- If you reject the change order, the Change Order Detail page shows Status-Rejected and the requester and implementation group receive a notification.
- If the change order is incomplete, the Change Order Detail page shows Status-Incomplete and the requester receives a notification.

More information:

[How the CAB Process Works](#) (see page 628)

[CAB Responsibilities](#) (see page 627)

[CAB Approvals](#) (see page 679)

Implement a Change Order

For a particular change order, an implementation group completes one or more work items. As a member of the implementation group, you report information about the change.

To implement a change order

1. Log in to CA SDM as a member of the implementation group and open the change order.

The Change Order Detail page appears.

2. On the Workflow Tasks tab, click the Implement Change Order link.

The Task List page appears.

3. On the Tasks Tab, click the Perform link.
The Perform Tasks page appears.
4. Follow the instructions for reviewing the change order and click Confirm.
The Change Order Detail page shows Status-Implementation in progress.
5. Implement the change. For example, if the change order stated to install AntiVirus patch on Exchange Server 1 and Exchange Server 2, complete the change order according to company guidelines and standards.
6. On the Workflow Tasks tab, click the Implement Complete Form link.
The Tasks page appears.
7. On the Tasks Tab, click the Perform link.
The Implementation Complete page appears.
8. On the Impl Complete tab, answer the questions to describe how your group implemented the change order and click one of the following options:

Complete

Specifies that all change order tasks completed successfully. Sets the change order Status to Implementation Complete and Closure Code to Successful.

Incomplete

Specifies that one or more items on the change order could not be completed.

The requester receives one or more notifications about change order completion or back out. The system also responds in one of the following ways:

- If the change order is complete, the Change Order Detail page shows the following:
 - Status-Implementation Complete
 - Closure Code-Successful.

The system notifies the implementation group to complete the Post Implementation Review.

- If the change order is incomplete, Change Order Detail page shows the following:
 - Status-Backed out
 - Closure Code-Unsuccessful.

The system also notifies the implementation group to determine whether the back out was successful.

When the back out is successful, the system notifies the implementation group to complete the Post Implementation Review. If the back out is unsuccessful, the system creates an incident and notifies the implementation group to complete the Post Implementation Review.

Complete the Post Implementation Review

When the work finished, the implementation group completes the Post Implementation Review (PIR) that describes the outcome of the change. As a member of the implementation group, you typically complete the PIR three to seven days following implementation of the change order. However, the Change Management Process Definition sets a default delay of ten seconds for the assignment of this task.

To complete the Post Implementation Review

1. Log in to CA SDM as a member of the implementation group and open the change order.

The Change Order Detail page appears.

2. On the Workflow Tasks tab, click the Post Implementation Review link.

The Task List appears.

3. On the Tasks tab, click the Perform link.

The PIR page appears.

4. On the PIR tab, answer the questions to describe the resolution and click Submit.

The Change Order Detail page shows Status-Closed.

ActivityNode Actor not found: Update Object -Service Desk r12

Symptom:

The CA Workflow Command window shows the following error:

ActivityNode Actor not found: Update Object -Service Desk r12

Solution:

1. Log in to the CA Workflow application.
2. Click the Actors Tab.
3. Select the Processes Folder.
4. Select File/New Actor.
5. Select Update Object –Service Desk r12 from the drop-down list.
6. Click OK.

As part of the Change Management Process Definition, the system uses the new actor.

Change Order Does Not Close

Symptom:

During the last step in the workflow, the change order does not close. The CA Workflow process instance shows the Running status and CA Workflow reports an Actor Fault error.

Solution:

When closing the change order, verify that you set all required fields for example, the Root Cause.

CAB Console and Reporting

The CAB Console is a dashboard that facilitates online and quick approvals of change orders that require CAB approval. The Change Manager and other CAB members use the console to view details about a change order (and its associated Workflow tasks and configuration items) and approve or reject the change order. The CAB Console lets team members review, approve or reject a change order, and continue to the next change order quickly. For change order requests, the CAB can deal with the request by satisfying it directly or by escalating/referring the request to an appropriate group.

The Change Manager can use CA SDM built-in summary and detail reporting options to accomplish the following:

- Report approved and rejected change orders
- Report change orders awaiting approval

To print or view summary and detail reports, first select the records you want to include in the report. You can select specific records for a report using the search feature of the list pages.

For example, from the Change Order list you can enter search criteria to create a list of change orders awaiting CAB approval that you can then use to generate a report.

Note: For more information about generating summary and detail reports in CA SDM, see the Generate Reports information in the *Online Help*.

More information:

[Manage CAB Groups](#) (see page 677)

[Assign CAB Groups](#) (see page 678)

[CAB Approvals](#) (see page 679)

[Change CAB Console Properties](#) (see page 679)

[Change Management Reporting](#) (see page 680)

Manage CAB Groups

You can create and manage CAB groups with members appropriate for the change orders under consideration. The CAB can include members from the application team, development manager, component owner, QA, support, and any additional parties deemed necessary.

Note: Before you implement a CAB group, configure the appropriate contacts for your business structure.

To create a CAB group

1. On the Administration tab, select Groups.
The Group Search page appears.
2. Click Create New.
The Create New Group page appears.
3. Complete the fields as appropriate.
4. Click Save.
The CAB Group appears on the Group List page.

To assign members to the CAB group

1. On the Group Detail page, select the Members tab.
 2. Click Update Members.
The Contact Search page displays.
 3. Enter the search criteria to display the desired contacts and click Search.
The Members Update page displays, listing the contacts that matched the search criteria.
 4. From the list on the left, select the contacts you want to assign to this group. To select multiple items, hold down the Ctrl key while clicking the left mouse button.
 5. When you have selected all the contacts you want, click the selection button ($\geq$).
 6. Click OK.
The selected contacts move to the Members list on the right.
- The Group Detail page displays, with the selected contacts listed on the Members tab.

To add the CAB group to a change category

Note: Make sure that the Category_Defaults option is installed, so that the CAB field on change orders defaults to a CAB group.

1. On the Administration tab, select CA SDM, Change Orders, Categories.

The Change Category List appears.

2. Select a category from the list.

The Update Change Category page appears.

3. Select the appropriate CAB group from the CAB field.
4. Complete other fields as appropriate.
5. Click Save.

The Change Category Detail page displays a successful save message.

6. Click Close Window.

The CAB group is associated with the change category and the change order.

More information:

[Change Order and Issue Categories](#) (see page 266)

[Area Defaults](#) (see page 382)

Assign CAB Groups

You can assign a CAB group to review a change order before it is implemented.

To assign a CAB Group to a change order

1. Select the desired change order on the change order list.

The Change Order Detail page appears.

2. Click Edit.
3. Select CAB in the Details area.

The Groups List page appears.

4. Select a CAB group.
5. Click Save, Close Window.

The CAB group is assigned to the change order.

CAB Approvals

The Change Advisory Board (CAB) must approve a change order before implementing the change order if the CAB Approval field is set to YES. During the approval process, when the change manager clicks Approve or Reject, the change order status is changed to Approved or Rejected, respectively.

Note: If you want to use different status values, you can update the Approve or Reject button code using Web Screen Painter.

More information:

[Change CAB Console Properties](#) (see page 679)

Change CAB Console Properties

Using Web Screen Painter, you can change the CAB Console properties that appear on web forms in the CAB Console. For example, you can do the following:

- Rename the Approve and Reject buttons. These buttons are properties on the `orderwf_approval_console.htmlpl` (CA Workflow tasks) and `order_approval_console.htmlpl` (change orders) web forms.
- Customize the status values of the Approve and Reject buttons by changing the 'REJ' or 'APR' values.
 - The 'Reject Task' button calls a function `approve_reject('REJ')`.
 - The 'Approve Task' button calls a function `approve_reject('APR')`.

Important! You can only associate an active transition with a button. Do not deactivate a predefined status transition that is associated with a button otherwise the predefined status transition no longer functions.

To change CAB Console properties

1. In Web Screen Painter, open the CAB Console form that you want to change.
Web Screen Painter opens and displays the form.
2. On the Design tab, right-click the control you want to change, and select Properties.
The Properties - *control* page appears.

3. Change the properties you want by entering a new value for each one.

Changes take effect as soon as you click outside the field, and when you close the Properties page.

The Web Screen Painter displays a brief summary of the significance of a property in a note that appears at the bottom of the Properties form when you select the property.

Note: For information about Web Screen Painter, see the *Implementation Guide*.

Example: Customize the CAB Console for Change Workflow Tasks

This example shows how to customize the task status using Web Screen Painter.

1. In Web Screen Painter, open the `orderwf_approval_console.html` web form.

Web Screen Painter opens and displays the form.

2. On the Design tab, right-click the Reject Task button, and select Properties.

The Properties - Button page opens.

3. Find the function `approve_reject('REJ')`.

'REJ' is the status code for the Task status Reject.

4. Enter a new value for 'REJ'.

Changes take effect as soon as you click outside the field, and when you close the Properties page.

The Web Screen Painter displays a brief summary of the significance of a property in a note that appears at the bottom of the Properties form when you select the property.

Change Management Reporting

The Change Manager with appropriate privileges can use BusinessObjects InfoView to accomplish the following:

- Report change volume by operating system, change category, group, implementer, risk, status, implementation date, affected configuration items (CIs), and changes originating from incident or problem tickets.
- Report successfully implemented changes grouped by change category, urgency, priority, impact, % successful vs. total for the specified period, and the change requestor's group.

- Report failed changes grouped by category, urgency, priority, impact, reason for back out, % failed vs. total for the specified period, and the change requestor's group.
- Report the total number of change request's grouped by change category Change Coordinator, Change Manager, risk level, priority, and affected CIs for a specific time period.

You can navigate to the predefined reports in the left pane of the InfoView window to view, schedule, modify, or run the report or to view the history and properties for a report.

Note: For more information about using BusinessObjects InfoView, see the CA Business Intelligence information in the *Online Help*.

Risk Assessment

Risk assessments let you identify, evaluate, and quantify the risks of change orders that belong to change categories, prior to modifying a system or service in your environment. You create risk surveys to evaluate risks, and associate the surveys with change categories. When a user creates a change order and specifies a change category, the survey associated with that category is available for completion and submission.

The risk survey lists a series of single or multiple choice questions. Each answer has a weightage value. When creating a change order, the user selects the appropriate answers and submits the survey. The evaluated risk level is based on weightages of answers selected by the user.

More information:

[View Default Risk Surveys](#) (see page 682)

[Deploy a Risk Survey Example](#) (see page 683)

How to Implement the Risk Survey

Implement the risk survey as follows:

1. Establish risk levels for your organization.
2. Create or select a risk survey from the default list.
3. Create or modify risk survey questions and answers.
4. Modify risk ranges for the risk survey.
5. Associate the risk survey with a change category.
6. After you associate a risk survey with a change category, the Risk Survey button appears for the requester when they save a change order using the specified change category.
7. View the evaluated risk based on the risk survey results.
8. (Optional) Override the evaluated risk value from the Activities menu.

Note: For detailed information about creating and modifying risk surveys, see the *Online Help*.

View Default Risk Surveys

CA SDM provides default risk surveys that you can associate with change categories.

To view the default risk surveys

1. Navigate to Service Desk, Change Orders, Risk Survey.
2. Select General from the Risk Survey Name column.
The General Risk Survey Detail page appears..
3. Click View Survey.

The risk survey appears, listing questions, answers and weightage values for each answer. This survey applies to general changes within your organization.

4. Close the survey.
5. The General Risk Survey Detail page appears; close it.

Deploy a Risk Survey Example

This example demonstrates how to deploy a default risk survey on your system using a default change category.

1. Navigate to Service Desk, Change Orders, Categories.

The Change Category List appears.

2. Select Add.IT.Other from the Symbol column.

The Add.IT.Other Category detail page appears.

3. Click Edit.

The Add.IT.Other Update Change Category page appears.

4. Click Risk Survey.

The Risk Survey Template Search appears.

5. Search for the risk survey. In this example, search for General and select the risk survey to add it to the detail form.

The Risk Survey field populates with General.

6. Save and close the window.

If a user creates a change order using the Add.IT.Other category, the Risk Survey button will appear when they save the change order.

How to Access a Risk Survey Directly from a URL

You can provide a dynamic URL to allow a user to access a Change Order Risk Survey directly through a URL from within a task in CA Workflow. You can append this dynamic URL with the *KEEP.UsingURL=1* parameter to indicate that the Risk Survey is being accessed directly with a URL.

Create the dynamic URL within CA Workflow as follows:

1. Open your CA Workflow Process Definition and add activity node for a user to complete the Risk Survey.
2. Within the activity node, add a dynamic URL using the appropriate syntax. Verify to append the *KEEP.UsingURL=1* parameter, such as in the following example:

```
http://<hostname>:CA Portal/CAisd/pdmweb.exe?CNT_ID=<ID of contact>+CRID=<ID of Change Order>+OP=DO_RISK_SURVEY+KEEP.UsingURL=1
```

hostname

Specifies the hostname of the CA SDM server.

port

Specifies the port where CA SDM is installed.

ID of Contact

Specifies the internal ID of the contact completing the Risk Survey.

ID of Change Order

Specifies the internal ID of the Change Order associated with the Risk Survey.

The following example displays a sample URL:

```
http://hostname:8080/CAisd/pdmweb.exe?CNT_ID=21A7CC606A3011DEA39AA8010000A800+CRID=400009+OP=DO_RISK_SURVEY+KEEP.UsingURL=1
```

Important! If you do not add the *KEEP.UsingURL=1* parameter to the dynamic URL, the user gets an error after submitting the risk survey, and the risk level is not calculated.

3. Save the CA Workflow Process Definition.
4. Create a Change Category and associate a Risk Survey and the CA Workflow Process Definition to the Change Category.
Save your changes.
5. Open a Change Order with this Change Category and save it to instantiate the CA Workflow Process.
6. As a part of the CA Workflow Process, a user is assigned to a task to complete the Risk Survey. The user clicks the URL, as previously outlined.
7. The user logs in to CA SDM using the URL and is brought directly in context of the Risk Survey for the Change Order.

8. The user submits the survey and a message indicating success or failure displays.
If the submission is successful, the risk level is calculated and updated in the change order.

Impact Explorer

Impact Explorer is an advanced tool for managing and controlling change within an organization. Impact Explorer allows the Change Manager to explore CIs that are attached to a change order and also interact directly with an attached CI or its child CIs.

Impact Explorer provides these advantages:

- Displays all CIs that are attached to a change order
- Displays all child and peer-to-peer relationships for attached CIs
- Provides the ability to attach any related CI to the change order
- Displays a Descendants List of successively related CIs for any attached CI
- Provides the ability to launch CMDB Visualizer for a CI

More information:

[Launch Impact Explorer](#) (see page 686)

[Explore Attached CIs](#) (see page 686)

[View a CI in Impact Explorer](#) (see page 687)

[Add a Related CI to a Change Order](#) (see page 687)

[Display the CI Descendants List](#) (see page 688)

[Launch CMDB Visualizer from Impact Explorer](#) (see page 688)

[Configuring Impact Explorer](#) (see page 688)

Launch Impact Explorer

You can access Impact Explorer from the Config. Items tab of any change order.

To launch Impact Explorer

1. Open a Change Order Detail page.
2. Select the Config. Items tab.
3. Click Impact Explorer.

The Impact Explorer page appears.

Note: If you close the Change Order Detail page, the Impact Explorer page also closes.

Note: If Impact Explorer is launched from multiple change orders, multiple Impact Explorer pages are displayed.

Explore Attached CIs

The Impact Explorer tree contains a node for each CI that is attached to a change order. A plus (+) symbol indicates that a CI has at least one child CI.

Note: If more than 100 CIs are attached to the change order, only the first 100 are displayed.

To display the relationships for an attached CI, click the plus (+) symbol for the CI. The related CIs appear in the tree. In addition to the related CI names, the tree also displays relationship types in brackets.

Note: If the change order has more than 100 attached CIs, only the first 100 are displayed. To view the next 100 CIs, click More...

View a CI in Impact Explorer

The Impact Explorer tree displays a node for each CI that is linked to a change order. A plus (+) symbol indicates that a CI has at least one child CI.

To view a CI

1. Click a CI node in the left pane.

The Configuration Item Detail page appears in the right pane.

Note: You can also display the CI Detail page in the right pane by right-clicking on the CI and selecting View from its context menu.

2. Click the change order node at the top of the left pane.

The Change Order Detail page appears.

Note: If the change order has more than 100 attached CIs, only the first 100 are displayed. To view the next 100 CIs, click More...

Add a Related CI to a Change Order

While exploring CI relationships, you may want to attach a related CI to the change order.

To add a related CI to a change order

1. Click the Config. Items tab for a change order.

The Configuration Items List page appears.

2. Click Impact Explorer.

The Impact Explorer tree displays attached CIs.

3. Click any plus (+) for an attached CI node

CIs related to the node appear.

4. Right-click it a related CI and select Add to Change Order from the context menu.

The new attached CI appears in the Config. Items tab for the change order.

Display the CI Descendants List

For any CI, you can display lines of related CIs (named *descendants*). In addition to basic information about each CI, the CI Descendants List displays relationship levels, with the originating CI as level 1. CI children begin at level 2, and their children at 3, and so on. If a CI is encountered more than once among the relationships traced, only its closest relationship level is displayed.

Example: Multiple Descendant Paths for One CI

Due to multiple relationships, a descendant search finds the same related CI at level 2 and also at level 4. The descendant CI is displayed only at level 2.

To list descendants for a CI

1. From a Change Order Detail page, select the Config. Items tab.
2. Click Impact Explorer to display attached CIs.
3. Expand a CI node if needed.
4. Right-click a CI and select List Descendants.

The CI Descendants List is displayed.

Launch CMDB Visualizer from Impact Explorer

Impact Explorer provides the capability to launch CMDB Visualizer from an attached or related CI.

To launch CMDB Visualizer from Impact Explorer

1. From a Change Order Detail page, click Impact Explorer.
The Impact Explorer page displays.
2. Right-click any CI and select Launch Visualizer from its context menu.

CMDB Visualizer is launched with the CI as the focus.

Configuring Impact Explorer

An Administrator can configure the Impact Explorer as follows:

- Specify the number of child CIs displayed for an attached CI
- Specify the number of relationship levels displayed in the CI Descendants List
- Suppress the display of child CIs

The number of child CIs displayed is configurable by adding a `NX_IMPACT_EXPLORER_MAX_CHILD_NODES` setting to the `NX.env` configuration file. The default is 100.

Example: Configure the Display of Child CIs

Using the following setting, Impact Explorer only displays ten children at a time for an attached CI:

```
@NX_IMPACT_EXPLORER_MAX_CHILD_NODES=10
```

The default depth of the CI Descendants List is nine (9) levels. The depth is configurable by adding a `NX_IMPACT_EXPLORER_MAX_LEVELS` setting to the `NX.env` configuration file.

Example: Configure the Number of Descendant Levels

Using the following setting, the CI Descendants List displays the originating CI and only its relationally adjacent child CIs:

```
@NX_IMPACT_EXPLORER_MAX_LEVELS=2
```

Use the following Options Manager option to suppress the display of a certain type of child CI:

```
IMPACT_EXPLORER_EXCLUDE_HIER=boolean
```

If the `IMPACT_EXPLORER_EXCLUDE_HIER` option is installed and set to Yes, child CIs that are related through the CMDB Relationships tab of a CI Detail page are not displayed in the Impact Explorer tree or in the CI Descendants List.

Note: Child CIs that are related through CA CMDB are still displayed.

Chapter 17: Managing Reports

This section contains the following topics:

[CA Business Intelligence Reports](#) (see page 691)
[Reporting Scenarios](#) (see page 692)
[Reporting Components](#) (see page 692)
[Reporting Data Flow Diagram](#) (see page 695)
[Display Web-Based Reports in InfoView](#) (see page 696)
[Security and Authorization](#) (see page 697)
[How to Point an Existing CA Business Intelligence Server to a CA SDM Server](#) (see page 702)
[How to Set Up Data Partitions Security for Reporting](#) (see page 704)
[Replicated Database for Offline Reporting](#) (see page 706)
[Role-Based Reports](#) (see page 706)
[Web-Based Reports](#) (see page 715)
[Key Performance Indicators](#) (see page 717)
[Ad Hoc Reports](#) (see page 728)
[Example Ad Hoc Reports](#) (see page 729)
[Dashboard Reports](#) (see page 734)
[Write CA Business Intelligence Reports](#) (see page 735)

CA Business Intelligence Reports

CA Business Intelligence is a web-based component that packages Business Objects technology, integrated with CA SDM and a variety of common data sources, such as SQL Server, Oracle, and Open Database Connectivity (ODBC).

CA Business Intelligence uses BusinessObjects Enterprise as the default reporting system. Predefined reports are provided for CA SDM, Knowledge Management, and Support Automation.

With CA Business Intelligence reporting, users can do the following:

- Customize existing reports
- Drill down on Business Objects report data to show the data beneath charts and summarized groups

- Export instances of reports to different output formats
- Build ad hoc reports
- Publish new reports and distribute them to authorized users
- Schedule reports to run at specified times
- Use dashboard reporting to monitor daily operations for all CA SDM ticket types.

Note: Access to Reporting is restricted through CA SDM Data Partitions Security.

Reporting Scenarios

CA Business Intelligence integrated with BusinessObjects Enterprise supports the following reporting scenarios:

- **Role-based reporting**—From the CA SDM Reports tab, authorized users can view reports defined for their role, then click the InfoView button to manage their personal reports in BusinessObjects InfoView. CA SDM uses the InfoView interface to collect, organize, and present information in report formats. In InfoView, predefined reports are grouped in public folders.
- **Web-based reporting**—Web-based reports are predefined reports in CA SDM, Knowledge Management, CMDB, and Support Automation. They are developed with either Web Intelligence or Crystal Reports. The reports are accessed in InfoView and can be used as models for defining site-specific reports.
- **Ad hoc reporting**—Ad hoc reports are created and administered from InfoView using a Web Intelligence plugin-based interface. You can store and manage reports in a personal workspace (My Folders). Ad hoc reporting is intended for users who want to create basic reports easily without writing queries.
- **Dashboard reporting**—Dashboard reporting lets you monitor daily operations for all CA SDM ticket types (request/incident/problem, change order, or issue) in InfoView. Each report contains analytics about the top performers working on active tickets, so you can monitor their progress. You can work with individual predefined dashboard reports or use the corporate dashboard to view all CA SDM daily operations in a single view.

Reporting Components

BusinessObjects Enterprise and its associated tools, coupled with BusinessObjects Crystal Reports XI are the backbone of the CA BI architecture.

Although Crystal reports are delivered as the primary component of CA Business Intelligence, the report creation and maintenance tool, Crystal Reports XI, is not delivered as part of CA Business Intelligence. Crystal Reports XI is a separately licensed product that can be purchased from BusinessObjects Enterprise, and used in conjunction with CA Business Intelligence.

Note: Microsoft Access predefined reports are no longer developed or provided with CA SDM.

You can use the following components to administer, monitor, and configure the CA Business Intelligence reporting environment:

- **CA SDM Database/ Domsrvr / ODBC Driver**—Report data is stored in a SQL Server or Oracle CA SDM database. BusinessObjects reporting applications (Crystal Reports and Web Intelligence) access the database using an ODBC driver that connects directly with the CA SDM object engine (domsrvr). All CA SDM security, including data partition and tenancy restrictions, is automatically applied to reports, but not to your reporting environment. You can set up your reporting environment so it works with existing data partitions in CA SDM.

Note: For information on how to establish data partitions security for your reporting environment, see [How to Set Up Data Partitions Security for Reporting](#) (see page 697).

- **Central Management Server (CMS)**—The central repository that stores all objects used in every reporting process.
- **Central Management Console**—An administrative component that provides access to all BusinessObjects administration functions. Using the CMC, you can deploy reports and assign user access and folder permissions for InfoView. User permissions, roles, and authentication options must be established for your reporting environment using the CMC.

Note: User permissions, roles, and authentication options must be established for your reporting environment before you use CA Business Intelligence. For more information on defining security, see [Security and Authentication](#) (see page 697).

- **BusinessObjects Universe**—Describes the classes (tables) and objects (columns) which are used in reports. The CA SDM universe is installed and configured during the installation. At the completion of the installation, the universe connection is assigned to various groups and users in CA SDM.
- **Designer**—A BusinessObjects component that lets you modify the CA SDM Universe, which is a metalayer between CA SDM schema, and BusinessObjects reporting tools. The Import/Export Wizard facilitates object population or extraction within the CMS.
- **Default Predefined Reports**—Predefined reports are web-based CA SDM and Knowledge Management reports developed with either Web Intelligence or Crystal Reports. The reports can be used as models for defining site-specific reports.

- **InfoView**—BusinessObjects InfoView is a web interface that allows authorized CA SDM users to interact with web-based predefined reports by viewing, running, and scheduling report types including, but not limited to, Web Intelligence and Crystal Reports. Reports are contained in folders in the public section in InfoView.
- **Ad Hoc Reports**—Ad hoc reports let you create and administer reports using a Web Intelligence plugin-based interface. This tool is intended for users who want to create basic reports easily without writing queries.


Note: For ad hoc reporting usage examples, see [Ad Hoc Reports](#) (see page 728).

- **Dashboard Reports**—Dashboard reports let you monitor daily operations for all CA SDM ticket types (request/incident/problem, change order, or issue) in BusinessObjects InfoView. Each report contains analytics about the top performers working on active tickets, so you can monitor their progress.
- **CA SDM Reports Tab**—Authorized users can view their role-based reports on the CA SDM Reports tab, then click the InfoView button on this tab to manage their personal reports in InfoView.

Note: For information on how to manage role-based reports and display new reports on the Reports tab, see [Web-Based Reports](#) (see page 715).

Reporting Data Flow Diagram

The following diagram illustrates component data flow and execution for CA Business Intelligence integrated with BusinessObjects Enterprise:



Display Web-Based Reports in InfoView

CA SDM uses the BusinessObjects InfoView interface to collect, organize, and present information in report formats.

Important! For all reports, you *must* update the Date fields to generate the reports. Other fields are optional.

Follow these steps:

1. From the Reports tab in CA SDM, click the InfoView button.

The InfoView home page appears.

2. From the Header panel, select Document List.

The Document List displays the CA SDM reports, administrative tools, samples, and so on in public folders. When you select a folder, its contents appear in the details panel. The folders that you can access on the Document List depend on the rights that are granted to you by your administrator.

For example, to view the reports for CMDB, navigate to CA Service Desk/CMDB. The list of reports displays on the details panel.

3. Double-click on a report.

The report form appears. The form contains the fields for your input to generate the selected report.

4. Enter the information relevant to generate the selected report.

Example: If you open the CIs Added in Time Range Report you need to enter the time range within which you want to view the added CIs, the CI Tenant Names, and so on.

Important! If you are using the Oracle database and want to generate All Change Impact Report, CIs Relationships Report, and Root Cause Analysis Report, then click Oracle (These Reports will work on Oracle Only) from the CMDB folder to access the reports. If you generate these reports from the SQL Server folder, the reports are not generated and an error message is displayed. Similarly, if you are using the SQL database and want to generate All Change Impact Report, CIs Relationships Report, and Root Cause Analysis Report, then click SQL Server (These Reports will work on SQL Server Only) from the CMDB folder to access the reports.

Note: It is mandatory to fill all the fields to generate a report.

5. Click OK.

The report appears.

Note: BusinessObjects InfoView includes a User's Guide that describes how to use BusinessObjects InfoView. To access the User's Guide, click the help icon in InfoView.

Security and Authorization

The default security configuration for BusinessObjects Enterprise is performed through the CA Business Intelligence configuration. The configuration determines the security policies of folders, universes, universe connections and tools. It also provides methods for adding users and mapping them to groups, and setting some preference options.

Specifically, the CA Business Intelligence reporting configuration involves:

- Setting up security
- Deploying reports
- Deploying universes
- Configuring Web Intelligence settings

At the completion of installation, the universe connection is assigned to various groups and users in CA SDM.

The administrator can log on to the BusinessObjects CMC and modify the default settings at any time. Users are authorized access to InfoView folders based on the CA Business Intelligence group to which they belong.

Groups and Users

The Groups listed in the following table will be added to the Central Management Server (CMS) during the CA Business Intelligence configuration. They relate to CA SDM roles with the same names. During the configuration phase, an optional check box was available to indicate whether sample users are added to the CMS. If it was selected, your default CMS configuration includes one sample user for each group. In the CMC, you can use these sample users as models when you establish user permissions and authentication options for your reporting environment.

Group Name	User Name
Change Manager	Change Manager User
Customer Service Manager	Customer Service Manager User
Knowledge Manager	Knowledge Manager User
Knowledge Analyst	Knowledge Analyst User
Service Desk Manager	Service Desk Manager User
Incident Manager	Incident Manager User
Problem Manager	Problem Manager User

CA SDM Data Partitions in Infoview

Consider the following information about the relationships between CA SDM data partitions and InfoView:

- The analyst connects to CA Business Intelligence through the Reports tab in CA SDM or in InfoView with the access type of a default reporting role.
- The CA SDM login credentials *must* match the InfoView credentials. The administrator sets an authentication method such as secLDAP or secEnterprise. CA SDM logs in to CA Business Intelligence, which then logs in to CA SDM through the ODBC.
- The CA Business Intelligence analyst login associates with the CA SDM login. CA Business Intelligence uses this CA SDM login and the reporting role associated with the access type of the login. If the analyst does not have an associated CA SDM login, CA Business Intelligence uses the system default. The system default is defined in the ODBC parameters in the Universe Designer and the reporting role of the access type.
- To restrict a user from viewing a report, the administrator can disable access to the report or the folder that contains it.

Universe and Universe Connections

The default configuration also includes the security policy for accessing the CA SDM universe and universe connections. To allow full access, all of the default groups are defined to have Full Control access.

Group Name	Access Level
Change Manager	Full Control
Customer Service Manager	Full Control
Knowledge Analyst	Full Control
Knowledge Manager	Full Control
Service Desk Manager	Full Control
Incident Manager	Full Control
Problem Manager	Full Control

Report Folders

The default configuration comes with a set of folders containing predefined reports for CA SDM and Knowledge Management. Each CA SDM group is configured to have access to a subset of these folders.

Folder Name	Group Name	Access Level
Aggregate	■ Change Manager	■ View
	■ Customer Service Manager	■ View
	■ Knowledge Manager	■ No Access
	■ Knowledge Analyst	■ No Access
	■ Service Desk Manager	■ Full Control
	■ Incident Manager	■ View
	■ Problem Manager	■ View
Asset	■ Change Manager	■ View on Demand
	■ Customer Service Manager	■ View
	■ Knowledge Manager	■ No Access
	■ Knowledge Analyst	■ No Access
	■ Service Desk Manager	■ Full Control
	■ Incident Manager	■ View
	■ Problem Manager	■ View
Change Order (includes all sub-folders)	■ Change Manager	■ Full Control
	■ Customer Service Manager	■ No Access
	■ Knowledge Manager	■ No Access
	■ Knowledge Analyst	■ No Access
	■ Service Desk Manager	■ View
	■ Incident Manager	■ View
	■ Problem Manager	■ View

Folder Name	Group Name	Access Level
Issue (includes all sub-folders)	■ Change Manager	■ No Access
	■ Customer Service Manager	■ Full Control
	■ Knowledge Manager	■ No Access
	■ Knowledge Analyst	■ No Access
	■ Service Desk Manager	■ No Access
	■ Incident Manager	■ No Access
	■ Problem Manager	■ No Access
Request (includes all sub-folders)	■ Change Manager	■ No Access
	■ Customer Service Manager	■ No Access
	■ Knowledge Manager	■ No Access
	■ Knowledge Analyst	■ No Access
	■ Service Desk Manager	■ Full Control
	■ Incident Manager	■ View
	■ Problem Manager	■ No Access
Knowledge Management	■ Change Manager	■ No Access
	■ Customer Service Manager	■ View
	■ Knowledge Manager	■ Full Control
	■ Knowledge Analyst	■ View
	■ Service Desk Manager	■ Schedule
	■ Incident Manager	■ View
	■ Problem Manager	■ View
Survey	■ Change Manager	■ No Access
	■ Customer Service Manager	■ No Access
	■ Knowledge Manager	■ View
	■ Knowledge Analyst	■ No Access
	■ Service Desk Manager	■ Full Control
	■ Incident Manager	■ View
	■ Problem Manager	■ View

Folder Name	Group Name	Access Level
Incident and Problem Management	■ Change Manager	■ No Access
	■ Customer Service Manager	■ No Access
	■ Knowledge Manager	■ View
	■ Knowledge Analyst	■ No Access
	■ Service Desk Manager	■ Schedule
	■ Incident Manager	■ Full Control
	■ Problem Manager	■ Full Control

Access Levels

The default configuration includes the following access levels for groups and users:

No Access

Sets all permissions to Not Specified.

View

Allows the user to see the folder, report or universe. If the report contains data, the user can open and interact with it. If the report does not contain data, the user cannot refresh the report. By default, the user can edit the report and save to a personal folder and refresh it there. You can explicitly prevent users from copying corporate documents to personal folders by setting an individual right that denies "Copy Objects to another folder".

Schedule

Allows a user to schedule a report but not refresh it in real time.

View On Demand

Allows a user to refresh a report in real time. When the report is a Web Intelligence document, the user also needs View On Demand access to the Universe and Universe connection to perform the refresh.

Full Control

Allows a user to create new reports within a folder, modify existing reports or delete items.

Advanced

When the preceding access levels do not meet your needs, we provide more granular access by choosing advanced.

When a user or group's access level is set to Advanced, more granular control of rights is available than those assigned through the selection of View, Schedule, View On Demand or Full Control.

BusinessObjects folders use inherited security. You receive the same authority in lower level folders as that of the top folder level that was assigned to you or your group, unless overriding restrictions are applied at lower levels. Default authorizations are provided at the folder level and group level. Users will inherit the rights of their group for all objects in the folder and subordinate folders.

CA Business Intelligence is installed with two groups: Administrators and Everyone. The Everyone group is assigned an Access Level of Schedule which allows scheduling and viewing of all report objects.

More information:

[Report Folders](#) (see page 699)

How to Point an Existing CA Business Intelligence Server to a CA SDM Server

If you have an existing CA Business Intelligence server, you can point it to a CA SDM server as follows:

1. [Create a New ODBC Data Source](#) (see page 702).
2. [Configure the Universe](#) (see page 703).
3. [Export the Universe](#) (see page 704).
4. Launch InfoView to verify the connection.

Create an ODBC Data Source

You can create an ODBC data source using the ODBC Data Source Administrator.

To create an ODBC Data Source

1. Navigate to Start, All Programs, Administrative Tools, Data Sources (ODBC).
The ODBC Data Source Administrator opens.
2. Select the System DSN tab and click Add.
The Create New Data Source page appears.
3. Select DataDirect OpenAccess and click Finish.
The DataDirect OpenAccess ODBC Setup page appears.
4. Specify `casd_servername` in the ODBC Name field.

5. Click Advanced.

The Advanced page appears.

6. Click Add.

The Open Access Database Setup page appears.

7. Complete the following fields:

- **Name.** Specify *casd_servername*.

- **IP Address.** Specify the *servername* or IP address.

Note: Ping *servername* to determine its IP address to use the IP address instead of *servername*.

- **Port.** Specify 1706.

- **Type.** (Optional) Specify the database used on the server. This field is only used as reference information.

8. Click OK.

9. Select *casd_servername* from the Database drop-down list.

10. Click OK.

The ODBC data source is created.

Configure the Universe

After you create the ODBC data source, you need to configure the universe by establishing a connection between CA SDM and CA Business Intelligence.

To configure the universe

1. Launch Designer and login as administrator.

2. Click File, Import.

The Import Universe page appears.

3. Browse and select your CA SDM universe and click OK.

The Universe successfully imported message appears.

4. Click File, Parameters.

The Universe Parameters page appears.

5. Click Edit.

The Edit CA SDM connection page appears.

6. Specify the username and password of the CA SDM privileged user, such as ServiceDesk.

7. Select `casd_servername` from datasource name and click Next.
8. Click Test Connection to verify that the CA SDM universe communicates with CA Business Intelligence.
9. Click Next, Next, Finish.
The universe is configured.

Export the Universe

After you configure the universe, you need to export it from the Designer.

To export the universe

1. Click File, Export.
The Export Universe dialog appears.
2. Click Continue.
3. Click OK.
The universe is exported.

How to Set Up Data Partitions Security for Reporting

Access to CA Business Intelligence reporting is restricted through CA SDM Data Partitions Security. All CA SDM security, including data partition and tenancy restrictions, is automatically applied to reports.

Data partition security for your specific reporting environment is not applied during the configuration phase. You set up your reporting environment so it works with existing data partitions in CA SDM.

Use the following tools to accomplish this task:

- **BusinessObjects Central Management Console (CMC)**—Lets you administer BusinessObjects Enterprise user authentication and permissions.
- **BusinessObjects Designer**—Lets you modify the universe for CA SDM, Knowledge Management, CMDB, and Support Automation, which is a metalayer between CA SDM schema, and reporting tools.
- **CA SDM Security and Role Management**—Lets you set up data partitions security to determine what data users can access.

Add the CA SDM Privileged User to CMC

The universe connection is configured by default to use the Service Desk Privileged User and Password when accessing data. This user should be added to the CMC as a new CA Business Intelligence user. There are two purposes for adding this user. First, you need this user if you plan to set up data partition security for reporting and second, it allows you to use this user when initially testing reports from the CA SDM Reports tab. The Reports tab requires a user who is defined to both CA SDM and CA Business Intelligence.

Use the instructions provided in the *Implementation Guide* to add your CA SDM users to the CMC and configure the CA SDM Privileged User account.

Define Universe Database Credentials

The universe must use the database credentials associated with the BusinessObjects user account.

Note: Verify that you sign on to the BusinessObjects Designer using the Privileged User account and not the Administrator User account.

To define universe database credentials

1. From the Start menu, browse to All Programs, BusinessObjects XI, BusinessObjects Enterprise, Designer.
The Designer window appears.
2. Log on to the Designer.
The Import Universe dialog appears.
3. Select File, Import.
The Import Universe dialog appears.
4. Select the CA SDM, Knowledge Management, CMDB, or Support Automation Universes folder from the drop-down list.
Note: If this is the first time you are using the Designer, select Browse, CA Universes.
5. Verify the file path for the import folder in the Import Folder box.
6. Click OK.
The universe window appears.
7. Select File, Parameters.
The Universe Parameters dialog appears.
8. On the Definition tab, click Edit.
The Login Parameters dialog appears.

9. Select the Use database credentials associated with Business Objects user account check box.
10. Click Next, Test Connection, and step through the universe connection dialogs that appear.
11. Click OK to finish.
12. Select File, Export.

The Export Universe dialog appears.

13. Select the universe that you want to use from the Domain drop-down list.
14. Select Everyone from the Groups list
15. Click OK to export universe changes.

Establish Data Partitions

In Security and Role Management, create a data partition constraint that will restrict the database record access for all report users assigned to the data partition.

Note: For information on establishing data partitions constraints, see [Data Partitions Setup](#) (see page 168).

Replicated Database for Offline Reporting

To manage potential performance issues that may affect the reporting components installed with CA SDM, you can create a replicated database for offline reporting purposes.

Note: For more information about creating a replicated database for offline reporting, see the sample documentation and scripts delivered in the NX_ROOT\samples\reporting directory.

Role-Based Reports

CA SDM presents role-based reports in two reporting frames on the Reports Tab. Each frame provides graphical views that let users drill down on report data to show the data beneath charts and summarized groups. You can manage role-based reports and display new BusinessObjects reports on the Reports tab.

The Reports List page contains details of reports that are available for use. You click the Reports List icon that appears in the selected frame to display this page.

Define Role-Based Reports for the Role

You can manage the report web forms that display on the Reports List page when a user assigned to this role is logged into the system.

To define a role-based report for the role

1. From the Administration tab, navigate to Security and Role Management, Role Management, Role List.

The Role List page appears. The following default roles are available for Reporting:

- Change Manager
- Customer Service Manager
- Knowledge Manager
- Knowledge Analyst
- Service Desk Manager
- Incident Manager
- Problem Manager

2. Select *one* of the Reporting roles from the list.

The Role Detail Form page appears. This page contains the following tabs:

Authorization

Allows you to define the authorization level assigned to the role.

Function Access

Defines role access to each CA SDM functional area.

Web Interface

Customizes the web interface for the role by defining the web pages and online help content the users can access.

Knowledge Management

Specifies the Knowledge Management privileges for the role

KT Document Visibility

Specifies which document statuses the role is allowed to view (for example, draft, retired, and published).

Tabs

Defines the tabs that appear when a user assigned to this role is logged in to CA SDM.

Report Web Forms

Defines the report web forms that are available to this role.

Go Resources

Specifies which record types appear in the "Go" drop-down list for the role. On the Role Detail page, select the Report Web Forms tab.

3. Click the Report Web Forms tab.

The Reports Web Form List page appears. This page contains details of reports available for use.

4. Click Update Web Forms.

The Web Form Search page appears.

5. Enter the search criteria to display the web forms and click Search.

The Web Forms Assigned Update page appears, listing the forms that matched the search criteria.

6. From the list on the left, select the web forms you want to display for this role. To select multiple items, hold down the CTRL key while clicking the left mouse button.

7. When you have selected all the forms you want, click Select Button.

The selected forms move to the Web Forms Assigned list on the right.

8. Click OK.

The Role Detail page appears, with the selected web forms listed on the Report Web Forms tab.

Display New Reports on the Reports Tab

When a report is approved for use within CA SDM, it is moved in to the public section in InfoView making it available to authorized users. To add the report to CA SDM, the administrator must perform some additional steps.

You can display new reports created in CA Business Intelligence on the CA SDM Reports tab. Use the following tools to accomplish this task:

- Web Screen Painter
- Role Management

Before you begin, do the following:

- Install and configure CA Business Intelligence so it works with CA SDM.

Note: For installation information, see the *Implementation Guide* and *CA Business Intelligence Installation Guide*.

- Established user permissions, roles, and [authentication](#) (see page 697) options for your reporting environment.

Step 1: Create Two Web Form Records to Call the New Reports

By default, the Reports tab contains two reporting frames that let authorized users display new reports. In this step, you will create two web form records and define the URLs that point to the new reports.

To create two web form records to call the reports

1. From the Administration tab, navigate to Security and Role Management, Role Management, Web Forms.

The Web Forms List page appears.

2. Click Create New.

The Create New Web Form page appears.

3. Fill in the following fields:

Web Form Name

Specify a name that identifies the web form. This is a required field.

Example: Asset List report

Record Status

Specify active.

Code

Specify a unique code value that identifies the web form to the system. After the code is defined, it cannot be changed.

Example: asset_list

Note: Make a note of the value you enter in the Code field.

Type

Select Report.

Description

(Optional) Enter a brief description of the web form.

Resource

Specify the URL that calls the new report.

Example: Open the Asset List web form from the Web Form list in Role Management and specify the URL that appears in this form.

`$BOServerURL?sPath=[Home],[Public+Folders],[CA+Reports],[CA+Service+Desk],[Asset]&sDocName=Asset+List&sViewer=html`

4. Click Save.

The web form definition is saved and the Web Form Detail page appears.

5. Repeat steps 1-3 to create a second web form record that will call the second report URL. Make a note of the value entered in the Code field.

Step 2: Create a Mutiframe Page and Assign the New Reports

In Web Screen Painter, create a multiframe web page with two vertical frames called a Frameset. Then, assign the web form records created in step 1 to this Frameset.

Note: You can create this page with any number of frames. Be aware that adding additional frames will limit visibility on the Reports tab.

To create a multiframe page and assign the reports

1. In Web Screen Painter, select File, New.
The New Form dialog box appears.
2. Fill in the interface and form group fields as appropriate.
3. Select multiframe.template from the File Name list.
4. Click New.
The multiframe.html window appears.
5. Select Controls, Insert Frameset.
The Insert Frameset dialog box appears.
6. Do not change the default settings and click OK.
Two vertical frames appear in the Frameset.
7. Right-click the left vertical frame and select Properties from the short-cut menu.
The Properties - Frameset dialog box appears.
8. Select the blank field next to the web_form_name attribute, then click the (...) button.
The Enter Web Form Name dialog box appears.
9. Specify the code value of the first report.
Example: asset_list
10. After the report web form name is found, click Validate.
11. Right-click the right vertical frame, specify the code value of the second report, then validate the report.
12. Save the multiframe.html page. Make a note of this file name.
Example: report_mframe.html

13. Select File, Publish.

Publishing makes the form available to all CA SDM users. For the Analyst role, the file is available in the following location: Program Files/CA/Service Desk/Site/mods/www/html/web/analyst/report_mframe.html.

Note: When searching for multiframe web forms in CA SDM, navigate to Security and Role Management, Role Management, Web Forms on the Administration tab. The Code column specifies the web_form_name on the Properties tab for a multiframe form in Web Screen Painter.

Step 3: Create a Start Page for the Reports Tab

In Role Management, create a start page to call the multiframe web page created in Step 2.

To create a start page for the Reports tab

1. From the Administration tab, navigate to Security and Role Management, Role Management, Web Forms.

The Web Forms List page appears.

2. Click Create New.

The Create New Web Form page appears.

3. Complete the following fields:

Web Form Name

Specifies a name that identifies the web form. This field is required.

Example: start page

Record Status

Specifies active or inactive.

Code

Specifies a unique code value that identifies the web form to the system. After the code is defined, it cannot be changed.

Example: start_page

Type

Specifies HTML.

Description

(Optional) Specifies a brief description of the web form.

Resource

Specifies the URL that calls the new report.

Example: Select the Administrator Reports Multiframe record on the Web Form List page. Specify the URL that appears in this form. At the end of this path, specify the new multiframe page (report_mframe.html).

`$cgi?SID=$SESSION.SID+FID=123+OP=DISPLAY_FORM+HTML=report_mframe.html`

Click Save.

The new start page record appears on the Web Form List page.

Step 4: Create a Reports Tab Record and Assign the Start Page

In Role Management, create a new Reports tab and specify the start page created in step 3.

To create a tab and assign the start page that calls the reports

1. Select Security and Role Management, Role Management, Tabs on the Administration tab.

The Tab List page appears.

2. Click Create New.

The Create New Tab page appears.

3. Complete the following fields:

Tab Name

Specify the name that identifies the tab within the administrative interface. For example, the tab name appears on the Tab List page.

Example: Reports Tab

Code

Specify the code that identifies this tab to the system. Once the code is defined, it cannot be changed.

Example: reports_tab

Record Status

Specify active.

Display Name

Specify the name that appears on the tab's graphical presentation in the user interface.

Example: Reports Tab

Starting Page

Specify the initial web form that appears in the main window when a user selects this tab.

Example: start page

4. Click Save.

The form appears in the Web Form List page.

Step 5: Assign the Tab Record to a Role Record

In Role Management, assign the Reports tab to the desired role listed on the Role Detail page. When a user assigned to this role logs in to the system, they will see the new Reports tab in the web interface.

To assign the tab record to a role record

1. Select Security and Role Management, Role Management, Role List on the Administration tab.

The Role List page appears.

2. Select the desired role from the Role list.

The Role Detail page appears.

3. On the bottom of this page, select Tabs, then Update Tabs.

The new tab appears on the Tab List.

Create a Start Page for the InfoView Button

You have the option of including the InfoView button that opens BusinessObjects InfoView in a new window on the Reports tab. This option is controlled through the Starting Page option that appears on the Reports tab record.

To create a start page for the InfoView button

1. From the Administration tab, navigate to Security and Role Management, Role Management, Web Forms.

The Web Forms List page appears.

2. Click Create New.

The Create New Web Form page appears.

3. Complete the following fields:

Web Form Name

(Required) Specifies a name that identifies the web form.

Example: InfoView page

Record Status

Specifies active or inactive.

Code

Specify a code value that identifies the web form to the system.

Example: info_view_page

Type

Specifies HTMLPL.

Description

(Optional) Describes the web form.

Resource

Specifies the URL that calls the reporting frames.

Example: Select the Administrator InfoView record on the Web Form List page. Specify the URL that appears in this form. At the end of this path, specify the new Reports tab start page (start_page).

```
$cgi?SID=$SESSION.SID+FID=123+OP=DISPLAY_FORM+HTMLPL=show_report_frames.html+K  
EEP.report_form=start_page
```

Click Save.

The InfoView start page appears on the Web Form List page.

Step 7: Assign the InfoView Button Start Page to the Reports Tab Record

In Role Management, assign the InfoView button start page created in step 6 to the Reports tab record.

To assign the Infoview button start page record to the Reports tab record

1. Select the new Reports Tab web form record from the Tab List page in Role Management.

The detail page appears.

2. Select Edit.

The Update Tab page appears.

3. Specify the new InfoView button start page in the Starting Page field.

4. Click Save.

When a user opens the Reports tab, the InfoView button appears in the top right corner of the form.

Web-Based Reports

CA Business Intelligence installs a set of web-based, predefined reports for CA SDM and Knowledge Management. They are developed with either BusinessObjects Enterprise Web Intelligence or Crystal Reports. The reports can be used as models for defining site-specific reports.

These reports are contained in folders that are automatically deployed to the CA Business Intelligence reporting server after installation.

Security can be assigned to folders and documents to specify whether they can be accessed globally, by specific roles, or by individuals.

Note: BusinessObjects InfoView includes a User's Guide that describes how to use BusinessObjects InfoView. To access the User's Guide, click the help icon in InfoView. For general information on using predefined reports, see the *Online Help*.

More information:

[BusinessObjects InfoView Interface](#) (see page 715)

[Navigate to Reports](#) (see page 715)

[InfoView Preferences](#) (see page 716)

[Schedule Reports](#) (see page 716)

[Data Analysis Setup](#) (see page 717)

[Publish and Distribute Reports](#) (see page 717)

BusinessObjects InfoView Interface

From the CA SDM Reports Tab, analysts and managers can work with reports defined for their role or manage their personal reports by clicking the InfoView button. From InfoView, users can access Crystal Reports and Web Intelligence Documents, and other objects, and organize them to suit their preferences.

The features that are available in InfoView vary by content type, but in general, the user can view information in their web browser, export it to other business applications (such as Microsoft Excel), and save it to a specified location.

Note: The folders and objects that users see in InfoView are dependent on their group (role) assignment.

Navigate to Reports

You can navigate the folders in the left-hand pane of the InfoView window to view, schedule, modify, or run the report or to view the history and properties for a report.

To view a report

1. Use *one* the following methods to open InfoView:
 - From the CA SDM Reports tab, click the InfoView button.
 - From the Start menu, browse to All Programs, BusinessObjects XI, BusinessObjects Enterprise, BusinessObjects Enterprise Java InfoView.The InfoView home page appears.
 2. In the left pane, navigate the folder structure-Public Folders/CA Reports/CA SDM.
 3. Click the folder that corresponds to the type of report you want to view.
 4. Click the report name for the type of information you want to see.
- The report appears in BusinessObjects InfoView.

InfoView Preferences

Users can set general preferences to specify that InfoView should start with one of their personal dashboard pages designed with Web Intelligence's My InfoView tool. They can also set their personal Web Intelligence and Crystal Reports viewing preferences.

Schedule Reports

InfoView supports scheduling of reports. For ad hoc reporting, scheduled reports are stored in the My Folders section. If a report is set to "refresh on open" the system will access the database to obtain the latest information each time the end user views the report. Users can schedule reports to run at certain times.

Administrators can define calendars to reflect appropriate scheduling dates. For example, a user can select a calendar using the "When" option (for timing and frequency) that will display available business days and ignore non-working days.

Note: For more information on scheduling reports, see BusinessObjects documentation.

With scheduling, users can:

- Specify the timing and frequency of the schedule (now, hourly, daily, and weekly, for example).
- Specify the destination, such as an inbox, file location or email recipient.
- Indicate which inboxes the report should be sent to.
- Specify the output, such as Web Intelligence, Crystal Report, Microsoft Excel, and Adobe Acrobat.

- Specify caching options. By default, the document results are stored on the Output File Repository Server. Users can choose to have the system cache the report on the Web Intelligence Report Server by selecting a cache format and locale.
- Select a server group to process the request. If a defined event is selected, the object will run only when the additional condition or event occurs.

Note: If CA SDM data partitions are used to manage data restrictions in InfoView's public folder section, scheduled reports contained in public folders must be defined for each user.

Data Analysis Setup

An interactive viewer provides an extensive set of tools for adjusting how the report data is viewed. You can drill to various levels of detail, such as from the group, to the assignee, or to the actual requests.

The interactive viewers lets you do the following:

- Change the appearance of the report by presenting the data in a different "block" format, such as a pie chart. For example, by right-clicking in the report and selecting "Turn table to," report presentation features appear.
- Perform different tasks, such as sorting, creating report breaks, calculating, filtering, and ranking the report.
- Share information with other users and groups.

Publish and Distribute Reports

InfoView users can save report data in Excel, PDF or CSV format, and then distribute it to a file location, an inbox, email address, or FTP site.

When a report is approved for use within CA SDM, it is moved in to the public section making it available to authorized users. Security can be assigned to folders and documents to specify whether they can be accessed globally, by specific roles or by individuals. Security is administered in the BusinessObjects Central Management Console.

Key Performance Indicators

Key Performance Indicators (KPIs) are metrics you can use to identify areas of your service management environment that can require administrative attention or configuration tuning.

Use the CA SDM web interface to configure your KPI definitions. The data they produce is stored in the CA SDM database and is available for producing [web-based reports](#) (see page 715).

Note: In addition to defining KPI queries, you can configure the KPI daemon to retrieve CA SDM ticket data whenever a ticket is opened, closed, or certain fields are modified.

By defining and monitoring a planned set of KPIs, you can measure progress toward your organization's performance-related goals, and gather valuable data for driving strategic decisions about your IT environment.

Note: For information about the KPI database tables, see the *Technical Reference Guide*.

KPI Types

CA SDM supports three KPI types:

- **System KPIs** are installed with the product. You can customize them to meet your needs, but you cannot create new System KPIs.
- **Stored Query KPIs** call a stored query to retrieve metrics from the database. You can create custom Stored Query KPIs, or modify the predefined examples installed with the product.
- **SQL KPIs** allow you to incorporate a complete SQL query directly into the KPI. You can create custom SQL KPIs, or modify the predefined examples installed with the product.

Notes:

- For details about configuring all three KPI types, see the *Online Help*.
- KPIs, whether predefined or user-created, cannot be deleted. When a KPI is no longer needed, you can deactivate it by setting the Record Status to inactive.

Predefined KPIs

Several KPIs of each type are installed with CA SDM. You may find it useful to examine them while reading this guide.

You can use the predefined Stored Query and SQL KPIs with their original definitions, or as models for your custom definitions.

You can use the predefined System KPIs with their original definitions, or modify some of their fields to meet your needs.

Note: Most of the predefined KPIs are set to inactive status by default. To view them in the web interface, filter the KPI List page to display inactive KPIs.

KPI Daemon

The KPI daemon manages the retrieval, organization, and storage of KPI metric data. The daemon runs continuously, and collects data at specified intervals from various system resources.

When the refresh time of a KPI is reached, the KPI daemon interacts with other system components as follows:

- **System KPI**—Sends a request to a target daemon (webengine, domsrvr, bpvirtodb, or db_agents) to retrieve count or duration data.
- **Stored Query KPI**—Sends a request to the domsrvr to collect count data.
- **SQL KPI**—Sends a request to the domsrvr to collect count, sum, max, or duration data.

When the KPI daemon receives the requested data, it stores the resulting metrics in the database.

Note: The calculated duration is based on a change to values in a ticket in real time, not in business hours.

System KPIs

System KPIs enable you to collect metric data related to the operation of CA SDM processes.


The following process types are supported:

- **domsrvr**—The CA SDM Object Manager (server process). The Object Manager also caches various records and tables for clients.
- **bpvirtodb**—The BOP Virtual Database server enables operation of multiple Object Managers in the CA SDM environment. All Object Managers running on primary or secondary servers connect to the virtual database, which arbitrates their access to database agents. For example, when retrieving a new range of ticket reference numbers, the virtual database ensures that only one Object Manager at a time accesses the table containing the reference numbers. The virtual database also performs caching of database information for Object Managers.
- **db_agents**—Database agents perform SQL queries, and handle other interactions between CA SDM and the database management system (DBMS). Database Agents adhere to the CA SDM database schema, and translate the SQL code to the form required by the particular DBMS (for example, Oracle).

- **webengine**—The CA SDM component that connects web browsers to an Object Manager through a pdmweb cgi running on a Microsoft IIS or Apache Tomcat web server. There must be a web engine for WSP on the primary server so WSP Schema Designer can write schema files. Web engines are the true client of an Object Manager for user client web browsers. Web engines cache .html web forms for connected users. You can manipulate the cache using the pdm_webcache utility and see client connection statistics using the pdm_webstat utility.

Note: For more information about these processes, see the *Implementation Guide*.

The following diagram illustrates data flow for System KPIs.



Each system KPI produce one of the following metric types:

- Count
- Duration

System KPI Examples

- This example produces a count of the update requests sent to BOP Virtual Database:
updateCt
- This example reports how long requests for updates, inserts, and deletes to BOP Virtual Database took to complete:
virtdbUpdateResponseDt

Stored Query KPIs

With Stored Query KPIs you can generate reports based on count metrics retrieved from the CA SDM database.

CA SDM provides a set of predefined stored queries. Many of them are useful just as they are. You can also customize them to meet your needs, or use them as models for writing your own queries.

All Stored Query KPIs have a Count metric type.

Note: You can use stored queries to produce counter fields for your web interface Scoreboard, or KPI metrics, or both. To use a stored query in a KPI definition, the query must be enabled for KPI usage. For more information, see [Stored Query Setup](#) (see page 313).

Stored Query KPI Examples

- This example produces a count of the open change orders at the assignee's location:

@cnt.location.name Changes
- This example produces a count of the workflow tasks that will violate an SLA before midnight of the current day:

Issue Workflow Tasks that will Violate an SLA today

SQL KPIs

SQL KPIs provide more flexibility than Stored Query KPIs. With SQL KPIs you can write your own queries to the CA SDM database to produce the following types of metrics:

- Sum
- Count
- Max
- Duration

Note: Your SQL code must comply with SQL92 syntax.

The following considerations apply to SQL KPIs:

- To see the SQL code for these example queries, open the query definitions from the KPI List page.

Note: For instructions, see the *Online Help*.

- The @ symbol is not supported in SQL KPIs. Instead, use syntax as shown in the following example and an attribute alias:

```
SELECT * FROM cr WHERE assignee_last_name = "Smith"
```

For this example, you could use a predefined attribute alias. The Attribute Alias Detail page appears as follows:

```
" Object Name      = cr
" Alias Name       = assignee_last_name
" Status          = Active
" Alias Value      = assignee.last_name
```

Examples: SQL Query KPI

- This predefined example produces the sum of the estimated costs of all pending change order workflow tasks:

Est Cost Sum of Pending Change Workflow Tasks

- This predefined example produces a count of the active change orders that have been closed and reopened:

Count of Reopened Change Orders

KPI Fields

The following table describes the KPI fields. The Sys, S.Q, and SQL columns indicate the KPI types each field belongs to (System, Stored Query, or SQL).

Note: All fields are required unless specified as optional.

Field	Sys.	S.Q.	SQL	Description
KPI Name	X	X	X	Identifies the display name for the KPI. Not editable.
Type	X	X	X	Defines the record as a System, Stored Query, or SQL KPI. Not editable.
Keep Existing Version	X	X	X	Specifies that the current version of the KPI record will be retained if the record is updated. Editable when the @NX_ALWAYS_KEEP_KPI_VERSIONS setting in the NX.env is set to No.
Version	X	X	X	Identifies the version number of the KPI record. The version number is incremented automatically each time the record is updated. Not editable.

Field	Sys.	S.Q.	SQL	Description
Record Status	X	X	X	Specifies whether the KPI is active (gathering data) or inactive (not gathering data). Editable only by using the "Edit in List" feature on the KPI List page.
Metric Type	X	X	X	Specifies the type of calculation the KPI will perform: ■ Stored Query KPIs are always Count. ■ System KPIs can be Count or Duration. ■ SQL KPIs can be Count, Sum, Max, or Duration. Editable for SQL type KPI only.
Refresh Time	X	X	X	Specifies the time interval for retrieving KPI metrics from the database. The value is editable. Refresh time is specified in HH:MM:SS format.
Process Type	X			System KPI (see page 719) metrics can be derived from the following CA SDM processes: ■ domsrvr—Object Manager ■ bpvirtddb—BOP Virtual Database Server ■ db_agents—Database agents ■ webengine—The CA SDM web client
System Name	X			The internal name of the System KPI. Not editable.
User Context		X	X	(Optional) Specifies the userid of a CA SDM contact. The contact's role and tenant assignments are used to determine data partitioning for the metrics produced by the KPI. The value is editable.
Stored Query		X		The name of the stored query called by this KPI. The value is editable.
SQL Query			X	The SQL code for the query. The value is editable.
Description	X	X	X	(Optional) Provides a detailed description of the KPI. The value is editable.

Retrieve Ticket Data

To allow reporting on how long tickets remain in the various processing states, you can configure the KPI daemon to retrieve CA SDM ticket data whenever a ticket is opened or closed, and whenever any of the following fields values are changed:

- Active
- Assignee
- Area or Category
- Group
- Impact
- Organization
- Priority
- Root Cause
- Status
- Service Type

To enable ticket retrieval, install the KPI Ticket Data Table option available in Options Manager KPI folder.

The collection of data from new and updated tickets is enabled. The ticket data is written into the `usp_kpi_ticket_data` database table, and is available for generating web-based reports.

Note: For instructions, see the *Online Help*.

The following considerations apply to ticket retrieval:

- The KPI daemon may take up to 30 minutes to populate the ticket information to the `usp_kpi_ticket_data` table.
- The `support_lev` field allows tracking of Service Type processing when the `classic_sla_processing` option is installed. The `classic_sla_processing` option enables the Service Type processing from CA SDM version 6.0 and earlier.
- Enabling this feature may result in degraded CA SDM performance.

More information:


[NX.env File](#) (see page 727)

KPI Ticket Data Flow

The KPI Ticket Data Table feature works on a trigger mechanism. When the `kpi_ticket_data_table` option is installed, the KPI daemon begins monitoring CA SDM ticket objects to track the following events:

- Open new ticket
- Check in changes to one or more monitored fields
- Close ticket

When one of these events occurs, a `POST_CI` trigger fires. The trigger sends a `BPMMessage` with a trigger attribute list to KPI daemon, and adds the returned data to the `usp_kpi_ticket_data` table, as illustrated in the following data flow diagram.



KPI Ticket Data Table Record Types

There are five operation types in the `usp_kpi_ticket_data` table:

- INIT (0)
- NO_INIT_REC (1)
- REOPEN (2)
- UPDATE (3)
- CLOSE (4)

The type value (integer value shown in parentheses) is stored in the operation field. The end_time of the previous record is stored in the prev_time field. The ktd_duration field stores the duration, which is calculated as end_time minus prev_time. The previous record is defined in the following rules:

1. The INIT record indicates a ticket has been created.
2. An UPDATE record is created for each update to a ticket. It gets prev_time from a previous UPDATE record to the same ticket field if a previous UPDATE record exists. Otherwise, it gets prev_time from the INIT, NO_INIT_REC or REOPEN record, depending on which record has an end_time value closest in time to the UPDATE record.
3. The CLOSE record is the same as UPDATE.
4. If a ticket is reopened, a single REOPEN record is created to show the Active field value. The REOPEN record gets prev_time from a previous CLOSE record for the same Active field.
5. If a ticket was created before turning on the Options Manager option, it may have records for updates but without an INIT record. The program in this case creates a NO_INIT_REC record and gets prev_time from the first non-INIT record it encounters.

Note: For complete documentation of the table fields, see the *Technical Reference Guide*.

Troubleshooting

You can diagnose problems you might experience with KPI usage.

Verify that the KPI Daemon is Running

To verify that the KPI daemon is running, run the pdm_status command-line utility as follows:

```
pdm_status
```

Examine the output for the following:

```
KPI Daemon myserver Running myserver 3568 Wed Feb 06 07:23:53
```

NX.env File

Review the \$NX_ROOT/NX.env file to verify that the basic KPI configuration is correctly set.

If you have installed the KPI Ticket Data Table option, the NX.env file includes the following:

```
#####
# NX_KPI_TICKET_DATA_TABLE
# Enable the collection of changed fields from ticket objects
# (like Requests, Issues and Change Orders) to write record entries
# into the usp_kpi_ticket_data table.
#####
@NX_KPI_TICKET_DATA_TABLE=Yes

#####
# NX_KPI_FUZZ_TIME
# Specifies a tolerable delay for each KPI within which kpi_daemon
# sends a request to retrieve KPI data when this KPI's refresh time
# is timeout
#####
@NX_KPI_FUZZ_TIME=20
```

Note: The default value of @NX_KPI_FUZZ_TIME is 20 seconds. You can modify this value within the range of 0 to 40 seconds. If you set this variable to a value higher than 40, 40 seconds will be used.

```
#####
# NX_ALWAYS_KEEP_KPI_VERSIONS
# The keep_version attribute in Kpi table is displayed on the KPI
# Detail Edit form as a checkbox. The NX_ALWAYS_KEEP_KPI_VERSIONS
# option specifies that checkbox is read-only and is always checked.
#####
@NX_ALWAYS_KEEP_KPI_VERSIONS=Yes
```

Turn on KPI Daemon Tracing

Use the pdm_logstat utility to turn on trace logging to monitor KPI daemon activity.

To turn on KPI daemon tracing, run the following command:

```
pdm_logstat -f cache_table_mgr.c VERBOSE
```

To turn off KPI daemon tracing, run the following command:

```
pdm_logstat -f cache_table_mgr.c
```

Examine the `stdlog.0` file for entries pertaining to KPI daemon activity.

The following example shows normal KPI daemon connections to the host computer and the domsrvr:

```
02/06 05:42:14.58 garbo-2k3-bb kpi_daemon 2432 SIGNIFICANT kpi_daemon.c 117 KPIDaemon ready for action!
02/06 05:42:14.80 garbo-2k3-bb domsrvr 792 SIGNIFICANT connmgr.c 2314 Connecting client kpi_daemon:garbo-2k3-bb
02/06 05:42:14.81 garbo-2k3-bb kpi_daemon 2432 SIGNIFICANT main.c 220 KPI daemon connected with domsrvr
```

The following example indicates an error in a KPI named `webSessionCT`:

```
01/16 13:24:46.74 jed web:local 2152 ERROR sys_kpi.c 96 Invalid KPI metric type: 1382180215 (KPI name:webSessionCT)
```

The following examples show KPI daemon activity:

```
02/06 07:23:50.47 garbo-2k3-bb kpi_daemon 3568 TRACE cache_table_mgr.c 427 Cache Table manager created Kpi_Obj (KPI dob
record_id:9003) (KPI type:2)
02/06 07:23:50.53 garbo-2k3-bb kpi_daemon 3568 TRACE cache_table_mgr.c 427 Cache Table manager created Kpi_Obj (KPI dob
record_id:9103) (KPI type:1)
```

Note: For information about KPI types, see [KPI Ticket Data Table Record Types](#) (see page 725).

Ad Hoc Reports

Ad hoc reports are developed in InfoView using Web Intelligence. You can create ad hoc reports.

Note: For more information about ad hoc reporting, see the BusinessObjects Enterprise documentation. To access the documentation, click the help icon in InfoView.

Web Intelligence Interface

In InfoView, Web Intelligence provides a custom reporting tool and it is intended for users who want to create basic reports easily without writing queries. Web Intelligence uses predefined report models and templates that manage data connections, querying, and data relationships, so you need only drag and drop data fields onto a template to create tabular or matrix reports. You can use all InfoView features that are described in the *InfoView User Guide*.

A Web Intelligence document references a BusinessObjects object named a universe to create reports. You build new reports on BusinessObjects universes. CA Business Intelligence provides the CA SDM universe with the product.

How Ad Hoc Reporting Works

Ad hoc reporting works as follows:

1. You start the ad hoc report by first selecting the CA SDM universe in InfoView.
2. The universe maps to the CA SDM database containing corporate business information. When you connect to a universe, Web Intelligence automatically launches the document editor selected on the Web Intelligence Document Preferences page in InfoView.

Note: For information on setting preferences, see the *InfoView User's Guide*.

3. Using the universe, you build queries to retrieve data to use in a report from objects grouped in folders called classes. Each class can also contain one or more subclasses. Subclasses contain objects that are a further subcategory of the objects in the upper level of the class. Classes and Objects are presented in a tree structure (hierarchy) in the Data tab.

Note: For instructions on how to build queries for Web Intelligence documents, see the BusinessObjects Enterprise documentation.

4. After you have built your query, by adding the required objects to the Result Objects pane, and defined query filter properties, you are ready to run the query. When you run a query, the universe asks the database to retrieve the data that corresponds to the demands of each of the objects in the query. You run a query by clicking Run Query on the toolbar.

Important! When selecting objects for your query, always choose objects, measures and filters from only one universe base class and its subclasses. If you include objects, measures and filters from multiple universe base classes in a query, you might receive unexpected results and poor performance.

Example Ad Hoc Reports

You can work with Web Intelligence and query examples to create ad hoc reports using the default CA SDM universe. When you create a report, make sure that you save your report at regular intervals. If the Web Intelligence connection session times out, you will lose your report modifications.

Note: For information about changing the session time-out value, see the *Implementation Guide*.

Example: View All Open Priority 1 and 2 Requests for All Users

In this example, you specify prompt values that gather additional data before generating the report.

Example: View all open priority 1 and 2 requests for all users

1. Click New, Web Intelligence Document from the menu bar.
The New Web Intelligence Document page appears.
2. Select the CA SDM universe in the Universe column.
The Web Intelligence document appears along with the universe information in the left pane.
Note: To improve navigation, close the Attached Service Type folder in the left pane.
3. Locate and then expand the Request, Request Detail folders in the left pane. Use the scroll bar to locate folders.
4. Select each of the following objects, then drag and drop them from the left pane to the top right pane under Results Objects.
 - Summary
 - Priority Symbol
 - Customer Combo Name
5. Select the following objects, then drag and drop them to the bottom right panel under Query Filters. Close the Request Detail folder when you are finished.
 - Status Symbol
 - Priority Symbol
6. Expand the Request Filters folder, then drag and drop the Customer Combo Name with Userid filter to the bottom right panel under Query Filters.
7. In the Query Filters pane, click the In list arrow on the Status Symbol filter, and select the Equal to operator.
8. Select Value(s) from list from the Operand Type menu.
The List of Values dialog box appears.
9. In the Status Symbol list, select Open, then click the green arrow.
The value appears in the Value(s) Selected field.
10. Click OK.
11. In the Query Filters pane, click the In list dropdown on the Priority Symbol filter, and select the In List operator.

12. In the Priority Symbol list, multi-select 1 and 2, then click the green arrow.

The values 1 and 2 appear in the Value(s) Selected field.

13. Click OK.

14. Click Run Query on the toolbar to generate the report.

Note: Since the Customer Combo Name with Userid quick filter uses the Prompt function, the Prompts dialog box appears. This dialog allows you to gather additional data before the report is generated.

15. In the Prompts window, select the first option ALL, then click the green arrow button.

All user names appear in the Select Customer list.

16. Click Run Query to generate the report.

The report appears in a new window.

17. Double-click the report's title and type Priority 1 & 2 Open Requests for All Users in the text box. You can change the font size to 16 or other sizes.

18. To change other aspects of the report, click the Properties tab.

19. To sub-group the data by Customer Last Name, right-click any data cell in the Customer Last Name column, and select Set as Section.

20. Click Save and specify a location using the Save As option.

Example: View All Open Requests that Do Not Include a Status of Work In Progress

In this example, you specify the Count reporting function to return the total number of open requests.

Example: View all open requests that do not have a status of work in progress

1. Click New, Web Intelligence Document from the menu bar.

The New Web Intelligence Document page appears.

2. Select the CA SDM universe in the Universe column.

The Web Intelligence document appears along with the universe information in the left pane.

Note: To improve navigation, close the Attached Service Type folder in the left pane.

3. Locate and then expand the Request, Request Detail folders in the left pane. (Use the scroll bar to navigate folders.)

4. Select each of the following objects, then drag and drop them from the left pane to the top right pane under Results Objects.
 - Status Symbol
 - Ref Num
 - Summary
5. Select the Status Symbol object, then drag and drop it to the bottom right panel under Query Filters. Close the Request Detail folder.
6. Expand the Request Filters folder, then drag and drop the Active filter to the bottom right panel under Query Filters.
7. In the Query Filters pane, click the In list arrow on the Status Symbol object, and select the Not Equal to operator.
8. Select Work in Progress from the Operand Type menu.
9. Click Edit Report on the main toolbar.

The Report Viewer opens.
10. In the report, select the Ref Num column. This action enables the Insert Sum command on the Reporting toolbar.
11. Click the Insert Sum command, then select Count from the menu.
12. In the report, right-click the =[Status Symbol] cell, then select Set as Section from the context menu.
13. Click Refresh Data on the main toolbar to run the report.
14. Click Save.

The report displays the total number (count) of requests that *do not* have a status of Work in Progress.

Example: View All Closed Requests in the Last 30 Days for Users Whose Last Name Begins with "C"

In this example, you specify prompt values for the Close Date and Customer Last Name that return all closed requests based on these values.

Example: View all closed requests in the last 30 days for users whose last name begins with "C"

1. Click New, Web Intelligence Document from the menu bar.

The New Web Intelligence Document page appears.

2. Select the CA SDM universe in the Universe column.

The Web Intelligence document appears along with the universe information in the left pane.

Note: To improve navigation, close the Attached Service Type folder in the left pane.

3. Locate and expand the Request, Request Detail folders in the left pane. Use the scroll bar to navigate folders.
4. Select each of the following objects, then drag and drop them from the left pane to the top right pane under Results Objects.
 - Ref Num
 - Close Date
 - Customer Last Name
 - Summary
5. Select each of the following objects, and drag them to the bottom right panel under Query Filters.
 - Status Symbol
 - Close Date
 - Customer Last Name
6. Click the In list arrow on the Status Symbol filter and select Value(s) from list from the Operand Type menu.

The List of Values dialog box appears.
7. Select the following values in the Status Symbol list:
 - Closed
 - Closed-Unresolved
 - Close Requested
 - Problem-Closed
 - Problem-Fixed
 - Resolved
8. Click the green arrow.

The values appear in the Value(s) Selected field.
9. Click OK.
10. In the Query Filters pane, click the In list arrow on the Close Date object and select the Between operator.
11. Click the In list arrow and select Prompt from the first and second Operand Type menus.

12. Click the In list arrow on the Customer Last Name object and select the Matches pattern operator.
13. Click the In list arrow and select Prompt from the Operand Type menu.
14. Click the Prompt Properties icon that appears next to the In list arrow.
The Prompt dialog box appears.
15. In the Prompt text field, enter a pattern for Customer Last Name.
16. Click OK.
17. Click Run Query on the toolbar to generate the report.

Note: Since the Close Date and Customer Last Name objects use the Prompt function, the Prompts dialog box appears. This dialog allows you to gather additional data before the report is generated.

18. In the Prompts window, select or type the prompt values for each prompt as follows:
 - Enter Close Date (Start): Specify a date that is 30 days earlier than today's date.
 - Enter Close Date (End): Specify today's date.
 - Enter a pattern for Customer Last Name: Specify C%.
19. Click Run Query to generate the report.
The report appears in a new window.
20. Double-click the report's title and type View All Closed Requests in Last 30 Days for Users Whose Last Name Begins with "C" in the text box. You can change the font size to 16 or other sizes.
21. To change other aspects of the report, click the Properties tab.
22. Click Save and specify a location using the Save As option.

Dashboard Reports

You can use dashboard reporting to monitor daily operations for all CA SDM ticket types (request/incident/problem, change order, or issue), Knowledge Management, Support Automation, and CMDB in BusinessObjects InfoView. Each report contains analytics about the top performers working on active tickets, so you can monitor their progress.

With dashboard reporting, you can:

- View summary and detail information about active tickets by priority, analyst, category, or group.
- Find out how many tickets were resolved within a particular time frame and much more.
- Edit, print, track, and save reports in other formats, such as .xls and .pdf.

You can work with individual predefined dashboard reports or use the corporate dashboard to view all CA SDM daily operations in a single view. The corporate dashboard shares vital information about daily operations for all ticket types (request/incident/problem, change order, or request) by priority, analyst, category, or group.

Note: When you work with dashboard reporting, you can use the InfoView features that are described in the Working with Dashboard and Analytics section of the *BusinessObjects InfoView User's Guide*. To access the User's Guide, click the help icon in InfoView.

View Dashboards and Reports in InfoView

The administrator account can log on and view the reports in BusinessObjects InfoView. The folders and reports that users see in InfoView depend on the account they log on to and the rights granted to the account.

To verify the dashboards reports in InfoView

1. Log in to CA SDM.
2. From the Reports tab, click the InfoView button.
InfoView appears in your web browser.
3. On the InfoView toolbar, select Document List.
Folders and objects appear in the navigation area.
4. Select Public Folders, CA Reports, CA SDM Dashboards.
The following folders appear: Change Order, Incident, Issue, Problem, Request, Configuration, and Knowledge.
5. Select the appropriate report.
The report appears in the right pane.
6. Double-click the title to display a report.
The report is open for viewing.

Write CA Business Intelligence Reports

You can write CA Business Intelligence reports for CA SDM.

More information:

[CA SDM ODBC Driver](#) (see page 736)

[Write SQL for BusinessObjects Reports](#) (see page 737)

[PDM Functions](#) (see page 737)

[Attribute Aliases](#) (see page 740)

[pdm_isql Interactive SQL](#) (see page 740)

CA SDM ODBC Driver

Business Objects reporting applications (Crystal Reports and Web Intelligence) access data using an ODBC driver that connects directly with the CA SDM object engine named the domsrvr.

This connection provides a number of benefits:

- SELECT statements used by BusinessObjects reports reference objects and attributes using their CA SDM names (in other words, their Majic names). For example, a SELECT statement for a report on contacts could be written:
`SELECT combo_name FROM cnt WHERE last_name LIKE 'smith%'`
- The CA SDM ODBC driver maps attributes like combo_name and objects like cnt to their corresponding DBMS name or names.
- All CA SDM security, including data partition and tenancy restrictions, is automatically applied to reports. All connections between BusinessObjects and CA SDM are associated with a CA SDM contact, and the CA SDM ODBC edits input SELECT statements to enforce the security restrictions associated with the end-user reporting role. BusinessObjects does not connect directly to the database.
- The CA SDM Attribute Alias feature "flattens" or de-normalizes the CA SDM database. Attribute aliases are additional attributes in CA SDM objects that allow a query to reference attributes in joined tables without an explicit join, allowing the base table to be used as if it were a reporting view.
- The CA SDM ODBC driver supports date literals in queries, and automatically translates values in CA SDM internal date format to a standard DBMS date.

Important! The CA SDM ODBC driver is supported only for BusinessObjects reporting (Crystal and Web Intelligence). CA SDM does not provide a standalone ODBC client, and does not recommend use of the ODBC driver with applications other than BusinessObjects.

Write SQL for BusinessObjects Reports

All SQL statements used by BusinessObjects reports are processed by the CA SDM ODBC driver. The ODBC driver supports standard SQL 92 SELECT statements with the following changes and extensions:

- CA SDM object names are used in place of DBMS table names, and CA SDM attribute names are used in place of DBMS column names.
- A CA SDM attribute alias name can be used anywhere in the query where a column name is valid. The ODBC driver replaces the attribute alias reference with one or more joins.

Note: For more information, see [Attribute Aliases](#) (see page 740).

- CA SDM DERIVED attributes (such as combo_name) can be used in the selection list only. They are not supported in any other part of the query, including in the WHERE clause.

Note: Many Combo Name With Userid objects have been provided in the universe, such as the Customer Combo Name With Userid object used as a filter in the sample ad hoc report provided in the *Administration Guide*. These objects allow Combo Name to be used as filter prompts in ad hoc queries with Web Intelligence, to overcome the limitation of including Combo Name in the WHERE clause. They present both Combo Name and Userid in the filter prompt, but use only the selected Userids in the resultant SQL query.

- Queries can contain date literals in either of the forms:
d'yyyy-mm-dd hh:mm:ss xm' (where xm is either am or pm)
ts'yyyy-mm-dd hh:mm:ss'

These literals can be used anywhere in the query. The ODBC driver automatically converts them to CA SDM internal date format (the number of seconds from midnight January 1, 1970).

PDM Functions

To assist in working with specialized CA SDM features and data types, the ODBC driver extends SQL to support a number of additional query functions. All driver-supported functions begin with the string "Pdm", and are known as PDM functions as described in the following table:

PDM Functions	Description
PdmAddDays([date,] count)	When used with one argument, adds the number of days in its argument to today's date and returns the result. When used with two arguments, adds the number of days in its second argument to the value of the date column specified in its first argument and returns the result. This function may be used anywhere in the query

PDM Functions	Description
PdmAddMonths([date],[count])	When used with one argument, adds the number of months in its argument to today's date and returns the result. When used with two arguments, adds the number of months in its second argument to the value of the date column specified in its first argument and returns the result. The single argument form can be used anywhere in the query. The two-argument form can be used only in the selection list.
PdmDay([date])	When used with no arguments, returns the current day as an integer. When used with one argument, returns the day associated with the value of the date column specified in its argument. The zero argument form can be used anywhere in the query. The one-argument form can be used only in the selection list.
PdmDownTime(slaName, workshift, startDate, endDate)	Calculates the downtime between two dates under the specified SLA and workshift. This function can be used only in the selection list.
PdmMonth([date])	When used with no arguments, returns the current month as an integer from 1 to 12. When used with one argument, returns the month associated with the value of the date column specified in its argument. The zero argument form can be used anywhere in the query. The one-argument form can be used only in the selection list.
PdmMonthName([date])	When used with no arguments, returns the localized name of the current month ("January", "February", and so on). When used with one argument, returns the localized name of the value of the date column specified in its argument. The zero argument form can be used anywhere in the query. The one-argument form can be used only in the selection list.
PdmDay([date])	When used with no arguments, returns the current day as an integer. When used with one argument, returns the day associated with the value of the date column specified in its argument. The zero argument form can be used anywhere in the query. The one-argument form can be used only in the selection list.
PdmSeconds(date)	Returns the value of the date column specified in its argument in its raw form as the number of seconds from midnight January 1, 1970. This function can be used only in the selection list. The argument is required.
PdmString(column)	Returns the string equivalent of value of the column specified in its argument. This function can be used with UUID, date, or string columns. It can be used only in the selection list.

PDM Functions	Description
PdmToday()	PdmToday() [timeAdj [, day [, month [, year]]]]) Evaluates to today's date (in seconds from 1/1/1970), adjusted according to the arguments: timeAdj: -1—adjust time to beginning of day (0:00:00); +1—adjust time to end of day (23:59:59) day: negative—adjust date by number of days specified positive—set day to absolute value specified (or to last day of month, whichever is less) month: negative—adjust date by number of months specified positive—set month to absolute value specified (or to December (12), whichever is less) year: negative—adjust date by number of years specified positive—set year to absolute value specified Adjustments are applied in the order year, month, day. A zero or omitted argument is ignored.
PdmYear([date])	When used with no arguments, returns the current year as a four-digit integer. When used with one argument, returns the year associated with the value of the date column specified in its argument. The zero argument form can be used anywhere in the query. The one-argument form can be used only in the selection list.

Attribute Aliases

Attribute aliases are additional attributes in CA SDM objects that reference data from joined tables using Majic dotted join syntax (where the syntax `srelname.attrname` is a reference to attribute `attrname` in the table referenced by foreign key `srelname`). A large number of predefined attribute aliases are provided with CA SDM Release 12.7, with names that typically are the same as the corresponding Majic join, with underscores replacing the dots that indicate the join. For example, the following SELECT statement might be used for a report that lists information about the request assignees:

```
SELECT ref_num, assignee_combo_name, assignee_organization_name
FROM cr WHERE customer_last_name LIKE 'smith%'
```

The CA SDM ODBC driver automatically builds joins as required to access the tables referenced by attribute aliases. A user in the CA SDM administrator role can easily add new attribute aliases online, providing a column-at-a-time way to extend the view corresponding to an object.

To access the Attribute Alias table, select the Administration tab, and browse to CA SDM, Codes, Attribute Aliases.

pdm_isql Interactive SQL

A command line utility, `pdm_isql`, is provided with CA SDM to allow interactive entry of SQL SELECT statements. SELECT statements entered into `pdm_isql` are sent to the CA SDM ODBC driver, allowing you to test SQL SELECT statements outside of BusinessObjects.

To use `pdm_isql`

1. Ensure that `$NX_ROOT/bin` is in your path.

2. Enter the command:

```
pdm_isql
```

3. At the `pdm_isql` prompt, enter the command:

```
connect username*password@casd_hostname
```

(Where *username* and *password* are valid CA SDM login credentials, and the host name is the host name of a CA SDM server with a web engine.)

4. Enter SQL select statements followed by a semicolon.

Chapter 18: Generating Reports in CA SDM

This section contains the following topics:

[Generate Reports](#) (see page 741)

[Database Views](#) (see page 741)

[Reports Method Setup](#) (see page 745)

[Reports Formatting](#) (see page 745)

[Column Sort Order Modification](#) (see page 746)

[Summary and Detail Reports](#) (see page 746)

[Analysis Reports](#) (see page 747)

Generate Reports

CA SDM provides a variety of built-in reporting capabilities and options including the ability to do the following:

- Print forms for assets and individual change orders, issues, incidents, problems, and more.
- Generate summary or detail reports for lists of change orders, issues, incidents, problems, requests, configuration items.
- Generate analysis reports.

Note: For more information about how to use the CA SDM reports, see *Generating Reports* in the *Online Help*. In addition to the reporting features that are built into CA SDM, you can create and generate custom reports based on the tables and views in your CA SDM database. For more information about custom reports, see the *Implementation Guide*.

Database Views

The CA SDM database is a repository for data entered and used to operate your service desk. CA SDM provides several database views and enables you to create customized reports of the database.

These views are of two types:

- Basic
- Advanced

More information:

[View Field Descriptions](#) (see page 899)

Basic Views

The basic views are based on CA SDM's tables. These views show data as long as the implementation makes use of Requests, Incidents, Problems, Change Orders, Issues, Assets, or Contacts / Groups.

Using the basic views, you can design and produce many reports, including:

- Detailed and summarized lists of open priority 1 requests sorted by request area.
- Detailed and summarized lists of change orders assigned to the level 1 group that have been open for more than 24 hours, sorted by priority.
- Counts of issues open for more than a specified number of days, sorted by priority, assigned group, or priority and assigned group.

The following basic views are provided:

View Name	Description
View_Contact_Full	All contacts, including their contact type, location, organizations, and service types
View_Contact_to_Environment	All contacts and their environment (assets)
View_Group	All the groups defined in the database
View_Group_To_Contact	All contacts (including managers) within their group assignments
View_Request	All requests, including their service types, severities, urgencies, categories, assets, impact numbers, assignees by name, customers by name, groups, statuses, and priorities
View_Act_Log	The request activity log information, including activity type and analyst name
View_Request_to_Act_Log	All requests with their activity logs (this view joins View_Request and View_Act_Log)
View_Request_to_Properties	Requests and their properties (this may not include all requests, especially if no properties are assigned to them)

View Name	Description
View_Change	All change orders, including their statuses, priorities, categories, organizations, affected end users by name, requesters by name, assignees by name, groups, service types, and impact numbers
View_Change_Act_Log	The change order activity log information
View_Change_to_Change_Act_Log	All change orders with their activity logs (this view joins View_Change and View_Change_Act_Log)
View_Change_to_Properties	Change orders and their properties (this may not include all change orders, especially if no properties are assigned to them)
View_Change_to_Change_WF	Change orders and their workflow tasks (this may not include all change orders, especially if no workflow tasks are assigned to them)
View_Change_to_Assets	Change orders and their assets (this may not include all change orders, especially if they have no assets)
View_Change_to_Requests	Change orders with basic information about attached requests (this view joins View_Change and the Call_Request table, which may not include all change orders—especially if no requests are attached to them)
View_Issue	All issues, including their statuses, priorities, categories, organizations, affected end users by name, assignees by name, groups, service types, and so on
View_Issue_Act_Log	The issue activity log information
View_Issue_to_Issue_Act_Log	All issues with their activity logs (this view joins View_Issue and View_Issue_Act_Log)
View_Issue_to_Properties	Issues and their properties (this may not include all issues, especially if no properties are assigned to them)
View_Issue_to_Issue_WF	Issues and their workflow tasks (this may not include all issues, especially if no workflow tasks are assigned to them)
View_Issue_to_Assets	Issues and their assets (this may not include all issues, especially if they have no assets)

Advanced Views

The advanced views are based on CA SDM's audit_log table.

Note: You must install the audit_ins and /or the audit_upd options in the Options Manager and restart the system before using these views.

Using the advanced views, you can report on the duration of a ticket (that is, request, change order, or issue) in a particular state. For example, you can report on the following:

- Show the length of time, between opening and assignment to L2 for requests opened since January 1, 2002
- Show the length of time, issues remained priority 3 before escalating to priority 2 for issues opened since January 1, 2002.

The following advanced views are provided, which are views of the audit_log table in the CA SDM database that specifically query for the status, priority, group, or assignee attributes:

View Name	Description
View_Audit_Status	All tickets sorted by the time in each status (does not include tickets that have no status)
View_Audit_Priority	All tickets sorted by the time in each priority
View_Audit_Group	All tickets sorted by the time in each group assignment (does not include tickets that have no group)
View_Audit_Assignee	All tickets sorted by the time in each individual assignment (does not include tickets that have no assignee)

Important! You must install audit logging options using the Options Manager, to view data in the advanced views. The descriptions for the Audit Options audit_ins and audit_upd in the Online Help provide more information.

More information:

[Options Manager Usage](#) (see page 317)

Reports Method Setup

Report methods let you specify where report results are to be sent when you select a report. Examples are the summary and detail reports available on the Reports menu and the reports available from the Analysis menu. Several predefined report methods are provided, and you can modify them.

Note: For information about how to define report methods, see the *Online Help*.

Reports Formatting

Using the `dataent.fmt` file found in `$NX_ROOT/fig/cfg` (UNIX) or `installation_directory\fig\cfg` (Windows) you can customize various data formats on the reports you print. You can view and modify this file using any text editor (Windows users should use WordPad to edit the file).

The following lines in the file control some of the date and time formats:

```
default_date = "long_date"
short_date = "M/D/YY// Enter a date as MM/DD/YY"
long_date = "MM/DD/YYYY// Enter a date as MM/DD/YY"
default_tod = "hour_12"
hour_12 = "h:mm:ss a(am,pm)// Enter a time of day as hh:mm:ss am/pm"
hour_24 = "HH:mm:ss// Enter a time of day as 00:00:00 - 23:59:59"
hms_12 = "h:mm:ss a(am,pm)// Enter a time of day as hh:mm:ss am/pm"
hms_24 = "HH:mm:ss// Enter a time of day as 00:00:00 - 23:59:59"
default_date_time = "date_time12"
date_time12 = "M/DD/YYYY h:mm:ss a(am,pm)// Enter a date and time as MM/DD/YY hh:mm:ss am/pm"
date_time24 = "M/DD/YYYY HH:mm:ss// Enter a date and time as MM/DD/YY 00:00:00-23:59:59"
```

The following is an example of how you might change these lines to support European dates and times:

```
default_date = "long_date"
short_date = "D/M/YY// Enter a date as DD/MM/YY"
long_date = "DD/MM/YYYY// Enter a date as DD/MM/YYYY"
default_tod = "hour_24"
hour_12 = "h:mm:ss a(am,pm)// Enter a time of day as hh:mm:ss am/pm"
hour_24 = "HH:mm:ss// Enter a time of day as 00:00:00 - 23:59:59"
hms_12 = "h:mm:ss a(am,pm)// Enter a time of day as hh:mm:ss am/pm"
hms_24 = "HH:mm:ss// Enter a time of day as 00:00:00 - 23:59:59"
default_date_time = "date_time24"
date_time12 = "M/DD/YYYY h:mm:ss a(am,pm)// Enter a date and time as MM/DD/YY hh:mm:ss am/pm"
date_time24 = "M/DD/YYYY HH:mm:ss// Enter a date and time as MM/DD/YY 00:00:00-23:59:59"
```

Note: If you are using `pdm_extract` to export data from the CA SDM database to another system, and want to extract data to use the date format specified in the `dataent.fmt` file, use the `-d` flag when you invoke `pdm_extract`.

Column Sort Order Modification

Administrators can modify the sort order of columns that display in reports using Web Screen Painter.

Users can sort data in the follows orders:

- Ascending (A to Z, zero to 9, or earliest to latest date)
- Descending order (Z to A, 9 to zero, or latest to earliest date)

Note: For more information, see the *Web Screen Painter Help*.

Summary and Detail Reports

CA SDM has built-in summary and detail reporting options. You can select specific records for a report using the search feature of the list windows. For example, from the Request List window you can enter search criteria to create a list of requests that you can then use to generate a report.

To print or view summary and detail reports, you must first select the records you want to include in the report. You can also print a single page detail report for each record by selecting Print Form from the File menu within any detail page. To print a report for a newly created record, you must first save the record.

Note: For more information about using reports, see the *Online Help*.

Analysis Reports

The CA SDM Administration Tab also provides the following built-in analysis reports for a global and detailed perspective on the service desk process:

- Request / Issue Reports
- Request Area / Issue Category Reports
- Request Area Priority / Issue Category Priority Reports

You can display the analysis report for today, for the past thirty days, and year-to-date.

More information:

[Generate Request or Issue Reports](#) (see page 747)

[Generate Request Area or Issue Category Reports](#) (see page 747)

[Generate Request Area Priority or Issue Category Priority Reports](#) (see page 748)

Generate Request or Issue Reports

Request or Issue reports provide statistics, such as number of requests or issues opened, number of requests or issues closed, average time open, and average time until closed, for a specified period. This report is sorted by date, oldest to newest.

To generate Request or Issue reports from the web interface

1. Select the Administration tab.
2. Expand the Service Desk and Analysis nodes.
3. Select Request or Issue.
4. Select the appropriate report based on the time frame you want.

Generate Request Area or Issue Category Reports

Request Area or Issue Category reports provide the number of requests or issues opened in the specified period for each request area or issue category. This report is sorted alphabetically by request area or issue category.

To generate Request Area or Issue Category report from the web interface

1. Select the Administration tab.
2. Expand the Service Desk and Analysis nodes.

3. Select Request or Issue.
4. Select the appropriate report based on the time frame you want.

Generate Request Area Priority or Issue Category Priority Reports

Request Area Priority or Issue Category Priority reports provide the number of requests or issues opened by priority in the specified period for each request area or issue category. This report is sorted by priority (highest to lowest) and then alphabetically by request area or issue category.

To generate a Request Area Priority or Issue Category Priority report from the web interface

1. Select the Administration tab.
2. Expand the Service Desk and Analysis nodes.
3. Select Request or Issue.
4. Select the appropriate report based on the time frame you want.

Chapter 19: Managing Knowledge

This section contains the following topics:

[Knowledge Management](#) (see page 749)

[Find Procedures for Knowledge Management](#) (see page 750)

[Knowledge and Multi-Tenancy](#) (see page 750)

[How to Set Up a Knowledge Base](#) (see page 751)

[How to Use Documents in the Knowledge Base](#) (see page 753)

[Knowledge Search](#) (see page 758)

[Forums](#) (see page 759)

[Knowledge Tree Documents](#) (see page 759)

[Knowledge Documents Schedule](#) (see page 760)

[Access Export/Import](#) (see page 769)

[Web Services](#) (see page 782)

Knowledge Management

Knowledge management refers to the concept of finding, organizing, and publishing knowledge. Knowledge management captures information quickly and efficiently and then delivers this information to a user or group. The information that is captured and made available for retrieval is referred to as a knowledge base.

Users access a knowledge base by using a search engine. Knowledge Management lets you create and manage content that resides in a knowledge base. You set category and document permissions to use groups or roles. Knowledge Management helps you provide customers with solutions to complex issues. Effective knowledge management quickly delivers solutions to customers through a process that is user-friendly and easy to navigate.

To manage knowledge effectively, you do the following:

- Create a meaningful hierarchy of content.
- Identify gaps in existing knowledge.
- Perform updates and maintenance to help ensure relevance of content.
- Measure the value of available content.

More information:

[Configure Knowledge Management Data Partition Constraints for Role-Based Permissions](#) (see page 172)

Find Procedures for Knowledge Management

CA SDM provides step-by-step procedures for Knowledge Management in this guide and in the *Online Help*.

To find step-by-step procedures for Knowledge management in the *Online Help*

1. Log in to CA SDM.

The CA SDM main page appears.

2. Do any of the following actions:

- Click Help, Help from the main menu and navigate the *Online Help* to Administration, Knowledge Management.

The Knowledge Management topics appear and you can navigate the hierarchy to locate the information you want.

- Click Help, Help on This Window from the main menu.

The *Online Help* appears and displays a help topic for the page for which you want information.

Knowledge and Multi-Tenancy

Knowledge documents and knowledge categories can be specific to a tenant, shared among multiple tenants, or shared as public. When you search for a knowledge document, the results are limited to documents viewable by the tenant associated with the ticket. FAQ search results are tenant-specific. Ratings for public tenant documents must be kept separately for each tenant.

You can administer the following settings at the tenant level:

- Action Content
- Document Templates
- Knowledge Approval Process
- Recommended Documents
- Knowledge Management ROI
- Reports

All other Knowledge Management system settings are administered globally, at the service provider level.

Note: Multi-tenancy is configured in Options Manager using the Administration tab.

How to Set Up a Knowledge Base

Effective knowledge management consists of more than the development of a large knowledge base. Knowledge management also involves the implementation of processes and procedures for maintaining relevant and accurate content.

To set up a knowledge base, do the following tasks:

1. Analyze how the knowledge base will be used and the scope of it.
2. Consider who the customers are and what information they will likely need.
3. Develop a strategic plan that identifies the potential issues that can result in the need for customer support. The problems you identify can guide you in determining what content must reside in your knowledge base.
4. Execute the development of your knowledge base. Create, organize, maintain, and secure this data. You can perform all these functions by using Knowledge Management.

Import Sample Knowledge Data

Sample Knowledge data from Knowledge Broker, PCHowTo, and Right Answers is provided for your use.

To use the sample Knowledge data, import the data to the Knowledge Management database.

Note: For more information about KT sample data, see the *Implementation Guide*.

Knowledge Base Monitoring

You can monitor the efficiency of the knowledge base using the following reporting tools. These tools let you view statistics on the usefulness of your documents and their effectiveness in solving problems.

Knowledge Report Card

Lists statistics for documents you have created. Each user has an individual Knowledge Report Card.

Web-Based Reports

Displays metrics that describe how knowledge is meeting user needs. Some of the most commonly used features include:

- Listing the most frequently accessed documents.
- Displaying user searches that did not return any results.

Knowledge Base Re-Index

The following changes made to search settings require that you re-index the knowledge base using the Knowledge Re-index utility:

- After importing knowledge
- After changing parse settings
- After mass deletions
- When search failures occur

Re-indexing is necessary because existing documents do not reflect any changes made to the synonyms, noise words, special terms, and other research parameters until they have been re-indexed. All new synonyms, noise words, and special terms must be re-indexed to help ensure that keyword searches of the knowledge base are current and accurate.

You run Knowledge Re-index from a command line. `pdm_k_reindex.exe` is the executable file for Knowledge re-index.

Note: Re-indexing the documents in the knowledge base can be a time-consuming operation, depending on the size of your database. We recommend that you run the Knowledge Re-index utility after all changes have been added.

Index and De-Index Queue Settings for Batch and Instant Processing

Indexing and De-Indexing both run a batch process to include a predefined number of documents at one run. These batch processes are used for performance optimization. If more documents are included in the batch, system performance increases.

The number of documents you can process is limited; the limit depends on the size of the documents and the linked attachments. The document size is calculated based on the pure text and its attachments. Image and format elements are not calculated.

Note: You can limit the size of attachments by navigating to Attachments Library, Repositories on the Administration tab and editing the repository to set the File Limit Size (KB).

The recommended batch maximum size is between 2-12 MB (per the EBR_MAX_INDEX_BATCH_SIZE parameter of the NX.env file and the average document size).

- If the average size of your document (including attachments) is approximately 0.1 MB, keep the default setting in NX.env:

```
@EBR_MAX_INDEX_BATCH_SIZE=128
@NX_EBR_INDEX_QUEUE_TIMEOUT=10
@NX_EBR_REINDEX_QUEUE_TIMEOUT=1
@NX_EBR_INDEX_QUEUE_ONLINE=Yes
@NX_EBR_NON_KD_INDEX_QUEUE_ONLINE=Yes
```

This setting means that one batch processes 128 documents, that batch executions have ten second intervals, and when in reindex, the wait interval between two batches is one second.

- If the average size of your document (including attachments) is approximately 0.5 MB, keep the default setting in NX.env

```
@EBR_MAX_INDEX_BATCH_SIZE=25
@NX_EBR_INDEX_QUEUE_TIMEOUT=10
@NX_EBR_REINDEX_QUEUE_TIMEOUT=10
@NX_EBR_INDEX_QUEUE_ONLINE=No
@NX_EBR_NON_KD_INDEX_QUEUE_ONLINE=No
```

This setting means that one batch processes 25 documents, that batch executions have ten second intervals, and when in reindex, the wait interval between two batches is ten seconds.

How to Use Documents in the Knowledge Base

Knowledge documents provide you with information about knowledge that is stored in the knowledge base. Creating quality knowledge requires input from several individuals. Each individual is responsible for performing specific tasks throughout various stages in the lifecycle of a knowledge document. Knowledge documents reside in the knowledge base and are managed as part of the following ongoing process:

1. Identify content to include in the knowledge base.
2. Create a knowledge document.

Knowledge documents are placed into categories that some organizations assign to owners.

Note: When a document is created or updated, it is placed in an owner Inbox. Until items are published, the items in the Inbox do not appear as resolutions and are not added to the knowledge base.

3. Revise the document.

After a document arrives in an Inbox folder on the scoreboard, users can perform their assigned tasks by modifying their documents according to their assigned roles.

All users with full (read/write) permission to the document can modify the document. The current owner has full permissions to the document, but not necessarily have explicit write permissions. Users can create versions or roll back to a previous version when a problem with the document is found.

4. Submit the document.

In addition submission from the customer or employee self-service interface, knowledge can also be submitted from CA SDM. This option lets the analyst submit a new resolution from an existing ticket. This option also provides a link between a problem and its resolution, and can help other users with similar problems find a resolution.

5. Publish the document.

After a document has passed through the complete approval process cycle, it can be published. A document that has been published becomes part of the viewable knowledge base on the start date, which is the current date by default. The document is only viewable by groups that have been granted access rights to read it. A user with full permissions can edit a published document.

6. Evaluate and decide whether to perform the following tasks:

- **Unpublish the document**—The owner of the document, a knowledge manager, or administrator disallows publishing the document for editing purposes.
- **Create a Rework Version**—Users with full editing permissions can create a rework-draft version of a published document while it remains online and available for viewing and searches. A rework version starts as a copy of the document that is replaced in the knowledge base after it is verified and republished.
- **Retire the document**—The owner of the document, a knowledge manager, or a system administrator can retire the document from the knowledge base.

The tasks that are performed and who performs each task can be defined to meet the approval process structure that exists in an organization.

Note: You can track knowledge activities from the Event Log tab on contact detail pages. For example, if an end user views a knowledge document, the event log updates to display the action performed and a link to the document. The event log also tracks the duration the user had the document opened.

Knowledge Submission from CA SDM

If you acquire knowledge that you think should be added to the knowledge base, you can submit it for consideration. In addition to submitting knowledge from the self-service interface, you can submit it from CA SDM. An analyst can submit a new resolution from an existing ticket, which provides a link between a problem and its resolution. In addition, it can help other users with similar problems find a resolution. In an organization that can have hundreds of open tickets simultaneously, this can save valuable time.

All knowledge documents must be assigned to a knowledge category. If the Incident/Request Area in CA SDM matches the knowledge categories in Knowledge Management, the category is automatically selected for the knowledge submission.

Knowledge Submission from Self-Service

By default, any user that is logged in to Knowledge Management can submit knowledge for consideration from the Submit Knowledge page. This page lets a customer submit knowledge without contacting their local service desk representative. After you submit knowledge, it passes through a publishing process where the knowledge is reviewed and edited before it is added to the knowledge base.

It is important to type information into every field that appears in the Submit Knowledge page. Pay special attention when completing the Summary field. Typically, this field contains a succinct overview of the document that you are submitting.

Document Attributes

Setting document attributes helps you manage documents in your knowledge pool. Document attributes can be updated to assign a new subject expert or owner to the document. You can also specify the date the document becomes available in the knowledge base and the date when it expires. By selecting different document templates, you can modify the appearance of each document.

Document Permissions

You can set, view, and edit permissions for a document. These permissions can be assigned to different groups of individuals. When setting permissions, you can decide to inherit permissions from the primary category of a document or specify new permissions. By default, documents inherit their permissions from their primary category. This default handles access permissions at the category level rather than the document level.

Important! When creating a knowledge document, verify that document permissions include users that can be later assigned on the document through the approval process. When a group is assigned to a document, users from that group may not have the permission to view the document. If the document is assigned to a specific user, default data partition constraints allow the user to view the document.

Resolution Editing

You can create knowledge documents with Knowledge Management using the HTML Editor. This feature lets you modify the Resolution field of a knowledge document in *What You See Is What You Get* (WYSIWYG) form. Using the HTML Editor improves the authoring process by enabling you to build documents without typing HTML code, format text, insert graphics and tables, and create links to other documents. These actions speed up the writing process and also provide an easy method of editing a document for organizational purposes.

The HTML Editor lets you do the following:

- **View HTML Code**—The Design Mode and Source Mode tabs on the HTML Editor toolbar lets you switch between different views.
- **Format Text**—You can type Document text directly into the HTML Editor window. You can then format the text using the toolbar and menu options.
- **Insert Graphics**—You can insert a graphic into your HTML text by selecting an image from the images library or by uploading your own graphic.
- **Insert Tables**—In a knowledge document, you can display large amounts of information neatly to users. You can use tables to format such information in rows and columns.
- **Insert Hyperlinks to Other Documents**—You can insert hyperlinks to knowledge documents when creating or changing a knowledge document. In this manner, a knowledge document can be linked to other relevant documents or URLs. When the HTML Editor is used in Design Mode, hyperlinks appear as blue, underlined text.
- **Insert Action Content**—You can insert Action Content (a live URL) into the Resolution field of a knowledge document which, when clicked by the end user, creates an incident, or performs some other action.

Document Publication Preparation

After you submit a knowledge document, you prepare it for publication by setting its attributes and permissions. Assigning controls such as settings and categories are important to help organize documents in an efficient way.

More information:

[How to Use Documents in the Knowledge Base](#) (see page 753)

[Document Attributes](#) (see page 755)

[Document Permissions](#) (see page 755)

Document Publication

After a document has passed through the complete approval process cycle, it can be published. A document that has been published becomes part of the viewable knowledge base on the start date, which is the current date by default. The document is only viewable by groups that have been granted access rights to read it.

When a knowledge document is published, the user is not typically permitted to modify the document unless it is unpublished first, and then modifications can be made. During this time, the knowledge document is offline and unavailable to users. The owner of the document, a knowledge manager, or a system administrator can unpublish the document using the Rework button and the Unpublish check box. Unpublishing a document returns it to draft status. An administrative user can then select the next step in the workflow process.

Version Documents

Using the document versioning capabilities of Knowledge Management, an analyst with editing privileges can create a *Rework-Draft* version of a document. A rework version starts as a copy of the document that is replaced in the knowledge base after it is verified and republished. The need to unpublish the document first is avoided.

Users with editing privileges can also perform the following versioning tasks:

- Save a draft version of a document.
- Roll back to a previous version if a problem with the current version occurs (Draft or Published). The Versions tab on the Update Document page is where users can select different versions for rollback to replace the current version.
- Track the number of document versions that are saved, deleted, and archived in the knowledge base.

The Knowledge Documents Inbox on the CA SDM Scoreboard is the repository for documents in all statuses including saved and assigned draft and rework-draft documents.

How to Manage Document Versions

Administrators can set up and manage document versions by performing the following steps:

1. Identify who can edit published documents and create Rework-Draft versions. The role being used for a particular contact record controls editing privileges.
2. Define an approval process template that groups tasks or steps to complete during the document lifecycle. By default, a built-in approval process template allows users to create documents.
3. Determine whether to use the document approval process. Analysts who are permitted to bypass the approval process can identify which tasks they want to start with when they create a Rework-Draft version.
4. Create archive and purge rules for document version maintenance.

Document Expiration

When a published document reaches its expiration date, the product typically retires it (that is, removes the document from the knowledge base and the approval process).

Document Archive and Purge

You can manage the size of your knowledge pool by using Archive and Purge, which removes old and unused documents automatically. Archive and Purge improves the efficiency of knowledge searches by returning only current documents.

Archive and Purge runs as a background process and automatically removes inactive records in the CA SDM database according to rules that you configure. These rules act on CA SDM objects at specific time intervals.

Knowledge Search

Knowledge Management provides users with the following options for retrieving knowledge:

Category browsing

Finds solutions based on categories. In each category, additional subcategories can narrow the search results to a set of solutions that are likely to be most relevant to an issue.

Knowledge tree

Asks a series of questions to guide the user to possible solutions. The responses lead the individual to a solution that appears to be most relevant.

Bookmark

Provides access to frequently viewed documents that are included in a bookmark list.

Forums

Forums let you communicate about existing issues. Using forums you can share documents globally or among predefined groups that work together in knowledge-sharing and brainstorming over existing challenges.

Forums broaden the scope of knowledge contributions by allowing communication on general questions, usability tips, and so on. You can create a forum from the Knowledge tab and from a service desk ticket.

Knowledge Tree Documents

The Knowledge Tree Designer is a visual tool used to create knowledge trees quickly and easily. A knowledge tree is a representation of expert knowledge in a particular area. The Knowledge Tree Designer lets analysts and knowledge engineers build detailed trees that guide the user through a series of questions and potential answers until they reach a resolution. This tool eliminates the need for specialized scripting or programming skills and can only involve an analyst working with a subject expert to create a tree design.

Knowledge trees can be complex in design. Therefore, it is recommended that you construct a diagram to map the design before creating a knowledge tree. The diagram must contain a series of question, possible responses, and associated resolutions. The hierarchy of questions and responses in the diagram you construct is the foundation on which the knowledge tree is built.

After you have mapped out a diagram of your knowledge tree, you build the tree using the Knowledge Tree Designer.

Knowledge Documents Schedule

The Knowledge Documents Schedule provides Change Managers, Knowledge Managers, and Level 2 Analysts a high-level view of key events in the knowledge management lifecycle. The following events can be scheduled for a Knowledge Document:

- Submission
- Publication
- Review
- Expiration

The Knowledge Management Schedule tab displays a similar calendar format as the Change Calendar, but instead of showing start times and end times, it only shows specific dates.

[Configuring scheduling views](#) (see page 639) involves setting macros statements.

The Knowledge Document Schedule can also be exported to iCalendar format.

Note: When exporting schedules on some calendar programs, selecting the Open option instead of Save causes the file to import incorrectly. To avoid this issue on Knowledge Management and Change Order schedules, select the Save option instead of Open. After you save the exported file, import it through the calendar program interface.

Knowledge Schedule Filter

The Knowledge Schedule Filter displays the following fields:

Tenant

Filters the search by tenant. This field displays for the privileged user in a multi-tenancy installation.

Event Type

Filters the search by the following event types:

- Submission
- Review
- Publication
- Retirement

Schedule Start Date

Specifies the date for the beginning of a range for filtering the history to show only entries for a specified time frame.

Schedule End Date

Specifies the date to specify the end of a range for filtering the history to show only entries for a specified time frame.

Schedule Timezone

Specifies a timezone to view your search results.

Note: If no timezone is selected, the events are displayed in your current timezone.

Owner

Defines the name of the contact assigned to maintain the document. Enter the name of the contact in "last name, first name" format or open the Contact Search dialog so you can locate and select a contact.

Assignee

Defines the name of the contact assigned to handle the record. Enter the name of the contact in "last name, first name" format or open the Contact Search dialog so you can locate and select a contact.

Subject Expert

Defines the name of a contact with expertise in the document's subject matter. Enter the name of the contact in "last name, first name" format open the Contact Search dialog so you can locate and select a contact.

Status

Select one of the following document statuses to perform your search:

- Draft
- Published
- Retired

Category

Defines a category by which to filter documents retrieved. Enter the name of the category in the box or open the Category Search window. When you click Search, the product returns only documents associated with the specified category.

Initial View

Select the view of the Knowledge Management Calendar you want to see:

Month

Displays a calendar for the full month that includes the earliest implementation start date specified in the Start Date field. The calendar shows abbreviated information about each change within the range (change number, start and end time, and first affected CI).

Week

Displays seven consecutive days in a single column, beginning with the earliest implementation start date specified in the Start Date field. The calendar includes summary information about each change within the specified range (change number, start and end time, summary description, assignee, group, and the first ten affected CIs).

Day

Displays a view similar to the week view, except that it shows only the day specified in the Start Date field.

***n* Days**

Displays a view similar to the week view, except that it continues for the specified number of days.

List

Displays a standard CA SDM list page.

Note: You can click the More icon to display the Additional Search Arguments field. This field is intended only for expert users who understand SQL and Majic and can use it to specify search arguments that are not available in the standard search filter fields. You can enter a SQL WHERE clause in this field to specify an additional search argument.

Knowledge Schedule Views

The Knowledge Documents Schedule has the following views:

Month

Displays a calendar for the full month that includes the earliest implementation start date specified in the Schedule Start Date field. The calendar shows abbreviated information about each change within the range (change number, start and end time, and first affected CI).

The Knowledge Documents Schedule has similar functionality to the Change Orders Month view with the following exceptions:

- Knowledge events only have a date (no time) and event type groups them, similar to change orders that are grouped for each time period and change type.

- Events types have the following default predefined colors, using the schedGroup macros:
 - Submission—Black
 - Review—Green
 - Publication—Blue
 - Retirement—Red

Note: If you see a review date referring to the past, it indicates that the review was probably not done.

Week

Displays seven consecutive days in a single column, beginning with the earliest implementation start date specified in the Schedule Start Date field. The calendar includes summary information about each change within the specified range (change number, start and end time, summary description, assignee, group, and the first ten affected CIs).

Day

Displays a view similar to the week view, except that it shows only the day specified in the Schedule Start Date field.

***n* Days**

Displays a view similar to the week view, except that it continues for the specified number of days.

List

Displays a standard CA SDM list page.

Scheduling View Configuration

You configure the monthly and weekly scheduling views by specifying pdm_macro statements in the <head> section of the HTML forms defining the schedule. We recommend using the Source View of Web Screen Painter to edit these forms.

Any form that displays a schedule must contain the following:

- A schedConfig macro
- At least one schedAttr macro
- At least one schedGroup macro

The configuration macros are in a separate source file referenced by a `pdm_include` statement in the main source file. This file lets you configure your schedule without modifying the main source file.

For example, the configuration macros for the Change Calendar form `list_chgsched.html` are in a file named `list_chgsched_config.html`. For the Knowledge Lifecycle Schedule, you can modify the `list_kdsched_config.html` using the same macros.

You can find `list_chgsched_config.html` and `list_kdsched_config.html` in the following directory:

`$NX_ROOT\bopcfg\www\html\web\analyst\`

schedConfig Macro—Configure Schedule

The `schedConfig` macro specifies that a form contains a schedule and provides basic configuration information. The following values are valid macro arguments:

autosearch=1|0

Specifies whether the schedule form reloads data from the server when the user selects a view outside the currently selected date range. Setting the value to 1 (default) causes the form to search automatically when the user selects a view with one or more days outside the date selection range of the search filter. Setting the value to 0 requires the user to press the Search button to initiate a search.

defaultView=0|1|7|30|99

Specifies the default view for the search filter as 0 (list), 1 (day), 7 (week), 30 (month), or 99 (n-day).

The specification for `defaultView` affects only the initial display of the search filter. After the schedule displays, CA SDM automatically keeps the filter view selection aligned with the current view.

Default: 30

firstday=0|1|2|3|4|5|6|7

Specifies the first weekday on the monthly view as a number between 0 (Sunday) and six (Saturday).

Default: 0

export=xxx|0

Specifies the code name of the template used for exporting in iCalendar format. Setting the value to 0 indicates the export feature and button are disabled.

Default: ChangeSchedule

legend=1|2|0

Specifies the location of the schedule legend showing the name and formatting of the groups on the schedule. You can set the value to 1 to position the legend above the schedule, or 2 to position the legend below the schedule. Set the value to 0 to disable the legend.

Default: 2

maxGroups=0/n

Specifies the maximum number of groups to be displayed in a single cell of the calendar month view.

If there are more than maxGroups scheduled for a single day, CA SDM displays only the first maxGroups-1, and replaces the last with a "...nn more changes" hyperlink that the user can mouseover or click to see the full list. Set the value to 0 to disable this feature and allow an unlimited number of events in a calendar cell.

Default: 4

nday=(n,n,...)

Specifies selections for the drop-down list for the n-day view.

The specification is a list of day counts that are to be included in the drop-down list, or 0 to indicate that the n-day drop-down list is omitted from the schedule. The first value specified is the default for the drop-down list.

Default: (3,7,14,28)

round=(hr,min)|0

Specifies whether schedule start and end dates are rounded when collecting change orders or knowledge documents into groups. Specify round=0 to disable rounding.

By default, schedule start and end date groups objects. All CA SDM dates include a time, and without rounding, objects scheduled as little as a minute apart would be in separate groups. Rounding determines the group after adjusting the start date to an earlier hour or minute and the end date to a later hour or minute.

The value of round specifies either an hour or a minute (but not both). Times are rounded to the nearest multiple of the value specified, for example:

round=(0,15) rounds to the nearest quarter hour

round=(0,30) rounds to the nearest half hour

round=1 rounds to the nearest hour

round=12 rounds to the nearest half day (12:00 AM or PM)

round=24 rounds to the nearest day

Default: (0,15)

timefmt=24hr|([am],[pm])

Specifies the format of times on the calendar views of the schedule.

The default value of 24hr specifies that times are displayed in 24 hour format (0:01 - 23:59). The alternative value of (am,pm) specifies a suffix for either morning or afternoon times, or both.

Note: All schedConfig arguments are optional.

schedAttr Macro—Specify a Stored Attribute

The schedAttr macro specifies an attribute stored for each item selected for the list. Stored attributes are available for hover information on the monthly view, for the detailed or summary information in views other than the monthly view, and in the setSchedEvents() JavaScript function. The following values are valid macro arguments:

attr=xxxx

(Required) Specifies an attribute from the object on the schedule, such as a change order or Knowledge Document. Dotted attributes are permitted. The keyword attribute name C inn can be used on the Change Calendar to specify that first nn CIs associated with the change order are included with the information stored.

Note: This argument is the only required argument for the schedAttr macro.

attrRef=.COMMON_NAME|xxxx

Stores the attribute of the referenced table stored for an SREL attribute (ignored for non-SREL attributes). The attribute name specified must be prefixed with a dot.

Default: .COMMON_NAME

label=

Displays a label for the attribute on the n-day view.

Default: the Majic DISPLAY_NAME of the attribute

ident=1|0

Specifies whether the attribute is an identifier for the object (such as a reference number of a change order). Identifier attributes are displayed without a label in front of the group name in hover information and the n-day view.

Default: 0

detail=1|0

Specifies whether the attribute is included in the detail information shown on views other than the monthly view. Detail information is the information shown when the Summary Only check box on the view is not selected.

Default: 1

hoverInfo=1|0

Specifies whether the attribute is included in the hover information pop-up displayed on the monthly view when the mouse pointer hovers over a group, or the user presses Alt+Right Arrow when focus is on the group.

Default: 0

summary=1|0

Specifies whether the attribute is included in the detail information shown on views other than the monthly view. Detail information is the information shown when the Summary Only check box on the view is not selected.

Default: 0

Note: CA SDM displays attributes in summary, detail, or hover information in the same order as their schedAttr macros.

schedGroup Macro—Specify an Event Group

The schedGroup macro specifies the name and color coding of a group of items. The monthly view aggregates all items in a group into a single event. Views other than the monthly show individual items in the format for the group to which they belong. The following optional values are valid macro arguments:

grpname=xxx

(Required) Specifies the name of the group. The macro automatically assigns a number to the group and assigns the number to a JavaScript variable with a name of the form schedGroup_XXX, where XXX is the name of the group. This variable can be used in the setSchedEvents() JavaScript function to create an event belonging to the group.

Note: This argument is the only required argument for the schedGroup macro.

label=xxx

Specifies a label for the group. If specified, the label is displayed in all views.

legend=xxx|0

Displays a description of the group for the legend that appears at the bottom of the schedule. Groups appear in the legend if at least one example of the group exists in the current view. Specifying 0 causes the group always to be excluded from the legend.

Default: 0

color=black|color

Specifies the color of text in items of this group. You can specify color in CSS format, either a valid web color or a hexadecimal value preceded by a pound sign.

Example: Enter either "#FF0000" or "red" for red.

Default: black

bgcolor=white|color

Specifies the background color of items of this group. You can specify bgcolor in CSS format, either a valid web color or a hexadecimal value preceded by a pound sign.

Example: Enter either "#FF000" or "red" for red.

Default: white.

style=normal|bold|italic

Specifies the style of text of this group in the normal, bold, or italic style.

Default: normal

The setSchedEvents() JavaScript Function

The setSchedEvents() JavaScript function creates events in the schedule. Modify this function when you want to view any new group objects. The predefined group objects appear by default.

CA SDM calls setSchedEvents() once for each object (change order or knowledge document) selected by the schedule search filter. The function creates events for the object by calling a second function, schedEvent(), and passing the group ID, start date, and end date of the event.

The function can create any number of events (including zero) for an object. The default setSchedEvents() function for the Change Calendar (list_chgsched.html) creates one event for each change order and groups change orders by change type. This function is coded as follows:

```
1.      function setSchedEvents( chg )
2.      {
3.      var grpnum;
4.      switch( chg["chgtype"] - 0 ) {
5.          case 100: grpnum = schedGroup_std; break;
6.          case 300: grpnum = schedGroup_emer; break;
7.          default: grpnum = schedGroup_norm; break;
8.      }
9.      chg.schedEvent( grpnum, chg["sched_start_date"], chg["sched_end_date"] );
10.     }
```

The case parameter specifies the change type ID. To list the case IDs, see [Create a Change Type](#).

The function has a single argument of a JavaScript object containing the attributes specified by schedAttr macros. The switch statement in lines 4-8 examines the chgtype attribute of the change order, and assigns the appropriate group number from one of the schedGroup_xxxx variables defined by previous schedGroup macros. On line 9, it calls the schedEvent() function to create an event in the schedule, passing the group number previously assigned and the schedule start and end dates. The dates are available in the argument object because they were specified in earlier schedAttr macros.

Access Export/Import

The Knowledge Export/Import tool migrates and synchronizes data between different Knowledge Management systems and enables export/import of Knowledge Documents. Use this tool to define export/import transactions.

Note: If multi-tenancy is installed, all tenanted and public categories that were imported from a previous system display as public in the new multi-tenancy system, but imported documents can be tenant-specific.

To access the KEIT page, browse to Knowledge, Documents, Export/Import on the Administration tab.

The Export/Import appears and contains the following elements:

Export Transactions

Displays a list of exports.

Export/Import Templates

Displays a list of templates that you can edit.

Import

Displays a list of packages available for import.

Import Transactions

Displays a list of imports.

Note: When importing or exporting knowledge from the web interface, Knowledge Management uses the default Command Line Utility Role, and not the current role the user is logged into CA SDM. The default Command Line Utility Role is set on the access type for the user.

How to Export/Import Knowledge

Export/Import templates let you define an export/import transaction.

To export or import knowledge, do the following steps:

1. Do any of the following:
 - Create an Export/Import template to create an Export package.
 - Use an existing Export/Import Template to create an Export package.
2. Export the data to another system by doing any of the following:
 - Click Export on the Template Detail page.
 - In the Export Transaction page, right-click a transaction and click Retry on the shortcut menu.
 - [Use the pdm ket utility](#) (see page 780)
3. Import knowledge data from the package by doing any of the following:
 - Copy the knowledge package to a targeted CA SDM installation and place it in the \$NX_ROOT/site/keit/import directory.
 - On the Knowledge Import Packages List page, right-click a Knowledge Package, and click Import on the shortcut menu.
 - In the Import Transaction page, right-click a transaction and click Retry on the shortcut menu.
 - [Use the pdm kit txt utility](#) (see page 781)

Note: The knowledge package name must start with "package_" prefix for it to be listed in the Knowledge Import page.

Important! The r11.0 import utility has been renamed pdm_kit_txt.exe to allow importing r11.0 text files. This utility does not support any of the r12.0 import enhancements.

More information:

[Export/Import Packages](#) (see page 773)

Export Transactions

You can sort transactions by package, template and status. All KEIT transactions are recorded in the following log files:

stdlog

Logs CA SDM information.

keitstat.log

Provides statistical information on the KEIT operation.

keitinfo.log

Provides detailed information on the KEIT operation.

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

The Export Transactions page contains the following fields:

Package

Displays the export package name.

Run By

Displays the user who ran the export.

Start Time

Displays the start time of the transaction.

End Time

Displays the end time of the transaction.

Document Count

Displays the number of documents exported.

Failure Count

Displays the number of failed exports.

Progress (%)

Displays the transaction's progress.

Status

Displays the transaction's status.

Import Transactions

The Import Transactions page lets you sort transactions by package, template, and status by clicking any of the columns.

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

The page contains the following fields:

Package

Displays the import package name.

Run By

Displays the user who ran the import.

Start Time

Displays the start time of the transaction.

End Time

Displays the end time of the transaction.

Document Count

Displays the number of documents imported.

Failure Count

Displays the number of failed imports.

Progress (%)

Displays the transaction's progress.

Status

Displays the transaction's status.

More information:

[Show Import Settings](#) (see page 777)

Export/Import Documents

To migrate and synchronize data between different Knowledge Management systems, use the Knowledge Export/Import Tool (KEIT) in Knowledge Administration, which lets you export and import knowledge documents.

You can create KEIT templates to define an export/import transaction.

Note: To import r11.0 knowledge data, use the pdm_kit-txt utility.

Export/Import Packages

The following syntax names export packages:

- Server name
- Date
- Start time

You can define package directories by editing the *keit_path* option under Options Manager, Knowledge. Packages are stored in the following default directories:

- Export Packages
KEIT_PATH\export
- Import Packages
KEIT_PATH\import

To add attributes for export/import availability, add the following attribute names to the NX.env file:

@NX_KEIT_AVAILABLE_FIELDS

Adds the attribute name.

@NX_KET_ADDL_FIELDS

Adds the attribute name, and if the new attribute is a SREL to another object, adds the attribute name used to synchronize the source and target systems.

Example: Add Attributes

```
@NX_KEIT_ADDL_FIELDS = STATUS_ID.STATUS,DOC_TEMPLATE_ID.TEMPLATE_NAME
```

The export and import directories contain the following files:

keit_config.xml

Contains the configuration file for an export or import package.

data.xml

Contains the data file containing values of knowledge documents.

rep.xml

Contains the repositories that are referenced in the data.xml file.

images

Specifies the image files embedded in knowledge documents.

View Export/Import Templates

The Knowledge Export/Import Tool (KEIT) templates are used to define export/import transactions.

Note: The default export/import template is named All-Docs and does not contain the STATUS_ID field. This template creates the data.xml file and imports documents with a draft status. If you manually add the STATUS_ID field to the All-Docs template before the export operation, the original document status is preserved upon importation.

To view Knowledge Export/Import templates

1. Click the Administration tab.
The Administration console appears.
2. Click Knowledge, Documents.
Click Export/Import, Export/Import Templates.
The Knowledge Export/Import Templates List appears.
3. Click a template in the Name column.
The KEIT template opens.

Note: You can export list results to Excel for use outside of CA SDM by clicking the Export button on the List page.

Search Export and Import Templates

You can search for a template using the Knowledge Export/Import Template Settings page.

Note: If multi-tenancy is installed, select the appropriate tenant from the drop-down list. The public (shared) option creates the object for all tenants. A tenant drop-down list appears in the search filter. If you select <empty> in this drop-down list, the search is public. A tenant column also appears on the list page.

To search Knowledge Export/Import templates

1. On the Administration tab, browse to Knowledge, Documents, Export/Import, Export/Import Templates.
The Knowledge Export/Import Templates List appears.
2. Click Search and enter your search parameters.
The Search Results appear.

3. Click Show Filter.
The filter appears.
4. Modify your search using the following options:
 - Package**
Searches by package name.
 - Earliest Modified Date**
Searches by earliest modified date.
 - Latest Modified Date**
Searches by latest modified date.
 - Additional Search Arguments**
Provides additional search arguments.
The search results appear.
5. Click a template in the search results.
The template opens.

Create an Export/Import Template

You can create and modify a template using the Knowledge Export/Import Template Settings page in the Administration tab.

To create a Knowledge Export/Import template

1. On the Administration tab, browse to Knowledge, Documents, Export/Import, Export/Import Templates.
The Knowledge Export/Import Templates List appears.
2. Click Create New.
The Create New Export/Import Template page appears.
3. Complete the following fields:
 - Template Name**
Identifies the name of the template.
 - Description**
Provides a brief description of the template.

4. Complete the appropriate fields on the following tabs:

- [Export Fields](#) (see page 776)
- [Export Filter](#) (see page 777)
- [Import Settings](#) (see page 777)

Click Save.

The template is created.

See Also

[Export Transactions](#) (see page 770)

[Import Transactions](#) (see page 771)

Show Export Fields

To show the export fields, create a Knowledge Export/Import Template (KEIT).

The Export Fields tab appears and contains the following fields:

Available

Displays Knowledge Document fields that are available for export.

Note: If you want to preserve status of the documents, such as Draft, for the export and import process, add STATUS_ID to the Exported column.

Exported

Specifies document fields to export.

Export Attachments

Exports Knowledge Document file attachments.

You can use the arrows in the Select Document Attributes section to add document attributes export.

Note: If you edit the export/import template using Firefox, the list size does not change after you use the arrows in the first edit. Close and reopen the template in edit mode, and using the arrows again causes changes to the list size. This behavior does not occur when you use Internet Explorer.

Important! Exported fields in the KEIT have a higher priority than when the fields are selected as defaults in the Import Settings tab. If you select a default export field in the Import Settings tab, it is not processed as the default field.

Show the Export Filter

To show the export filter, create a Knowledge Export/Import Template (KEIT).

The Export Filter tab appears and contains the following fields:

Category

Opens the Knowledge Category page. Use this option to add categories to the list on the Export Filter tab.

Remove Category

Removes categories from the list on the Export Filter tab.

Clear Category

Clears the category list on the Export Filter tab.

Include child categories

Exports documents from child categories of exported categories.

Include secondary categories

Exports secondary categories of exported documents.

Include all documents linked to selected categories

Exports all documents linked to selected categories.

Note: Enabled only if you select Include secondary categories.

Additional Filter

Provides an additional WHERE clause.

Show Import Settings

To show import settings, create a Knowledge Export/Import Template (KEIT).

The Import Settings tab appears and contains the following fields:

Error Threshold (%)

Stops the import process (sets the status to Failed) if the percent of errors exceeds the specified number.

Override published documents

Allows imported documents to override published documents of the destination server.

Override documents in all non-published statuses

Allows imported documents to override non-published documents of the destination server.

Use default values when overriding documents

Uses default values on overridden documents for defined fields.

Index documents immediately

Indexes the document after the import process.

Status

Select Draft, Published or Retired.

Priority

Sets the priority level of the import settings.

Template

Selects one of the following default templates:

- Built in - Knowledge Document
- Built in - Knowledge Tree
- Built in - Quick Editing

Owner

Allows you to set the owner by opening the Contact Search page.

Author

Allows you to set the author by opening the Contact Search page.

Subject Expert

Allows you to set the subject expert by opening the Contact Search page.

Assignee

Allows you to set the assignee by opening the Contact Search page.

Expiration date

Opens the date helper.

Review Date

Opens the date helper.

Product ID

Allows you to set the product ID by opening the Product Search page.

Asset

Opens the Configuration Item Search page.

Root Cause

Opens the Root Cause Search page.

Service Desk Priority

Sets the Service Desk priority of the import settings.

Severity

Sets the severity level of the import settings.

Impact

Sets the impact level of the import settings.

Urgency

Sets the urgency level of the import settings.

Edit an Export/Import Template

You can edit export and import templates from the Knowledge Export/Import Templates List page.

To edit an export or import template

1. Navigate to Knowledge, Documents, Export/Import, Export/Import Templates.

The Knowledge Export/Import Templates List appears.

2. Select a template.

The Template Detail page appears.

3. Click Edit.

Modify the appropriate fields:

Template Name

Identifies the name of the template.

Description

Provides a brief description of the template.

4. Modify the appropriate fields on the following tabs:

- [Export Fields](#) (see page 776)
- [Export Filter](#) (see page 777)
- [Import Settings](#) (see page 777)

5. Click Save.

Your edits are saved.

More information:

[Import Transactions](#) (see page 771)

pdm_ket Utility—Knowledge Export Tool

The pdm_ket utility exports knowledge based on a template created from the web interface from a source computer to a knowledge package.

Attachments and their links are exported to data.xml during the export. Before you import the attachments, manually move them to following directory on the destination server:

\$NX_ROOT/site/attachments/default

Use this utility to do the following:

- Create configuration file based on the related knowledge export template.
- Export data with UUID of document, and content such as title, summary, problem, resolution and other document attributes like owner, status, and so on.
- Export all unique images used by the exported documents (always exported).
- Export all unique attachments used by the exported documents (EXP_ATTMENT field).

The files are copied from the remote repository to the local package folder.

The utility is invoked as follows:

```
pdm_ket -n <template name> [-h] [-v]
```

-n <template name>

Defines the name of the template used for export (case sensitive).

-h

(Optional) Displays help on the utility.

-v

(Optional) Enables extensive logging (bop_logging) of program events. This option is commonly used for internal problem solving.

Example: Using pdm_ket to Export Knowledge Using the my_template Template.

```
pdm_ket -n my_template
```

The pdm_ket utility can be scheduled using a third-party scheduler for exporting Knowledge Documents.

pdm_kit Utility—Knowledge Import Tool

The pdm_kit utility imports data to the destination server according to the settings in the configuration file of a package.

Important! The r11.0 import utility has been renamed to pdm_kit_txt.exe to allow importing r11.0 text files. This utility does not support any of the Release 12.7 import enhancements.

The referenced data previously referenced by the pdm_ket utility gets the real UUID or ID value of the destination server. When running pdm_kit utility a new userid parameter is applied. The pdm_kit utility works as follows:

1. Imports documents by replacing the userid value (for contacts) or referenced name (for fields like asset) with appropriate UUID of the destination server.
2. Imports images.
3. Imports attachments.
4. Uploads files from the local package folder to the remote repositories.

Note: If editing published documents is disabled, then the imported document is created as a rework version.

The utility is invoked as follows:

```
pdm_kit [-h] -f -u [-v]
```

-h

(Optional) Displays help for the utility on the interface.

-f

Specifies the path to the package

-u

Specifies the default user.

-v

(Optional) Enables extensive logging (bop_logging) of program events. This option is commonly used for internal problem solving.

Example: Using pdm_kit to Import a Package

```
pdm_kit -f c:\package_path -u ServiceDesk
```

Allow Users to Export and Import Knowledge

You can allow analysts to grant export and import permissions by managing your role list on the Administration tab.

To grant export/import permissions

1. Navigate to Security and Role Management, Role Management, Role List.

The Role List appears.

2. Select a role with the Analyst Interface Type.

The Role Detail page appears.

3. Click Edit.

The Update Role page appears.

4. Click the Knowledge tab.

Do any of the following:

- Select the Allow Export check box.
- Select the Allow Import check box.
- Select both check boxes.

5. Click Save.

The Role Detail page reappears.

6. Review your changes in the Knowledge tab.

Close the Role Detail page.

You can also restrict users from exporting/importing by clearing the check boxes on the Knowledge tab of the Role Detail page.

Web Services

Knowledge can be accessed through SOAP web services. Various methods are available, permitting the search, retrieval, creation, and updating of documents, and a range of other operations.

Note: For more information about SOAP web services, see the *Implementation Guide*.

Chapter 20: Administering Knowledge Management

This section contains the following topics:

[Knowledge Administration](#) (see page 783)
[Find Procedures for Knowledge Administration](#) (see page 784)
[Knowledge Management Roles and Functions](#) (see page 784)
[How to Manage Role Privileges and Document Visibility](#) (see page 794)
[Action Content](#) (see page 795)
[Document Approval Process](#) (see page 799)
[Automated Policies](#) (see page 806)
[Knowledge Document Control](#) (see page 810)
[Knowledge Categories](#) (see page 825)
[Reports and Metrics](#) (see page 835)
[Search](#) (see page 837)
[CA SDM Integration Options](#) (see page 852)
[Solution Survey](#) (see page 861)
[Knowledge Management System Settings](#) (see page 864)

Knowledge Administration

You can configure administrative settings for Knowledge Management. You can set [self-service Knowledge options](#) (see page 786) to meet the needs of users and to create an efficient and effective environment for managing and delivering knowledge. Use the Administration tab to set system options. You can set knowledge permissions based on the group or role of a contact. You apply settings that help to conform the functionality and use of Knowledge Management.

Important! Re-indexing documents in the knowledge base and running the calculation for Automated Policies and the Knowledge Report Card can be time-consuming. We recommend that you run these operations during non-business hours or when your system is less busy.

Find Procedures for Knowledge Administration

CA SDM provides step-by-step procedures for knowledge administration in this guide and in the *Online Help*.

To find step-by-step procedures for knowledge administration in the *Online Help*

1. Log in to CA SDM.

The CA SDM main page appears.

2. Do any of the following actions:

- Click Help, Help from the main menu and navigate the *Online Help* to Administration, Knowledge Administration.

The Knowledge Administration topics appear and you can navigate the hierarchy to locate the information you want.

- Click Help, Help on This Window from the main menu.

The *Online Help* appears and displays a help topic for the page for which you want information.

Knowledge Management Roles and Functions

Knowledge Management is designed for a wide variety of users, from administrators and knowledge managers, who maintain the product to customers, and employees, who use the system to find solutions to their problems. Although one person can fill multiple roles, the following roles are the basic user roles found in Knowledge Management:

- **Customer**—An external end user who performs basic self-service tasks.
- **Employee**—An internal end user who performs basic self-service tasks.
- **Knowledge Analyst**—A user that is responsible for one or more steps within the knowledge management process. This user interacts with service desk analysts to create and maintain a quality solution base.
- **Knowledge Manager**—A supervisor for the Knowledge Analyst. This role handles knowledge document reassignments and escalations, and manages day-to-day administrative aspects of the solution, including creating the category structure, defining the approval process, managing noise words, special terms, synonyms, and other settings and options that are more dynamic in nature than what the Knowledge Management Administrator controls.

- **Administrator**—The administrator who has access to all functionality in the CA SDM and Knowledge Management products within a single role. This role is typically used when implementing CA SDM to help ensure that all users and roles are set up properly and for a CA SDM environment that has a single person performing all administration tasks.
- **Knowledge Management Administrator**—An administrator who is responsible for configuring and monitoring the knowledge management process. This role includes creating the category structure, defining the approval process, and configuring default search and security settings.

Different levels of access are associated with each role in the CA SDM environment. These levels help define the tasks that each role performs.

See Also

[Knowledge Management User Interfaces](#) (see page 785)

[Knowledge Management Configuration and Management Functions](#) (see page 786)

[Self-Service Knowledge Options](#) (see page 786)

[Documents and Users](#) (see page 791)

[How to Manage Role Privileges and Document Visibility](#) (see page 794)

Knowledge Management User Interfaces

The following user interfaces help you manage knowledge:

- **Self-Service**—In the self-service interface, customers and employees using CA SDM can access knowledge documents and submit knowledge for further consideration. Customers can search, browse, or use bookmarks to locate and view knowledge documents.
- **Knowledge Documents**—In the knowledge documents interface, accessed from the Knowledge Documents node on the CA SDM Scoreboard, all users of the system can view their Inbox and follow-up comments. The *administrator* manages their unassigned/unindexed documents, automated document lifecycle policies, and user forums.
- **Knowledge Management**—In the knowledge management interface, accessed from the Knowledge tab in CA SDM, the *knowledge analyst* or *knowledge manager* can find, organize, and publish knowledge. They can also filter the documents displayed using advanced options, and sort the results by relevance, modified date and much more.
- **Knowledge Administration**—In the knowledge administration interface, accessed from the Knowledge node on the CA SDM Administration tab, the *administrator*, *knowledge manager*, or *knowledge management administrator* can set system options. Settings can help conform the functionality and use of Knowledge Management.

Knowledge Management Configuration and Management Functions

You can perform the following configuration and management functions in Knowledge Management:

- Create "action content" (a live action URL) that you can insert into the Resolution field of a document.
- Set up the approval process and define the knowledge document lifecycle process.
- Set up automated policies that automate certain tasks in the knowledge document approval process.
- Set up document options related to comments, submitting knowledge, templates, and document settings.
- Create templates that control how a document displays information.
- Manage the default Knowledge Management search engine and configure noise words, special terms, and synonyms used to perform keyword and natural language searches.
- Create "recommended documents" that display in the self-service interface when users search for knowledge solutions.
- Manage the knowledge category structure to make document access easier.
- Set up the Knowledge Report Card and general system settings.
- Define surveys that collect and analyze customer feedback about knowledge document performance.
- Manage integration of Knowledge Management into CA SDM, including field mapping and request and issue search configuration.

Self-Service Knowledge Options

Customers and employees using CA SDM have access to knowledge documents through the self-service interface. Customers can search, browse, or use bookmarks to locate and view knowledge documents. Giving customers access to knowledge documents and services permits customers to find answers themselves and reduces the pressure on resources.

In the self-service user interface, employees, and customers can view the following knowledge solution options:

- **Search for Knowledge Solutions**—Employees and customers can enter keywords and phrases in a search field that retrieves a list of knowledge solutions. You can configure this option in the following location: Administration, Knowledge, Search, Search Settings.

- **View Top Solutions**—Employees and customers can display a list of top solutions in the self-service interface. The following Administrator setting: Administration, Knowledge, System, General Settings, Top Solutions—Number of Documents to Display, determines the number of documents to display.
- **Prompt for Knowledge Survey**—After opening a knowledge document, the user is able to read the document and access a series of survey questions. One of these questions lets the customer indicate through a prompt whether they feel that their problem has been solved. You can configure this option in the following location: Administration, Knowledge, Solution Survey, Survey Settings.
- **Suggest Knowledge**—Employees and customers can, where permitted, view a list of knowledge suggestions when they create a ticket in the self-service interface. You can configure this option in the following location: Administration, Knowledge, Service Desk Integration, Suggest Knowledge.

Advanced Search

In Advanced Search Settings, customers and employees can use the following options to refine a search for the solution to a problem:

Knowledge Management Search

Searches for certain keywords, which serve as preliminary matches.

Natural Language Search

Searches for words and accounts for word proximity, word order, and relevance.

Search results are listed by relevance. Relevance is determined by the specified search criteria (expressed as EXCELLENT, GOOD, and so on). Documents with the highest relevance (EXCELLENT) are listed first.

Each result can include a title that appears as a link, a summary of the document, and additional statistics relevant to the document, such as Relevance Rating, Document ID, Modify Date, FAQ Rating, and Hits Received.

Depending on how the system administrator configures Knowledge Management, users can open an incident, rate a document, and provide comments when a knowledge document is open.

Knowledge Document Information

Each knowledge document contains key document fields that provide you with information. Depending on the document template used, different fields appear or they appear differently in the document.

A title identifies the document. The document also includes the following information:

Summary

Summarizes the document, briefly describes the problem.

Problem

Describes the problem.

Resolution

Describes how to resolve the problem in plain text or HTML format.

See Also

Lists documents that are related to the current document. If you click a document hyperlink, a separate page opens that contains the related document.

Attachments

Lists files that are attached to the document and can be downloaded. Attachments provide additional information about the document.

Related Categories

Lists categories that are related to the document. If you click a category hyperlink, the Search Tools page updates to display that category.

Related Tickets

Links to requests, incidents, problems, issues, and documents that have been opened as a result of the document or that have been solved by using the document.

Properties

Indicates additional document properties selected in the document template. By default, the modify date and document ID are displayed.

Comments

Lists comments from users for the document. Included with the comments are the contact name, email address, and the date.

Rate & Comment

Provides follow-up comments and feedback about whether the document was helpful in answering your question. You can rate the usefulness of the document, based on the following questions:

- Did this document solve your problem?
- How helpful was this forum?

You can also add a comment, and assign it to another analyst for follow-up, using one the following comment types:

- Broken Link
- Candidate for Publication
- Candidate for Retirement
- Incorrect Information
- Missing Information
- Recommend New Content
- Review
- Solution Does Not Work

The Page Options section lets you take the following courses of action:

- Edit
- Add/Remove Bookmark
- Subscribe
- Rate & Comment
- Email
- New Forum
- New Incident
- New Incident based on this Document
- Printable Version

More information:

[Create a Comment Type](#) (see page 812)

[Configure Self-Service Policies](#) (see page 860)

Announcements

Announcements are a quick way of providing solutions to a problem. Announcements provide a solution for a widespread problem and to answer calls to the service desk.

Current announcements display in the right pane of the main CA SDM window when you first log in to the system.

Sort Documents

You can sort documents in various order.

To sort documents

1. On the Knowledge tab, select a tree node from the category list in the left pane.
2. Expand the node and click a category.

The Document List page appears.

3. Sort the documents using the Order by list in the following orders:

- FAQ rating
- Number of hits
- Recently modified documents
- Solution count

The documents are listed in the order you selected.

Document Browsing

You can categorize knowledge documents to let customers browse for information based on frequently asked questions (FAQs). By selecting a knowledge category, the related subcategories and knowledge documents appear. You can select the document you would like to view or select a subcategory to narrow your search further.

To enhance self-service capabilities, a dynamic list of the most frequently used documents appear.

Note: Users can specify criteria about an item of interest and the search engine finds the matching knowledge documents and displays them on the search results page as a set of "recommended document" links. The search query can be expressed as a keyword or set of words (phrase) that identify the desired concept that one or more documents may contain. For more information, see [Create New Recommended Documents](#) (see page 848).

Document Bookmarking

You can bookmark frequently accessed documents. After a document has been bookmarked, it is added to the Bookmark list. This list can help you to locate documents quickly that you frequently view. After you have added a document to a bookmark list, a Remove button appears in the Bookmark field. You can remove bookmarks from the list when they are no longer frequently accessed.

The My Bookmarks folder stores links to the most frequently referenced documents. When you click My Bookmarks in the Category pane on the Knowledge tab, the list of bookmarked documents displays in the Knowledge Document List pane.

Incidents and Problems

Customers sometimes encounter problems that cannot be solved simply by searching for knowledge. Not all problems have a direct solution in the knowledge base. When a customer has a problem that cannot be solved, they can create an incident that is sent to an analyst for further processing. The incident describes the problem, and it can also be based on an existing knowledge document. The more information that is added to an incident, the easier it is for the analyst to solve.

Many ITIL-defined activities are supported in Knowledge Management, including the following activities:

- Incident Management
 - Knowledge searches can help find known errors during further incident investigation and diagnosis
 - Incident categorization
- Problem Management
 - Accessing information about known errors, and helping with problem matching to obtain the resolution when the problem has occurred before
 - Maintaining and providing access to information about workarounds
 - Recording information about procedures, work instructions, diagnostic scripts, and known errors (while supporting rich content, editing tools, measurement, and a definable approval process for the development of resolutions)
 - Problem analysis (through the linkage and analysis of incidents)

Documents and Users

The knowledge documents interface, accessed from the Knowledge Documents node on the CA SDM Scoreboard, lets all users of the system view their Inbox and follow-up comments. The administrator manages their unassigned/unindexed documents, automated document lifecycle policies, and user forums.

Document Inbox

The Knowledge Documents Inbox (and Group Inbox) on the CA SDM Scoreboard contains the documents that are assigned to you or your group. The Inbox is the central repository for documents in all statuses including saved and assigned draft and rework-draft documents.

The Inbox is the day-to-day task container and it is an important tool for managing the approval process. When a document is created or updated, it is placed in an owner Inbox. Items that appear in an Inbox require the attention of the user as part of the publishing process. Until they are published, items in the Inbox may not appear as resolutions and are not added to the knowledge base. Regularly monitor your Inbox for new documents.

Display Follow-up Comments

The Follow-Up Comments Inbox on the CA SDM Knowledge Documents Scoreboard contains the comments that are assigned to you or your group.

You can view summary information for your assigned follow-up comments on the Comment List page.

To display your follow-up comments, browse to Knowledge Documents, Follow-Up Comments.

Document Attributes

Setting document attributes helps you manage documents in your knowledge pool. Document attributes can be updated to assign a new subject expert or owner to the document. You can also specify the date the document becomes available in the knowledge base and the date when it expires. By selecting different document templates, you can modify the appearance of each document.

Document Permissions

You can set, view, and edit permissions for a document. These permissions can be assigned to different groups of individuals. When setting permissions, you can decide to inherit permissions from the primary category of a document or specify new permissions. By default, documents inherit their permissions from their primary category. This default handles access permissions at the category level rather than the document level.

Important! When creating a knowledge document, verify that document permissions include users that can be later assigned on the document through the approval process. When a group is assigned to a document, users from that group may not have the permission to view the document. If the document is assigned to a specific user, default data partition constraints allow the user to view the document.

Knowledge Categories

Assign each document to a primary category. For example, any knowledge related to Microsoft Word must be added to the Microsoft Word category. In addition, Knowledge Management lets you associate a document with multiple secondary categories and other documents. In this way, a document can be classified under many different applicable categories and can provide for more successful search results. After you create a document link, a *see also* link appears when viewing either of the linked documents. The *see also* link lets you go directly from one linked document to the other.

Attachment Permissions

You manage attachment permissions in the Attachments Library, Repositories node in the Administration tab. You can set up repository folder permissions based on the needs of your company.

Note: A document can have different permissions than the attachments linked to the document.

Add a Document Attachment

When working with a knowledge document, you can attach supplemental files or URLs to the document. These supplemental files or URLs provide easy access to additional information related to the knowledge documents. Attached files are stored in a repository.

You use the Attachments tab to manage file and URL attachments in the current document. In documents that are created using the default Built-In Knowledge Document or Built-In Knowledge Tree template, links to attachments display under the "Attachments" heading.

Note: The Attachments tab displays only for published documents. For unpublished documents, the Attachment tab displays Repositories, Files, and Attachments panes, and a list of URLs and files currently attached to the document.

To add a file attachment to a document

1. Open the document for editing.
2. Click the Attachments tab.
3. Enter the URL you want to attach.
4. Do either of the following:
 - Click Attach File to add a file from your system.
 - Click Attach File From Library to select the file name you want to attach from the Repositories list.

The name displays on the Attachments tab.

View Document History

You can review a record of all the actions performed on a document.

To view document history, click the History tab in the Update Document page.

An event list appears showing the details of each event. Information displayed on the History tab exists for reference purposes only and cannot be modified.

Document Notifications

In CA SDM, you can set up a list of users, known as contacts, which can be notified if certain events occur. The notification process informs individuals of the changing status of a document, keeping them up to date on its progress. Only users that appear in this list can be sent the notifications that have been designated for the document.

Document Changes

After a document arrives in an Inbox folder on the scoreboard, users can perform their assigned tasks by modifying and saving their documents according to their assigned roles (for example, an analyst is responsible for checking the technical content of the document, while another analyst checks the formatting).

All users with full permission to the document can modify the document. The current owner has full permissions to the document but may not have explicit write permissions. Only the owner can change the owner of the document (on the Attributes tab).

How to Manage Role Privileges and Document Visibility

You can set up Knowledge Management security permissions by managing role privileges for users in your environment. These permissions let you establish what users can access when they want to view or create knowledge, and how users are authenticated when they log in to the system.

You can manage role privileges and document visibility for Knowledge Management by doing the following:

1. From the Administration tab, select Security and Role Management, Role Management, Role List.

The Role List page appears.

2. Select the role you want such as Knowledge Manager.

The Role Detail and Update Role pages display the following tabs:

Knowledge Management

Specifies the Knowledge Management privileges for the role.

KT Document Visibility

Specifies which document statuses the role is allowed to view, such as draft, retired, and published.

3. Complete the tab pages and save changes.

Security and role privileges are defined.

Note: For more information about setting up security and defining roles, see the *Online Help*.

Action Content

You can create "action content" (a live URL) that can be inserted into the Resolution field of a knowledge document that, when clicked by the end user, creates an incident, or performs some other action. Using action content, a substantial degree of definition and classification can be achieved without the user even realizing it.

The steps to insert a live "action content" link into a document are simple and no coding is required. The Knowledge Management HTML editor handles the generation of the HTML code.

Note: Action content is primarily used for interaction with external applications.

More information:

[View Action Content](#) (see page 796)

[Create Action Content \(Action URL\)](#) (see page 797)

[Create Action Content \(Internal HTML\)](#) (see page 797)

[Edit Action Content](#) (see page 798)

[Search for Action Content](#) (see page 799)

View Action Content

You can view the details of the live action URLs that let you define actions for your end users. Knowledge Management lets you link knowledge documents with automated tasks, so you can embed Self Service scripts into knowledge documents.

To view Action Content

1. On the Administration tab, browse to Knowledge, Action Content.

The Action Content List appears.

2. The page displays action details and execution options.

Name

Displays the name of the action content.

Code

Specifies the HTML code that links the action content to the document.

Action URL

Specifies the URL link that performs some type of action, such as opening a website or a form.

Search

Searches for items by name.

Show Filter

Searches for items in the Comment List page. You can complete one or more of the search fields and click Search. The Action Content List page displays the items that match your criteria.

Create New

Creates an action content link.

3. Search or action content that you can use with Knowledge Management and Support Automation.

The search results appear.

Create Action Content (Action URL)

You can create Action Content for a knowledge documents. These action URLs can launch a live website that is accessible to all users of your system. You can also link action URLs to automated tasks on your server, and you can embed these scripts into knowledge documents, tree documents, document templates, and knowledge forums.

To create action content (Action URL)

1. On the Administration tab, browse to Knowledge, Action Content.

The Action Content list page appears.

2. Click Create New.

The Create New Action Content page displays.

3. Complete the fields. The following fields require explanation:

Code

Specifies a unique identifier for this action content item.

Use Internal HTML

[Creates an internal link](#) (see page 797) in the application that dynamically passes information, such as the user name and session ID, from knowledge documents into third-party applications. Do not select this option.

Action URL

Specifies a URL that links to a web page, template, or automated script, for example: `http://www.ca.com`.

Click Save.

The action content is created.

Create Action Content (Internal HTML)

You can create Action Content for a knowledge documents. These action URLs can launch self service scripts and also third-party applications. This link is generated dynamically and automatically transfers attributes about the user, such as the user name, to the target application (`http://www.ca.com?USERNAME=BBB`, for example). User attributes are specified in an HTML file.

To create action content with an internal HTML file

1. Create an internal HTML file that passes data to the target application.

Note: The `act_content_sample.html` file is available in the following location:
`NX_ROOT\bopcfg\www\html\default`.

2. Save the HTML file in the following location:
`NX_ROOT\site\mods\www\html\default` directory.

3. From the Administration tab, navigate to Knowledge, Action Content.

The Action Content list page appears.

4. Click Create New.

The Create New Action Content page appears.

5. Complete the following field:

Code

Specifies a unique identifier for this action content item.

6. Select the Use Internal HTML check box.

7. Specify the appropriate HTML file in the Action URL field, for example:
act_content_sample.html

Note: You can use Support Automation scripts by using the format *SA_SCRIPT=[Self Service Script ID]* in the URL.

8. Click Save.

The action content is created. When the user clicks this Action Content link within a knowledge document, attributes about the user, such as the user name, are dynamically passed on to the target application.

Edit Action Content

You can edit action content that has already been created on the Action Content List page.

To edit action content

1. From the Administration tab, select Knowledge, Action Content.

The Action Content List appears.

2. Select the item you want to edit.

The Action Content Detail page appears.

3. Click Edit.

The Update Action Content page appears.

4. Edit in the fields as appropriate.

5. Click Save.

The action content is updated.

Search for Action Content

You can enter search criteria to filter the Action Content List to display only the items you want to see. You can also search for individual items to view or edit.

To search for action content

1. From the Administration tab, select Knowledge, Action Content.

The Action Content List page opens.

2. Click Show Filter.

The search filter appears.

3. Complete one or more of the search fields:

Name

Identifies this action content item.

Status

Indicates whether this item is Active or Inactive.

Code

Specifies a unique identifier for this action content item.

Knowledge Management URL

Specifies that the Action URL links to the internal web page (.html).

4. (Optional) Click More.

Additional fields are displayed, letting you further restrict the items that display in the Action Content List.

5. Click Search.

The Action Content List page displays the items that match your search criteria. You can select an item to view or edit it.

Note: You can export list results to Excel for use outside of CA SDM by clicking the Export button on the List page.

Document Approval Process

For administrators who want to control the management of their knowledge base, the ability to customize the document editing and approval process is essential. You can design Approval Process templates that specify how, when, and by which an employee document can be modified and published to the public. Approval Process templates can designate various approval processes best suited to your business environment. The approval process you implement can be modified over time to become simpler or more complex.

The Approval Process Manager lets you define Approval Process templates. By default, the Built-in Approval Process template is used. However, you can create a template or edit an existing template. When creating an Approval Process template, you define statuses and add tasks to the template. The approval process involves a series of tasks that are performed on a knowledge document. The owner that is assigned a task in the Approval Process template performs each task.

The following statuses are the various states the document is associated with during the stages of the approval process:

Draft

Specifies a new document.

Published

Specifies a document that has passed through the complete approval cycle and becomes part of the viewable knowledge base.

Rework-Draft Version

Specifies a rework version of a copy of the document that is replaced in the knowledge base after it is verified and republished.

Retired

Specifies a document that has reached its expiration date. You can also create your own statuses, which you later associate with tasks.

Approval Process Manager

Knowledge Administrators can use the Approval Process Manager to perform these actions:

- Determine which groups can read a knowledge document and which groups can write (or edit) the knowledge document.
- Identify the tasks in an approval process template that determine the lifecycle of documents created with the template.
- Define the various statuses the document can be associated with during the approval process.

Important! When creating a knowledge document, verify that document permissions include users that can later be assigned to the document through the approval process. Users from that group do not necessarily have the permission to view the document. If the document is assigned to a specific user, default data partition constraints let the user view the document.

Define an Approval Process for Document Editing

Knowledge administrators can specify who can edit documents before the approval process and after publishing.

Note: Users with full (read/write) permissions can edit published documents.

To define an approval process for document editing

1. On the Administration tab, browse to Knowledge, Approval Process Manager, Approval Process Settings.

The Approval Process Settings page appears.

2. Specify who can edit documents *before* the documents are published. Select one of the following options:

Documents may be edited by a task assignee, an owner or users with the appropriate Access Type views

Specifies that the following contacts can edit documents:

- A contact assigned to the current task
- A contact specified as an owner of the document for the current task
- A knowledge manager
- A system administrator

Documents may be edited by users with full permissions

Specifies that any user with write permissions to the document can edit it.

3. Specify who can edit documents *after* the documents are published. Select one of the following options:

User with full permissions may edit documents after they have been published

Specifies that a user with full permissions can edit published documents.

User with full permissions can change published document's attributes

Specifies that any user with write permissions to the document can change only attributes of published documents such as configuration items or products.

Document must be unpublished before editing is allowed

Specifies that the user must unpublish a document before editing it.

Click Save.

The approval process is defined.

Create an Approval Process Template

The tasks in an approval process template define the lifecycle of documents created with the template.

Note: If multi-tenancy is installed, select the appropriate tenant from the drop-down list. The public (shared) option creates the object for all tenants.

To create an approval process template

1. From the Administration tab, select Knowledge, Approval Process Manager, Approval Process Templates.

The Approval Process Template List appears.

2. Click Create New.

The Approval Process Template Detail page appears.

3. Enter a name for the template and a description.

4. Click Save.

The Approval Process Template Detail page displays additional fields.

5. Select the task you want to perform when you create a rework version of the document using this template.

6. Select the task you want to perform when you unretire a document that was created using this template.

7. Click Insert Task to create a task to add to the template.

The Create New Task page displays.

8. Complete the fields. The following fields require explanation:

Task

Names the task.

Assignee

Assigns a person to the task. You can click the search icon to select a name.

Note: You can [add an alternate list of assignees](#) (see page 803) after the task is created.

Click Save.

The approval process template is created.

More information:

[Add Alternate Assignees to a Task](#) (see page 803)

[Edit an Approval Process Template](#) (see page 803)

[Search for an Approval Process Template](#) (see page 804)

Add Alternate Assignees to a Task

You can add alternate assignees when you edit a task.

To add alternate assignees to a task

1. Select the task you want to edit from the Approval Process Template list.
The Task Detail page appears.
2. Edit the fields as appropriate.
3. Click the Add button on the Assignee List.
The Create New Assignee page appears.
4. In the Assignee field, enter the name of the person that you want to assign to the task, or click the search icon to select the name.
5. Repeat Step 4 as necessary to create a list of alternate assignees.
6. Click Save.

The alternate assignees are added to the task.

Edit an Approval Process Template

You can edit an approval process template.

To edit an approval process template

1. From the Administration tab, select Knowledge, Approval Process Manager, Approval Process Templates.
The Approval Process Template List appears.
2. Select the template you want to edit and click its name.
The Approval Process Template Detail page appears.
3. Click Edit.
The Update Approval Process Template page appears.
4. Edit the fields as appropriate.
5. Click Save.

The approval process template is updated.

Search for an Approval Process Template

You can enter search criteria to filter the Approval Process Template List to display only the items you want to see. You can also search for individual items to view or edit.

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

To search for an approval process template

1. From the Administration tab, select Knowledge, Approval Process, Approval Process Templates.

The Approval Process Template List page appears.

2. Click Show Filter.
3. Type in the first few characters of the template name you want.
4. (Optional) Click More.

Additional fields are displayed, letting you to restrict the items that display in the Approval Process Template List.

5. Click Search.

The Approval Process Template List page displays the items that match your search criteria. You can select an item to view or edit it.

Note: You can export list results to Excel for use outside of CA SDM by clicking the Export button on the List page.

Document Status Definitions

You can add and delete user-defined document statuses, and modify the names and descriptions of predefined document statuses.

More information:

[Create a Document Status](#) (see page 805)

[Search Document Statuses](#) (see page 805)

[Edit Document Status](#) (see page 806)

Create a Document Status

You can create a document status for Knowledge documents in your CA SDM system.

To create a document status

1. From the Administration tab, select Knowledge, Approval Process Manager, Document Statuses.

The Document Status List appears.

2. Click Create New.

The Document Status Detail page appears.

3. Enter a name for the status and a description.

4. Click Save.

The document status is created.

Search Document Statuses

You can enter search criteria to filter the Document Status List to display only the items you want to see. You can also search for individual items to view or edit.

To search for a document status

1. From the Administration tab, select Knowledge, Approval Process Manager, Document Statuses.

The Document Status List page appears.

2. Click Show Filter.

3. Type in the first few characters of the status name you want to find.

4. (Optional) Click More.

Additional fields are displayed, letting you restrict the items that display in the Document Status List.

5. Click Search.

The Document Status List page displays the items that match your search criteria. You can select an item to view or edit it.

Note: You can export list results to Excel for use outside of CA SDM by clicking the Export button on the List page.

Edit Document Status

You can edit a document status.

To edit a document status

1. From the Administration tab, select Knowledge, Approval Process Manager, Document Statuses.

The Document Status List appears.

2. Select the status you want to edit and click its name.

The Document Status Detail page appears.

3. Click Edit.

The Update Document Status page appears.

4. Edit the Status and Description fields as appropriate.

5. Click Save.

The document status is updated.

Automated Policies

Administrators can automate certain tasks in the knowledge document approval process based on document lifecycle policies and actions they define. By automating tasks, end users who are searching for solutions can solve problems faster, and they can do so without contacting other individuals, which provides a benefit to the organization.

Automated policies work with events and macros.

More information:

[View Automated Policies](#) (see page 807)

[How to Set Up Automated Policies](#) (see page 808)

[Create an Automated Policy](#) (see page 808)

[Edit an Automated Policy](#) (see page 809)

[Schedule Automated Policies](#) (see page 809)

[View Document Lifecycle Policy Reports](#) (see page 810)

View Automated Policies

An automated policy describes the condition by which documents are flagged for correction and promoted to publication or retirement throughout the various stages of the document lifecycle process. For example, you can specify the "fix broken links" default policy that matches documents found in the knowledge base with broken links. The task of fixing the problem can be assigned to an analyst.

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

To view automated policies, select the Administration tab, Knowledge, Automated Policies, Policies.

The Automated Policies List page displays the details of the automated policies. You can edit the default policies, or define your own. Each policy is associated with the following components:

Stored Query

Contains a set of action macros that execute when the policies identify and match a document during processing. After processing, the stored query condition event displays a role-based Knowledge Management lifecycle policy report on the CA SDM Scoreboard.

The administrator is responsible for monitoring the reports and for providing feedback and recommendations to the appropriate document editors.

Action Macro

Contains code that lets users set a flag, increase the priority, or perform some other action. You can modify the macros that appear on the Macro list, or define your own.

Note: You can export list results to Excel for use outside of CA SDM by clicking the Export button on the List page.

More information:

[View Document Lifecycle Policy Reports](#) (see page 810)

[Schedule Automated Policies](#) (see page 809)

How to Set Up Automated Policies

Administrators can set up the Automated Policies feature by performing the following steps:

1. (Required) A batch process must be defined in the Automated Policies Scheduler that executes on the server to present the data required to view the Knowledge Management Lifecycle Policy reports. This action step also applies to the Knowledge Report Card.

Note: For more information about automated policies, see the *Implementation Guide*.

2. (Required) For security and role management, define the stage by which users can view and search on documents during their lifecycle on the Knowledge Document Visibility tab in Role Management.
3. On the Automated Policy List page, you can edit the default policies, or define your own.

Note: For new policies, the administrator must include the "Disregard Life Cycle Policies" field in the stored query; otherwise, it does not appear on the Attributes tab.

Create an Automated Policy

You can create an automated policy that activates when an action occurs, such as when a document is published or retired from the knowledge base.

Note: If multi-tenancy is installed, select the appropriate tenant from the drop-down list. The public (shared) option creates the object for all tenants.

To create an automated policy

1. On the Administration tab, browse to Knowledge, Automated Policies, Policies.
The Create New Automated Policy page appears.

2. Enter a name and description for the policy in the appropriate fields.
3. Enter a stored query name or select one using the search icon.
4. Click Add Action.

The Macro List page appears.

5. On the Macro List page, select one of the predefined action macros, or define your own (click Create New).

The action macro appears in the Action Information list on the Create New Automated Policy page.

Note: You can delete an action macro: right-click the name and select Delete from the shortcut menu.

6. Click Save.

The new policy appears in the Automated Policy list.

7. (Optional) Right-click the title to edit a policy.

The selected policy appears in the Policy Update page and you can edit it.

Edit an Automated Policy

You can update an automated policy.

To edit an automated policy

1. On the Administration tab, browse to Knowledge, Automated Policies, Policies.

The Automated Policy List appears.

2. Right-click the policy to edit, and select Edit from the shortcut menu.

The Update Automated Policy page appears.

3. Edit the entry fields as appropriate.

4. Click Save.

The automated policy is updated.

Schedule Automated Policies

You can specify the date and time at which CA SDM performs the calculation and executes batch processing of the policies.

To schedule automated policies

1. Select the Administration tab, browse to Knowledge, Automated Policies, Scheduling.

The Automated Policies page appears.

2. Complete the following fields:

Last Updated

Specifies the Run Calculation check box.

Schedule

Specifies a date either in the text box or in the Calendar. You can select the time interval at which CA SDM performs the calculation and runs the policies.

Click Save.

The policies are processed at the specified date and time.

View Document Lifecycle Policy Reports

You can view the document lifecycle policy reports that Automated Policies generate.

To display the reports, select the Service Desk tab, Knowledge Documents, Automated Policies, Policies.

When you select a policy, the following summary information appears on the Document List page:

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

Title

Displays the title of the document that is either flagged for correction or promoted for publication or retirement.

Policy/Actions

Displays the policy name and action content defined for the policy.

Attributes

Displays the properties of the document that the policy affects.

Knowledge Document Control

When a user opens a knowledge document, they can add comments, submit knowledge, and view their knowledge tree documents. They can also specify the content and appearance of documents by selecting a template.

Administrators can control how users interact with documents by performing the following tasks:

- Define the [comment types](#) (see page 811) that appear in knowledge documents and drop-down lists in the end-user interface.
- Specify document [settings](#) (see page 864) related to comments, submitting knowledge, and viewing documents.
- Create [document templates](#) (see page 816) that specify the content and appearance of documents.
- [Export](#) (see page 770) (or import) knowledge content from another resource or system.
- Manage document [versioning](#) (see page 757) capabilities.
- Manage the [Knowledge Documents Schedule](#) (see page 760).

See also

[Comment Types](#) (see page 811)

[Document Templates](#) (see page 816)

[How to Create Knowledge Document Links](#) (see page 823)

Comment Types

If the analyst notices a typo or problem with the content in a document, they can add a comment that flags the document for correction, then assign the problem to another analyst for follow-up.

Administrators can define the comment types that appear in various list views within the end user interface.

View Comment Types

When a user opens a knowledge document, they can add a comment that flags the document for correction, promotion, or retirement, and much more. The Comment Type List page lets you manage the details of the comment types.

To display this page, select the Administration tab, Knowledge, Documents, Comment Types.

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

The Comment Type List page appears and displays the following columns:

Name

Displays the list of comment types that appear in documents and drop-downs lists within the user interface. You can edit the following default comment types, or define your own:

- Broken Link
- Candidate for Publication
- Content Not Clear
- Difficult to Find
- General Comment

- Incorrect Information
- Missing Information
- Recommend New Content
- Review
- Solution Does Not Work

Show in User View

Displays the comment in various list views within the user interface.

Follow-up Required

Specifies if the user is required to respond to this type of comment.

Time to Complete (Days)

Defines the number of days by which the user must follow up on this type of comment.

Status

Indicates whether the comment type is active or inactive.

This page contains the follow buttons:

Search

Searches for items by name.

Show Filter

Searches for a comment in the Comment List page. Complete one or more of the search fields and click Search. The Comment Type List page displays the comments that match your criteria.

Clear Filter

Clears the filter.

Create New

Creates a comment type.

Note: You can export list results to Excel for use outside of CA SDM by clicking the Export button on the List page.

Create a Comment Type

You can define the comment types that appear in various list views within the end user interface on the Create New Comment Type page.

Note: If multi-tenancy is installed, select the appropriate tenant from the drop-down list. The public (shared) option creates the object for all tenants.

To create a comment type

1. On the Administration tab, select Knowledge, Documents, Comment Types.

The Comment Type List page appears.

2. Click Create New.

The Create New Comment Type page appears.

3. Complete the fields as appropriate. The following fields require explanation:

Time to Complete (Days)

Defines the number of days by which the user must follow up on this type of comment.

Show in User View

Displays the comment in various list views within the user interface.

Follow-up Required

Specifies if the user is required to respond to this type of comment.

Click Save.

The new comment type appears on the Comment Types List page.

Edit a Comment Type

You can update a comment type that has already been created on the Comment Type List page.

To edit a comment type

1. From the Comment Type list, do one of the following:

- Select a comment, and select Edit from the Comment Type Detail page.
- Right-click a comment, and select Edit from the shortcut menu.

The Comment Type Detail page appears.

2. Complete the fields as appropriate.
3. Click Save.

The updated comment type appears on the Comment Type List page.

Document Notifications

In CA SDM, you can set up a list of users, known as contacts, which can be notified if certain events occur. The notification process informs individuals of the changing status of a document, keeping them up to date on its progress. Only users that appear in this list can be sent the notifications that have been designated for the document.

Set Up a Follow-Up Comment Notification

Several default activity notifications listed on the Activity Notification List page let you set up notifications to users when a follow-up comment is assigned to them.

The Follow-Up Comments queue on the scoreboard is the repository for assigned and unassigned follow-up comments.

To set up a follow-up comment notification

1. On the Administration tab, browse to Notifications, Activity Notifications.

The Activity Notification List page appears.

2. Select *one* of the following activity notifications:

- Follow-Up Comment Assigned
- Follow-Up Comment Closed

The Activity Notification Detail page appears.

3. Click Edit.

The Update Activity Notification page appears.

4. Change the fields as appropriate.

5. Click Save.

The Activity Notification Detail page appears.

6. Click Close window.

The modified activity notification appears in the Activity Notification List when you redisplay the list.

Specify Document Settings

If you are a system administrator you can specify settings related to comments, submitting knowledge, and viewing knowledge tree documents.

To specify document settings

1. On the Administration tab, select Knowledge, Documents, Document Settings.

The Document Settings page appears.

2. Complete the following fields:

Knowledge Tree Document Viewing

Specifies the viewing mode in which knowledge tree documents open. Select Open in Tree Mode (default) to open the knowledge tree directly or Open in Document Mode to open the document in document view. If you open in Document Mode, click Display to show the knowledge tree.

Comments

Specifies if users can submit comments for documents and view document comments. Select one of the following options:

- **Allow comment submission and comment viewing (default)**—Displays a Comment field at the bottom of an open document so users can submit comments for the document. Users can view comments already associated with the open document.
- **Allow comment submission but not comment viewing**—Displays a Comment field on the right of an open document so users can submit comments for the document. Users cannot view comments already associated with the open document.
- **Allow neither comment submission nor comment viewing**—Denies users the ability to submit or view comments. The Comment field does not display in an open document.

Submit Knowledge

Defines the repository for user-submitted documents. The product populates the list with the names of repositories defined on the Attachment Library pane.

Consider the following information:

- When an analyst creates a document in a category that has an owner and assigns the document to the category owner:

Category owner becomes the assignee and owner of that document.

Category owner should receive a document assigned notification.
- When an analyst creates a document in a category that has no owner and assigns the document to the category owner:

No one becomes the assignee or owner of that document.

Submit KD notification should send as defined in the administration.
- When an analyst creates a document and does assign the document to the category owner:

That user becomes the assignee and owner.

A notification does not send.
- When an employee or customer creates new documents, CA SDM takes the action as stated in the first two points.

Maximum Resolution Size

Defines the maximum size (in characters) that the Resolution field in a document can contain.

Limits: The maximum characters allowed is 256000.

Default: 32768.

Duplicate Document Avoidance

Forces a search for similar documents when the user creates a knowledge document.

Notification before Expiration

Defines the number of days before the document expires and a notification is sent.

Default: 7

Note: This value only applies to documents that have been updated or created in CA SDM. If you migrate documents from CA SDM r11.2, and the expiration dates are set before the migration, this option does not apply, unless or until you update the document after the migration.

Click Save.

The changes apply when you open a knowledge document or a knowledge tree document.

Document Templates

You can use document templates to control the format and default content of knowledge documents. Every knowledge document uses a document template to define its properties and appearance when opened. By default, a built-in template is associated with new knowledge documents.

The Template Editor lets you do the following:

- Design a document template that can later be associated with a document on the Contents tab of the Document Editor page.
- Modify the built-in template and other templates.

By editing templates, you can create templates to associate with documents. You can select the Properties options and edit the Header and Body sections using the HTML Editor.

- Update document templates from a previous release to support [knowledge relationships](#) (see page 822), such as parent-child links.

Create a Document Template

A document template specifies the content and appearance of documents in the knowledge base. You can apply the default templates:

- Built In—Knowledge Document
- Built In—Knowledge Tree
- Built In—Quick Editing

The product uses the default templates when you create knowledge documents and knowledge tree documents unless you create document templates and associate them with your documents.

Note: While creating a knowledge document, select the appropriate tenant from the drop-down list. The public (shared) option creates the object for all tenants.

To create a document template

1. On the Administration tab, browse to Knowledge, Documents, Document Templates.

The Document Templates List page appears.

2. Click Create New.

The Create New Document Template page appears.

3. Complete the following fields:

Template

(Required) Defines a unique name for the template.

Detail

Displays the static content that appears in documents created using the current template. If you select the HTML Source option, you can edit HTML code for the body directly in the Body. If you select the Quick View option, the Body field is read-only and displays the static body content as it appears at run time.

4. (Optional) Click Set Default Values.

The Default Values Template page appears.

You can set default values when creating a document. If you change the template on an existing document, there is no impact.

5. (Optional) Hide the Title, Summary, Problem, or Resolution from appearing in your template.

You can hide these fields when you want a simple document, such as a question and answer template.

6. Click Edit Detail.

The HTML Editor page appears and you can specify the static content and layout of documents that use the template. You can edit the code using the toolbar to insert placeholder tags.

Important! You can remove knowledge relationships, such as parent, child, and related links, by deleting the {TAG_PARENT} and {TAG_RELATED} tags from the template.

7. Click OK.

The Detail field shows the updated content.

8. (Optional) Click Quick View.
Shows the content as it appears in documents based on the template.
 9. (Optional) Click HTML View.
Shows the content as HTML code.
 10. Click Save.
The Document Template Detail page appears.
 11. Click Close Window.
New documents that use the template show the new content and layout.
The name of the template appears in the Template List when you display the list again.
- Note:** Changes to the document using the template display after starting a new session by logging into the system.

Edit a Document Template

You can edit the name, layout, and content of a template using Knowledge Management.

To update a document template

1. On the Administration tab, browse to Knowledge, Documents, Document Templates.
The Document Templates List pane appears.
2. Select a template name.
The Update Document Template page appears.
3. Click Edit Detail.
The HTML Editor opens.
4. Change the name, content, or layout of the template. You can change the following fields:

Template

Defines a unique name for the template. This field is required.

Detail

Displays the static content that appears in documents created using the current template. If you select the HTML Source option, you can edit HTML code for the body directly in the Body field. If you select the Quick View option, the Body field is read-only and displays the static body content as it appears at run time.

5. Click Edit Detail to specify the static content and layout of documents that use the template.

The HTML Editor window opens.

6. Edit the code using the toolbar to insert placeholder tags. Click OK.

The Detail field shows the updated content.

7. (Optional) Click Quick View.

The content appears as a document based on the template.

8. (Optional) Click HTML View.

The content appears as HTML code.

9. Click Save.

The Document Template Detail page appears.

10. Click Close Window.

New documents that use the template show the new content and layout.

A change to the name of the template appears in the list when you refresh the list.

List Document Templates

The Document Template List page lets you create and manage document templates, which specify the content and appearance of documents in the knowledge base.

The following default templates install with the product:

- Built In—Knowledge Document
- Built In—Knowledge Tree
- Built In—Quick Editing

The product uses the default templates when you create knowledge documents and knowledge tree documents unless you create document templates and associate them with your documents.

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

To list document templates, select Administration, Knowledge, Documents, Document Templates from the Administration tab.

The Document Template List page appears and includes the following columns:

Template Name

Lists document templates currently defined in the product. Select a template name to open the Update Document Template window. Right-click a template name to open the Template Name shortcut menu, which contains commands for working with the selected template.

Default

Displays a mark to indicate that a template is used by default for new documents or when a specified template for a document is deleted. You can set a template as the default, by right-clicking the template name and selecting Set as Default from the shortcut menu.

Note: When you create a document or delete a template associated with a document, the product associates the default template with it (unless you specify another template on the Attributes tab of the Create New Document window or the Update Knowledge Document window).

This page contains the following fields:

Template

Defines the name of the template for which to search. This field only displays when you click Show Filter.

Additional Search Arguments

Defines additional criteria by which to search. This field only displays when you click a More link in the Knowledge Search pane.

This page contains the following buttons:

Clear Filter

Returns all filter fields on the pane or window to their default values.

Create New

Opens the Create New Document Template window so you can define a new document template.

Search

Initiates a search for items that match the specified criteria. When you specify no criteria, the product returns all appropriate items (for example, folders/documents, contacts, templates, noise words, or permission groups).

Show Filter/Hide Filter

Displays or hides fields with which you can filter a search for items on the current window, pane, or dialog.

Note: You can export list results to Excel for use outside of CA SDM by clicking the Export button on the List page.

Filter the Document Template List

The default templates are used when you create knowledge documents and knowledge tree documents, unless you create document templates and associate them with your documents. The Document Template List page lets you create and manage document templates, which specify the content and appearance of documents in the knowledge base. You can filter templates from this page.

To filter the document template list

1. Click the Administration tab.
The Administration page appears.
2. Click Knowledge, Documents, Document Templates.
The Document Template List appears.
3. Do *any* of the following:
 - Select a template from the Template Name column.
The template opens.
 - Search for a template using the search function.
The template search results appear.
 - Select Show Filter to filter your search.
The template search is filtered.

Update Document Templates to Display Knowledge Relationship Links

You can update knowledge document templates from a previous release of CA SDM to display knowledge relationships. You modify the template by adding the {TAG_PARENT} and {TAG_RELATED} tags, allowing documents using the template to display document relationships, such as parent-child links and tenant.

To modify templates to display knowledge relationships

1. Open the knowledge document template.
The Update Document Template page appears.
2. Click Edit.
3. Select the HTML Source view.
Note: You can also click Edit Detail open the HTML Editor and add the {TAG_PARENT} and {TAG_RELATED} tags.
4. Select *Select Template Placeholder* on the Edit Detail dropdown to add tags to the document template.
The tags are added to the template.
5. Save and close the template.
The template is updated to display knowledge relationships.

How to Create Knowledge Document Links

You can maintain your Knowledge Management environment by creating document links. Knowledge relationships let you create document hierarchies and manage changes to your knowledge documents.

You can create knowledge relationships as follows:

1. Define knowledge document templates to display or hide the parent-child relationships in view mode.
2. Create or modify a knowledge document.

Note: You can only add links to an unpublished knowledge document by default. If the document is already published, open the document in edit or rework mode to create document links. You can modify the permissions for documents before the approval process and after publishing.

3. Create *any* of the following document links from the Related Knowledge tab, as appropriate to your Knowledge Management environment:

- Link the document as a non-hierarchical See Also.

You can link knowledge documents to other existing documents.

- Link the document as a parent or child.

You can link knowledge documents in parent-child relationships.

Note: If you modify a parent document, an alert appears saying that child documents can be affected.

- Link the global document to a tenanted document.

You can link knowledge documents to a single tenant, or to multiple tenants. For example, a child document can have a different tenant than the parent document.

The History tab updates when you create document links to help track changes.

4. Save the document and verify that links appear when opening the document in User View.

You can verify that the document links appear according to the permissions you established.

Create a Parent-Child Link

You can create parent-child relationships for knowledge documents. Use parent-child links to organize related documents for users in your environment.

To create a parent-child link

1. Create a knowledge document or open a document in edit mode.

The knowledge document appears.

2. Select the Related Knowledge tab.

The Categories page appears.

Note: You can also create links from the Find Similar tab. Use search criteria to locate documents.

3. Locate the document you want to link in the Categories (X) pane.

Right-click the document in the Documents pane and select the *Link this as Parent* or *Link this as Child* option to create the link between the documents.

The document link is created.

Create a See Also Link

You can create a See Also relationship for knowledge documents in your environment. Use See Also relationships to manage knowledge hierarchies in your environment.

To create a See Also link

1. Create a knowledge document or open a document in rework mode.

The knowledge document appears.

2. Select the Related Knowledge tab.

The Categories page appears.

Note: You can also create links from the Find Similar tab.

3. Locate the document you want to link in the Categories (X) pane.

Right-click the document in the Documents pane and select the option *Link this Document as See Also*.

The document link is created.

Remove Document Links

You can remove links from an unpublished knowledge document, such as a parent-child relationship. If the document is already published, open the document in edit or rework mode to remove document links.

To remove a document link

1. Open the knowledge document.
The document appears.
2. Click Edit.
The Update Document page appears.
3. Select the Related Knowledge tab.
The Categories page appears.
4. Navigate to the Document Links pane.
5. Select option *Remove this Link* from context menu so as to remove the link.
The document links are removed.

Knowledge Categories

Knowledge documents are arranged into *knowledge categories*. Knowledge engineers, knowledge managers, and administrators can manage the categories. Each of these individuals uses Knowledge Management to create, copy, and modify categories. However, only knowledge managers and administrators can delete categories. When you construct the category structure in Knowledge Categories, you are creating the hierarchical structure that service desk employees, customers, and analysts use to navigate to relevant documents.

Assign each document to a primary category. For example, any knowledge related to Microsoft Word must be added to the Microsoft Word category. In addition, Knowledge Management lets you associate a document with multiple secondary categories and other documents. In this way, a document can be classified under many different applicable categories and can provide for more successful search results.

The category structure performs the following functions:

- Organizes knowledge solutions into manageable groups.
- Makes it easier to assign access rights.

- Makes it possible to search for solutions using FAQ/Browse.
To use category browse functionality from an incident, the incident area and the knowledge category must match. When you create an incident based on a knowledge document, the document category sets the incident area; therefore, the category and the area always match.
- Creates a document link—A *see also* link appears when viewing either of the linked documents. The *see also* link lets you go directly from one linked document to the other.

Create a Knowledge Category

For each category, you can define properties that identify attributes or qualities to be associated with the ticket and create a workflow that identifies all the individual tasks required to fulfill the ticket.

You can use categories to specify default values for certain fields in tickets, or automatically associate a level of service to tickets by assigning a default service type to categories. Whenever an analyst assigns a category to a ticket, all the information you associate with the category is automatically associated with the ticket.

Note: If you are using multi-tenancy, a tenant drop-down list appears in the Knowledge Document search filter. If you select <empty> in this drop-down list, the search is public. A tenant column also appears on the list page.

To create a category

1. On the Administration Tab, browse to Knowledge, Knowledge Categories.
The Knowledge Categories page appears.
2. Right-click the category under which you want to create the category. Select New Category from the shortcut menu.
The Create New Category page opens to the Content tab.
3. Complete the [fields](#) (see page 828) as appropriate.
4. Click Permissions.
The Permissions tab appears.
5. Select one of the following permissions options for the category:

Inherit from Parent

Specifies that the new category has the same permission settings as its parent category.

Note: The Inherit from Parent option is not available if you select the TOP category before opening the Create Category page.

Control by Group

Specifies category permissions for groups to have read or write access to the category.

Control by Role

Specifies category permissions for roles to have read or write access to the category.

Note: If you change controls, such as changing the category permission from group to role, a warning appears that previous permissions are deleted for that category.

Grant Write Permission to Everyone

Specifies that all users have write access to the category. Write access indicates that you can edit or delete this category.

Note: The Grant Read Permission to Everyone check box is automatically selected if you select the Grant Write Permission to Everyone check box.

Grant Read Permission to Everyone

Specifies that all users have read access to the category. Read permission indicates that you can view the category, but you cannot edit or delete it. Users with administrative rights can edit a folder even if their associated permission group cannot. If a user belongs to multiple permission groups with varying levels of access to the category, the user gets the highest available access level (for example, if one group has read-only access and the other write access, the user gets write access).

Note: The Grant Read Permission to Everyone check box is automatically selected if you select the Grant Write Permission to Everyone check box.

Important: When you grant permissions for Everyone, the access by role or group is the same. If you selected Everyone and Control by Role, after you save the permissions, the Control by Group becomes selected.

6. (Optional) Specify the read-write permissions to specific groups or roles from the Available and Selected lists.

You use this option to manage which groups or roles have read or write access to the category. You can select one or more permission groups or roles from the Available Groups/Roles list, and then use the Add and Remove buttons to move the selected groups or roles to the Groups/Roles with Write Permission and Groups/Roles with Read Permission lists.

7. Click Save.

The Category Detail page appears.

8. Click Close Window.

The Knowledge Categories pane refreshes to include the new category.

More information:

[Category Fields](#) (see page 828)

Category Fields

Title

Names the category.

Description

Describes the category.

Category Owner

Indicates the person responsible for the category. When a contact is defined as the owner of a category, the contact has a link on the Knowledge Report Card named "My Categories," from which they can view statistics for that category and the documents it contains. This person is also the default owner for new documents in the category when the user who creates the documents is not an analyst, or an analyst creates the documents with 'Assign to Category Owner' selected.

Documents Template

Defines the document template to use for all documents associated with this category. The <empty> option means that none have been defined, but by default, the predefined template is used.

Approval Process Template

Defines the default template to use for the approval process for all documents associated with this category. The approval process template defines the workflow steps a document must go through before it is published. The default is <empty>, which indicates the application default template is used.

Allow forums to be created in this category

Specifies whether analysts can create forums within this category.

Request/Incident/Problem Area

Designates a Request/Incident/Problem area that your administrator defines to designate an area of responsibility. You can click the search icon to select from the available areas.

Issue Category

Designates an Issue Category that your administrator defines to designate an area of responsibility. You can click the search icon to select from the available areas.

Modify a Category

You can modify a category in your Knowledge Management environment. Categories determine the contents of change orders and issues after their creation. For each category, you can define properties that identify attributes or qualities to be associated with the ticket and create a workflow that identifies all the individual tasks required to fulfill the ticket.

You can use categories to specify default values for certain fields in tickets, or automatically associate a level of service to tickets by assigning a default service type to categories. Whenever an analyst assigns a category to a ticket, all the information you associate with the category is automatically associated with the ticket.

To modify a category

1. On the Administration Tab, browse to Knowledge, Knowledge Categories.
The Knowledge Categories page appears.
2. Right-click the category to modify, and select Edit Category from the shortcut menu.
The Update Category page opens to the Content tab.
Note: You cannot delete or modify the Top category.
3. Update one or more of the [fields](#) (see page 828) as appropriate.
4. Click Permissions.
The Permissions tab appears.
5. Select one of the following permissions options for the category:

Inherit from Parent

Specifies that the new category has the same permission settings as its parent category.

Note: The Inherit from Parent option is not available if you select the TOP category before opening the Create Category page.

Control by Group

Specifies category permissions for groups to have read or write access to the category.

Control by Role

Specifies category permissions for roles to have read or write access to the category. If you select Everyone and Control by Role, after you save the permissions, the Control by Groups becomes selected.

Note: If you change controls, such as changing the category permission from group to role, a warning appears that previous permissions are deleted for that category.

6. (Optional) Specify the read-write permissions for the category.

You use this option to manage which groups or roles have read or write access to the category. When you select this option, the page refreshes to include the Grant Write Permission to Everyone and Grant Read Permission to Everyone [check boxes](#) (see page 832).

7. Click Save.

The modified category appears in the Knowledge Categories List page.

Delete a Category

You can remove a category from the Knowledge Category structure. When you delete a category, you can specify whether to remove the subcategories of the selected category, and all associated document links.

Note: This feature is only available to system administrators and knowledge managers.

To delete a category

1. On the Administration Tab, browse to Knowledge, Knowledge Categories.

The Knowledge Categories page appears.

2. Right-click the category to delete, and select Delete Category from the shortcut menu.

The Delete Category page appears.

Note: You cannot delete or modify the Top category.

3. (Optional) Do *one* of the following:
 - Select the Include Subcategories check box to delete all subcategories of the selected category.
 - Clear the Include Subcategories check box to move all subcategories from the selected category to its nearest available parent category.

4. (Optional) Do *one* of the following:
 - Select the Include Documents check box to delete documents that reside in the selected category. When you select the Include Documents check box, the following options display:
 - **Delete documents linked by primary category only**—Deletes only documents for which the selected category is identified on the Knowledge Categories pane as the primary category of the document.
 - **Delete all documents linked to the category**—Deletes all documents in the selected category.
 - Select the Include Documents check box to relocate documents from the selected category to its nearest available parent category.

Click OK.

The product deletes the selected category and the Knowledge Category pane refreshes.

Move a Category

You can move a category, its subcategories, and all associated document links from their current location to another category.

Note: This feature is only available when Knowledge Management is installed.

To cut and paste a category

1. On the Administration Tab, browse to Knowledge, Knowledge Categories.
The Knowledge Categories page appears.
2. Right-click the category to move, and select Cut Category from the shortcut menu.
The product puts the selected category, its subcategories, and all associated document links in memory.
3. Right-click the category into which to paste the cut information, and select Paste Category from the shortcut menu.
The cut category moves from its original location to beneath the selected category.
The Knowledge Categories pane refreshes to display the new category structure.

Copy a Category with Document Links

You can put a copy of a category, its subcategories, and all associated document links into another category without removing the selection from its original location.

To copy and paste a category with document links

1. On the Administration Tab, browse to Knowledge, Knowledge Categories.
The Knowledge Categories page appears.
2. Right-click the category to copy, and select Copy Category with Document Links from the shortcut menu.
The product puts the selected category, its subcategories, and all associated document links in memory.
3. Right-click the category into which to paste the copied information, and select Paste Category from the shortcut menu.
The copied information appears beneath the selected category. The Knowledge Categories pane refreshes to display the new category structure.

Copy a Category without Document Links

You can place a copy of a category and its subcategories into another category without removing the selection from its original location or copying associated document links.

To copy and paste a category without document links

1. On the Administration Tab, browse to Knowledge, Knowledge Categories.
The Knowledge Categories page appears.
2. Right-click the category to copy, and select Copy Category from the shortcut menu.
The product puts the selected category and its subcategories in memory.
3. Right-click the category into which to paste the copied information, and select Paste Category from the shortcut menu.
The copied information appears beneath the selected category. The Knowledge Categories pane refreshes to display the new category structure.

Manage Category Permissions

You can manage category permissions in your Knowledge Management environment by granting specific read/write permissions based on group or role. You can specify who can view, edit, delete, or add subcategories to a category.

Note: You can use empty write permissions for documents and categories. Empty write permissions are set when you do not select write permissions. Only privileged users can modify categories and documents with empty write permissions.

To manage category permissions

1. On the Administration tab, browse to Knowledge, Knowledge Categories.

The Knowledge Categories page appears.

2. Right-click the category under which you want to create the category, and select New Category from the shortcut menu to manage the permissions of a new category.

The Create New Category page displays the Content tab.

3. Click Permissions.

The Permissions tab appears.

4. Select one of the following permissions options for the category:

Inherit from Parent

Specifies that the new category has the same permission settings as its parent category.

Note: If you select the Top category before opening the Create Category page, the Inherit from Parent option is not available.

Control by Group

Specifies category permissions for groups to have read or write access to the category.

Control by Role

Specifies category permissions for roles to have read or write access to the category.

Note: If you change controls, such as setting the permission from group to role, a warning appears that says existing permissions are deleted.

Grant Write Permission to Everyone

Specifies that all users have write access to the category. Write access indicates that you can edit or delete this category.

Note: The Grant Read Permission to Everyone check box is automatically selected if you select the Grant Write Permission to Everyone check box.

Grant Read Permission to Everyone

Specifies that all users have read access to the category. Read permission indicates that you can view the category, but you cannot edit or delete it. Users with administrative rights can edit a folder even if their associated permission group cannot. If a user belongs to multiple permission groups with varying levels of access to the category, the user gets the highest available access level (for example, if one group has read-only access and the other write access, the user gets write access).

Note: The Grant Read Permission to Everyone check box is automatically selected if you select the Grant Write Permission to Everyone check box.

Important: When you grant permissions for Everyone, the access by role or group is the same. If you selected Everyone and Control by Role, after you save the permissions, the Control by Group becomes selected.

5. Grant the appropriate read/write permissions to specific groups or roles from the Available and Selected lists.

You can select one or more permission groups or roles from the Available Groups/Roles list, and then use the Add and Remove buttons to move the selected groups or roles to the Groups/Roles with Write Permission and Groups/Roles with Read Permission lists.

Note: All groups or roles added to the permission list are automatically added to the read permission list. When you select the Grant Read Permission to Everyone check box, you do not need to add groups or roles to the read permission list.

6. Click Save.

The Category Detail page appears.

7. Click Close Window.

The modified category appears in the Knowledge Categories pane.

Reports and Metrics

You can monitor the efficiency of the knowledge base using the following reporting tools:

- Knowledge Report Card
- Predefined web-based reports for Knowledge Management
- Role-based web forms

These tools let you view statistics on the usefulness of your documents and their effectiveness in solving problems.

More information:

[Knowledge Report Card](#) (see page 835)

[Web-Based Reports](#) (see page 836)

[Role-Based Report Web Forms](#) (see page 837)

Knowledge Report Card

The Knowledge Report Card displays information about the knowledge contribution from each end user and provides feedback about which knowledge documents are most effective. You can use the information to improve the processes of creating knowledge documents and providing the best support to end users in your environment.

Define Knowledge Report Card Statistics

Use the Knowledge Report Card page to define the schedule at which the product calculates and sends notifications about the Knowledge Report Card and to define the content of Knowledge Report Card notification emails.

Note: Rework versions and retired documents are not presented when the Knowledge Report Card calculation is run.

To define Knowledge Report Card statistics

1. On the Administration tab, select Knowledge, Knowledge Report Card.

The Report Card page opens.

2. Complete the following fields as appropriate.

Last Updated

Runs the Report Card calculation.

Default: Deactivated

Note: When the calculation is not run, and the statistics to present the data are not collected, the following message appears when the Knowledge Report Card command is specified from the View menu on the Knowledge tab: "Please Run Report Card Calculation."

Schedule

Schedules the Report Card.

- **Report Card Calculation will next run on xxx and will run every xxx**—Specifies the frequency with which to recalculate the Report Card statistics.
 - **Report Card Email Notifications will be sent on xxx and will be sent every xxx**—Specifies the frequency with which to send Report Card notifications.
- Default:** Never
- **Report Card Email Should Display Statistics for the Past xxx**—Specifies the amount of time for which the Report Card notification contains information. This field is only available when you select a value other than Never from the Report Card Email Notifications will be Sent Every xxx list.

Default: 365 days

3. Click Save.

The Knowledge Report Card statistics are defined.

Web-Based Reports

CA Business Intelligence installs a set of predefined Knowledge Management reports. These reports are automatically deployed to the BusinessObjects reporting server after the CA SDM installation.

The reports are developed with either BusinessObjects Web Intelligence or Crystal Reports. Authorized users can display reports on the CA SDM Reports Tab.

Role-Based Report Web Forms

You can define the report web forms that display when an authorized manager or analyst clicks the Report List icon on the Reports tab.

Report web forms are managed through the Role List in the following location: Administration, Security and Role Management, Role Management.

Search

The Search feature enables administrators to perform the following tasks:

- Manage the Knowledge Management search engine and define the settings used to manage noise words, special terms, and synonyms that are excluded or included in searches.
- Define settings used to parse documents.
- Define default search settings.
- Create "recommended documents" that display in the search results. Dynamic FAQ listing is used to push (bubble up) the recommended documents to users.

Note: A document can have different permissions than the attachments linked to the document.

See also

[KT Search Engine](#) (see page 837)

[Recommended Documents](#) (see page 847)

[Set Up Default Search Settings](#) (see page 851)

KT Search Engine

After you install CA SDM, the KT search engine is configured as the default. Searches of the knowledge base are limited to knowledge documents. The search engine is located on the Administration tab, Knowledge, Search, KT Search Engine node.

The KT Search Engine node lets you manage the following options:

- Noise Words
- Special Terms
- Synonyms

More information:

[Use the Knowledge Management Search](#) (see page 838)

Use the Knowledge Management Search

After you install CA SDM, the Knowledge Management search engine is configured as the default search engine. Searches of the knowledge base are limited to knowledge documents.

You can define accessibility and defaults to all knowledge sources based on a user role. By default, knowledge documents are searchable for all user roles.

Follow these steps:

1. Click the Administration tab.
The Administration console appears.
2. Click Options Manager, Search Engine.
The Option List appears.
3. Click ebr_version.
The Options Detail page appears.
Click Edit. The Update Options page appears.
4. Select KT Search Engine.
5. Click Save, Refresh.
The Options Detail page is updated with your selection.
6. Click Close Window.
7. Recycle CA SDM services.

See also

[Noise Words, Synonyms, and Special Terms](#) (see page 839)

[Create Noise Words](#) (see page 839)

[Edit a Noise Word](#) (see page 840)

[View Noise Words](#) (see page 840)

[Create a Special Term](#) (see page 841)

[Edit a Special Term](#) (see page 842)

[View Special Terms](#) (see page 842)

[Create a Synonym](#) (see page 843)

[View Synonyms](#) (see page 843)

[Edit a Synonym](#) (see page 844)

[Parse Settings](#) (see page 844)

Noise Words, Synonyms, and Special Terms

You can define words (synonyms, noise words, and special terms) that affect natural language and keyword searches performed in CA SDM. Adding or deleting the following terms has a significant effect on search results that are returned to the user:

Noise Words

Specifies words that do not generally contribute to the search process and can, therefore, be ignored. For example, prepositions such as a, an, the, or, and to are often identified as noise words. Use the Noise Words List page to specify words that the product can ignore in documents and queries without affecting search results.

Special Terms

Specifies a term that you want identified as a single word during the search process, although it can consist of several words or contain special characters. For example, words that have a nonalphanumeric character, such as the forward slash (/) in TCP/IP, the hyphen (-) in dial-up, or the underscore (_) in LOCAL_SERVER can be added as special terms. In your evaluation of which words to define, consider valid words that can be divided during the search process because they have a nonalphanumeric character.

Synonyms

Specifies a word that has the same meaning as another word. You can add synonyms so when a user searches for a particular word, and a corresponding synonym for that word exists in your knowledge base, the information can be found. You can define several synonyms for the same word. The system automatically creates reverse synonyms from the keywords you define. For example, if you define computer as a synonym for the word PC, PC automatically becomes a synonym for the word computer. Use the Synonyms List page to specify keyword/synonym pairs that the product uses interchangeably when parsing documents and queries. These keyword/synonym pairs can improve search results.

Note: After you create, modify, or delete noise words, special terms, synonyms, or parse settings, use the Knowledge Re-Index utility provided with the product to re-index the knowledge base.

Create Noise Words

Use the Create New Noise Word page to add a new word to the list of words that the product can ignore in documents and queries without affecting search results.

To create a noise word

1. On the Administration tab, select Knowledge, Search, KT Search Engine, Noise Words.

The Noise Words List page appears.

2. Click Create New.
3. The Create New Noise Word page appears.
4. Enter the word you want to define as a noise word in the Noise Word FIELD.

Note: You cannot define a noise word that exists as a synonym or keyword.

5. Click Save

The new noise word appears on the Noise Words List page. You can click the Edit button to update the new noise word.

Note: After you create, modify, or delete noise words, special terms, synonyms, or parse settings, use the Knowledge Re-Index utility provided with the product to re-index the knowledge base.

Edit a Noise Word

You can edit a noise word that has already been created on Noise Words List page.

Note: After you create, modify, or delete noise words, special terms, synonyms, or parse settings, use the Knowledge Re-Index utility provided with the product to re-index the knowledge base.

To edit a noise word

1. Right-click the desired noise word in the Noise Words list, and select Properties from the shortcut menu.

The Update Noise Words page opens.

2. Enter your edits in the Noise Word field.

Note: You cannot define a noise word that exists as a synonym or keyword.

3. Click Save.

The updated noise word appears on the Noise Word List page.

Note: After you define, modify, or delete noise words, special terms, synonyms, or parse settings, use the Knowledge Re-Index utility provided with the product to re-index the knowledge base.

View Noise Words

You can view the summary information for each noise word on the Noise Words List page.

To display summary information for each noise word, select the Administration tab, then Knowledge, Search, KT Search Engine, Noise Words.

The Noise Words List page appears and contains the following fields:

Noise Word

Defines a noise word for which to search. This field only displays when you click Show Filter.

Additional Search Arguments

Defines additional criteria by which to search. This field only displays when you click a More link in the Knowledge Search page.

Noise Words List

Displays words to filter out of search requests. The list displays when you click the Search button. If you entered a word in the Noise Word field, the list only displays the specified word. If you entered nothing in the Noise Word field, the list displays all noise words defined for the product.

From the Noise Words List page, you can perform the following tasks:

- [Create Noise Words](#) (see page 839)
- [Edit a Noise Word](#) (see page 840)

Create a Special Term

Use the Create New Special Term page to specify a word or phrase that the product treats as single keywords when parsing documents and queries.

To create a special term

1. On the Administration tab, browse to Knowledge, Search, KT Search Engine, Special Terms.
The Special Terms List page appears.
2. Click Create New.
The Create New Special Term page appears.
3. Enter the word or phrase you want to define as a special term in the Special Term field.
4. Click Save.

The Create Special Term page closes and the Special Term Detail page opens so you can review the word or phrase you added. You can use the Edit button to update the new term. The new special term appears on the Special Terms List page.

Edit a Special Term

You can edit a special term.

To edit a special term

1. Right-click the term you want to edit in the list, and select Edit from the short-cut menu.

The Update Special Term page appears.

2. Enter the word or phrase you want to define as a special term in the Special Term field.
3. Click Save.

The Update Special Term page closes and the Special Term Detail page appears so you can review the word or phrase you added. The updated term appears on the Special Terms List page.

View Special Terms

You can view the summary information for each special term.

To view special terms, select the Administration tab, then Knowledge, Search, KT Search Engine, Special Terms.

The Special Terms List page appears, and you can use the following fields to modify the terms that appear by default, or define your own:

Special Term

Defines a special term for which to search. This field only displays when you click Show Filter.

Additional Search Arguments

Defines additional criteria by which to search. This field only displays when you click a More link in the Knowledge Search pane.

Special Terms

Displays words or phrases containing spaces or characters (such as hyphens, slashes, and so on) that the product treats as single keywords when parsing queries. If you entered a word in the Special Term field before clicking Search, the list only displays the specified word. If you entered nothing in the Special Term field, the list displays all special terms defined for the product.

From the Special Terms List page, you can perform the following tasks:

- [Create a special term](#) (see page 841)
- [Edit a special term](#) (see page 842)

Create a Synonym

Synonyms are words or phrases that have the same or nearly the same meaning as the keyword with which they are associated.

If you define a new complex synonym (that is, a synonym containing multiple words separated by spaces or other delimiters), also create an identical special term so that the product can treat the synonym as a single entity. For example, if you define "Computer Associates" as a synonym for "CA", also define "Computer Associates" as a special term.

Note: You cannot define a synonym or keyword that exists as a noise word.

To create a synonym

1. On the Administration tab, browse to Knowledge, Search, KT Search Engine, Synonyms.

The Synonyms List page appears.

2. Click Create New.

The Create New Synonyms page appears.

3. Enter the word or phrase you want to define as a synonym in the Synonyms field.

4. Click Save.

The Create Synonyms page closes and the Synonyms Detail page opens so you can review the word or phrase you added. The new synonym appears on the Synonyms List page.

View Synonyms

You can view the summary information for each synonym.

To view synonyms, select the Administration tab, then Knowledge, Search, KT Search Engine, Synonyms.

The Synonyms List page appears, and you can modify the keywords that appear by default, or define your own using the following fields:

Keyword

Defines a keyword for which to search. This field only displays when you click Show Filter.

Synonym

Defines a synonym for which to search. This field only displays when you click Show Filter.

Additional Search Arguments

Defines additional criteria by which to search. This field only displays when you click a More link in the Knowledge Search pane.

Synonyms List

Displays keyword/synonym pairs defined in the product. For each keyword displayed in the Keyword column, the Synonym column displays one or more synonyms.

From the Synonyms List page, you can perform the following tasks:

- [Create a synonym](#) (see page 843)
- [Edit a synonym](#) (see page 844)

Edit a Synonym

You can edit a synonym.

To edit a synonym

1. Right-click the term you want in the list, and select Edit from the short-cut menu.
The Update Synonym page appears.
2. Enter the word or phrase you want to define as a special term in the Synonym field.
3. Click Save.

The Update Synonym page closes and the Synonym Detail page opens so you can review the word or phrase you added. The updated term appears on the Synonym List page.

Parse Settings

When you publish a document to the knowledge base, the product parses the information in the Title, Summary, Problem, and Resolution fields of the document into keywords. When a user searches the knowledge base, the product compares keywords from the user query with the keywords parsed from the knowledge base to produce a result list. The Parse Settings page lets you define the settings used to parse documents in the knowledge base.

Note: The Parse Settings feature is only available with the default KT search engine.

Define Parse Settings

When you publish a document to the knowledge base, the product parses the information in the Title, Summary, Problem, and Resolution fields of the document into keywords. When a user searches the knowledge base, the product compares keywords from the user query with the keywords parsed from the knowledge base to produce a result list.

To define the settings used to parse documents in the knowledge base, browse the Administration tab in Knowledge, Search, Parse Settings.

The Parse Settings appears, and you can use the following fields to define settings:

Maximum Search Keywords

Defines the maximum number of keywords to extract when the product parses the search text.

Default: 20

Note: The valid range is 1-100, so that a CA SDM Knowledge administrator can change the value within this range, based on search needs and parameters of a specific Knowledge database. Use a lower number of search keywords for faster performance.

Language

Specifies the language type to use for parse processing. Select one of the following settings:

English

Performs certain types of processing specific to the English language (for example, de-pluralizing search terms) during a search, if applicable.

Other European

Performs only European specific processing during the search.

Korean

Performs only Korean specific processing during the search.

Other Far East

Performs processing for other far east languages during the search.

Note: When you are in a Chinese, Japanese, or Korean operating environment, verify that you understand the available parsing approaches, and limitations of [Multi-Byte Character Set \(MBCS\) languages](#) (see page 847), before implementing your Knowledge Management system to help ensure that user expectations are set appropriately.

Valid Character Range

Defines the range of alphanumeric characters to consider valid when parsing the Title, Summary, Problem, and Resolution fields in a document. The product treats any other characters as separators.

Note: When you select Yes from the Recognize Special Terms list, the product does not parse words and phrases defined as special terms.

Default: a-z, which indicates that the alphabetic characters a through z are valid characters for parsing.

The Valid Character Range field contains the appropriate letters that parsing uses. The letters not presented in the Valid Character Range are removed.

The recommended values for different languages are as follows:

Language	Valid Character Range
German	a-zäöüß
Spanish	a-záéíóúñ
French	a-zàâçêëëïîôû
Portuguese Brazil	a-zàáãçéêíóú
Italian	a-zàèéîïù
Simplified Chinese	a-z
Japanese	a-z
Traditional Chinese	a-z
Korean	a-z

Note: Japanese contains the "a-z" range plus a list of Katakana valid characters, excluding punctuation marks.

Remove Similar Words

Specifies whether the product removes structurally similar keywords from the groups used in a search. You can select one of the following settings:

Yes

Removes structurally similar keywords from the search criteria.

Note: When you select Yes, the product also removes similar words when you save or publish the document. This setting can impact whether a document is searchable if the Remove Similar Words field was set to Yes. The similar word may have not been indexed and used in the later search and retrieval of the document.

No

Leaves structurally similar keywords in the search criteria.

Default: No

Remove Noise Words

Specifies whether the product removes noise words when parsing the Title, Summary, Problem, and Resolution fields in a document. You can select one of the following settings:

Yes

Removes noise words from the search criteria.

No

Leaves noise words in the search criteria.

Default: Yes

Recognize Special Terms

Specifies whether the product considers special terms as single entities or as multiple words when parsing the Title, Summary, Problem, and Resolution fields in a document. You can select one of the following settings:

Yes

Processes special terms as single entities in the search criteria.

No

Processes the words that comprise special terms as separate entities in the search criteria.

Default: Yes

Multi-Byte Character Set Search Limitations

Make sure that you understand the available parsing approaches, and limitations of MBCS languages, before implementing your Knowledge Management system to help ensure that user expectations are set appropriately. This limitation of the product impacts search features using Japanese, Chinese, or Korean language text within the system. The word parsing mechanism used by the search mechanism is controlled on the [Parse Settings](#) (see page 844) page.

For the English, Other European, and Korean settings, the product assumes that punctuation, "white space," or both characters separates words. This assumption allows the document text to be broken into specific words, and allows noise words to be ignored and application of known synonyms and special terms to search terms.

Alternatively, when the Far East language setting is selected, the parsing routine uses a character-by-character parsing approach to accommodate some Far East language text approaches of not using white-space delimiters between words. This setting tells the parser to assume that each character is treated as a full word. The setting applies to all text to be searched. Because the language setting changes the way that the search parsing works, the entire search index must be recreated if the language setting is changed to or from Far East.

Recommended Documents

CA SDM users can specify criteria about an item of interest and the search engine finds the matching knowledge documents and display them on the search results page as a set of "recommended document" links. The search query can be expressed as a keyword or set of words (phrase) that identify the desired concept that one or more documents can contain.

The list of documents that meet the search criteria is sorted, and ranked (from highest to lowest) to place the most relevant documents first in the search results. Using recommended documents helps users reduce the time required to find the desired information.

To provide a set of matching documents that are sorted according to some criteria quickly, the search engine collects data through the condition type (phrase, keywords or category), which the administrator configures on the Create Recommended Documents page.

More information:

[Create Recommended Documents](#) (see page 848)

[Edit a Recommended Document Condition](#) (see page 850)

[View Recommended Documents](#) (see page 850)

[Search Recommended Documents](#) (see page 851)

Create Recommended Documents

Administrators can create recommended documents that users can find when they specify criteria about an item of interest.

Note: If multi-tenancy is installed, select the appropriate tenant from the drop-down list. The public (shared) option creates the object for all tenants.

To create a recommended document

1. On the Administration tab, browse to Knowledge, Search, Recommended Documents.

The Recommended Documents List appears.

2. Click Create New.

The Create New Recommended Documents page appears.

3. Complete the following fields as appropriate.

Knowledge Document

Specifies a knowledge document or click the search icon to open the Knowledge Document Search page.

Condition Type

Specifies a [condition type](#) (see page 849) by which the search engine sorts and matches the document.

- If the condition type is Full Match, Exact Phrase, or Keywords, a text field appears. You can enter a phrase or keyword that identifies the concept you want for the document to contain.
- If the condition type is Knowledge Category, the Knowledge Category link appears. You can specify a knowledge category to associate with this document.

Status

Defines the status of this record as active or inactive.

Click Save.

The new recommended document is saved to the knowledge base and appears on the Recommended Documents List page.

Condition Type Field

The search engine locates documents by the following condition types:

Full Match

Searches for documents by the search phrase entered in the search text. A match occurs only when the search engine finds all the same words in a phrase.

Exact Phrase

Searches for documents by the exact phrase entered in the search text. A match occurs only when the search engine finds the exact set or sequence of words in a phrase.

Keywords

Searches for documents by the keywords entered in the search text. A match occurs only when the search engine finds all of the keywords.

Knowledge Category

Searches for documents by knowledge category. A match occurs only when a user navigates to a category configured for recommended documents.

Edit a Recommended Document Condition

You can update a recommended document condition.

To edit a recommended document condition

1. On the Administration tab, browse to Knowledge, Search, Recommended Documents.

The Recommended Document List page appears.

2. To edit a condition, right-click the title in the Condition column.

The Update Recommended Document page appears.

3. Complete the fields as appropriate.

4. Click Save.

The updated condition appears on the Recommended Document list.

View Recommended Documents

You can view the summary information for each recommended document.

To view recommended documents, select Search, and then Recommended Documents on the Administration tab.

The Recommended Document List page appears and lists the following columns:

Condition

Indicates the condition by which the search engine sorts and matches document.

Knowledge Document

Displays the documents in the result set.

Author

Defines an author of the knowledge document.

Modify Date

Displays the date on which the document was last modified.

From the Recommended Document List page, you can perform the following tasks:

- [Create a recommended document](#) (see page 848)
- [Edit a recommended document condition](#) (see page 850)
- [Search recommended documents](#) (see page 851)

Search Recommended Documents

You can use the search function to filter the Recommended Documents List to show only the items you want to see.

Note: If multi-tenancy is installed, the list page displays Tenant and Public Data settings in the search filter. Public Data can be Excluded or Included with Tenant data; Only searches for public objects exclusively. On detail pages, select the appropriate tenant from the list. If you select <empty>, the object is public.

To search for a recommended document

1. On the Administration Tab, browse to Knowledge, Search, Recommended Results.
The Recommended Documents List appears.
2. Click Show Filter.
3. Complete the fields as appropriate. The following fields require more explanation.

Type

Specifies a condition type by which the document is matched and displayed as a recommended document link.

Phrase/Keywords

Specifies search query information by which to identify the desired concept that the document contains.

Click Search.

The Recommended Documents List populates with all items that match your search criteria.

Set Up Default Search Settings

You can set up search options to be used as the default options that appear when users search for knowledge using the search field.

Note: These search options are overwritten by any personal search settings users define in the Preferences window, or any additional search options in the Knowledge Search pane on the Knowledge tab or in the Knowledge Categories pane on the Administration tab.

To set up default search options

1. From the Administration tab, navigate to Knowledge, Search, Search Settings.
The Search Options page displays.

2. Select the following options as appropriate:

Recommended Results

Specifies the number of documents to display in the search results list.

Default Search Fields

Specifies which document fields to include by default in keyword searches. Select a check box to include the associated field in default searches. Clear a check box to exclude the associated field from default searches. The following document fields are available for searching:

- Title
- Summary
- Problem
- Resolution
- Attachments

Search Settings for All Sources

Specifies whether searches can include all knowledge sources. For example, knowledge categories and request areas.

Search Settings in a Ticket Context

Specifies whether searches can include all fields defined in a service desk ticket (incident, problem, issue, change order, or request).

- For these options, select *one* of the following match types:
 - **Any of the words (OR)**—Includes a document in the result set when it contains any of the words in the Search field. This option is the default selection.
 - **All of the words (AND)**—Includes a document in the result set only when it contains all the words in the Search field.

Click Save.

The default search settings are set up.

CA SDM Integration Options

The following configuration options are available for CA SDM integration:

Field Mapping

Specifies which CA SDM fields to populate with Knowledge Management information, and whether to overwrite existing information.

Issue Search Configuration

Selects the fields to search when you click the Search Knowledge button on a ticket.

Request Issue Search Configuration

Selects the fields to search when you click the Search button on a ticket.

Request/Incident/Problem Search Configuration

Selects the fields to search when you click the Search button on a ticket.

Suggest Knowledge

Selects the fields to search when you click the View Knowledge before Save button on a ticket.

More information:

[Define Field Mapping](#) (see page 853)

[Define Issue Search Configuration](#) (see page 855)

[Define Request/Incident/Problem Search Configuration](#) (see page 856)

[Knowledge Suggestions](#) (see page 857)

[Define Issue Categories](#) (see page 858)

[Define Request/Incident/Problem Areas](#) (see page 859)

[Configure Self-Service Policies](#) (see page 860)

Define Field Mapping

Administrators can use the Field Mapping page to specify which fields to populate with Knowledge Management information, and whether to overwrite existing information.

To define fields in service desk for Knowledge Management

1. On the Administration tab, browse to Knowledge, Service Desk Integration, Field Mapping.

The Field Mapping page appears.

2. Complete the following fields as appropriate:

Populate Service Desk Values from Knowledge Management

Specifies whether to use information from Knowledge Management to populate fields in service desk issues or requests.

- Select the check box to make fields in the Knowledge Management, Populate Empty Service Desk Values, and Overwrite Service Desk Values columns available so you can specify which Knowledge Management information to use to populate fields in service desk issues or requests.
- Clear the check box to make fields in the Knowledge Management, Populate Empty Service Desk Values, and Overwrite Service Desk Values unavailable. In this case, users must manually populate service desk issues or requests created from Knowledge Management.

Default: This check box is selected.

Service Desk

Identifies the fields in issues or requests that correspond to fields listed in the Knowledge Management column.

For each check box selected in the Populate Empty Service Desk Values column, information from the corresponding field in the Knowledge Management column populates the issue or request.

Knowledge Management

Identifies the Knowledge Management fields that correspond to the service desk fields listed in the Service Desk column.

For each check box selected in the Populate Empty Service Desk Values column, information from the corresponding field in the Knowledge Management column populates the issue or request.

The Knowledge Management column contains two drop-down lists:

- The first drop-down list corresponds to the Summary field in the Service Desk column and specifies the Knowledge Management field (Title, Summary, or Problem) with which to populate the Summary field in an issue or request.

Default: Summary.

- The second drop-down list corresponds to the Description field in the Service Desk column and specifies the Knowledge Management field (Title, Summary, or Problem) with which to populate the Description field in an issue or request.
- **Default:** Problem.

Populate Empty Service Desk Values

Specifies which empty fields in a service desk issue or request to populate with information from Knowledge Management.

- Select a check box to map information from the Knowledge Management field to the corresponding service desk field if that field currently contains no information.
- Clear a check box if you do not want to map information from the Knowledge Management field to the corresponding service desk field.

Default: The check boxes corresponding to the

Summary, Description, Product, Asset, and Request Area fields in service desk are selected.

Overwrite Service Desk Values

Specifies which fields in a service desk issue or request to overwrite with information from Knowledge Management.

- Select a check box to replace information in the service desk field with information from the corresponding Knowledge Management field.
- Clear a check box if you do not want to replace information in the service desk field with information from the corresponding Knowledge Management field.

These check boxes are only available when the corresponding check boxes in the Populate Empty Service Desk Values column are selected.

Default: All Overwrite Service Desk Values check boxes are cleared.

Click Save.

Field mapping is defined.

Define Issue Search Configuration

You can define the fields within an Issue to search in when you click the Search Knowledge button on a ticket. The fields you select are copied to the corresponding fields in the Search Filter on the Knowledge tab of the Issue Detail window. The population of the Search Filter fields from the ticket occurs when the Knowledge tab is selected or the Reset Filter button (on the Knowledge tab) is clicked.

To define the fields in an Issue to use for Knowledge Base searches

1. From the Administration tab, select Knowledge, CA SDM Integration, Issue Search Configuration.

The CA SDM Integration page displays.

2. Select the fields you want to be available for Knowledge Base searches.

- Summary
- Description
- Configuration Item
- Priority
- Category
- Root Cause
- Product

Note: You cannot select both the Summary and Description fields.

3. Select the "Automatically run search when the Knowledge tab of an issue is selected" option if you want to search the knowledge base automatically when the Knowledge tab on the detail page is selected.
4. Click Save.
Issue search is configured.

Define Request/Incident/Problem Search Configuration

You can define the fields within a request, incident, or problem to search in when you click the Search Knowledge button on a ticket. The fields you select are copied to the corresponding fields in the Search Filter on the Knowledge tab of the ticket detail page. The population of the Search Filter fields from the ticket occurs when the Knowledge tab is selected or the Reset Filter button (on the Knowledge tab) is clicked.

To define the fields in a request, incident, or problem to use for Knowledge Base searches

1. From the Administration tab, select Knowledge, CA SDM Integration, Request/Incident/Problem Search Configuration.

The CA SDM Integration page displays.

2. Select the fields you want to be available for Knowledge Base searches.

- Summary
- Description
- Configuration Item
- Severity
- Impact
- Urgency
- Priority
- Request Area
- Root Cause

Note: You cannot select both the Summary and Description fields.

3. Select the "Automatically run search when the Knowledge tab of a request is selected" option if you want to search the knowledge base automatically when the Knowledge tab on the detail page is selected.
4. Click Save.

The fields in a request, incident, or problem to use for Knowledge Base searches search is configured.

Knowledge Suggestions

Employees and customers can, where permitted, view a list of knowledge suggestions when they create a ticket in the self-service interface.

If a solution is found, and the ticket is not saved, the documents that were suggested can be credited through a self-service rating system in the document form. This rating system differs depending on the self-service policy settings defined on the Search Settings page.

The data retrieved can be used for reports, dashboards and also while searching the knowledge base, where users can filter the documents that have successfully resolved their tickets.

The benefits of self-service are in the form of fewer support calls and redundant tickets created, which translates into reduced operational costs.

The administrator must enable this feature before use and configure the appropriate issue categories and request areas for which knowledge is suggested in the self-service interface.

Define Issue Categories

You can define issue categories for which knowledge is suggested to customers and employees during ticket creation.

You can also mark the Suggest Knowledge feature as active or inactive. When you mark this feature as inactive, it is no longer available to customers and employees, but it remains available in the database for future use. If you decide to use this feature in the future, you can go back and mark it as active.

To define suggest knowledge issue categories

1. On the Administration tab, browse to Knowledge, Service Desk Integration, Suggest Knowledge, Issue Categories.

The Suggest Knowledge for Issue Categories page appears.

2. Select the **Do not suggest knowledge** option to mark the Suggest Knowledge feature as active.

Default: Inactive

If you mark this feature active, the following options appear:

By default knowledge will be:

Suggested

Indicates that you *do* want knowledge suggested for all issue categories except the definitions in the issue categories list.

Not Suggested

Indicates that you *do not* want knowledge suggested for all issue categories except the definitions in the issue categories list.

For all Issue Categories except the following:

Displays the list of request areas where knowledge is either suggested or not suggested to employees and customers in the self-service interface. The self-service user is not allowed to edit the request areas, they are read-only.

3. Click *one* of the following buttons:

Add

Adds the selected request area to the list.

Remove

Removes the selected request area from the list.

Remove All

Removes all request areas from the list.

Save

Saves the request area information to the knowledge base.

Issue categories are defined.

Define Request/Incident/Problem Areas

You can define request, problem, and incident areas for which knowledge is suggested to customers and employees during ticket creation.

You can also mark the Suggest Knowledge feature as active or inactive. When you mark this feature as inactive, it is no longer available to customers and employees, but it remains available in the database for future use. If you decide to use this feature in the future, you can go back and mark it as active.

To define suggest knowledge request/incident/problem areas

1. On the Administration tab, browse to Knowledge, CA SDM Integration, Suggest Knowledge, Request/Incident/Problem Areas.

The Suggest Knowledge for Request/Incident/Problem Areas page appears.

2. Select the **Do not suggest knowledge** option to mark the Suggest Knowledge feature as active.

Default: Inactive

If you mark this feature as active, some additional options appear:

By default knowledge will be:

Suggested

Specifies this option if you want knowledge suggested for all request/incident/problem areas except those defined in the Request Area list.

Not Suggested

Specifies this option if you *do not* want knowledge suggested for all request/incident/problem areas except those defined in the request area list.

For all Request/Incident/Problem Areas except the following:

Displays the list of request areas where knowledge is either suggested or not suggested to employees and customers in the self-service interface. The self-service user is not allowed to edit the request areas, they are read-only.

Click Save.

The suggest knowledge request/incident/problem areas are defined.

Configure Self-Service Policies

You can configure self-service policies that credit documents based on a set of user scenarios.

To configure self-service policies

1. On the Administration tab, browse to Knowledge, CA SDM Integration, Suggest Knowledge, Self-Service Configuration.

The Search Settings page appears.

2. Specify the appropriate policy settings that credits documents and save avoided tickets based on the following user scenarios:
 - User did not open any suggested document
 - User opened a suggested document
 - User accepted a suggested document as a solution to their problem
 - User searched for knowledge, opened the document, and exited

Click Save.

Self-service policy settings are saved.

Solution Survey

Solution Surveys let you collect and analyze customer feedback about Knowledge Document performance. You can modify the survey settings by selecting Knowledge, Solution Survey from the Administrative Interface. The survey appears on a published Knowledge Document and lets customers, guests and employees rate the effectiveness of the Knowledge Document.

This feature contains the following components:

- FAQ Settings
- Survey Settings

More information:

[Define FAQ Settings](#) (see page 861)

[Define Solution Survey Settings](#) (see page 863)

Define FAQ Settings

Use the FAQ Settings page to set parameters by which the product calculates the FAQ rating assigned to each document. The product bases the FAQ rating on the following criteria:

- How frequently the document was accessed in the past
- How helpful the document was to users
- How the document's effectiveness has decreased over time

By default, the document list page displays documents sorted in order of FAQ rating (that is, in order of usefulness). The most useful documents "bubble up" to the top of the document list. Over time, documents tend to move downward in the document list as users learn solutions to the problems.

To define FAQ Settings

1. Select Knowledge, Solution Survey, FAQ Settings from the Administration tab.

The FAQ Settings page appears.

2. Complete the following fields as appropriate:

Last Updated

Specifies whether to run the FAQ Rating Service and displays the date on which FAQ ratings were last updated.

- Select the Run the FAQ Rating Service check box to run the FAQ calculation service using the settings on this window.
- Clear the Run the FAQ Rating Service check box to turn off the FAQ calculation service.

Schedule

Defines the frequency at which the product updates FAQ ratings. This field contains the following components:

Perform the FAQ Calculation Every...

Specifies the time that elapses before the product updates the FAQ rating for documents.

Default: 1 day

From...

Specifies the time of day that the product should begin recalculating FAQ ratings.

Default: 00:00 (12:00 A.M.)

To...

Specifies the time of day that the product should stop recalculating FAQ ratings, regardless of whether the calculation is finished.

Default: 07:00 (7:00 A.M.)

Note: This setting initially takes effect the day after product installation. For example, if you install the product on April 19, 2008, the FAQ server runs for the first time on April 20, 2008.

Aging

Defines the number of times a document FAQ rating is recalculated before it reaches 0. Based on the specified value, the document FAQ rating decreases and eventually becomes 0, at which point it appears at the bottom of the document list (when the list is sorted by FAQ Rating).

Default: 180

For example, if the Aging value is 180 for a document with a rating of 4 (very helpful), the document's FAQ rating is 0 (zero) when the product has recalculated the FAQ rating 180 times.

Note: By default, the FAQ bubble-up calculation requires bu_trans data for the last 180 days, where 180 is the aging factor. So, if you change the aging factor for FAQ to more than 365 days, you should extend the archive rules for the bu_trans table accordingly.

Days New

Specifies the number of days that a newly created or imported document displays in the New Documents folder on the Knowledge tab.

Default: 5 days

Rating

Specifies the default rating (Not Helpful at All, Somewhat Helpful, or Very Helpful) for documents that users have opened but not rated.

Default: Somewhat Helpful

Click Save.

The FAQ settings are defined.

Define Solution Survey Settings

Use the Survey Settings page to configure how a knowledge document displays when accessed to solve a problem or answer a question.

To define solution survey settings

1. Select Knowledge, Solution Survey, Survey Settings from the Administration tab.

The Survey Settings page appears.

2. Complete the following fields as appropriate. Click Save.

The solution survey settings are defined.

Knowledge Management System Settings

You can set default information to display on the Knowledge tab (at logon) the format that categories display in the Knowledge Categories pane on the Administration tab, and the number of documents to list in the Top Solutions list on the Knowledge Management home page.

More information:

[Configure General Settings](#) (see page 864)

Configure General Settings

You can set the default information to display on the Knowledge tab at logon, the format in which categories display in the Knowledge Categories pane on the Administration tab, and the number of documents to list in the Top Solutions list on the Knowledge Management home page.

To configure general settings

1. Select Administration, Knowledge, System, General Settings from the left pane of the Administration tab.

The General Settings pane appears.

2. Complete the following settings as appropriate:

Search Tool Opening Screen

Specifies the information displayed by default on the Knowledge tab. You can select one of the following options:

- **Open with FAQ/Search**—Displays the Category, Knowledge Search, and Knowledge Document List panes.
- **Open with Knowledge Tree Document ID**—Displays the knowledge tree with the document ID specified in the field provided. You can return to the knowledge tree document itself, and then to the Category, Knowledge Search, and Knowledge Document List panes.

Default: Open with FAQ/Search

Category Viewing

Specifies the format in which document categories display in the Knowledge Categories pane on the Administration tab. You can select one of the following options:

- Display categories in tree view—Presents categories in a hierarchical tree structure in the Knowledge Categories pane. Categories expand to reveal associated subcategories. In this manner, you can view all the categories in the tree simultaneously.
- Display categories in list view—Presents categories in a list format in the Knowledge Categories pane. When you select a category, its subcategories display in a list. You can view only the current level of categories or subcategories at one time. Use the Up One Level link to return to the previous category level.

Note: If you have more than 250 categories under the top category, or under any category, use the Display Categories in List View option and not the tree view.

Top Solutions

Specifies the number of documents to list in the Top Solutions list on the CA SDM home page.

Default: 10

Include global data

Displays data from all tenants in Top Solutions.

Default: Enabled

Document Indexing Notifications

Sets a user to receive email notifications about status or when errors occur with document indexing. The user must have an email address in the ca_contacts table to receive these email notifications. Use the Notification Page for the assignee contact record for setting notification methods.

Important! It is critical that you set an Assignee to receive document indexing notifications in the Document Indexing Notifications section. An email address must be defined for this assignee on the Contact Notification page to enable email notifications.

Click OK.

General settings are configured.

Chapter 21: Administering Support Automation

This section contains the following topics:

[Automating Support in Your Environment](#) (see page 867)
[Support Automation Analyst Administration](#) (see page 870)
[Support Automation User Administration](#) (see page 875)
[Support Automation Activity Notification Administration](#) (see page 879)
[Support Automation Page Adaptations](#) (see page 880)
[Support Automation System Properties](#) (see page 882)
[Support Automation Queue Administration](#) (see page 883)
[Ticket Template Management](#) (see page 885)
[Administration Settings](#) (see page 885)
[How to Customize Support Automation Tools](#) (see page 887)
[Session Log Administration](#) (see page 893)
[Support Automation Reports](#) (see page 894)
[Resolve Tickets Using Live Assistance](#) (see page 894)

Automating Support in Your Environment

You can implement a support strategy using a combination of processes and tools. CA SDM provides the tools to administer live assistance, develop automated tasks, and deliver them through the various channels of support.

You use the associated processes to create and maintain an environment that does the following:

- Reduces the average support call duration
- Reduces overall support costs
- Increases resolution rates
- Provides more end-user satisfaction

Live Assistance

Live Assistance provides end-user support through the use of tools that enhance remote interaction between analysts and end users. You can use automated, predefined responses to communicate with the end user. You gather detailed information about an end-user computer and act to provide support.

You provide live assistance using the following interfaces:

Support Automation Analyst Interface

Lets analysts interact with end users and provide support during assistance sessions.

End-User Client

Lets the end user chat with the analyst, while the analyst provides support to their computer.

Support Automation Connectivity

The analyst and end user never communicate directly with each other. You do not require direct peer-to-peer connectivity between the two users. All data transfer is routed through the server, verifying that you can communicate even when the end-user computers are behind restrictive firewalls.

You can connect to end-user computers using the following connections:

Socket

Using a socket connection is the best way for you to connect. Socket connections are the fastest type of connection, with the least overhead, minimal latency, and the most efficient type of connection.

HTTP (or HTTPS)

Using an HTTP connection can result in better connections than a direct socket connection because corporate firewalls can block direct socket connections. HTTP connections generate a significant amount of network traffic overhead compared to direct socket connections. Due to HTTP overhead and processing on the server, the number of simultaneous sessions is much lower when most the connections to the server are over HTTP.

Proxy

Socket Proxy is a mode of operation for the Support Automation server to off-load some of the CPU-intensive operations. For example, encryption/decryption from the main server, and to provide a server component that can go into the DMZ (or similar zone) within the logical network topology.

Typically, you attempt to connect through the direct socket connection first, and, then connect through HTTP if the direct socket fails. However, you can specify custom connection settings on the client computer to alter this sequence.

Note: For details about configuring communication settings, see the *Online Help*.

End-User Client

The end user client connects end users to analysts in live assistance sessions. End users chat with analysts in WebChat, but if you must use the tools in the Support Automation Analyst Interface, CA SDM launches the client on the end-user computer. When the client launches, instructions appear for the end user specific to their web browser.

Server Load

In large deployments, high server load can degrade the application performance. For this reason, you can off-load some of the processing to one or more Socket Proxy servers as follows:

- Offload encryption and decryption of the incoming and outgoing data for all analysts or clients (connecting either through Direct Socket or through HTTP).
- Offload the processing of HTTP traffic from and to those clients connecting through HTTP to the Socket Proxy.

DMZ Server Component

In some network environments, if you allow direct socket access to the application servers that are running the Support Automation server application, it can be considered a security risk. In these environments, you can use Socket Proxy within the DMZ to provide a middleman component between the internet users and the application servers. Using Socket Proxy in this scenario offloads some of the processing from the primary server.

How Socket Proxy Works

The Socket Proxy works as follows:

1. On the configured external port, the Socket Proxy listens for incoming connections from analysts or end users.
2. The Socket Proxy establishes a peer connection to the Primary Server on the configured internal port, for every connection. These two connections are named the end user connection and the server connection, respectively.

3. The end user connections are encrypted and the Socket Proxy encrypts/decrypts any data coming in or going out over the end user connection.

Note: The server connections are not encrypted.

4. For each incoming data-packet, the protocol structure is verified and a checksum value validated before the data is passed on to the primary server through the server connection.
5. The primary server off-loads the encryption and decryption processing.
6. The Socket Proxy closes the matching peer connection once the end user or server connection closes.

Support Automation Analyst Administration

Support Automation analysts monitor and manage multiple end-user requests in live assistance sessions in your environment. Analysts use Support Automation tools to interact with end users and provide live assistance.

Analysts access the interface from a CA SDM ticket, such as an incident, or the Support Automation tab. You can manage access levels to set permissions for tools that the analysts can use. You can enable and disable Support Automation tools for specific tenants. If a tool is disabled for a tenant, analysts cannot use that tool in assistance sessions.

Important! The Support Automation Analyst Interface only runs on Windows. For more information about supported operating systems, see the *Release Notes*.

More information:

[How Analysts Launch Live Assistance](#) (see page 871)

[How to Configure Live Assistance for Analysts](#) (see page 872)

[How End Users Join Assistance Sessions](#) (see page 873)

[How Analysts Automate Support for End Users](#) (see page 874)

[How Analysts Provide Live Assistance](#) (see page 874)

How Analysts Launch Live Assistance

Analysts launch Live Assistance as follows:

1. The analyst does *any* of the following:
 - Logs into CA SDM and selects the Support Automation tab.
 - Opens a CA SDM ticket and selects the Support Automation tab.
2. CA SDM prompts the analyst to install the 32-bit Java Runtime Environment (JRE) version 1.6 or later, if it is not already installed. The Support Automation Analyst Interface does not support the 64-bit JRE.

Note: The Safari browser requires the 32-bit JRE 1.6.0_30 or later.

CA SDM launches the Support Automation Analyst Interface.

Note: The sa_login_session table creates a record every time an analyst launches the Support Automation Analyst Interface and when an end user launches the Web Client. For more information about the sa_login_session table, see the *CA SDM Technical Reference Guide*.

Configure Java Connection Options

You can configure Java connection options to resolve an issue where the Support Automation Analyst Interface cannot connect to the Support Automation server. This issue occurs when the default browser configuration of the analyst is not set up properly for your environment and fails to launch Live Assistance. You can edit the analyst connection settings from a browser or the Java Control Panel.

To configure connection options in your browser

1. Open your web browser, such as Internet Explorer.
The browser appears.
2. Click Tools, Internet Options.
The Internet Options dialog appears.
3. Configure the appropriate connection settings, such as a direct connection or proxy server.
4. Click OK.

The connection opens are configured.

To configure connection options in the Java Control Panel

1. Open the Java Control Panel with the following command:

```
javaws -viewer
```

The Java Control Panel appears.

2. On the General tab, click Network Settings.

The Network Settings dialog appears.

3. Configure the appropriate settings, such as a direct connection or proxy server.

Click OK.

The connection opens are configured.

How to Configure Live Assistance for Analysts

You configure the Support Automation Analyst Interface by establishing role and security permissions for analysts in your live assistance environment as follows:

1. Create or modify the appropriate [roles](#) (see page 876) and [Support Automation Access Levels](#) (see page 878) for analysts in your live assistance environment.

You update roles and Access Levels to give analysts permission to specific tools they can use in assistance sessions.

2. Establish and manage [queues](#) (see page 883) for your live assistance environment.

You create queues to route incoming live assistance requests appropriately.

3. Establish and manage [activity notifications](#) (see page 879) for your live assistance environment.

You create email notifications to alert analysts when they get an assistance session request in their queue.

4. Establish and manage [chat presets](#) (see page 893) for your live assistance environment.

You create chat presets that analysts use to send preconfigured responses to commonly asked questions and common situations.

5. Establish and manage automated tasks for your live assistance environment.

You create automated tasks to execute specific actions on the end-user computer.

Note: You can only create scripts and upload them to the server through the Automated Task Editor IDE.

How End Users Join Assistance Sessions

An end user requests assistance sessions from their CA SDM home page or by directly contacting the help desk, such as through telephone or email. The Support Automation analyst invites the end user to a session from the CA SDM ticket.

1. The end user does *any* of the following:
 - Requests Live Chat from the CA SDM Home Page.

The Live Chat Launch page appears asking the end user for the incident area and description. The end user clicks Continue, the session opens, and the end user is placed in the appropriate queue.
 - Clicks a link from an email notification and logs into the assistance session using their credentials and a session join code provided by the analyst.

If the end user joins from the email notification, they bypass the queue.
 - Clicks Join Analyst Now and provides the session join code to bypass the queue.

Note: When the Support Automation end-user client is launched, an executable is downloaded to launch the program. The end user starts it manually, however for security reasons there is a limited time to launch the executable. After the time expires, an error message appears on the end-user computer when they try to start the launcher executable.
2. The analyst provides live assistance to the end user through chat. The end user uses their web browser to chat with the analyst, or they can launch the agent executable.

Note: If the analyst cannot resolve the session using WebChat, they can invite the end user from the Analyst Session window to use the full tools available in the Support Automation Analyst Interface. The end user must accept the request to launch the client on their computer and for the analyst to use these tools.

How Analysts Automate Support for End Users

Analysts use Live Assistance tools to perform the following tasks on end-user computers:

- Host live chat sessions
- View file systems
 - Create, modify, rename, or delete files and directories
 - Copy and transfer files and folders to the end-user computer
- View system registries
 - Create, edit, or delete registry records
 - Export or import registry values from the end-user registry
- Capture end-user screenshots when connection quality is not sufficient for Remote Control assistance
- View the end-user computer desktop
- Remotely control the end-user computer
- Launch a program on the end-user computer
- Restart or shut down the end-user computer
- Run automated tasks

Note: For detailed information about how analysts use the Live Assistance tools, see the *Online Help*.

More information:

[Support Automation Access Level Administration](#) (see page 878)

How Analysts Provide Live Assistance

Analysts provide Live Assistance to end users by using the Support Automation Analyst Interface. Analysts handle end users routed to their queues, manage assistance sessions, and join sessions where they have permission.

1. The end user requests an assistance session from the CA SDM home page, or a ticket, such as an incident, request, or issue.

2. The end user joins a queue.

Note: If the end user joins the session using the link from the analyst email invitation, they bypass the queue.

You configure CA SDM request areas and issue categories for queue routing.

3. The analyst selects the end user from the queue.
4. The assistance session begins and the analyst provides live assistance.

Note: If the analyst launches the Support Automation Analyst Interface from the Support Automation tab, no CA SDM ticket is associated with the assistance session.
5. The analyst creates the ticket when closing the assistance session and define the status as *SA-Open* or *SA-Resolved*, or they transfer the assistance session to another queue.
6. The analyst closes the session and the end user receives an email notification with the Session Log.

Note: For details about how analysts handle queues and manage assistance sessions, see the *Online Help*.

Support Automation User Administration

System administrators and tenant administrators configure CA SDM contacts, role permissions, access levels, and privacy levels to define user permissions. The following lists users that use Support Automation.

System Administrator

Defines system-wide access to add, edit, and modify all Support Automation defaults and functions in the Administration tab. The system administrator sets up tenants and analysts, customizes Support Automation system properties, and performs system password resets.

Tenant Administrator

Defines administrative rights at the tenant level and does not grant access to create or edit other tenants or reset user passwords. The Service Provider tenant determines permissions.

Analyst

Defines rights for the user that provides live assistance to end users in your support environment.

End User

Defines rights for users that can request live assistance from analysts in your support environment, such as employees and customers.

How to Configure Support Automation Role Permissions

You can configure CA SDM roles to have Support Automation permissions. You can set role permissions by configuring Support Automation access levels for analysts and privacy levels for end users in your live assistance environment as follows.

1. Configure the appropriate access levels for analysts in your live assistance environment.

You create analyst access levels to manage analyst permissions in your system, such as enabling and disabling specific Support Automation Analyst Interface tools.

2. Configure the appropriate privacy levels for end users in your live assistance environment.

You create privacy levels to manage end-user access levels in your system.

3. [Assign](#) (see page 879) access levels to roles.

You assign end-user privacy levels and analyst access levels to roles in your environment.

Note: For detailed information about creating and modifying Support Automation access levels for analysts and configuring security levels for end users, see the *Online Help*.

Support Automation Anonymous and Registered Users

The Support Automation server accepts registered or anonymous users, depending on CA SDM permissions. If permitted, the guest user lets anonymous users log in to CA SDM. You can authenticate with the server to gain access to the following:

- Live Assistance
- Self-Service
- Automated Tasks Editor
- End-User Client

How to Configure Support Automation for Guest Users

You can configure your live assistance environment to allow guest users. You configure guest users for specific tenants or make them available to the entire system as follows:

1. Create an Access Type that allows [web authentication](#) (see page 877).
You select the "Open - always allow access" Validation Type.
2. [Assign](#) (see page 878) the Access Type to a contact.
You can create a guest contact for different tenants in your environment.
3. Notify the appropriate users in your environment about the anonymous login, using the following URL format:
`http://hostname?USRNAME=UserName`

Create an Access Type for Guest Users

You can create an access type that lets guest users log in to CA SDM without authenticating. If you are the service provider, you can create a guest access type for each tenant in your environment.

To create a guest access type

1. Click Security and Role Management, Access Types.
The Access Type List appears.
2. Click Create New.
The Create New Access Type page appears.
3. Complete the fields as appropriate, such as the symbol and description.
4. On the Web Authentication tab, select *Open-always allow access* from the Validation Type drop-down list.
5. Save the access type.
The access type is created.
6. (Optional) Right-click on the Access Type List page to refresh the list page.
The guest access type appears in the list.

Assign the Guest Access Type to a Contact

You can assign the guest access type to a contact after you create an access type that uses web authentication so that the contact can log in to CA SDM. You can notify the appropriate users in your environment about the guest contact after you assign the access type.

To assign the guest access type

1. Click Security and Role Management, Contacts.

The Contact Search page appears.

2. Click Create New.

The Create New Contact page appears.

Note: You can also modify a contact.

3. Select guest from the Contact Type drop-down list.
4. (Optional) Associate the contact with a tenant. You can also make the contact public.
5. Save the contact.

The contact is created and the guest access type is associated with the contact.

Support Automation Access Level Administration

You manage Support Automation access levels and assign them to CA SDM roles in your support environment. Support environments vary in size and structure, so your implementation of access levels can vary.

In a small support environment, there can be only one or two analysts categorized within a single access level, such as Analyst. In a larger support environment, the tenant administrator can set up many analyst access levels, each with different access and support privileges.

Important! If you are in a multi-tenancy environment, analysts that do not belong to the service provider only have write access to their own tenant and subtenants. You can give the analyst write access to other tenants and subtenants by updating the function access of the accessed tenant to include non-service provider tenants.

The following access levels are available:

Analyst

Specifies the contact type that provides live assistance to end users in your support environment. Access levels define which queues, automated tasks, and tools are available for the analyst to use.

End User

Specifies the contact type that receives live assistance from analysts, such as employee and customer.

You manage Support Automation access levels from the Administration tab.

Note: For detailed information about creating and modifying Support Automation access levels for analysts and end users, see the *Online Help*.

Assign an Access Level to a Role

You can assign Support Automation access levels to existing CA SDM roles in your environment.

To assign an access level to a role

1. Select Security and Role Management, Role Management, Role List from the Administration tab.

The Role List page appears.

2. Click the role that you want to assign the access level, such as Administrator.

The Role Detail page appears.

3. Click Edit.

The Update Role page appears.

4. On the Authorization tab, select the access level you created from the SA Access drop-down list, and click Save.

The Role Detail page appears. Verify the Support Automation access is assigned to the role.

Support Automation Activity Notification Administration

You can use activity notifications to manage Support Automation activities. You can customize how end users, analysts, and administrators track and receive notifications when an activity occurs, such as the end of an assistance session.

Configure any of the following default notifications, as appropriate to your environment:

Queue Entry Notification

Notifies the analyst when an end user joins an assistance session queue.

Notifies the analyst when the assistance session is transferred to another queue.

Analyst Notification

Notifies the analyst when end-user queue wait time expires. The event of expiration is recognized with a CA SDM Event conditional macro.

Invite End User to Assistance Session - Incident

Notifies the end user when the analyst invites them to an assistance session from an incident or request.

Invite End User to Assistance Session - Issue

Notifies the end user when the analyst invites them to an assistance session from an issue.

Session Ended Notification

Notifies the system when the assistance session ends.

Important! If you want to use Support Automation functionality with an external system such as Star, the System_SA_User contact is set to the *Session Ended Notification* rule by default.

Support Automation Page Adaptations

You can configure the appearance of Support Automation pages for your end users according to your business needs. The following adaptations let you control what the user sees before joining the assistance session and after the session completes:

- Customize the header and footer of all pages the end user sees.
You can change the HTML code of the header and footer, and you can also modify the location of the CSS file that contains the style definitions.
- Customize the localization needs of your environment.
You can set the default locale for all analysts and end users.
- Customize the page layouts that end users see when they log in and wait to join assistance sessions.
You can customize the pages that direct your end users to suit the needs of your business.

Note: For details about configuring Support Automation pages for analysts and end users, see the *Online Help*.

Branding Administration

You can customize the header and footer of end-user facing pages, such as self-serve. You can change HTML code of the header and footer and also modify the location of the CSS file that contains the style definitions.

You can view a list of branding records, one for each tenant at maximum. Tenants can create their own branding, or if branding is not defined for the tenant, they use the default system settings. You can also enable localization of branding and view a list of all localized brandings for enabled localizations.

Important! Branding customizations from CA Support Automation r6.0 SR1 eFix5 do not migrate to CA SDM Release 12.7 automatically. We recommend that you review the customized branding to verify that it corresponds to the CA SDM branding. If necessary, copy and paste the Header, Footer, and CSS URL data of each division to the corresponding tenant (or public) in CA SDM to migrate the branding data.

Localization Administration

Localization lets you support multiple languages simultaneously on the same install, by translating elements of the application to a particular language. These elements can include system messages, icons, and content. Information is presented in the best form for the end user, no matter their native language.

Localized versions of Support Automation meet the needs of your global environment. Each tenant can handle multiple languages, and the localizations are supported across multiple tenants. Each tenant shares the default system settings with its localization. You can configure what language the end user and analyst use before launching Live Assistance from the list of enabled languages. The analyst can use a different language than the end user within an assistance session.

You can edit localized text for disclaimers. You cannot create or remove localizations, but you can enable or disable localizations.

Note: The administrator interface is available in the default server localization only.

Page Layout Configuration

When end users log in to Live Assistance from CA SDM, they view a default launch page. Likewise, when they are initially placed on-hold in a queue, they again view a default page.

If no special settings are specified for tenants, they default public settings. You can configure the following public settings for queues that do not have their own specialized settings:

- Wait Page
- End-user Post-Launch Page
- In Session
- Post-Logout Page
- Exit Survey Page

Each page has its own detail page which is comprised of a text field for entering the URL and a check box for marking this page as external.

Support Automation System Properties

You can customize many of the ways in which Support Automation handles activities to differ from the default installation. You can use the default system option settings to customize Support Automation behavior. For example, you can customize some of the following Support Automation options:

- Enable or disable the link for live chat from CA SDM employees and customers home page.
- Enable or disable the link for Joining a session from CA SDM employees and customers home page.
- Enable or disable the option to create a CA SDM incident when the end user disconnects while waiting to be served on a Support Automation queue.

System properties are tenant optional. If a tenant has not defined its properties, Support Automation uses the public (shared) settings. The product installation creates the default public properties.

Note: For more information about configuring Support Automation system properties, see the *Online Help*.

Support Automation Queue Administration

You use queues to route assistance session requests to the most appropriate analyst. The end user can select a category, or enter a description of their computer problem, and their ticket (such as an incident) routes to the appropriate queue.

After the initial product installation, the default queue is named Support. You can set up several queues to facilitate the sorting and tracking of different support requests, according to your business needs. You can assign only one default queue per tenant. If you do not assign a default to a tenant queue, or if the default tenant queue is unavailable, the system uses the public default queue. You set the working hours per queue.

The system automatically determines where to place the end user by mapping queues to incident areas. If an area maps to a queue, the end user selects a category and is routed to the appropriate queue. Search capabilities are applied to the description of an incident or issue category to identify relevant queues, and the end user is routed to the best matched queue only.

Note: For more detailed information about customizing and mapping Support Automation queues, see the *Online Help*.

Queue Management

You configure queues to help end users receive the appropriate live assistance from analysts. You manage queues to improve how end users are routed to assistance sessions as follows:

- Customize queues for analysts and tenants in your live assistance environment.
You can activate or deactivate queues and specify tenant and analyst permissions.
- Assign a default queue.

You can route end users to the default queue when they enter queries that do not match the queues configured in your environment. You can also customize queues for tenants in your environment.

Note: If the default tenant queue is missing or unavailable, the public queue is used.

- Assign hours of operation for your queues.
You can manage queues based on the availability of users in your support environment, such as enabling Support Automation services during business hours.

Important! You can assign workshifts to both your Support Automation hours and individual live assistance queues. Different workshifts assigned to Support Automation hours and individual queues can cause conflicts for analysts and end users in your support environment.

How to Manage Queue Summaries

You can compose the fields that an analyst views in the Queue Summaries page. You can select from the following set of fields for the queue summary, but you cannot create fields:

- Queue
- Question
- Wait Time
- E-mail
- Company
- IP Address
- Owning Analyst
- Category
- Language
- Tenant

More information:

[Queue Management](#) (see page 883)

[How to Manage the Queue Hours](#) (see page 884)

How to Manage the Queue Hours

You can enable Support Automation for each queue for specific hours of the day, to accommodate the working hours of analysts as follows:

- Create a separate schedule for each queue and for all automated support services.

Note: These settings do not limit self-service functions.

- Define support hours for the server by a global open or close status. An entry for each hour of the week indicates a difference from the global status of the server.
- The server uses the first entry for each hour based on the rules you establish.

This action effectively merges the support hour definitions from the parent tenant (or public) settings. This action can have counter-intuitive results if a mix of 'default-closed' and 'default-open' is used in the hierarchy.

Ticket Template Management

You can specify which ticket templates are available for the Analyst User Interface. You can select the ticket templates for Incident/Requests and Issue ticket types.

You can define if the template is default or active. When you create a ticket template, you can select from existing CA SDM templates. The default template must be active.

You can have only one ticket template as a default per tenant.

Note: For more detailed information about customizing ticket templates, see the *Online Help*.

Administration Settings

You can define and configure the following Support Automation settings:

- Message Routing Servers
- End-user Permissions (Privacy Levels)
- Support Automation hours

How to Configure Support Automation Settings

You can configure the Support Automation settings for your support environment according to your business needs. The settings let you control administrative functionality, such as message routing servers, privacy levels, and Support Automation hours.

- Configure message routing servers to improve performance during assistance session.

You can connect end users to the preferred local server of the analyst, or if the connection is unsuccessful, connect the session to the main default server.

- Define the privacy levels for end users in your environment.

You can enable specific Support Automation Analyst Interface tools, based on privacy levels used in your environment.

- Set your Support Automation hours to specific hours of operation or as open at all times.

You can direct live assistance requests when the support desk is closed, such as opening a web page informing the end user of the support desk hours and additional support options.

Note: For more information about configuring Support Automation settings, see the *Online Help*.

Message Routing Servers

You can use Message Routing Servers (MRS) to manage multiple Remote Control servers, based on the geographical location of the local server. Using MRS helps improve performance during assistance sessions. When you enable MRS, the Support Automation Analyst Interface and end-user client attempt to connect to the analysts preferred (local) server for sharing. If the connection is unsuccessful, the sharing session falls back to the main default server. The Live Log records which MRS you use during the assistance session.

You can create, update, remove, enable, or disable a message routing server object.

Note: For more information about configuring message routing servers, see the *Online Help*.

Support Automation Privacy Levels

Privacy levels are used to set what actions are allowed to be performed on different end users to protect user privacy. Privacy levels are associated with CA SDM roles. Three default permissions exist: High, Medium, and Low, but more can be defined if necessary. You can add, update, and delete the privacy levels that are available to the end user.

You can set the privacy level name and its description. You can define which functions for specified tools (File explorer, Remote registry, Run program, and so on) are enabled for this privacy level.

Note: For more information about configuring Support Automation privacy levels, see the *Online Help*.

Support Automation Hours

You can set Support Automation to operate at specific hours or to operate at all times. You manage these hours of operation based on the needs of end users in your support environment. End users cannot access Support Automation functionality when it is closed. You can make quick changes to several workshifts in a single step.

Important! You can assign workshifts to both your Support Automation hours and individual live assistance queues. Different workshifts that are assigned to Support Automation hours and individual queues can cause conflicts for analysts and end users in your support environment.

Note: For more information about configuring Support Automation hours, see the *Online Help*.

How to Customize Support Automation Tools

You can configure the Support Automation tools for your support environment according to your business needs. The settings let you control administrative functionality, such as automated tasks, chat presets, default credentials, and disclaimers.

- Customize the automated tasks list and classifications that Support Automation analysts use to provide support for end users.
- Configure chat presets for common responses to commonly asked questions and situations.
- Configure default credentials that let you gain access to an end-user computer.
- Customize disclaimers that end users see before launching self-service tasks.

Note: For more information about configuring Support Automation tools, see the *Online Help*.

Automated Tasks

An *automated task* is a collection of steps that define an automated process that the analyst or end user follows. Automated task steps are routines written in VBScript or JavaScript that carry out specific actions on the analyst or end-user computer. You can create new automated tasks and task steps using the Automated Task Editor. Common routines include gathering telemetry information, diagnosing problems, and implementing resolutions.

When you execute an automated task, the log updates. This log is both linked and accessed from the assistance session log. Entries in the automated task log consist of a number of timestamped text entries. The entries are created by calling `Functions.LogMessage()` or `WScript.Echo()`.

More information:

[How to Implement Automated Tasks](#) (see page 889)

How to Configure Automated Tasks

You install and configure the Automated Tasks Editor to manage automated tasks that Support Automation analysts use to provide support for end users. The end user can launch an automated task from a knowledge document and the self-service interface, or an analyst executes an automated task during an assistance session. Automated tasks provide analysts with detailed information about an end-user computer. You create self-service automated tasks that interact with the end user and process their input. These tasks can change the file system, registry, download install software, and so on. You configure automated tasks as follows:

1. Install the Automated Tasks Editor.

You launch the installer from the following location on the installation media:

`casd.nt\SAScriptWriter`

Note: You can also copy the installer and deploy it to the appropriate users in your support environment.

The Automated Task Editor is installed.

2. Open the Automated Tasks Editor.

The Automated Tasks Editor installation creates a shortcut on your desktop.

3. Set the following connection parameters:

- a. Click Tools, Server.

The Server Configuration dialog appears.

- b. Enter your hostname and port.

Default Port: 8070

- c. Enter the user name and password of a user with read/write access to the Automated Task Editor, such as a Support Automation Analyst.

- d. Click Test.

- e. Click OK.

4. Create automated tasks and upload them to your server.

You can upload public tasks or assign them to specific tenants and subtenants.

Important! Only roles from the Service Provider tenant with the Update Public flag enabled can upload tasks and libraries to the server. All task library content and static content are stored as public data.

How to Implement Automated Tasks

You can use automated tasks to assist end users in your support environment. Automated tasks complete specific actions on the end-user computer, without the need for the analyst or end user to complete the process. These scripts can help you gather telemetry information, diagnose computer problems, and implement resolutions.

To implement automated tasks, do the following:

1. Identify opportunities for support automation.

You identify common problems that end users encounter, and decide that you can automate some solutions to reduce your support costs.

2. Research automation of the potential solutions.

You research resolutions to common problems and gather data about diagnostic processes you plan to use.

3. Design tasks to automate end-user support.

You design the end-user experience you want for each task with the Automated Task Editor.

4. Implement and test the automated tasks.

You test the automated tasks to verify that they resolve common problems encountered in your support environment and help reduce support costs.

5. Deploy and monitor the automated tasks.

You deploy the automated tasks to end users in your environment by allowing analysts to use them in assistance sessions, or you can attach scripts to knowledge documents.

Important! If you are in a multi-tenancy environment and want to allow analysts to upload task and library content, their role must have the Update Public option enabled.

Note: CA Technologies can provide training in the creation of automated tasks and components, such as automated task step templates and libraries, which you can use in your environment. For more information about developing automated tasks, contact *CA Technologies Services*.

Automated Tasks Administration

You can create automated tasks and associate them with the server. You need read/write access to all the automated task-related tables to use the Automated Tasks Editor. You can perform user management functions with the application, such as assigning automated tasks to roles and tenants.

Note: If you are a Service Provider analyst and have access to multiple tenants, you can select the tenant context of any task update operation against the server. You can also assign an automated task as public.

Automated Task Deployment

After you author and test an automated task in the Automated Task Editor, you can deploy it to a classification on the CA SDM server for use in assistance sessions or self-service. Some settings that are not part of the automated task definition are set when you deploy the task.

You can upload or download automated tasks directly to the server from the Automated Task Editor. You select the appropriate tenant when creating automated task classifications. If you upload scripts to the server, select the classification or tenant where you have access, or set the script as public. You can also import automated tasks from XSDF files and export them to XSDF files.

Important! If you are in a multi-tenancy environment and want to allow analysts to upload task and library content, their role must have the Update Public option enabled.

Upload an Automated Task

You can upload automated tasks that you created in the application. When you select a task, all dependent content automatically uploads, such as libraries and static content.

To upload an automated task

1. Open the Automated Task Editor.
The Support Automation Task Editor appears.
2. Select the automated task you want to upload.
3. Select the classification where you want to upload the task.

Note: If you are a privileged user in a multi-tenancy environment, select the appropriate tenant when uploading the automated task, or make the task public.

4. On the toolbar, select the Upload Automated Task icon.

The Support Automation Task Editor uploads the task to the server.

Edit an Automated Task

You can download automated tasks from the server and edit them in the Automated Task Editor. Any content that is newer in the version than the existing content on the server is imported into the database and made available to administrators of that tenant.

To edit an automated task

1. Open the Automated Task Editor.
The Support Automation Task Editor appears.
2. On the toolbar, select the Upload Automated Task icon.
The Open Automated Task from Server pane appears.
3. Select the automated task that you want to download.
Note: If you are a privileged user in a multi-tenancy environment, you can edit public or tenant-specific automated tasks.
4. Click Open Task.
The Support Automation Task Editor downloads the task to the client and opens it in the application.
The download creates a text file on the computer of the task author that contains the dependent content.

Automated Task Credentials

You can execute an automated task that you need administrative privileges to run, such as performing a software installation, by doing the following:

- Define Default Credentials for the tenant and configure the automated task to use the default credentials in automated task administration.
- Define Default Credentials for the tenant and select that you use the Execute As dialog in the Assistance Session window.
- Define automated task credentials for the task and specify that you use the Execute as dialog in the Assistance Session window.
- Specify the credentials in the Execute As dialog in the Assistance Session window.

Role Assignment

You assign the appropriate roles (for example, Analyst, Administrator) to use the automated task. You assign individual automated tasks to selected roles to limit the automated tasks to only those analysts in the assigned role. You can manage the different skill sets or levels of experience within the support organization.

You can manage the roles at a task level, selecting certain commonly executed tasks, such as diagnostics to run automatically when the end user enters a session.

Assign an Access Level to a Role

You can assign Support Automation access levels to existing CA SDM roles in your environment.

To assign an access level to a role

1. Select Security and Role Management, Role Management, Role List from the Administration tab.

The Role List page appears.

2. Click the role that you want to assign the access level, such as Administrator.

The Role Detail page appears.

3. Click Edit.

The Update Role page appears.

4. On the Authorization tab, select the access level you created from the SA Access drop-down list, and click Save.

The Role Detail page appears. Verify the Support Automation access is assigned to the role.

Chat Preset Administration

You can create common responses to commonly asked questions and situations. Instead of repeatedly typing the same information, you can save a response and use it in other chat sessions.

You can send the presets to end users at the beginning of each session automatically, such as a greeting. You can also automatically populate the presets with information specific to the current session, such as the analyst name.

You can use the following types of presets in an assistance session:

Chat Preset

Identifies a commonly used text response to an end-user question.

URL Preset

Identifies a commonly used URL that the end user can access.

You can localize each chat preset. The chat preset is synchronized with the end-user localization so that the end user receives correct localized presets.

How to Manage Chat Presets

You can use predefined responses to commonly asked questions and situations.

1. Understand the typical assistance sessions in your organization.
2. Create chat and URL presets based on the needs of your organization
3. Specify whether to make the presets available as tenant specific or public.
4. (Optional) Duplicate the preset tree structure to manage localized environments.
5. Send chat and URL presets during an assistance session.

Note: For more information about configuring Chat Presets, see the *Online Help*.

Default Credentials

You can run automated tasks on the end-user computer even if the end user would not typically have system access rights to perform such activities. If the current end user does not have administrative rights to view system information about their computer, you can run a restricted automated task using default credentials to gain access.

Note: For more information about configuring default credentials, see the *Online Help*.

Disclaimers

When end users launch self-service tasks, they are presented disclaimer text that they must agree to before they can continue.

You can create, update, and delete the disclaimer business objects.

Note: For more information about configuring Support Automation disclaimers, see the *Online Help*.

Session Log Administration

The session log lets you view all actions that the analyst performed during an assistance session, such as the tools used and chat details (excluding whispers). You can print or Email the session log to the end user. End users can also view and save the log, but they cannot modify the log.

View the Session Log

All actions you perform during the assistance session, such as chat dialog (excluding Whisper conversations), automated task results, and using a specific Support Automation Analyst Interface tool update the Session Log.

To view the session log

1. Open the active session view.
The Active Session page appears.
2. Select the Session Log from the toolbar or Sessions menu..
The Session Log page appears.
3. Click Refresh Now.
The page refreshes.
4. (Optional) Select the Auto-Refreshing check box.
5. (Optional) Click Save Log to Disk.
The save dialog appears. You can save the session log locally as an HTML file.

Support Automation Reports

CA Business Intelligence installs a set of predefined Support Automation reports. CA SDM automatically deploys these reports to the BusinessObjects server during the installation.

The Support Automation Administrator and Support Automation Analyst roles can use BusinessObjects InfoView to view detailed and summary information about the following:


- Analyst login metrics
- Assistance Sessions metrics
- Queue entries metrics
- Automated task execution metrics
- Tool usage per ticket category
- Real-time data about end users placed in support queues, active assistance sessions, and active (logged-in) analysts.

Note: For more information about using BusinessObjects InfoView, see the CA Business Intelligence information in the *Online Help*.

Resolve Tickets Using Live Assistance

An *assistance session* lets a Service Desk Analyst provide Live Assistance to End Users in CA SDM to resolve tickets. You view details in a CA SDM ticket about an End User that has a computer problem. You chat with the user and can invite the user to an assistance session. Use the Support Automation functionality in CA SDM to resolve tickets using Live Assistance. For example, the End User creates a ticket about a network connection problem with a software application.

The following diagram explains how a Service Desk Analyst resolves a ticket using Live Assistance:



Perform these steps to provide Live Assistance to resolve a CA SDM ticket:

1. [Receive a Ticket Request](#) (see page 896).
2. [Invite the End User to an Assistance Session](#) (see page 896).
3. [Resolve the Ticket with a Chat Session](#) (see page 897) or [Provide Live Assistance](#) (see page 897).
4. [End the Assistance Session and Close the Ticket](#) (see page 898).

Receive a Ticket Request

A user submits a ticket that describes an issue with their computer. View the ticket details from an email alert or the CA SDM Scoreboard. For example, the user cannot use a software application that they configured to synchronize with their network.

Follow these steps:

1. Log in to CA SDM and select Scoreboard, My Queue.
2. Open a ticket.
3. Review the ticket description.
4. (Optional) Add a comment to the ticket that asks the user to provide more information about the software application they configured incorrectly.

Invite the End User to an Assistance Session

Invite the user to an assistance session to resolve the ticket. For example, you require information about network settings of the end-user computer.

Follow these steps:

1. On the Support Automation tab of the ticket Detail page, click Invite End User.
2. Enter a greeting for the End User.
3. Click Launch.

The Support Automation Analyst Interface appears and you wait for the user to join the assistance session.

Resolve the Ticket with a Chat Session

Initiate a chat session and the user receives an email notification from CA SDM. This email contains a link to the assistance session and the ticket details. The link opens the Support Automation end-user client on their computer. Chat with the user to determine how to resolve their ticket.

Note: If the End User cannot join the assistance session from the URL, email the user to click Join Analyst Now on the Self-Service home page.

Follow these steps:

1. Click the Chat tab on the Support Automation Analyst Interface.
2. Complete any of the following actions:
 - Select a question, statement, or URL from the Chat Presets drop-down list.
 - Enter text in the chat window and click Send.

For example, you ask the End User about the network ports they specified during the software application configuration.
 - Push a specific URL to the End User.

The End User manually opens the link in their browser.

Provide Live Assistance

Use the Support Automation Analyst Interface to provide Live Assistance by performing different actions on the end-user computer. Perform actions such as running diagnostic scripts, browsing the file system, and controlling the end-user computer remotely. For example, a chat with the user determines that you can resolve the software application synchronization issue by using Live Assistance.

Follow these steps:

1. On the Support Automation Analyst Interface, click the appropriate tool tab to resolve the ticket:
 - Automated Tasks—Select an automated task from the left pane and click Execute. For example, execute a script that configures network settings for the software application.
 - File Explorer—Browse the file system of the user. For example, browse the hard drive to locate a specific file in the installation directory of the software application.
 - File Transfer—Transfer a file between computers. For example, transfer a file from your computer to replace a corrupt file in the installation directory.
 - Remote Registry—Browse the end-user registry and modify a registry entry. For example, modify registry values for the software application.

- Remote System Tools—Execute a program on the end-user computer or force the computer to reboot. For example, execute the configuration interface of the software application.
 - Remote Control—Control the end-user computer remotely. For example, control the end-user computer to configure the software application.
 - Screenshot—Capture a screenshot of the end-user computer. For example, connection issues prevent remote control from operating successfully, so you guide the user after viewing screenshots.
2. Select Remote Control and configure the software application to synchronize with the company network.
 3. Verify with the user that the software application synchronizes successfully and you resolved the ticket.

End the Assistance Session and Close the Ticket

After you verify that the ticket has been resolved, update the ticket and close the assistance session.

Follow these steps:

1. Click End in the Support Automation Analyst Interface to close the session.
The user receives an email notification with the session log.
2. (Optional) Click Session Log to view chat logs and Support Automation tool results.
3. Click the ticket number on the Support Automation Analyst Interface. For example, click Incident 40.
The Incident 40 Detail page opens in CA SDM.
4. Click Edit.
5. Change the ticket status to SA-Resolved.
6. Click Save and Close.
The Activity Log for the ticket is saved and the Live Assistance process is complete.

Appendix A: View Field Descriptions

This section contains the following topics:

[View Field Descriptions](#) (see page 899)
[View Act Log](#) (see page 900)
[View Audit Assignee](#) (see page 901)
[View Audit Group](#) (see page 902)
[View Audit Priority](#) (see page 902)
[View Audit Status](#) (see page 903)
[View Change Act Log](#) (see page 903)
[View Change](#) (see page 905)
[View Change to Assets](#) (see page 910)
[View Change to Change Act Log](#) (see page 911)
[View Change to Change WF](#) (see page 912)
[View Change to Properties](#) (see page 914)
[View Contact Full](#) (see page 916)
[View Contact to Environment](#) (see page 919)
[View Group](#) (see page 920)
[View Group to Contact](#) (see page 920)
[View Issue](#) (see page 921)
[View Issue Act Log](#) (see page 926)
[View Issue to Assets](#) (see page 927)
[View Issue to Issue Act Log](#) (see page 928)
[View Change to Request](#) (see page 929)
[View Issue to Issue WF](#) (see page 933)
[View Issue to Properties](#) (see page 935)
[View Request](#) (see page 936)
[View Request to Act Log](#) (see page 942)
[View Request to Properties](#) (see page 942)

View Field Descriptions

You can use the field description information in the basic and advanced views that are supplied with CA SDM.

The following points apply to many of the tables:

- You must turn on audit logging, found in Administration, Options Manager, Audit Log, to see data in the advanced views.
Note: For more information about audit logging, see the *Online Help*.
- pdmtime refers to date/time fields that are in GMT format (the number of elapsed seconds since 1/1/1970).
- The terms change request and change order are used interchangeably.

More information:

[Generating Reports in CA SDM](#) (see page 741)

View_Act_Log

The following is a basic view of the request activity log table. Activity type and the analyst's full name are also listed in the view. The activity log table (act_log) was joined with the activity type table (act_type) and the contact table (ca_contact) to give the actual activity type of each activity log entry, and the analyst who performed the activity. Extracted fields from the joins that might be useful are located at the end of this list.

Field	Remarks
id	act_log.id: The unique identifier for this record in the act_log table.
persid	act_log.persid: The unique identifier for this record in the act_log table, preceeded by the object identifier (alg for act_log) and a colon.
call_req_id	act_log.call_req_id: Pointer to call request persid to which this activity belongs. act_log.call_req_id = call_req.persid.
last_mod_dt	act_log.last_mod_dt: The last modify date/time (pdmtime).
time_spent	act_log.time_spent: The duration of time spent on this activity, stored as the total number of seconds. For example, 80 = 1 minute, 20 seconds.
time_stamp	act_log.time_stamp: User modifiable date/time of activity (pdmtime).
system_time	act_log.system_time: The date/time of record creation (pdmtime).
analyst	act_log.analyst: The uuid pointer to the contact uuid to get the analyst who performed the activity. act_log.analyst = ca_contact.contact_uuid.
description	act_log.description: The text description of this activity, which can be modified by the user.
action_desc	act_log.action_desc: The text description of automated action, which cannot be modified by the user.
type	act_log.type: The text pointer to a record in the activity type table. For example, act_log.type = act_type.code.

Field	Remarks
knowledge_session	act_log.knowledge_session: An identifier for a particular session of a particular user.
knowledge_tool	act_log.knowledge_tool: An indicator of the knowledge tool used for the search, such as NLS_FAQ or EXPERT, etc.
internal	act_log.internal: An integer flag (1 or 0), which indicates if this log entry is intended for all to see or just for internal use.
activity_type	act_type.symActivity: The type derived from act_log.type = act_type.code.
analyst_lastname	View_Contact_Full.last_name: The analyst's last name, derived from act_log.analyst = ca_contact.contact_uuid.
analyst_firstname	View_Contact_Full.first_name: The analyst's first name.
analyst_middlename	View_Contact_Full.middle_name: The analyst's middle name.

View_Audit_Assignee

The following is an advanced view of the audit log where assignee is tracked. This view shows the duration of time between assignee changes for every request and change order. Requests or change orders that are changed from a particular assignee to a null assignee, and then from null assignee back to a particular assignee, do not have the duration of the null assignee listed in this view. This view lists the following fields for both requests and change orders. There may be more than one entry for each audobj_uniqueid (request or change order).

Field	Remarks
audobj_uniqueid	audit_log.audobj_uniqueid: The audit log object unique id representing the chg.id or the call_req.id.
from_val	audit_log.attr_after_val: The 'changed from' assignee value.
to_val	audit_log.attr_after_val: The 'changed to' assignee value.
from_time	audit_log.attr_from_time: The beginning time an assignee was assigned (pdmtime).
to_time	audit_log.attr_from_time: The ending time the same assignee was assigned (pdmtime).

View_Audit_Group

The following is an advanced view of the audit log where group is tracked. This view shows the duration of time between group changes for every request and change order. Requests or change orders that are changed from a particular group to null group, and then from null group back to a particular group, do not have the duration of the null group listed in this view. This view lists the following fields for both requests and change orders. There may be more than one entry for each audobj_uniqueid (request or change order).

Field	Remarks
audobj_uniqueid	audit_log.audobj_uniqueid: The audit log object unique id which represents chg.id or call_req.id.
from_val	audit_log.attr_after_val: The 'changed from' group value.
to_val	audit_log.attr_after_val: The 'changed to' group value.
from_time	audit_log.attr_from_time: The beginning time group was assigned (pdmtime).
to_time	audit_log.attr_from_time: The ending time same group was assigned (pdmtime).

View_Audit_Priority

The following is an advanced view of the audit log where priority is tracked. This view shows the duration of time between priority changes for every request and change order. This view lists the following fields for both requests and change orders. There may be more than one entry for each audobj_uniqueid (request or change order).

Field	Remarks
audobj_uniqueid	audit_log.audobj_uniqueid: The audit object unique id, represents a request call_req.id or change order chg.id.
from_val	audit_log.attr_after_val: The 'changed from' value of priority.
to_val	audit_log.attr_after_val: The 'changed to' value of priority.

Field	Remarks
from_time	audit_log.attr_from_time: The beginning time priority was in a particular state (pdmtime).
to_time	audit_log.attr_from_time: The ending time priority was in the same state (pdmtime).

View_Audit_Status

The following is an advanced view of the audit log where status is tracked. This view shows the duration of time between status changes for every request and change order. This view lists the following fields for both requests and change orders. There may be more than one entry for each audobj_uniqueid (request or change order).

Field	Remarks
audobj_uniqueid	audit_log.audobj_uniqueid: The audit object unique id, which represents a request call_req.id or change order chg.id.
from_val	audit_log.attr_after_val: The 'changed from' value of priority.
to_val	audit_log.attr_after_val: The 'changed to' value of priority.
from_time	audit_log.attr_after_time: The beginning time status was in a particular state (pdmtime).
to_time	audit_log.attr_after_time: The ending time status was in the same state (pdmtime).

View_Change_Act_Log

The following is a basic view of all change order activity logs. This is a view of the change request activity log table (chgalg) joined with the activity type table (act_type) and the contact table (ca_contact) to give more meaningful data, such as the actual activity type description and full name of the analyst who performed the activity.

Field	Remarks
id	chgalg.id: The unique identifier for this record in the chgalg table.

Field	Remarks
persid	chgalg.persid: The unique identifier for this record in the chgalg table, preceded by the object identifier (chgalg for chgalg) and a colon.
change_id	chgalg.change_id: The pointer to the change order id to which this activity belongs. chgalg.change_id = chgalg.id
last_mod_dt	chgalg.last_mod_dt: The last modify date/time (pdmtime).
time_spent	chgalg.time_spent: The duration of time spent on this activity, stored as the total number of seconds. For example, 80 = 1 minute, 20 seconds.
time_stamp	chgalg.time_stamp: The user modifiable date/time of activity (pdmtime).
system_time	chgalg.system_time: The date/time of record creation (pdmtime).
analyst	chgalg.analyst: The uuid pointer to the contact uuid to get the analyst who performed the activity. chgalg.analyst = ca_contact.contact_uuid
description	chgalg.description: The text description of this activity, which can be modified by the user.
action_desc	chgalg.action_desc: The text description of the automated action, which cannot be modified by the user.
type	chgalg.type: The text pointer to a record in the activity type table. chgalg.type = act_type.code
internal	chgalg.internal: The integer flag (1 or 0), which indicates if this log entry is intended for all to see or just for internal use.
knowledge_session	chgalg.knowledge_session: An identifier for a particular session of a particular user.
knowledge_tool	chgalg.knowledge_tool: An indicator of the knowledge tool used for the search, such as NLS_FAQ or EXPERT, and so on.
analyst_lastname	View_Contact_Full.last_name: The analyst's last name, derived from chgalg.analyst = ca_contact.contact_uuid.
analyst_firstname	View_Contact_Full.first_name: The first name of the analyst.

Field	Remarks
analyst_middlename	View_Contact_Full.middle_name: The middle name of the analyst.
activity_type	act_type.sym: The activity type referenced by chgalg.type = act_type.code.

View_Change

The following is a basic view of all change orders, listing the status, priority, category, organizations, the affected end user's full name, the requester's full name, the assignee's full name, the group name and ID, and so on. Here, the change request table (chg) was joined with many other tables to give some more meaningful data about the change order.

Field	Remarks
id	chg.id: The unique identifier for this record in the chg table.
persid	chg.persid: The unique identifier for this record in the chg table, preceded by the object identifier (chg for table chg) and a colon.
chg_ref_num	chg.chg_ref_num: The change order reference number, which is used by analysts and customers to refer to a particular change order.
description	chg.description: The long description of a change order, as dictated by an analyst or customer.
status	chg.status: The unique identifier of a change order status, which is a pointer to the chgstat table. chg.status = chgstat.code.
active_flag	chg.active_flag: The integer flag to determine whether this change record is active or not (1 or 0).
start_date	chg.start_date: The date the first task goes to a pending status (pdmtime).
open_date	chg.open_date: The change order creation date (pdmtime).
last_mod_dt	chg.last_mod_dt: The last modified date (pdmtime).
last_mod_by	chg.last_mod_by: The pointer to the contact uuid who was the last contact to modify this change order. chg.last_mod_by = ca_contact.contact_uuid.

Field	Remarks
close_date	chg.close_date: The date the change order was set to inactive (pdmtime).
resolve_date	chg.resolve_date: The date the change order was set to a status configured to indicate the change was resolved (pdmtime).
rootcause	chg.rootcause: A pointer to a record in the rootcause table, which represents the original situation which required this change order to be executed. chg.rootcause = rootcause.id.
est_total_time	chg.est_total_time: The estimated total time (pdmtime) it will take to complete this change.
actual_total_time	chg.actual_total_time: The actual total time (pdmtime) it took to complete this change.
log_agent	chg.log_agent: A binary unique identifier referencing the ca_contact table, referencing the person who was the change's original creator. chg.log_agent = ca_contact.contact_uid.
assignee	chg.assignee: The pointer to the contact uid who is currently assigned to the change order. chg.assignee = ca_contact.contact_uid.
organization	chg.organization: The pointer to internal organization uid, which represents the organization to whom this change order belongs. chg.organization = ca_organization.organization_uid.
group_id	chg.group_id: The pointer to contact uid, which represents the group currently assigned to the change order. chg.group_id = ca_contact.contact_uid
affected_contact	chg.affected_contact: The pointer to the contact uid, which represents the affected contact for this change order. chg.affected_contact = ca_contact.contact_uid
requestor	chg.requestor: The pointer to the contact uid, which represents the person who ordered the change. chg.requestor = ca_contact.contact_uid
category	chg.category: The pointer to the change category code to get the category into which this change falls. chg.category = chgcat.code

Field	Remarks
priority	chg.priority: The pointer to priority enum, which represents the priority into which this change falls. chg.priority = pri.enum
need_by	chg.need_by: The date which indicates when the affected_end_user needs to have the change completed (pdmtime).
est_comp_date	chg.est_comp_date: The estimated completion date (pdmtime) for this Change Order.
actual_comp_date	chg.actual_comp_date: The actual completion date (pdmtime) of this change order.
est_cost	chg.est_cost: The estimated cost of this change order.
actual_cost	chg.actual_cost: The actual cost to implement this change order.
justification	chg.justification: A text field which allows a requestor to document the reason(s) this change is required.
backout_plan	chg.backout_plan: A text field that allows an analyst to document a backout plan for this change.
impact	chg.impact: A pointer to an impact table record, which indicates the scope of resources that this change affects. chg.impact = impact.enum
parent	chg.parent: A pointer to another change request id, which allows creation of a hierarchy of change orders. chg.parent = chg.id
effort	chg.effort: A text field which explains the plan for implementing this change order.
support_lev	chg.support_lev: A pointer to a service desc record, which automates some constraints for which this change must be completed. chg.support_lev = srv_desc.code
template_name	chg.template_name: The name of and pointer to a change order template. chg.template_name = chg_template.template_name
sla_violation	chg.sla_violation: The integer to count the number of times slas attached to this "change has been violated".
predicted_sla_viol	chg.predicted_sla_viol: (r5.5) Neugent related technology field.
macro_predict_viol	chg.macro_predict_viol: (r5.5) Neugent related technology field.

Field	Remarks
created_via	chg.created_via: A pointer to a record in the interface table, which indicates from which interface the change order originated. chg.created_via = interface.id
call_back_date	chg.call_back_date: A date/time field (pdmtime), which indicates a future date/time the requestor is to be contacted.
call_back_flag	chg.call_back_flag: A boolean indicator displayed as a checkbox to the user, indicating whether or not to notify the analyst at the chg.call_back_date.
string1	This is a user-definable text field.
string2	This is a user-definable text field.
string3	This is a user-definable text field.
string4	This is a user-definable text field.
string5	This is a user-definable text field.
string6	This is a user-definable text field.
service_date	chg.service_date: The Date/ Time (pdmtime) that an outside vendor is expected to spend to service this change order.
service_num	chg.service_num: The text field to document an outside vendor's service or purchase order number.
product	chg.product: A pointer to a record in the product table, which indicates the product that is affected by this change. chg.product = product.id
actions	chg.actions: A big text field for documenting actions.
type_of_contact	chg.type_of_contact: A pointer to a record in the toc table, which indicates a general categorization of the affected_end_user's perspective of the change order. chg.type_of_contact = toc.id
reporting_method	chg.reporting_method: A pointer to a record in the repmeth table, which classifies the origination of the change order, and is selected by the person creating the change order. chg.reporting_method = repmeth.id
person_contacting	chg.person_contacting: A pointer to a record in the perscon table, which indicates the role of the affected_end_user or requestor. chg.person_contacting = perscon.id

Field	Remarks
status_name	chgstat.sym: The description of the status as seen by a user. chg.status = chgstat.code
priority_num	pri.sym: The description of the priority as seen by a user. chg.priority = pri.enum
category_name	chgcatsym: The name of the Change Category as viewed by a user. chg.category = chgcatsym.code
organization_name	ca_organization.org_name: The name of an organization as viewed by a user. chg.organization = ca_organization.organization_uuid
affected_end_user_lastname	ca_contact.last_name: The affected end user's last name. chg.affected_end_user = ca_contact.contact_uuid
affected_end_user_firstname	ca_contact.first_name: The affected end user's first name. chg.affected_end_user = ca_contact.contact_uuid
affected_end_user_middlename	ca_contact.middle_name: The affected end user's middle name. chg.affected_end_user = ca_contact.contact_uuid
requester_lastname	ca_contact.last_name: The requestor's last name. chg.requestor = ca_contact.contact_uuid
requester_firstname	ca_contact.first_name: The requestor's first name. chg.requestor = ca_contact.contact_uuid
requester_middlename	ca_contact.middle_name: The requestor's middle name. chg.requestor = ca_contact.contact_uuid
business	ca_organization.org_name: The name of the Requester's organization as seen by users. chg.requestor = ca_organization.organization_uuid
assignee_lastname	ca_contact.last_name: The assignee's last name. chg.assignee = ca_contact.contact_uuid
assignee_firstname	ca_contact.first_name: The assignee's first name. chg.assignee = ca_contact.contact_uuid
assignee_middlename	ca_contact.middle_name: The assignee's middle name. chg.assignee = ca_contact.contact_uuid
groupID	ca_contact.contact_uuid: A binary representation of the internal id used for the group assigned to this change order. chg.group_id = ca_contact.contact_uuid

Field	Remarks
group_name	ca_contact.last_name: The name of the group assigned to this change order. chg.group = ca_contact.contact_uuid
service_type	srv_desc.sym: The name of the service type applied to this change order. chg.support_lev = srv_desc.code
impact_num	impact.sym: The description of the impact as seen by users. chg.impact = impact.enum
product_sym	product.sym: The product description as seen by users. chg.product = product.id
type_of_contact_sym	toc.sym: The Type Of Contact description as seen by users. chg.type_of_contact = toc.id
rpting_method_sym	repmeth.sym: The Reporting method description as seen by users. chg.reporting_method = repmeth.id
person_contacting_sym	perscon.sym: The Person Contacting description as seen by users. chg.person_contacting = perscon.id

View_Change_to_Assets

The following list of fields is a basic view of change orders and their assets. The change request table (chg) is indirectly joined with the network resource table (ca_owned_resource) to get a list of each change order's assets. This view may not list all change orders, particularly those that have no assets.

Field	Remarks
View_Change.*	All fields listed in the View_Change view defined earlier in this document.
assetID	ca_owned_resource.own_resource_uuid: The binary field which serves as the internal, unchanging unique identifier for an asset record.
asset_serial_num	ca_owned_resource.serial_number: The serial number for an asset record.

Field	Remarks
asset_class	ca_resource_class.name: A short description of the class to which an asset belongs. ca_owned_resource.resource_class = ca_resource_class.id
asset_family	ca_resource_family.name: The family of assets to which this asset belongs. ca_owned_resource.resource_class = ca_resource_class.id AND ca_resource_class.family_id = ca_resource_family.id
asset_name	ca_owned_resource.resource_name: The network name by which this asset is known.

View_Change_to_Change_Act_Log

The following is a basic view of all change orders and the activity logs that go with them. This view joins the View_Change view with the Change Order Activity Log (chgalg) to give detailed information about change orders and their activity logs.

Field	Remarks
View_Change.*	This shows all fields listed in the View_Change view defined earlier in this document.
chgalg_id	chgalg.id: The unique identifier for this record in the chgalg table.
chgalg_persid	chgalg.persid: The unique identifier for this record in the chgalg table, preceded by the object identifier (chgalg for chgalg) and a colon.
change_id	chgalg.change_id: This is a pointer to change the order id to which this activity belongs. chgalg.change_id = chgalg.id
chgalg_last_mod_dt	chgalg.last_mod_dt: The last modify date/time (pdmtime).
time_spent	chgalg.time_spent: The duration of time spent on this activity, stored as the total number of seconds. For example, 80 = 1 minute, 20 seconds.
time_stamp	chgalg.time_stamp: The user modifiable date/time of activity (pdmtime).
system_time	chgalg.system_time: The date/time of record creation (pdmtime).

Field	Remarks
analyst	chgalg.analyst: The uuid pointer to contact uuid to get the analyst who performed the activity. chgalg.analyst = ca_contact.contact_uuid
chgalg_description	chgalg.description: The text description of this activity, which can be modified by the user.
action_desc	chgalg.action_desc: The text description of automated action, which cannot be modified by the user.
type	chgalg.type: The text pointer to a record in the activity type table. chgalg.type = act_type.code
internal	chgalg.internal: The integer flag (1 or 0), which indicates if this log entry is intended for all to see or just for internal use.
knowledge_session	chgalg.knowledge_session: This is an identifier for a particular session of a particular user.
knowledge_tool	chgalg.knowledge_tool: This is an indicator of the knowledge tool used for the search, such as NLS_FAQ or EXPERT, and so on.
chgalg_analyst_id	chgalg.analyst: This is the uuid pointer to contact uuid to get the analyst who performed the activity. chgalg.analyst = ca_contact.contact_uuid

View_Change_to_Change_WF

This view is a result of the View_Change view joined with the Workflow task table (wf) to give a basic view of change orders and their workflow tasks. This may not list all change orders, particularly when there are no workflow tasks assigned.

Field	Remarks
View_Change.*	This shows all fields listed in the View_Change view defined earlier in this document.
wf_id	wf.id: The unique identifier for a record in the wf table.
wf_persid	wf.persid: This is a unique identifier for this record in the wf table, preceded by the object identifier (wf for wf) and a colon.
del	wf.del: This is a boolean indicator. It specifies whether this record is to be displayed to the user.

Field	Remarks
object_type	wf.object_type: This is the factory name, which is used to identify the type of record (for example, chg) to which this workflow task is attached.
object_id	wf.object_id: This is the unique identifier used to identify the specific record to which this workflow task is attached. wf.object_id = chg.id
task	wf.task: This is an identifier, which references the type of task this record represents. wf.task = tscky.code
wf_template	wf.wf_template: This is an identifier, which references from which template this workflow task record was created. wf.wf_template = wftpl.id
sequence	wf.sequence: This is an integer, which indicates the order this particular workflow task record should be displayed and executed by CA SDM (for example, Ascending).
wf_status	wf.status: This is an identifier, which references a tsckstat record that indicates the current status of this workflow task. wf.status = tsckstat.code
group_task	wf.group_task: This is a Boolean, which indicates whether this task belongs to a group.
asset	wf.asset: This is a UUID (binary) identifier, which references a record in the ca_owned_resource table. wf.asset = ca_owned_resource.own_resource_uuid
creator	wf.creator: This is a UUID (binary) identifier, which references a record in the ca_contact table. It indicates the person who created this workflow task. wf.creator = ca_contact.contact_uuid
date_created	wf.date_created: This is the Date/timestamp this workflow task was created (pdmtime).
wf_assignee	wf.assignee: This is the UUID (binary) identifier, which references a record in the ca_contact table. It indicates the person who is currently assigned to this workflow task. wf.assignee = ca_contact.contact_uuid
done_by	wf.done_by: This is the UUID (binary) identifier, which references a record in the ca_contact table. It indicates the person who completed or approved this workflow task. wf.done_by = ca_contact.contact_uuid
wf_start_date	wf.start_date: This is the timestamp when the workflow task moved into an active status (pdmtime).

Field	Remarks
wf_est_comp_date	wf.est_comp_date: This is the timestamp (pdmtime) when users believe this task will be completed.
est_duration	wf.est_duration: This is the estimated duration for this workflow task.
completion_date	wf.completion_dat: This is the timestamp (pdmtime) when this workflow task was completed.
actual_duration	wf.actual_duration: This is the actual amount of time it took to complete this workflow task.
wf_est_cost	wf.est_cost: This is the estimated cost of this workflow task.
cost	wf.cost: This is the actual cost required to complete this workflow task.
wf_description	wf.description: This is the description of the workflow task.
wf_last_mod_dt	wf.last_mod_dt: This is the timestamp (pdmtime) when this workflow task was last changed.
wf_last_mod_by	wf.last_mod_by: This is the UUID (binary) unique identifier referencing a record in the contact table, which indicates the last person to make changes to this workflow task. wf.last_mod_by = ca_contact.contact_uuid

View_Change_to_Properties

This view is a result of the View_Change view joined with the Properties table (prp) to give a basic view of change orders and their assigned properties. This may not list all change orders, particularly when there are no properties assigned.

Field	Remarks
View_Change.*	This shows all fields listed in the View_Change view defined earlier in this document.
prp_id	prp.id: This is an integer unique identifier for the property record.
prp_persid	prp.persid: This is a unique identifier for this record in the wf table, preceded by the object identifier (prp for prp) and a colon.

Field	Remarks
object_type	prp.object_type: This is the factory name, which is used to identify the type of record (for example, chg) to which this property record is attached.
object_id	prp.object_id: This is the unique identifier used to identify the specific record to which this property is attached. prp.object_id = chg.id
sequence	prp.sequence: This is an integer that indicates the order for which this particular property record should be displayed by CA SDM (for example, Ascending).
property	prp.property: This is an identifier, which references a record in the prptpl table. It represents the template from which this property was created. prp.property=prptpl.code
value	prp.value: This is the value entered by the user in response to the prp_description and prp.label fields.
prp_last_mod_dt	prp.last_mod_dt: This is the timestamp (pdmtime) when this property was last modified.
prp_last_mod_by	prp.last_mod_by: This is a binary identifier, which references a record in the ca_contact table. It represents the person who last modified this record. prp.last_mod_by = ca_contact.contact_uid
required	prp.required: This is a Boolean indicating whether this property must have a prp.value before the record is saved.
sample	prp.sample: This is a text field, which displays example values to guide the user in typing the most useful value in prp.value.
prp_description	prp.description: This is a text field, which explains what kind of value should be entered in prp.value.
label	prp.label: This is a short description of what should be placed in the prp.value field.

View_Contact_Full

The following of fields is a basic view of all contacts. This view lists all fields in the `ca_contact` table, plus referenced fields such as short descriptions for contact type, location name, organization names, and service type for each contact. This view has already been joined with the `ca_location`, `ca_organization`, `srv_desc`, and `ca_contact_type` tables to get the actual names and symbols for some of the fields in the `ca_contact` table. The actual names and symbols are located at the end of this view's field list.

Field	Remark
<code>contact_uuid</code>	<code>ca_contact.contact_uuid</code> : A binary-unique identifier for each <code>ca_contact</code> record.
<code>middle_name</code>	<code>ca_contact.middle_name</code> : The middle name of this contact.
<code>alias</code>	<code>ca_contact.alias</code> : The alternate, often informal name for this contact.
<code>last_name</code>	<code>ca_contact.last_name</code> : This contact's surname.
<code>first_name</code>	<code>ca_contact.first_name</code> : This contact's formal first name.
<code>pri_phone_number</code>	<code>ca_contact.pri_phone_number</code> : The contact's primary phone number.
<code>alt_phone_number</code>	<code>ca_contact.alt_phone_number</code> : The contact's alternate phone number.
<code>fax_number</code>	<code>ca_contact.fax_number</code> : The contact's facsimile number.
<code>mobile_phone</code>	<code>ca_contact.mobile_phone</code> : The contact's mobile phone number.
<code>pager_number</code>	<code>ca_contact.pager_number</code> : The number for issuing a page to the contact's page.
<code>email_address</code>	<code>ca_contact.email_address</code> : The contact's email address.
<code>location_uuid</code>	<code>ca_contact.location_uuid</code> : A binary-unique identifier referencing a record in the <code>ca_location</code> table, which indicates the contact's static location. <code>ca_contact.location_uuid = ca_location.location_uuid</code>
<code>floor_location</code>	<code>ca_contact.floor_location</code> : A contact's floor number.
<code>pager_email_address</code>	<code>ca_contact.pager_email_address</code> : A contact's email address for their pager.

Field	Remark
room_location	ca_contact.room_location: The contact's specific room on the floor at the static location.
contact_type	ca_contact.contact_type: An unique integer-identifier, referring to a row in the ca_contact_type table, which indicates the general function of this contact within the service desk application. ca_contact.contact_type = ca_contact_type.id
inactive	ca_contact.inactive: A boolean indicator of the state of this record, determining its inclusion or exclusion from standard searches within service desk.
creation_user	ca_contact.creation_user: The userid of the contact who created this record. ca_contact.creation_user = ca_contact.userid
creation_date	ca_contact.creation_date: A timestamp (pdmtime), indicating the date and time this contact was created.
last_update_user	ca_contact.last_update_user: The userid of the contact who last updated this contact record. ca_contact.last_update_user = ca_contact.userid
last_update_date	ca_contact.last_update_date: A timestamp (pdmtime), indicating the date and time this record was last modified.
version_number	ca_contact.version number: The internal version indicator.
department	ca_contact.department: An integer-unique identifier, referencing a row in the ca_resource_department table, that indicates the contact's department. ca_contact.department = ca_resource_department.id
comment	ca_contact.comment: A freeform text comment field for analysts to document important facts that influence the handling of this particular contact.
company_uuid	ca_contact.company_uuid: A binary-unique identifier, which references a row in the ca_company table. It indicates this contact's affiliation with a company. ca_contact.company_uuid = ca_company.company_uuid

Field	Remark
organization_uuid	ca_contact.organizaiton_uuid: A binary-unique identifier which references a row in the ca_organization table. It indicates this contact's working organization. ca_contact.organizaiton_uuid = ca_organization.organization_uuid
admin_organization_uuid	ca_contact.admin_organization_uuid: A binary-unique identifier, which references a row in the ca_organization table. It indicates this contact's administrative organization. ca_contact.admin_organization_uuid = ca_organization.organization_uuid
alternate_identifier	ca_contact.alternate_identifier: A user-defined identifier, usually an entity used by human resources to uniquely identify this contact.
job_title	ca_contact.job_title: An integer-unique identifier, which references the ca_job_title table. It indicates the standardized job title for this contact. ca_contact.job_title = ca_job_title.id
job_function	ca_contact.job_function: An integer-unique identifier, which references the ca_job_function table. It indicates a standardized general description of the contact's job function. ca_contact.job_function = ca_job_function.id
mail_stop	ca_contact.mail_stop: Mail Stop.
cost_center	ca_contact.cost_center: An integer-unique identifier, which references the ca_resource_cost_center table. It indicates this contact's primary cost center. ca_contact.cost_center = ca_cost_center.id
userid	ca_contact.userid: The user identifier this contact will use for logging into service desk.
supervisor_contact_uuid	ca_contact.supervisor_contact_uuid: A binary-unique identifier, referencing a row in the ca_contact table, that creates a hierarchy of contacts to indicate the reporting structure of each contact. ca_contact.supervisor_contact_uuid = ca_contact.contact_uuid
exclude_registration	ca_contact.exclude_registration: An internal flag.
delete_time	ca_contact.delete_time: A timestamp indicating when the inactive flag was set to 1.

Field	Remark
contact_type_name	ca_contact_type.name: A short description of this contact's ca_contact.contact_type.
location_name	ca_location.name: A short description of this contact's static location.
organization	ca_organization.org_name: A short description of this contact's ca_contact.organizaiton_uuid.
admin_organization	ca_organization.org_name: A short description of this contact's organization name. ca_contact.admin_organization_uuid
service_type	srv_desc.sym: A short description of this contact's usp_contact.c_service_type. ca_contact.contact_uuid = usp_contact.contact_uuid AND usp_contact.c_service_type = srv_desc.code
state_sym	ca_location.state: An integer-unique identifier, referencing a row in the ca_state_provice table, that indicates the State, Province, or other artificially defined geographic region. ca_contact.location_uuid = ca_location.location_uuid AND ca_location.state = ca_state_province.id

View_Contact_to_Environment

The following list of fields is a basic view of contacts and their environment (assets). This view is a view of the contact table (ca_contact), but is also joined with the owned resource table (ca_owned_resource) to get a list of all assets associated with a contact. This view may be joined with the View_Contact_Full view to get the contact type, service type, organizations and location of each contact. Or, this view may be joined with the individual tables to get that same information.

Field	Remarks
ca_contact.*	The ca_contact fields that are defined in the View_Contact_Full view definition.
asset_uuid	ca_owned_resource.own_resource_uuid: A binary-unique identifier for an asset in the ca_owned_resource table.
asset_name	ca_owned_resource.resource_name: The designated network name of this asset.

View_Group

The following list of fields is a basic view of the contact table, but lists only group contacts. Contact type id = 2308, which is for group types. You may want to join this view with other CA SDM tables to get more meaningful data for reporting. For example, you can join it with the Location (ca_location) table to find the name and address for the group's location. You can also join this view with the Organization (ca_organization) table to get the functional and administrative organization names for the group.

Field	Remarks
ca_contact.*	The ca_contact fields that are defined in the View_Contact_Full view definition.

The following example shows how joining tables works, and how reporting fields are extracted. The field from one table (on the right) is joined (->) with a field from another table (on the left). To join properly between tables and views, you need to understand the differences in joins for the database. The field defined in the parentheses, as follows, is what you may want to use in your reports if the previous tables are joined with the View_Group view:

- View_Group.contact_type -> ca_contact_type.id (ca_contact_type.sym)
- View_Group.location_uuid -> ca_location.location_uuid (ca_location.location_name)
- View_Group.organization_uuid -> ca_organization.organization_uuid (ca_organization.org_name)
- View_Group.admin_organization_uuid -> ca_organization(2).organization_uuid (ca_organization(2). org_name)

View_Group_to_Contact

The following list of fields is a basic view of all group contacts (members). It also includes managers. Here, View_Group is joined with the Group_Member table, and then joined with the ca_contact table. All fields in View_Group are listed, as well as the first, middle and last name from the ca_contact table. The Group_Member manager flag is also listed. The group member manager flag is 1 or 0, which means that the member is either (yes - 1) a manager, or (no - 0) not a manager. Most of the information in this view pertains to the group itself, not the actual members. This view is used to find information on a particular group, including the names of its members.

Field	Remarks
View_Group.*	The View_Group fields that are defined in the View_Group definition.
member_lastname	ca_contact.last_name: The surname of the group member.
member_firstname	ca_contact.first_name: The formal first name of the group member.
member_middlename	ca_contact.middle_name: The middle name of the group member.
grpmem_manager_flag	ca_contact.manager_flag: The Group Member Manager (1 or 0) indicator.

View_Issue

The following is a basic view of all issues, listing the status, priority, category, organizations, the requester's full name, the assignee's full name, the group name and ID, and so on. Here, the issue table is joined with many other tables to give some more meaningful data about the issue.

Field	Remarks
id	issue.id: The unique identifier for this record in the issue table.
persid	issue.persid: The unique identifier for this record in the issue table, preceded by the object identifier (iss for table issue) and a colon.
issue_ref_num	issue.iss_ref_num: The Issue reference number, which is used by analysts and customers to refer to a particular issue.
description	issue.description: The long description of an issue as dictated by an analyst or customer.
status	issue.status: The unique identifier of an issue status, which is a pointer to the issstat table. issue.status = issstat.code
active_flag	issue.active_flag: The integer flag used to determine whether this issue is active (1 or 0).
start_date	issue.start_date: The date the first task goes to a pending status (pdmtime).

Field	Remarks
open_date	issue.open_date: The issue creation date (pdmtime).
last_mod_dt	issue.last_mod_dt: The last modified date (pdmtime).
last_mod_by	issue.last_mod_by: The pointer to the contact uuid, which indicates the last contact to modify this issue. issue.last_mod_by = ca_contact.contact_uuid
close_date	issue.close_date: The date the issue was set to inactive (pdmtime).
resolve_date	issue.resolve_date: The date the issue was set to a status configured to indicate the issue was resolved (pdmtime).
rootcause	issue.rootcause: A pointer to a record in the rootcause table, which represents the original situation that required this issue to be logged. issue.rootcause = rootcause.id
est_total_time	issue.est_total_time: The estimated total time (pdmtime) it will take to complete this issue.
actual_total_time	issue.actual_total_time actual: The total time (pdmtime) it took to complete this issue.
log_agent	issue.log_agent: A binary-unique identifier, this references the ca_contact table, which in turn references the person who was the issue's original creator. issue.log_agent = ca_contact.contact_uuid
assignee	issue.assignee: A pointer to the contact uuid who is currently assigned to the change order. issue.assignee = ca_contact.contact_uuid
organization	issue.organization: A pointer to the internal organization uuid, which represents the organization to whom this issue belongs. issue.organization = ca_organization.organization_uuid
group_id	issue.group_id: A pointer to the contact uuid, which represents the group currently assigned to the issue. issue.group_id = ca_contact.contact_uuid
affected_contact	issue.affected_contact: A pointer to the contact uuid, which represents the affected contact for this issue. issue.affected_contact = ca_contact.contact_uuid
requestor	issue.requestor: A pointer to the contact uuid, which represents the person who asked this issue to be logged. issue.requestor = ca_contact.contact_uuid

Field	Remarks
category	issue.category: A pointer to the issue category code to reference the category into which this issue falls. issue.category = isscat.code
priority	issue.priority: A pointer to priority enum, which represents the priority into which this issue falls. issue.priority = pri.enum
need_by	issue.need_by: The date, which indicates when the affected_end_user needs to have the issue completed (pdmtime).
est_comp_date	issue.est_comp_date: The estimated completion date (pdmtime) for this Issue.
actual_comp_date	issue.actual_comp_date: The actual completion date (pdmtime) of this issue.
est_cost	issue.est_cost: The estimated cost of this issue.
actual_cost	issue.actual_cost: The actual cost to implement this issue.
justification	issue.justification: A text field, which allows a requester to document the reason(s) this issue is required.
backout_plan	issue.backout_plan: A text field that allows an analyst to document a backout plan for this issue.
impact	issue.impact: A pointer to an impact table record, which indicates the scope of resources that this issue affects. issue.impact = impact.enum
parent	issue.parent: A pointer to another issue id, which allows creation of a hierarchy of issues. issue.parent = issue.id
effort	issue.effort: A text field which explains the plan for implementing this issue.
support_lev	issue.support_lev: A pointer to a service desc record, which automates some constraints under which this issue must be completed. issue.support_lev = srv_desc.code
template_name	issue.template_name: The name of and pointer to an issue template. issue.template_name = iss_template.template_name
sla_violation	issue.sla_violation: The integer to count the number of times slas attached to this issue have been violated.
predicted_sla_viol	issue.predicted_sla_viol: (r5.5) Neugent related technology field.

Field	Remarks
macro_predict_viol	issue.macro_predict_viol: (r5.5) Neugent related technology field.
created_via	issue.created_via: A pointer to a record in the interface table. This indicates from which interface the issue originated. issue.created_via = interface.id
call_back_date	issue.call_back_date: A date/time field (pdmtime), which indicates a future date/time the requester is to be contacted.
call_back_flag	issue.call_back_flag: A boolean indicator displayed as a checkbox to the user, to indicate whether to notify the analyst at the issue.call_back_date.
string1	This is a user-definable text field.
string2	This is a user-definable text field.
string3	This is a user-definable text field.
string4	This is a user-definable text field.
string5	This is a user-definable text field.
string6	This is a user-definable text field.
service_date	issue.service_date: The Date/ Time (pdmtime) that an outside vendor is expected to service this issue.
service_num	issue.service_num: The text field to document an outside vendor's service or purchase order number.
product	issue.product: A pointer to a record in the product table, which indicates the product that is affected by this issue. issue.product = product.id
actions	issue.actions: A big text field for documenting actions.
type_of_contact	issue.type_of_contact: A pointer to a record in the toc table, which indicates a general categorization of the affected_end_user's perspective of the issue. issue.type_of_contact = toc.id
reporting_method	issue.reporting_method: A pointer to a record in the repmeth table, which classifies the origination of the issue, and is selected by the person creating the issue. issue.reporting_method = repmeth.id

Field	Remarks
person_contacting	issue.person_contacting: A pointer to a record in the perscon table, which indicates the role of the affected_end_user or requester. issue.person_contacting = perscon.id
status_name	issstat.sym: The description of the status as seen by a user. issue.status = issstat.code
priority_num	pri.sym: The description of the priority as seen by a user. issue.priority = pri.enum
category_name	isscat.sym: The name of the issue category as viewed by a user. issue.category = isscat.code
organization_name	ca_organization.org_name: The name of an organization as viewed by a user. issue.organization = ca_organization.organization_uuid
affected_end_user_lastname	ca_contact.last_name: The affected End User's last name. issue.affected_end_user = ca_contact.contact_uuid
affected_end_user_firstname	ca_contact.first_name: The affected End User's first name. issue.affected_end_user = ca_contact.contact_uuid
affected_end_user_middlename	ca_contact.middle_name: The affected End User's middle name. issue.affected_end_user = ca_contact.contact_uuid
assignee_lastname	ca_contact.last_name: The assignee's last name. issue.assignee = ca_contact.contact_uuid
assignee_firstname	ca_contact.first_name: The assignee's first name. issue.assignee = ca_contact.contact_uuid
assignee_middlename	ca_contact.middle_name: The assignee's middle name. issue.assignee = ca_contact.contact_uuid
groupID	View_Group.contact_uuid: A binary representation of the internal id used for the group assigned to this issue. issue.group_id = ca_contact.contact_uuid
group_name	View_Group.last_name: The name of the group assigned to this issue. issue.group = ca_contact.contact_uuid
service_type	srv_desc.sym: The name of the service type applied to this issue. issue.support_lev = srv_desc.code
impact_num	impact.sym: The description of the impact as seen by users. issue.impact = impact.enum

Field	Remarks
product_sym	product.sym: The product description as seen by users. issue.product = product.id
type_of_contact_sym	toc.sym: The Type Of Contact description as seen by users. issue.type_of_contact = toc.id
rpmtg_method_sym	repmeth.sym: The Reporting method description as seen by users. issue.reporting_method = repmeth.id
person_contacting_sym	perscon.sym: The Person Contacting description as seen by users. issue.person_contacting = perscon.id
created_via_sym	interface.sym: issue.created_via = interface.id.
rootcause_sym	rootcause.sym: issue.rootcause = rootcause.id.

View_Issue_Act_Log

The following is a basic view of all issue activity logs. This is a view of the issue activity log table (issalg) joined with the activity type table (act_type) and the contact table (ca_contact) to give more meaningful data, such as the actual activity type and full name of the analyst who performed the activity.

Field	Remarks
id	issalg.id: The unique identifier for this record in the issalg table.
persid	issalg.persid: The unique identifier for this record in the issalg table, preceded by the object identifier (issalg for issalg) and a colon.
issue_id	issalg.issue_id: The pointer to issue id to which this activity belongs. issalg.issue_id = issalg.id
last_mod_dt	issalg.last_mod_dt: The last modify date/time (pdmtime).
time_spent	issalg.time_spent: The duration of time spent on this activity, stored as the total number of seconds. For example, 80 = 1 minute, 20 seconds.
time_stamp	issalg.time_stamp: The user modifiable date/time of activity (pdmtime).
system_time	issalg.system_time: The date/time of record creation (pdmtime).

Field	Remarks
analyst	issalg.analyst: A binary-unique identifier referring to the contact uuid to get the analyst who performed the activity. issalg.analyst = ca_contact.contact_uuid
description	issalg.description: The text description of this activity, which can be modified by the user.
action_desc	issalg.action_desc: The text description of the automated action, which cannot be modified by the user.
type	issalg.type: The text pointer to a record in the activity type table. issalg.type = act_type.code
internal	issalg.internal: The integer flag (1 or 0), which indicates if this log entry is intended for all to see or just for internal use.
knowledge_session	issalg.knowledge_session: An identifier for a particular session of a particular user.
knowledge_tool	issalg.knowledge_tool: An indicator of the knowledge tool used for the search, such as NLS_FAQ or EXPERT, and so on.
analyst_lastname	View_Contact_Full.last_name: The Analyst's last name, derived from issalg.analyst = ca_contact.contact_uuid.
analyst_firstname	View_Contact_Full.first_name: The analyst's first name.
analyst_middlename	View_Contact_Full.middle_name: The middle name of the analyst.
activity_type	act_type.sym: The activity type referenced by issalg.type = act_type.code.

View_Issue_to_Assets

The following list of fields is a basic view of issues and their assets. The issue table (issue) is indirectly joined with the owned resource table (ca_owned_resource), and other asset-related tables, to get a list of each issue's assets. This may not list all issues, particularly those that have no assets.

Field	Remarks
View_Issue.*	The View_Issue view which defines all fields listed in the View_Issue view.

Field	Remarks
assetID	ca_owned_resource.own_resource_uuid: The binary field which serves as the internal, unchanging unique identifier for an asset record.
asset_serial_num	ca_owned_resource.serial_number: The serial number for an asset record.
asset_class	ca_resource_class.name: A short description of the class to which an asset belongs. ca_owned_resource.resource_class = ca_resource_class.id
asset_family	ca_resource_family.name: The family of assets to which this asset belongs. ca_owned_resource.resource_class = ca_resource_class.id AND ca_resource_class.family_id = ca_resource_family.id
asset_name	ca_owned_resource.resource_name: The network name by which this asset is known.

View_Issue_to_Issue_Act_Log

The following is a basic view of all issues and the activity logs that go with them. This view joins the View_Issue view with the View_Issue_Act_Log view to give detailed information about issues and their activity logs. Actual data is at the end of the fields list.

Field	Remarks
View_Issue.*	Please refer to the View_Issue view defined earlier in this document.
issalg_id	issalg.id: The unique identifier for this record in the issalg table.
issalg_persid	issalg.persid: The unique identifier for this record in the issalg table, preceded by the object identifier (issalg for issalg) and a colon.
issue_id	issalg.issue_id: The pointer to issue id to which this activity belongs. issalg.issue_id = issalg.id
issalg_last_mod_dt	issalg.last_mod_dt: The last modify date/time (pdmtime).

Field	Remarks
time_spent	issalg.time_spent: The duration of time spent on this activity, stored as the total number of seconds. For example, 80 = 1 minute, 20 seconds.
time_stamp	issalg.time_stamp: The user modifiable date/time of activity (pdmtime).
system_time	issalg.system_time: The date/time of record creation (pdmtime).
analyst	issalg.analyst: The unique binary pointer to the contact uuid to get the analyst who performed the activity. issalg.analyst = ca_contact.contact_uuid
issalg_description	issalg.description: The text description of this activity, which can be modified by the user.
action_desc	issalg.action_desc: The text description of the automated action, which cannot be modified by the user.
type	issalg.type: The text pointer to a record in the activity type table. issalg.type = act_type.code
internal	issalg.internal: The integer flag (1 or 0), which indicates if this log entry is intended for all to see or just for internal use.
knowledge_session	issalg.knowledge_session: An identifier for a particular session of a particular user.
knowledge_tool	issalg.knowledge_tool: An indicator of the knowledge tool used for the search, such as NLS_FAQ or EXPERT, and so on.
issalg_analyst_id	issalg.analyst: The unique binary pointer to the contact uuid to get the analyst who performed the activity. issalg.analyst = ca_contact.contact_uuid

View_Change_to_Request

The following is a basic view of change orders that have assigned requests only. This view is a result of the View_Change view joined with the request table (call_req) to give details about the change order and its associated request.

Field	Remarks
View_Change.*	This shows all fields listed in the View_Change view defined earlier in this document.
cr_id	call_req.id: This is the unique identifier for this record in the call_req table.
ref_num	call_req.ref_num: This is the Request reference number, which is used by analysts and customers to refer to a particular Request.
cr_summary	call_req.summary: This is a brief description of the request for quick reference.
cr_persid	call_req.persid: This is a unique identifier for this record in the call_req table, preceded by the object identifier (cr for table call_req) and a colon.
cr_description	call_req.description: This is the long description of a request, as dictated by an analyst or customer.
cr_status	call_req.status: This is a unique identifier referencing a record in the cr_stat table. It indicates the status of this request: call_req.status = cr_stat.code
cr_active_flag	call_req.active_flag: This is the Integer flag used to determine whether this request record is active (1 or 0).
time_spent_sum	call_req.time_spent_sum: This is the derived total of all act_log records' time_spent fields, stored in seconds (i.e. 80 = 1 minute 20 seconds).
cr_open_date	call_req.open_date: This is the Request creation timestamp (pdmtime).
cr_last_mod_dt	call_req.last_mod_dt: This is the last modified timestamp (pdmtime).
cr_close_date	call_req.close_date: This is the Timestamp for when the request was set to inactive (pdmtime).
cr_log_agent	call_req.log_agent: This is a binary unique identifier referencing the ca_contact table. It references the person who was the request's original creator. call_req.log_agent = ca_contact.contact_uuid
cr_group_id	call_req.group_id: This is a binary unique identifier referencing a record in the ca_contact table. It represents the group currently assigned to the request. call_req.group_id = ca_contact.contact_uuid

Field	Remarks
cr_assignee	call_req.assignee: This is a binary unique identifier referencing a record in the ca_contact table. It represents the person currently assigned to the request. call_req.assignee = ca_contact.contact_uuid
customer	call_req.customer: This is a binary unique identifier referencing a record in the ca_contact table. It represents the affected end user for this request. call_req.customer = ca_contact.contact_uuid
charge_back_id	charge_back_id: This is a text field available for use as an indicator of accounting jargon for expensing this request to the appropriate cost center.
affected_rc	call_req.affected_rc: This is a binary unique identifier referencing a row in the ca_owned_resource table. It represents the asset to which this request applies. call_req.affected_rc = ca_owned_resource.own_resource_uuid.
cr_support_lev	call_req.support_lev: This is a pointer to a service desc record, which automates some constraints under which this request must be completed. call_req.support_lev = srv_desc.code
cr_category	call_req.category: This is a unique identifier referencing a record in the prob_ctg table,. It represents the category to which this request belongs. call_req.category = prob_ctg.persid
solution	call_req.solution: This is a pointer to call solution to get solution. call_req.solution = crsol.persid
cr_impact	call_req.impact: This is an integer unique identifier referencing a row in the impact table. It indicates the scope this request is affecting. call_req.impact = impact.enum
cr_priority	call_req.priority: This is an integer unique identifier referencing a record in the pri table. It indicates how analysts will prioritize the work associated with this request. call_req.priority = pri.enum
urgency	call_req.urgency: This is an integer unique identifier referencing a row in the urgncy table. It documents the user's feeling of urgency for having this request resolved. call_req.urgency = urgncy.enum

Field	Remarks
severity	call_req.severity: This is an integer unique identifier referencing a row in the severity table. It indicates the severity of the consequences of this unresolved request. call_req.severity = sevrtty.enum
extern_ref	This specifies an associated ticket.
last_act_id	This is the id of the last activity.
cr_ticket	This is a pointer to a trouble ticket to get the associated ticket.
cr_parent	call_req.parent: This is a persid pointer to another request persid, which facilitates creation of a hierarchy of change orders. call_req.parent = call_req.persid
cr_template_name	call_req.template_name: This is a text value, which indicates this request is designated for and can be chosen from a list as a template for other similar requests. cr_template.template = call_req.persid
cr_sla_violation	call_req.sla_violation: This is an integer, which counts number of times slas attached to this request have been violated.
cr_predicted_sla_viol	call_req.predicted_sla_viol: (r5.5) Neugent related technology field.
cr_created_via	call_req.created_via: This is an integer pointer to a record in the interface table. It indicates from which interface the change order originated. call_req.created_via = interface.id
cr_call_back_date	call_req.call_back_date: This is a timestamp field (pdmtime), which indicates a future date/time the affected_end_user is to be contacted.
cr_call_back_flag	call_req.call_back_flag: This is a Boolean indicator displayed as a checkbox to the user, indicating whether to notify the analyst at the call_req.call_back_date.
event_token	call_req.event_token: This is used by CA NSM for message matching.
type	call_req.type: This is a text field referencing a record in the crt table. It indicates the ITIL type of this request. call_req.type = crt.code
cr_string1	This is a user-definable string.
cr_string2	This is a user-definable string.

Field	Remarks
cr_string3	This is a user-definable string.
cr_string4	This is a user-definable string.
cr_string5	This is a user-definable string.
cr_string6	This is a user-definable string.
change	call_req.change: This is an integer unique identifier, referencing a row in the chg table. It indicates the change order that was created as a result of this request. call_req.change = chg.id.

View_Issue_to_Issue_WF

This view is a result of the View_Issue view joined with the workflow task table (isswf) to give a basic view of issues and their workflow tasks. This may not list all issues, particularly if there are no workflow tasks assigned to them.

Field	Remarks
View_Issue.*	Please refer to the View_Issue definition earlier in this document for a description of each field.
wf_id	isswf.id: This is a unique identifier for a record in the isswf table.
wf_persid	isswf.persid: This is a unique identifier for this record in the isswf table, preceded by the object identifier (isswf) and a colon.
del	isswf.del: This is a boolean indicator of whether this record is to be displayed to the user.
object_type	isswf.object_type: This is the factory name, which is used to identify the type of record (for example, iss) to which this workflow task is attached.
object_id	isswf.object_id: This is a unique identifier used to identify the specific record to which this workflow task is attached. isswf.object_id = issue.id
task	isswf.task: This is an identifier, which references the type of task this record represents. isswf.task = tscky.code
wf_template	isswf.wf_template: This is an identifier, which references from which template this workflow task record was created. isswf.wf_template = wftpl.id

Field	Remarks
sequence	isswf.sequence: This is an integer that indicates the order for which this particular workflow task record should be displayed and executed by CA SDM (for example, Ascending).
wf_status	isswf.status: This is an identifier, which references a tskstat record. It indicates the current status of this workflow task. isswf.status = tskstat.code
group_task	isswf.group_task: This is a boolean, which indicates whether this task belongs to a group.
asset	isswf.asset: This is a unique binary identifier, which references a record in the ca_owned_resource table. isswf.asset = ca_owned_resource.own_resource_uuid
creator	isswf.creator: This is a unique binary identifier, which references a record in the ca_contact table. It indicates the person who created this workflow task. isswf.creator = ca_contact.contact_uuid
date_created	isswf.date_created: This is the date/timestamp that this workflow task was created (pdmtime).
wf_assignee	isswf.assignee: This is a unique binary identifier, which references a record in the ca_contact table. It indicates the person who is currently assigned to this workflow task. isswf.assignee = ca_contact.contact_uuid
done_by	isswf.done_by: This is a unique binary identifier, which references a record in the ca_contact table. It indicates the person who completed or approved this workflow task. isswf.done_by = ca_contact.contact_uuid
wf_start_date	wf_start_date: This is the timestamp when the workflow task moved into an active status (pdmtime).
wf_est_comp_date	isswf.est_comp_date: This is a timestamp (pdmtime) indicating when users believe this task will be completed.
est_duration	isswf.est_duration: This is the estimated duration for this workflow task.
completion_date	isswf.completion_date: This is a timestamp (pdmtime) indicating when this workflow task was completed.

Field	Remarks
actual_duration	isswf.actual_duration: This is the actual amount of time it took to complete this workflow task.
wf_est_cost	isswf.est_cost: This is the estimated cost of this workflow task
cost	isswf.cost: This is the actual cost required to complete this workflow task.
wf_description	isswf.description: This is a description of the workflow task.
wf_last_mod_dt	isswf.last_mod_dt: This is the timestamp (pdmtime) indicating when this workflow task was last changed.
wf_last_mod_by	isswf.last_mod_by: This is the unique binary identifier, which references a record in the contact table. It indicates the last person to make changes to this workflow task. isswf.last_mod_by = ca_contact.contact_uid

View_Issue_to_Properties

This view is a result of the View_Issue view joined with the issue properties table (issprp) to give a basic view of issues and their assigned properties. This may not list all issues, particularly if there are no properties assigned to them.

Field	Remarks
View_Issue.*	Please refer to the View_Issue definition earlier in this document, for a description of each field.
prp_id	issprp.id: This is an integer-unique identifier for the property record.
prp_persid	issprp.persid: This is a unique identifier for this record in the prp table, preceded by the object identifier (prp) and a colon.

Field	Remarks
sequence	issprp.sequence: This is an integer that indicates the order for which this particular property record should be displayed by CA SDM (for example, Ascending).
label	issprp.label: This is a short description of what should be placed in the issprp.value field.
value	issprp.value: This is a value entered by the user in response to the prp_description and issprp.label fields.
prp_last_mod_dt	issprp.last_mod_dt: This is the timestamp (pdmtime) indicating when this property was last modified.
prp_last_mod_by	issprp.last_mod_by: This is a binary identifier, which references a record in the ca_contact table. It represents the person who last modified this record. issprp.last_mod_by = ca_contact.contact_uuid
required	issprp.required: This is a boolean, indicating whether this property must have a issprp.value before the record is saved.
sample	issprp.sample: This is a text field, which displays example values to guide the user in entering the most useful value in issprp.value.
owning_iss	issprp.owning_iss: This is a unique identifier used to identify the specific record to which this property is attached. issprp.object_id = issue.persid
prp_description	issprp.description: This is a text field, which explains the type of value that should be entered in issprp.value.

View_Request

The following is a basic view of all requests. Here, the Request table has been joined with other CA SDM tables to give you more specific information, such as the request service type, severity, urgency, category, and priority. There is some additional information about the request listed, as well. All fields from the Request (call_req) table are selected. The extracted fields that are a result of the joined tables are listed at the end of this fields list.

Field	Remarks
id	call_req.id: This is the unique identifier for this record in the call_req table.

Field	Remarks
persid	call_req.persid: This is a unique identifier for this record in the call_req table, preceded by the object identifier (cr for table call_req) and a colon.
ref_num	call_req.ref_num: This is a Request reference number, which is used by analysts and customers to refer to a particular Request.
summary	call_req.summary: This is a brief description of the request for quick reference.
description	call_req.description: This is the long description of a request, as dictated by an analyst or customer.
status	call_req.status: This is a unique identifier referencing a record in the cr_stat table. It indicates the status of this request. call_req.status = cr_stat.code
active_flag	call_req.active_flag: This is an integer flag to determine whether this request record is active (1 or 0).
open_date	call_req.open_date: This is the Request creation timestamp (pdmtime).
time_spent_sum	call_req.time_spent_sum: This is the derived total of all of the act_log records' time_spent fields, stored in seconds (i.e. 80 = 1 minute 20 seconds).
last_mod_dt	call_req.last_mod_dt: This is the last modified timestamp (pdmtime).
close_date	call_req.close_date: This is the timestamp when the request was set to inactive (pdmtime).
resolve_date	This is the date for when the request was resolved (pdmtime).
rootcause	This is a pointer to the rootcause.id.
log_agent	call_req.log_agent: This is a binary-unique identifier, referencing the ca_contact table. It references the person who was the request's original creator. call_req.log_agent = ca_contact.contact_uuid

Field	Remarks
assignee	call_req.assignee: This is a binary-unique identifier, referencing a record in the ca_contact table. It represents the person currently assigned to the request. call_req.assignee = ca_contact.contact_uuid
group_id	call_req.group_id: This is a binary-unique identifier, referencing a record in the ca_contact table. It represents the group currently assigned to the request. call_req.group_id = ca_contact.contact_uuid
customer	call_req.customer: This is a binary-unique identifier, referencing a record in the ca_contact table. It represents the affected end user for this request. call_req.customer = ca_contact.contact_uuid
charge_back_id	charge_back_id: This is a text field available for use as an indicator of accounting jargon for expensing this request to the appropriate cost center.
affected_rc	call_req.affected_rc: This is a binary-unique identifier, referencing a row in the ca_owned_resource table. It represents the asset to which this request applies. call_req.affected_rc = ca_owned_resource.own_resource_uuid.
support_lev	call_req.support_lev: This is a pointer to a service desc record, which automates some constraints under which this request must be completed. call_req.support_lev = srv_desc.code.
category	call_req.category: This is a unique identifier, referencing a record in the prob_ctg table. It represents the category to which this request belongs. call_req.category = prob_ctg.persid
solution	call_req.solution: This is a pointer to call solution to get solution. call_req.solution = crsol.persid
impact	call_req.impact: This is an integer-unique identifier, referencing a row in the impact table. It indicates the impact scope of the request. call_req.impact = impact.enum
priority	call_req.priority: This is an integer-unique identifier, referencing a record in the pri table. It indicates how analysts will prioritize the work associated with this request. call_req.priority = pri.enum

Field	Remarks
urgency	call_req.urgency: This is an integer-unique identifier, referencing a row in the urgency table. It indicates the user's feeling of urgency for having this request resolved. call_req.urgency = urgncy.enum
severity	call_req.severity: This is an integer-unique identifier, referencing a row in the severity table. It indicates the severity of the consequences of this unresolved request. call_req.severity = sevrty.enum
extern_ref	This is an external reference to an associated ticket.
last_act_id	This identifies the id of the last activity.
cr_ticket	This is a pointer to the trouble ticket to get the associated ticket.
parent	call_req.parent: This is a ersid pointer to another request persid, which facilitates creation of a hierarchy of change orders. call_req.parent = call_req.persid
template_name	call_req.template_name: This is a text value, which indicates this request is designated for and can be chosen from a list as a template for other similar requests. cr_template.template = call_req.persid
sla_violation	call_req.sla_violation: This is an integer, which counts the number of times that slas attached to this request have been violated.
predicted_sla_viol	This specifies that a request has been predicted by neugents to likely violate SLA.
macro_predicted_violation	This indicates that the request has been predicted by neugents to likely violate SLA.
created_via	call_req.created_via: This is an integer pointer to a record in the interface table that indicates from which interface the change order originated. call_req.created_via = interface.id

Field	Remarks
call_back_date	call_req.call_back_date: This is a timestamp field (pdmtime), which indicates a future date/time that the affected_end_user is to be contacted.
call_back_flag	call_req.call_back_flag: This is a boolean indicator, displayed as a checkbox to the user, indicating whether to notify the analyst at the call_req.call_back_date.
event_token	call_req.event_token: This is used by CA NSM for message matching.
sched_token	call_req.sched_token: This is used by CA NSM for message matching.
type	call_req.type: This is a text field referencing a record in the crt table. It indicates the ITIL type for this request. call_req.type = crt.code
string1	This is a user-definable string.
string2	This is a user-definable string.
string3	This is a user-definable string.
string4	This is a user-definable string.
string5	This is a user-definable string.
string6	This is a user-definable string.
problem	This is an ITIL problem.
incident_priority	This is an ITIL incident priority.

Field	Remarks
change	call_req.change: This is an integer-unique identifier, referencing a row in the chg table. It indicates the change order that was created as a result of this request. call_req.change = chg.id.
service_type	srv_desc.sym: This indicates the actual Service Type. call_req.support_lev = srv_desc.code
severity_num	sevrty.sym: This is the actual Severity number. call_req.severity = sevrty.enum
urgncy_num	urgncy.sym: This indicates the actual Urgency number. call_req.urgency = urgncy.enum
category_name	prob_ctg.sym: This is the actual Request Area (problem category). call_req.category = prob_ctg.id
asset	ca_owned_resource.resource_name: This is the actual Asset name. call_req.affected_rc = ca_owned_resource.own_resource_uuid
impact_num	impact.sym: This is the actual Impact number. call_req.impact = impact.enum
assignee_lastname	ca_contact.last_name: This is the actual Assignee last name. call_req.assignee = ca_contact.contact_uuid
assignee_firstname	ca_contact.first_name: This is the actual Assignee first name. call_req.assignee = ca_contact.contact_uuid
assignee_middlename	ca_contact.middle_name: This is the actual Assignee middle name. call_req.assignee = ca_contact.contact_uuid
customer_lastname	ca_contact.last_name: This is the actual last name of the Affected End User. call_req.customer = ca_contact.contact_uuid
customer_firstname	ca_contact.first_name: This is the actual first name of the Affected End User. call_req.customer.ca_contact.contact_uuid

Field	Remarks
customer_middlename	ca_contact.middle_name: This is the actual middle name of the Affected End User. call_req.customer = ca_contact.contact_uuid
group_name	View_Group.last_name: This is the actual Group name.
GroupID	View_Group.contact_uuid: This is the actual Group key ID.
status_name	cr_stat.sym: This is the actual status.
priority_num	pri.sym: This is the actual Priority number.

View_Request_to_Act_Log

The following is a basic view of all requests with their activity logs. The View_Request view is joined with the View_Act_Log view to give more detailed information about each activity per request.

Field	Remarks
View_Request.*	Please refer to the field defined in the View_Request section earlier in this document.
View_Act_Log.*	Please refer to the fields defined in the View_Act_Log section earlier in this document.

View_Request_to_Properties

The following is a basic view of call requests and their properties. This view lists everything from the call request (call_req) table and the request property (cr_prp) table.

Field	Remarks
View_Request.*	Please refer to the fields defined in the View_Request section of this document.

Field	Remarks
crprp_id	cr_prp.id: This is an integer-unique identifier for the property record.
crprp_persid	cr_prp.persid: This is a unique identifier for this record in the cr_prp table, preceded by the object identifier (cr_prp) and a colon.
sequence	cr_prp.sequence: This is an integer that indicates the order for which this particular property record should be displayed by CA SDM (for example, Ascending).
label	cr_prp.label: This is a short description of what should be placed in the cr_prp.value field.
value	cr_prp.value: This is a value entered by the user in response to the prp_description and cr_prp.label fields.
crprp_last_mod_dt	cr_prp.last_mod_dt: This is the timestamp (pdmtime) identifying when this property was last modified.
crprp_last_mod_by	cr_prp.last_mod_by: This is a binary identifier, which references a record in the ca_contact table. It represents the person who last modified this record. cr_prp.last_mod_by = ca_contact.contact_uuid
required	cr_prp.required: This is the boolean, indicating whether this property must have a cr_prp.value before the record is saved.
sample	cr_prp.sample: This is a text field, which displays example values to guide the user in entering the most useful value in cr_prp.value.
owning_cr	cr_prp.owning_cr: This is the unique identifier used to identify the specific record to which this property is attached. cr_prp.object_id = call_req.persid
crprp_description	cr_prp.description: This is a text field, which explains the type of value that should be entered in cr_prp.value.

Appendix B: RFC 2251 LDAP Result Codes

This section contains the following topics:

[LDAP Return Codes](#) (see page 945)

[LDAP Server Return Codes](#) (see page 945)

[LDAP Client Return Codes](#) (see page 950)

[LDAP-Associated RFC Standards](#) (see page 952)

LDAP Return Codes

LDAP has a set of operation result codes that may be generated by the LDAP server in response to various LDAP requests. These codes indicate the status of the protocol operation and are categorized by server or client return code categories.

LDAP Server Return Codes

The following table lists the server return codes:

Hex	Decimal	Description
0x00	0	LDAP_SUCCESS Indicates the requested client operation completed successfully.
0x01	1	LDAP_OPERATIONS_ERROR Indicates an internal error occurred. The server is unable to respond with a more specific error and is also unable to properly respond to a request. It does not indicate that the client has sent an erroneous message.
0x02	2	LDAP_PROTOCOL_ERROR Indicates that the server has received an invalid or malformed request from the client.
0x03	3	LDAP_TIMELIMIT_EXCEEDED Indicates that the operation's time limit specified by either the client or the server has been exceeded. On search operations, incomplete results are returned.
0x04	4	LDAP_SIZELIMIT_EXCEEDED Indicates that in a search operation, the size limit specified by the client or the server has been exceeded. Incomplete results are returned.

Hex	Decimal	Description
0x05	5	LDAP_COMPARE_FALSE Does not indicate an error condition. Indicates that the results of a compare operation are false.
0x06	6	LDAP_COMPARE_TRUE Does not indicate an error condition. Indicates that the results of a compare operation are true.
0x07	7	LDAP_AUTH_METHOD_NOT_SUPPORTED Indicates that during a bind operation the client requested an authentication method not supported by the LDAP server.
0x08	8	LDAP_STRONG_AUTH_REQUIRED Indicates one of the following: ■ In bind requests, the LDAP server accepts only strong authentication. ■ In a client request, the client requested an operation, such as delete, that requires strong authentication. ■ In an unsolicited notice of disconnection, the LDAP server discovers the security protecting the communication between the client and server has unexpectedly failed or been compromised.
0x09	9	Reserved.
0x0A	10	LDAP_REFERRAL Does not indicate an error condition. In LDAPv3, indicates that the server does not hold the target entry of the request, but that the servers in the referral field may.
0x0B	11	LDAP_ADMINLIMIT_EXCEEDED Indicates that an LDAP server limit set by an administrative authority has been exceeded.
0x0C	12	LDAP_UNAVAILABLE_CRITICAL_EXTENSION Indicates that the LDAP server was unable to satisfy a request because one or more critical extensions were not available. Either the server does not support the control or the control is not appropriate for the operation type.
0x0D	13	LDAP_CONFIDENTIALITY_REQUIRED Indicates that the session is not protected by a protocol, such as Transport Layer Security (TLS), which provides session confidentiality.
0x0E	14	LDAP_SASL_BIND_IN_PROGRESS Does not indicate an error condition, but indicates that the server is ready for the next step in the process. The client must send the server the same SASL mechanism to continue the process.

Hex	Decimal	Description
0x0F	15	Not used.
0x10	16	LDAP_NO_SUCH_ATTRIBUTE Indicates that the attribute specified in the modify or compare operation does not exist in the entry.
0x11	17	LDAP_UNDEFINED_TYPE Indicates that the attribute specified in the modify or add operation does not exist in the LDAP server's schema.
0x12	18	LDAP_INAPPROPRIATE_MATCHING Indicates that the matching rule specified in the search filter does not match a rule defined for the attribute's syntax.
0x13	19	LDAP_CONSTRAINT_VIOLATION Indicates that the attribute value specified in a modify, add, or modify DN operation violates constraints placed on the attribute. The constraint can be one of size or content (string only, no binary).
0x14	20	LDAP_TYPE_OR_VALUE_EXISTS Indicates that the attribute value specified in a modify or add operation already exists as a value for that attribute.
0x15	21	LDAP_INVALID_SYNTAX Indicates that the attribute value specified in an add, compare, or modify operation is an unrecognized or invalid syntax for the attribute.
	22-31	Not used.
0x20	32	LDAP_NO_SUCH_OBJECT Indicates that the target object cannot be found. This code is not returned on the following operations: ■ Search operations that find the search base but cannot find any entries that match the search filter. ■ Bind operations
0x21	33	LDAP_ALIAS_PROBLEM Indicates that an error occurred when an alias was dereferenced.
0x22	34	LDAP_INVALID_DN_SYNTAX Indicates that the syntax of the DN is incorrect. However, if the DN syntax is correct, but the LDAP server's structure rules do not permit the operation, the server returns the following: LDAP_UNWILLING_TO_PERFORM

Hex	Decimal	Description
0x23	35	LDAP_IS_LEAF Indicates that the specified operation cannot be performed on a leaf entry. (This code is not currently in the LDAP specifications, but is reserved for this constant.)
0x24	36	LDAP_ALIAS_DEREF_PROBLEM Indicates that during a search operation, either the client does not have access rights to read the aliased object's name or dereferencing is not allowed.
	37-47	Not used.
0x30	48	LDAP_INAPPROPRIATE_AUTH Indicates that during a bind operation, the client is attempting to use an authentication method that the client cannot use correctly. For example, either of the following causes this error: ■ The client returns simple credentials when strong credentials are required. ■ The client returns a DN and a password for a simple bind when the entry does not have a password defined.
0x31	49	LDAP_INVALID_CREDENTIALS Indicates that during a bind operation, one of the following occurred: ■ The client passed either an incorrect DN or password. ■ The password is incorrect because it has expired; intruder detection has locked the account, or some other similar reason.
0x32	50	LDAP_INSUFFICIENT_ACCESS Indicates that the caller does not have sufficient rights to perform the requested operation.
0x33	51	LDAP_BUSY Indicates that the LDAP server is too busy to process the client request at this time, but if the client waits and resubmits the request, the server may be able to process it then.
0x34	52	LDAP_UNAVAILABLE Indicates that the LDAP server cannot process the client's bind request, usually because it is shutting down.

Hex	Decimal	Description
0x35	53	LDAP_UNWILLING_TO_PERFORM Indicates that the LDAP server cannot process the request because of server-defined restrictions. This error is returned for the following reasons: ■ The add entry request violates the server's structure rules. ■ The modify attribute request specifies attributes that users cannot modify. ■ Password restrictions prevent the action. ■ Connection restrictions prevent the action.
0x36	54	LDAP_LOOP_DETECT Indicates that the client discovered an alias or referral loop, and is thus unable to complete this request.
	55-63	Not used.
0x40	64	LDAP_NAMING_VIOLATION Indicates that the add or modify DN operation violates the schema's structure rules. For example: ■ The request places the entry subordinate to an alias. ■ The request places the entry subordinate to a container that is forbidden by the containment rules. ■ The RDN for the entry uses a forbidden attribute type.
0x41	65	LDAP_OBJECT_CLASS_VIOLATION Indicates that the add, modify, or modify DN operation violates the object class rules for the entry. For example, the following types of request return this error: ■ The add or modify operation tries to add an entry without a value for a required attribute. ■ The add or modify operation tries to add an entry with a value for an attribute which the class definition does not contain. ■ The modify operation tries to remove a required attribute without removing the auxiliary class that defines the attribute, as required.
0x42	66	LDAP_NOT_ALLOWED_ON_NONLEAF Indicates that the requested operation is permitted only on leaf entries. For example, the following types of requests return this error: ■ The client requests a delete operation on a parent entry. ■ The client request a modify DN operation on a parent entry.
0x43	67	LDAP_NOT_ALLOWED_ON_RDN Indicates that the modify operation attempted to remove an attribute value that forms the entry's relative distinguished name.

Hex	Decimal	Description
0x44	68	LDAP_ALREADY_EXISTS Indicates that the add operation attempted to add an entry that already exists, or that the modify operation attempted to rename an entry with the name of an entry that already exists.
0x45	69	LDAP_NO_OBJECT_CLASS_MODS Indicates that the modify operation attempted to modify the structure rules of an object class.
0x46	70	LDAP_RESULTS_TOO_LARGE Reserved for CLDAP.
0x47	71	LDAP_AFFECTS_MULTIPLE_DSAS Indicates that the modify DN operation moves the entry from one LDAP server to another and thus requires more than one LDAP server.
	72-79	Not used.
0x50	80	LDAP_OTHER Indicates an unknown error condition. This is the default value for NDS error codes which do not map to other LDAP error codes.

LDAP Client Return Codes

The following table lists the client return codes:

Hex	Decimal	Description
0x51	81	LDAP_SERVER_DOWN Indicates that the LDAP libraries cannot establish an initial connection with the LDAP server. Either the LDAP server is down, or the specified host name or port number is incorrect.
0x52	82	LDAP_LOCAL_ERROR Indicates that the LDAP client has an error. This is usually a failed dynamic memory allocation error.
0x53	83	LDAP_ENCODING_ERROR Indicates that the LDAP client encountered errors when encoding an LDAP request intended for the LDAP server.
0x54	84	LDAP_DECODING_ERROR Indicates that the LDAP client encountered errors when decoding an LDAP response from the LDAP server.

Hex	Decimal	Description
0x55	85	LDAP_TIMEOUT Indicates that the time limit of the LDAP client was exceeded while waiting for a result.
0x56	86	LDAP_AUTH_UNKNOWN Indicates that the ldap_bind or ldap_bind_s function was called with an unknown authentication method.
0x57	87	LDAP_FILTER_ERROR Indicates that the ldap_search function was called with an invalid search filter.
0x58	88	LDAP_USER_CANCELLED Indicates that the user cancelled the LDAP operation.
0x59	89	LDAP_PARAM_ERROR Indicates that an LDAP function was called with an invalid parameter value (for example, the ID parameter is NULL).
0x5A	90	LDAP_NO_MEMORY: Indicates that a dynamic memory allocation function failed when calling an LDAP function.
0B	91	LDAP_CONNECT_ERROR Indicates that the LDAP client has lost either its connection or cannot establish a connection to the LDAP server.
0x5C	92	LDAP_NOT_SUPPORTED Indicates that the client does not support the requested functionality. For example, if the LDAP client is established as an LDAPv2 client, the libraries set this error code when the client requests LDAPv3 functionality.
0x5D	93	LDAP_CONTROL_NOT_FOUND Indicates that the client requested a control that the libraries cannot find in the list of supported controls sent by the LDAP server.
0x5E	94	LDAP_NO_RESULTS_RETURNED Indicates that the LDAP server sent no results. When the ldap_parse_result function is called, no result code is included in the server's response.
0x5F	95	LDAP_MORE_RESULTS_TO_RETURN Indicates that more results are chained in the result message. The libraries set this code when the call to the ldap_parse_result function reveals that additional result codes are available.
0x60	96	LDAP_CLIENT_LOOP Indicates the LDAP libraries detected a loop. Usually, this happens when following referrals.

Hex	Decimal	Description
0x61	97	LDAP_REFERRAL_LIMIT_EXCEEDED Indicates that the referral exceeds the hop limit. The hop limit determines how many servers the client can hop through to retrieve data. For example, suppose the following conditions: ■ The hop limit is two. ■ The referral is to server D which can be contacted only through server B (1 hop) which contacts server C (2 hops) which contacts server D (3 hops) With these conditions, the hop limit is exceeded and the LDAP libraries set this code.

LDAP-Associated RFC Standards

The following table describes the LDAP-associated RFC standards available for your use:

RFC	Description
1274	The COSINE and Internet X.500 Schema
1275	Replication Requirements to provide an Internet Directory using X.500
1276	Replication and Distributed Operations extensions to provide an Internet Directory using X.500
1308	Executive Introduction to Directory Services Using the X.500 Protocol
1309	Technical Overview of Directory Services Using the X.500 Protocol
1430	A Strategic Plan for Deploying an Internet X.500 Directory Service
1488	The X.500 String Representation of Standard Attribute Syntaxes
1558	A String Representation of LDAP Search Filters
1617	Naming and Structuring Guidelines for X.500 Directory Pilots
1777	Lightweight Directory Access Protocol v2
1778	The String Representation of Standard Attribute Syntaxes
1779	A String Representation of Distinguished Names
1804	Schema Publishing in X.500 Directory
1823	The LDAP Application Program Interface
1959	An LDAP URL Format
1960	A String Representation of LDAP Search Filters
2044	UTF -8, a transformation format of Unicode and ISO 10646

RFC	Description
2164	Use of an X.500/LDAP Directory to support MIXER address mapping
2218	A Common Schema for the Internet White Pages Service
2247	Using Domains in LDAP/X.500 Distinguished Names
2251	Lightweight Directory Access Protocol (v3)
2252	Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions
2253	Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names
2254	The String Representation of LDAP Search Filters
2255	The LDAP URL Format
2256	A Summary of the X.500(96) User Schema for use with LDAPv3
2279	UTF-8, a transformation format of ISO 10646
2293	Representing Tables and Subtrees in the X.500 Directory
2294	Representing the O/R Address hierarchy in the X.500 Directory Information Tree
2307	An Approach for Using LDAP as a Network Information Service
2377	Naming Plan for Internet Directory-Enabled Applications
2531	Content Feature Schema for Internet Fax
2559	Internet X.509 Public Key Infrastructure Operational Protocols - LDAPv2
2587	Internet X.509 Public Key Infrastructure LDAPv2 Schema
2589	Lightweight Directory Access Protocol (v3): Extensions for Dynamic Directory Services
2596	Use of Language Codes in LDAP
2649	An LDAP Control and Schema for Holding Operation Signatures
2657	RFC 2657 - LDAPv2 Client vs. the Index Mesh
2696	LDAP Control Extension for Simple Paged Results Manipulation
2713	Schema for Representing Java(tm) Objects in an LDAP Directory
2714	Schema for Representing CORBA Object References in an LDAP Directory
2739	Calendar Attributes for vCard and LDAP
2798	Definition of the inetOrgPerson LDAP Object Class
2820	Access Control Requirements for LDAP
2829	Authentication Methods for LDAP
2830	Lightweight Directory Access Protocol (v3): Extension for Transport Layer Security
2849	The LDAP Data Interchange Format (LDIF) - Technical Specification

RFC	Description
2879	Content Feature Schema for Internet Fax (V2)
2891	LDAP Control Extension for Server Side Sorting of Search Results
3045	Storing Vendor Information in the LDAP root DSE
3062	LDAP Password Modify Extended Operation
3112	LDAP Authentication Password Schema
3296	Named Subordinate References in Lightweight Directory Access Protocol Directories
3377	Lightweight Directory Access Protocol (v3): Technical Specification
3384	Lightweight Directory Access Protocol (version 3) Replication Requirements

Appendix C: Reference Commands

This appendix provides detailed reference commands available for CA SDM.

This section contains the following topics:

[bop_sinfo--Display System Information](#) (see page 956)
[dbmonitor_nxd--Database Monitoring Daemon](#) (see page 957)
[pdm_backup--Write Database to ASCII File](#) (see page 958)
[pdm_cache_refresh--Refresh Database](#) (see page 960)
[pdm_configure--Open the Configuration Window](#) (see page 961)
[pdm_d_refresh--Start Failed Daemons](#) (see page 961)
[pdm_deref--Dereference ASCII Data](#) (see page 962)
[pdm_discimp -- Discovered Asset Import](#) (see page 966)
[pdm_discupd -- Discovered Asset Update](#) (see page 967)
[pdm_edit--Configure Server Processes](#) (see page 968)
[pdm_extract--Extract Data from Database](#) (see page 969)
[pdm_halt--Terminate Daemons or Stop Services](#) (see page 972)
[pdm_init--Start Daemons](#) (see page 973)
[pdm_key_refresh--Refresh Cached Key Information](#) (see page 973)
[pdm_lexutil--Modify CA SDM Lexicons](#) (see page 974)
[pdm_k_reindex--Knowledge Re-Index Utility](#) (see page 975)
[pdm_listconn--List Active Connections](#) (see page 978)
[pdm_load--Add, Update, and Delete Database Records](#) (see page 980)
[pdm_logfile--Change stdout Cutover Size](#) (see page 982)
[pdm_log4j_config Utility--Modify the log4j properties File](#) (see page 983)
[pdm_proctor_init--Start Proctor on Secondary Servers](#) (see page 987)
[pdm_replace--Replace a Database Table](#) (see page 988)
[pdm_rest_util--Manage the CA SDM RESTful Web Services Application](#) (see page 989)
[pdm_restore--Restore a Database](#) (see page 990)
[pdm_status--Show Status of Daemons or Processes](#) (see page 992)
[pdm_task--Set Environment Variables](#) (see page 992)
[pdm_text_cmd--Text API Command Line Interface](#) (see page 993)
[pdm_uconv--Convert Local Charset to UTF-8](#) (see page 996)
[pdm_userload--Add, Update, and Delete Database Records](#) (see page 998)
[pdm_webstat--Return Web Usage Statistics](#) (see page 1001)
[report--Generate Reports](#) (see page 1005)
[rpt_srv--Generate Reports](#) (see page 1006)
[uniconv--Start UNIX CA NSM Event Converter Daemon](#) (see page 1007)

bop_sinfo--Display System Information

The bop_sinfo utility displays information about a single Majic-defined object. You can run this utility on objects listed in the *Technical Reference Guide*.

Syntax

This command has the following format:

```
bop_sinfo [-s server] [-p] [-l] [-d] [-f] [-q] [-t] [-m] [-a] object [-h]
```

-s server

Specifies the server to query.

-p

Displays the producer information.

-l

Displays the domset list.

-d

Displays database object information, including the name and type of all attributes.

-f

Displays factory information, including the rel_attr and common_name.

-q

Displays the schema name of the associated table.

-t

Displays triggers on the object.

-m

Displays methods used by the object.

-a

Displays attribute details.

object

The name of the object to query

-h

Displays help on the utility.

Example: Display System Information for the dmn object

```
bop_sinfo -d dmn
```

Factory dmn

Attributes

id	INTEGER
producer_id	LOCAL STRING(20)
persistent_id	STRING(30)
sym	STRING(60) REQUIRED
delete_flag	SREL -> actbool.enum REQUIRED
desc	STRING(40)
tables	BREL <- dcon.dom_id {dom_id = ?}
last_mod	DATE
last_mod_by	SREL -> cnt.id
audit_userid	LOCAL SREL -> cnt.id

dbmonitor_nxd--Database Monitoring Daemon

The Database Monitoring daemon (dbmonitor_nxd) provides a mechanism to allow CA SDM cache of specific database tables to be refreshed when changes are made externally from CA SDM.

The main function of dbmonitor_nxd is to generate CHANGE notifications for changes in specified tables that did not occur through CA SDM. In order to perform this function, the monitor periodically queries the database, determines what was changed externally and then sends CHANGE notifications to the bpvirtb_nxd server. The bpvirtb_nxd server notifies all domsrvr servers of the change, which causes each domsrvr to update its cache of specific database objects and then notify all other processes that subscribe for changes in the specified tables.

This mechanism works well for the occasional external change in tables that are monitored. However, in cases where mass updates are made externally a storm of CHANGE notifications are broadcast which leads to many database queries from various CA SDM processes which significantly impacts CA SDM's performance.

In order to eliminate this impact on CA SDM performance, dbmonitor_nxd has been updated for this release of the product. The Monitor supports a command line interface that allows the user to start and stop the monitoring of specified tables.

Syntax

This command has the following format:

```
dbmonitor_nxd -c <command> -t <tables>
```

<command>

Enter start or stop.

<tables>

Specifies a table name or a comma delimited list of table names that must match one or more of the tables specified in the NX_DBMONITOR_TABLES environment variable.

Each request is sent to the dbmonitor_nxd daemon. The daemon takes the appropriate action and returns a message to the user indicating the action taken.

- If a start request is invoked for a table that is already started, no action is taken.
- If a stop request is made for a table that is already stopped, no action is taken.
- If monitoring is successfully stopped or started for a table, a log message is also written to stdlog.

Note: When the Monitor is paused for a table, all CA SDM processes that cache data from these tables may become out of date and no provision is made to update this cache.

For example, BOPLGIN caches Contact records (from the ca_contact and usp_contact tables) and this cache would not be updated if the Monitor was paused for the ca_contact table during the time external updates were loaded into the database. In the BOPLGIN case this will have little consequence because the essential Contact attributes cached in BOPLGIN are taken from the usp_contact table and not the ca_contact table.

Note: When the Monitor is paused for a table, Web Users will not be able to see changes in the table while viewing a detail form that were made externally while the Monitor was paused.

pdm_backup--Write Database to ASCII File

pdm_backup stops CA SDM and then writes one or more tables from a CA SDM database to an ASCII file. You can use this output file as the input file to pdm_restore. In addition to the contents of the database, pdm_backup backs up application configuration files.

If you have operating environment-specific backup tools, we recommend that you use them instead of pdm_backup. Because pdm_backup is a generic tool, it can be slow on some operating environment combinations.

Note: As part of its processing, pdm_backup first shuts down the daemons (UNIX) or services (Windows).

Syntax

This command has the following format:

```
pdm_backup [-d] [-g] [-v] -f filename [ALL | table1...tableN]
```

-d

Specifies to back up the database data only.

-g

Specifies that only non-database data be backed up. This means only windows (forms) and other non-database data is backed up.

-v

Specifies verbose mode, which writes comments about command progress to stdout.

-f filename

Specifies an output file.

ALL|table1...tableN

Specifies ALL files or the table or tables to write. If more than one table is specified, separate each with spaces.

- You can find table names in the CA SDM database schema file, `ddict.sch`, located in `$NX_ROOT/site` (UNIX) or `installation-directory\site` (Windows). `$NX_ROOT` or `installation-directory` is the directory where you installed CA SDM.
- If no table is specified, the entire CA SDM database is written, including window groups and menu registration files.

Restrictions

`pdm_backup` cannot be run while CA SDM is active.

Important! On UNIX, the `LIBPATH` must be set before running several CA SDM utilities. Use `pdm_task` to set the `LIBPATH` before running a utility. For example, input "`pdm_task pdm_clean_attachments ...`".

More information:

[pdm_restore--Restore a Database](#) (see page 990)

pdm_cache_refresh--Refresh Database

pdm_cache_refresh causes CA SDM daemons or processes to use data loaded into the database with pdm_userload or other database utilities, including non-CA SDM database tools.

Most CA SDM daemons and executables maintain a cache in dynamic memory of database records. This cache improves performance by eliminating the need to access the database when the required data is already available. CA SDM executables notify each other of database updates, so that the cache can always be kept up-to-date. However, there is no notification of updates from external utilities, such as pdm_userload, or third party database utilities. If new data is loaded into the database from one of these external sources, it is necessary to use the pdm_cache_refresh utility to notify executing modules that their cache must be refreshed from the database.

Syntax

This command has the following format:

```
pdm_cache_refresh [-f filename] [-t tablenameelist] [-d] [-v]
```

-f filename

Specifies a text file containing a list of database tables that have been modified externally. The text file consists of one or more lines, with each line containing one or more table names separated by spaces.

-t tablename

Specifies one or more tables that have been modified externally. If the list contains more than one table, the table names must be separated by semicolons, and the whole list must be enclosed in quotes. The tables are listed in the appendix "Data Element Dictionary" in the *Technical Reference Guide*.

For example, assume you have loaded location and site data into the CA SDM database with a third party utility. To tell CA SDM daemons to refresh their cache for these tables, issue the following command:

```
pdm_cache_refresh -t "Location;Site"
```

-d

Sends a message to the domsrvr, the CA SDM domain server. The domain server then reloads all selection lists, which causes any list windows displayed on a client to flash.

-v

Specifies verbose mode. The value 1 is the brief mode. The value 2 prints the progress messages to the log file.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

pdm_configure--Open the Configuration Window

pdm_configure opens a window containing the CA SDM configuration window. Use this window to set the CA SDM configuration after installing CA SDM, or to change the CA SDM configuration after CA SDM is running. Possible reasons to change the CA SDM configuration include the following:

- Changing the database type or server
- Reload default data
- Change passwords for CA SDM system accounts
- Rebuild schema files to incorporate changes
- Reconfigure as either Primary server or Secondary Server

Syntax

This command has the following format:

```
pdm_configure
```

Restrictions

pdm_configure can be run on a CA SDM server (Primary or Secondary) or a CA SDM Linux client. You must be the privileged user or root to run pdm_configure.

pdm_d_refresh--Start Failed Daemons

The pdm_d_refresh command line utility is used primarily for remote daemon configurations. It tells the daemon manager to try to start daemons that have failed to start ten times and that the daemon manager has flagged as "not runnable." This is often caused by starting the primary server when one of the following occurs:

- One or more of the secondary server(s) are not already started.
- A secondary server is stopped while the primary server is running.

The daemon manager tries to start remote daemons but if it cannot contact the secondary server, it tries ten times and stops. Running this utility flags all daemons as able to run and resets the restart counter. Then, the daemon manager tries to start all stopped daemons.

Syntax

This command has the following format:

```
pdm_d_refresh
```

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

pdm_deref--Dereference ASCII Data

pdm_deref processes ASCII-formatted input to exchange data found in one database table for data found in another database table. It can be used to create files compatible with pdm_userload from a non-CA SDM database or spreadsheet. It can also be used to create reports or output files for a non-CA SDM database or spreadsheet.

Important! Do not use pdm_deref if you are unfamiliar with SQL. The directories "export" and "import" in \$NX_ROOT/samples (UNIX) or installation-directory\samples (Windows) contain a standard set of specfiles for viewing.

Syntax

This command has the following format:

```
name pdm_deref -s specfile [-c|-e|-r] [-d] [-h] [-n] [-u] [-v] <filename
```

-s specfile

(Required) Specifies a file that defines which data is exchanged and the conditions under which it changes.

Specfile contains a list of SQL commands in the following format (note that "att" means attribute and "atts" means attributes):

Deref

```
{
input = <list of "from" atts from input file>
output = <list of "to" atts for output file>
rule = "SELECT <atts used to fill output atts> \
      FROM <table name> \
      WHERE <att from table name to match input 1> = ? \
      AND <att from table name to match input 2> = ? \
      OR <att from table name to match input 3> = ?"
}
```

-c

Produces comma-separated value (CSV) output, such as:

```
"field1","field2","field3"
```

The -c, -e, and -r output format options are mutually exclusive.

-e

Produces comma-separated value (CSV) output with embedded double quotes escaped by another double quote. For example:

```
"Text with a "quoted string" in it".
```

The -c, -e, and -r output format options are mutually exclusive.

-r

Produces left-justified output in the formats if the column labels are not supplied in the input file:

```
"label":"value"
```

or

```
"value"
```

This option is intended for use with line printers, for example:

```
Field_Name: Field Value
```

The -c, -e, and -r output format options are mutually exclusive.

-d

Produces diagnostic information.

-h

Displays help and usage information.

-n

Specifies that you do not want to treat 0 valued foreign keys as NULL. This argument should be used only under special circumstances when recovering a damaged database.

-u

Produces output without headers.

-v [1|2]

Specifies verbose mode. The value 1 is the brief mode. The value 2 prints the progress messages to the log file.

filename

(Optional) Specifies the ASCII-formatted input file to be processed. If omitted, stdin is used.

Restriction—Valid on UNIX only

Before using `pdm_deref` on UNIX, the `$NX_ROOT` environment variable must be set to the path of the CA SDM installation directory, and the `PATH` environment variable must include `$NX_ROOT/bin`.

Example: Using `pdm_deref` to Set Up Data for Input

This example shows how you can use `pdm_deref` to set up data for input.

Given the following data in the `ca_location` table:

```
id location_name_name
86873FA40BA4234A8CF7A418D7C8B2DB "Boulder NCC"
```

the following statement in the specfile:

```
Deref
{
input = place
output = location_uuid
rule = "SELECT id FROM ca_location WHERE location_name=?"
}
```

would change the following input:

```
last_name, first_name, place
{"Potts", "Elmore", "Boulder NCC"}
```

to the following output, which can be loaded into the ca.contact table with pdm_userload:

```
last_name, first_name, location_uuid
{"Potts", "Elmore", "86873FA40BA4234A8CF7A418D7C8B2DB"}
```

Example: Using pdm_deref to Set Up Data for Output

This example shows how you can use pdm_deref to set up data for output.

Given the following data in the ca_contact table:

```
id last_name first_name
"69499D5A2424884887E62EC9823F5E47" "Potts" "Elmore"
```

the following statement in the specfile:

```
Deref
{
  input = primary_contact_uuid
  output = firstname, lastname
  rule = "SELECT first_name, last_name FROM ca_contact
  WHERE id=?"
}
```

would change the following input:

```
location_name, primary_contact_uuid
{"Boulder NCC", "69499D5A2424884887E62EC9823F5E47"}
```

to the following output, which can be printed or sent to a spreadsheet:

```
location_name, firstname, lastname
{"Boulder NCC", "Elmore", "Potts"}
```

More information:

[pdm_extract--Extract Data from Database](#) (see page 969)

pdm_discimp -- Discovered Asset Import

Batch registration of non-CA SDM Discovered Assets. Use this to search the MDB for assets that were registered by other software products and register them as CA SDM assets, so they can be used in CA SDM.

Logic is similar to Discovered Assets dialog that can be launched from Asset Search/List web form. That is an interactive and batch process.

This program will query the ca_logical_asset, ca_asset, and ca_logical_asset_property tables, using various parameters, and attempt to register new CA SDM Assets from the discovered values.

Syntax

```
pdm_discimp [-l label] [-s serial number] [-t asset tag] [-n hostname] [-d dns name] [-m mac address] [-c asset class]
[-v] [-r] [-o object manager]
```

Asset selection criteria (use % for wild card):

l

Match this asset label.

s

Match this serial number.

t

Match this asset tag.

n

Match this hostname.

Asset property selection criteria (use % for wild card):

d

Match this dns name.

m

Match this mac address.

Other options:

c

Asset class to assign when registering new owned assets defaults to Discovered Hardware.

v

Verbose/diagnostic mode.

- r**
Register assets, otherwise runs in simulate mode.
- h**
Displays this information.
- o**
Object manager (domsrvr) to use for processing.

Note: If processing results in a blank Asset Label, the value found for the host name or DNS Name will be used as the Asset Label. Assets must have at least a Label and Asset Class to be registered for use in CA SDM.

Because of the structure of the MDB and CA SDM architecture limitations, two queries are performed to select the appropriate records to process. It is important to understand this as it could affect performance. The first query retrieves the rows from a join between the `ca_logical_asset` and `ca_asset` tables that match label, serial number, tag and hostname. Then for each resulting row, a query is performed against `ca_logical_asset_property` to match `dns_name` and `mac_address`. The asset from the first query is chosen for registration if the second query results in rows being returned.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use `pdm_task` to set the LIBPATH before running a utility. For example, input "`pdm_task pdm_clean_attachments ...`".

pdm_discupd -- Discovered Asset Update

Batch update of non-CA SDM Discovered Assets. Use this utility to update assets that were imported by the `pdm_discimp` command.

Syntax

```
pdm_discupd [-t] [-v] [-d domsrvr]
```

Where

- t**
Test
- v**
Verbose/diagnostic mode.

d

Object manager (domsrvr) to use for processing.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

pdm_edit--Configure Server Processes

pdm_edit is used to edit the startup template file, pdm_startup.tpl, on the CA SDM server. Some examples of when you need to edit the startup template file are to define object engine aliases, to set up the web engine to run on a secondary server, and to establish remote authentication on a secondary server.

For UNIX

UNIX users must change the permissions on the pdm_edit.pl file to "executable."

Syntax

On the CA SDM server installation, switch to the \$NX_ROOT/samples/pdmconf (UNIX) or *installation-directory*\samples\pdmconf (Windows) directory, and then enter the following command to start pdm_edit:

```
pdm_perl pdm_edit.pl
```

Note: pdm_edit is a Perl script, and we recommended that you use the version of Perl that comes with CA SDM, pdm_perl, to run it; however, other versions of Perl may also work.

Usage

After pdm_edit is running, you will be prompted for certain information, and the rest of the process is menu-driven. To get usage information, use the H menu option for help. You can also view the help in the file pdm_edit.README using any text editor.

Output

When you exit pdm_edit and save your changes, several files are created:

- pdm_startup.rmt stores your new configuration values
- alias_install.sh (UNIX) or alias_install.bat (Windows), if you defined Object Engine aliases
- pdm_startup.dat stores data that will be needed the next time you invoke pdm_edit--keep this file in the working directory with pdm_edit.pl

To put your changes into effect, follow these steps:

1. Create a backup of the pdm_startup.tpl that resides in the \$NX_ROOT/pdmconf (UNIX) or *installation-directory*\pdmconf (Windows) directory of the CA SDM primary server installation, and then use the newly created pdm_startup.rmt file to replace it.

Note: The pdm_startup.tpl file that you create from pdm_startup.rmt must be readable by the privileged user.

2. If you defined Object Engine aliases, run the alias_install script to reconfigure your primary server.
3. Run the configuration utility on the CA SDM primary server without making any changes. This will put your new configuration settings into effect the next time you start the CA SDM server.
4. If necessary, start all secondary CA SDM servers, and then start the primary server.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

More information:

[Start the Primary Server](#) (see page 51)

[Start a Secondary Server](#) (see page 50)

pdm_extract--Extract Data from Database

The pdm_extract command extracts data from specified CA SDM database tables or the entire CA SDM database, and outputs it as ASCII-formatted text.

Syntax

This command has the following format:

```
pdm_extract [-c|-e|-r] [-d] [-h] [-u] [-v] [-C] --B] [-f formatstring| ALL | table1 . . . TableN]
```

-c

Produces comma-separated value (CSV) output, such as:

```
"field1","field2","field3"
```

The -c, -e, and -r output format options are mutually exclusive.

-e

Produces comma-separated value (CSV) output with embedded double quotes escaped by another double quote. For example:

"Text with a "quoted string" in it"

The -c, -e, and -r output format options are mutually exclusive.

-r

Produces left-justified output in the formats if the column labels are not supplied in the input file:

"label": "value"

or

"value"

This option is intended for use with line printers, for example:

Field_Name: Field Value

The -c, -e, and -r output format options are mutually exclusive.

-d

Uses the date format found in the file \$NX_ROOT/fig/english/cfg/dataent.fmt (UNIX) or *installation-directory*\fig\english\cfg\dataent.fmt (Windows), which you can edit to suit your requirements.

-h

Displays help and usage information.

-u

Produces output without headers.

-v

Specifies verbose mode, which writes comments about command progress to stdout.

-C

Changes encoding from UTF-8 to another charset. The default output is UTF-8.

Example: To convert the output to JIS, you would run "-C iso-2022-jp"

Example: To encode to the operating system's native charset, use "DEFAULT" or "NATIVE".

-B

Suppresses the Byte Order Mark if the variable NX_ADD_UTF8_BYTE_ORDER_MARK is set.

The NX_ADD_UTF8_BYTE_ORDER_MARK option is a signature into a file. It allows editors that support UTF-8 to maintain the UTF-8 integrity of the file.

Note: This is only needed for non-ASCII data. If this is not installed, the default behavior omits the Byte Order Mark (BOM). If installed, set it to "1" or "Yes".

-f *formatstring*

Extracts specific records and fields according to *formatstring*, which is an SQL subset statement.

For a date after a period, use the following syntax:

```
pdm_extract -v -f "select id, ref_num from Call_Req where open_date >= DATE  
'2005-02-24'" > daterange1.txt
```

For a date range, use the following syntax:

```
pdm_extract -v -f "select id, ref_num from Call_Req where open_date >= DATE  
'2004-01-20' and open_date < DATE '2004-02-25'" > daterange2.txt
```

Note: Use single quotes around the date in the YYYY-MM-DD format.

The syntax for DATE is as follows:

DATE 'yyyymmdd'

yyyy = integer representing year (between 1970 and 2038)

mm = integer representing month

dd = integer representing day

Examples:

DATE '2005-01-18'

DATE '1999-12-25'

The syntax for TIMESTAMP is:

TIMESTAMP 'yyyymmdd hh.mm.ss[.nnnnnn][+|-][hh.mm]

yyyy = integer representing year (between 1970 and 2038)

mm = integer representing month

dd = integer representing day

hh = integer representing hour

mm = integer representing minutes

ss = integer representing seconds

nnnn = optional integer representing fractions of sec.

[+|-][hh.mm] = optional time zone interval.

Examples:

```
TIMESTAMP '1998-04-28 12:00:00.000000'  
TIMESTAMP '2004-10-17 18:30:45'  
TIMESTAMP '2005-03-21 12:00:12+08:00'  
TIMESTAMP '1999-05-10 09:12:23.005-03:30'
```

Note: The -d option is not needed, as it only affects the format of the output.

A command usage example follows:

```
pdm_extract -f "select * from Call_Req where open_date > TIMESTAMP '2004-01-12 12:00:00'"
```

In this example, all columns are being extracted from the Call_Req table where the open_date is after midnight 1/12/2004.

ALL

Extracts output from all tables in the database.

```
table1 . . . tableN
```

Extracts output from the specified tables. Table names must be separated by spaces.

The default format, if none is specified, is an ASCII file compatible with pdm_userload.

More information:

[pdm_deref--Dereference ASCII Data](#) (see page 962)

[pdm_replace--Replace a Database Table](#) (see page 988)

[pdm_userload--Add, Update, and Delete Database Records](#) (see page 998)

[rpt_srv--Generate Reports](#) (see page 1006)

pdm_halt--Terminate Daemons or Stop Services

pdm_halt cleanly terminates all CA SDM daemons (UNIX) or the System Server Service (Windows) on the system from which pdm_halt is executed. The pdm_halt utility usually takes about 30 seconds to complete. If it hangs for more than two minutes, press Ctrl+C to stop pdm_halt, and try again.

Syntax

This command has the following format:

```
pdm_halt [-w] [-a] [time]
```

-w

Waits for the daemons to stop.

-a

Stops all proctors defined in the pdm_startup file.

[time]

Specifies the number of seconds to wait until the command executes.

Restrictions

pdm_halt can be run on a CA SDM server or a CA SDM UNIX client. You must be the privileged user to run pdm_halt.

pdm_init--Start Daemons

Applies to UNIX only

pdm_init starts all CA SDM automatic processes on the system from which pdm_init is executed. These automatic processes are called *daemons* and run continually in the background while you work. Not all of the daemons are started; some are applicable only to certain operating systems. See the chapter "Managing Servers" for a list of daemons that this command can start.

Note: Use pdm_d_refresh to start daemons that failed to start the first time after remedying the problem that caused them not to start initially. In most cases you do not need to terminate the daemons running to start ones that initially failed.

Syntax

pdm_init

Restrictions

pdm_init can be run on a CA SDM server or a CA SDM UNIX client. You must be the privileged user to run pdm_init.

pdm_key_refresh--Refresh Cached Key Information

The pdm_key_refresh utility refreshes cached key information from the key control table in the database. If key control information is updated, this utility can be executed instead of stopping and restarting the system to force CA SDM to use the new keyid base value.

Important! Changing the key control table may cause data corruption. You should not attempt to change the key control without specific instructions from CA Technical Support.

Syntax

This command has the following format:

```
pdm_key_refresh
```

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

pdm_lexutil--Modify CA SDM Lexicons

The pdm_lexutil utility allows you to modify Service Desk lexicons, to add or delete words for the spell check dictionary.

Syntax

This command has the following format:

```
pdm_lexutil -a | -d [-f] [-l] wordlist
```

-a

Add words.

-d

Delete words.

-f

File or lexicon containing list of words to be added or deleted.

-l

Lexicon name.

Default: userdict.tlx

wordlist

Words to be added or deleted.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

pdm_k_reindex—Knowledge Re-Index Utility

The Knowledge Re-Index utility, `pdm_k_reindex.exe`, is located under the Knowledge Management installation directory.

Note: Re-indexing the documents in the knowledge base can be a time-consuming operation, depending upon the size of your database. We recommend that you run the Knowledge Re-Index utility after all the changes have been added.

To run Knowledge Re-Index, enter the following command at the command prompt:

```
pdm_k_reindex
```

Following are the options available with this command.

Interface:

-D

Defines the debug mode, such as printing to the command window.

-v

Defines the verbose mode, such as printing to the stdlog file.

-i

Does not create table indexes in the re-index table after re-indexing.

Note: Parameters with dash as a prefix, such as “-D”, should precede other parameters that do not have this prefix.

file:reindex.txt

Documents are re-indexed to the appropriate file.

+i

Creates indexes of the re-indexed table only, which is the search table after re-indexing. Old indexes are dropped before re-indexing.

+t

Switches names of search and re-index tables only.

Note: A “+” prefix denotes only this parameter applies.

sdtout

Defines the frequency of statistic appearing in the command window. By default the Knowledge Re-index utility provides statistics into the command window for every 1000 documents processed. However, sometimes statistics are required to be provided more often. Use the following parameter:

```
pdm_k_reindex -i sdtout:10
```

In this case, statistics display in the command window for every ten documents.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

When to Use pdm_k_reindex

Run the pdm_k_reindex utility when one or more of the following search settings were changed:

- Noise words
- Synonyms
- Special terms
- Language
- Remove Similar Words
- Remove Noise Words
- Valid Character Range
- Recognize Special Terms

An appropriate message appears on Knowledge node of Administration tab when a change occurs.

Re-Index Tracking

While the re-index is running, you can view the status of the process in the Re-Index Tracking section in the lower half of the page. Each field is described as follows:

Document #

Specifies the number of documents already processed.

Average Size (Words)

Specifies the size of the current documents, calculated by the number of words minus the number of noise words.

Rate (Docs/Sec)

Specifies the number of documents processed, per second.

Time from Start

Indicates the duration of the re-index process since start.

Time Remaining

Specifies the estimated amount of time remaining for the process, based on the current rate and the remaining number of documents.

Failures #

Represents the number of failed documents (maximum=100). When the maximum number of failures is reached, the administrator is prompted to either continue or cancel the process.

Note: If changes have been made to Noise Words, Special Terms, Synonyms or Parse Settings and you do not re-index, you are prompted the next time you access the Knowledge node of the Administration tab. Changes will take effect only after the Knowledge Re-index utility is run.

Import and Re-Indexing

When pdm_kit.exe is invoked from the command line, the pdm_kit utility imports the documents into the database. After pdm_kit is completed, assuming document indexing or re-indexing has not been disabled using the command-line options, the re-index utility (pdm_k_reindex.exe) is automatically invoked. The status and output of the re-indexing operation is automatically written to the "EBR_REINDEX.LOG" in the \$NX_ROOT\log directory.

Index and De-Index Queue Settings for Batch and Instant Processing

Indexing and De-Indexing both run a batch process to include a predefined number of documents at one run. These batch processes are used for performance optimization. If more documents are included in the batch, system performance increases.

The number of documents you can process is limited; the limit depends on the size of the documents and the linked attachments. The document size is calculated based on the pure text and its attachments. Image and format elements are not calculated.

Note: You can limit the size of attachments by navigating to Attachments Library, Repositories on the Administration tab and editing the repository to set the File Limit Size (KB).

The recommended batch maximum size is between 2-12 MB (per the EBR_MAX_INDEX_BATCH_SIZE parameter of the NX.env file and the average document size).

- If the average size of your document (including attachments) is approximately 0.1 MB, keep the default setting in NX.env:

```
@EBR_MAX_INDEX_BATCH_SIZE=128
@NX_EBR_INDEX_QUEUE_TIMEOUT=10
@NX_EBR_REINDEX_QUEUE_TIMEOUT=1
@NX_EBR_INDEX_QUEUE_ONLINE=Yes
@NX_EBR_NON_KD_INDEX_QUEUE_ONLINE=Yes
```

This setting means that one batch processes 128 documents, that batch executions have ten second intervals, and when in reindex, the wait interval between two batches is one second.

- If the average size of your document (including attachments) is approximately 0.5 MB, keep the default setting in NX.env

```
@EBR_MAX_INDEX_BATCH_SIZE=25
@NX_EBR_INDEX_QUEUE_TIMEOUT=10
@NX_EBR_REINDEX_QUEUE_TIMEOUT=10
@NX_EBR_INDEX_QUEUE_ONLINE=No
@NX_EBR_NON_KD_INDEX_QUEUE_ONLINE=No
```

This setting means that one batch processes 25 documents, that batch executions have ten second intervals, and when in reindex, the wait interval between two batches is ten seconds.

pdm_listconn--List Active Connections

The pdm_listconn utility can be used to list active connections for both clients and servers.

Syntax

This command has the following format:

```
pdm_listconn [-c] [-s] [-s -c] [-t nn] [proc1 [proc2...]]
```

The default command has the following format if no parameters are specified:

```
pdm_listconn -s -t 2
```

-c

Lists connections by client. The utility displays two lines for each client:

```
(n secs) client_type node|cproc
connected to alias (sproc) at time
```

n is the number of seconds it took for the client to respond.

client_type is "vbop," "animator daemon," or "web engine."

node is the IP address of the client's node.

cproc is the client's slump procname.

alias is the alias of the server the client is connected with.

sproc is the server's slump procname.

time is the connection time.

For example:

```
(0 secs) vbop client 141.202.211.34|vbop-0x40120000:anthill:0
        connected to CMD40120 on anthill (domsrvr) at 02/19/1999 10:44:16
```

-s

Lists connections by server (default). The utility displays two lines for each server:

```
(n secs) server alias (node|sproc) willingness willingness
count connected clients (use pdm_listconn -c -s to list clients by server)
```

n is the number of seconds it took for the server to respond.

node is the IP address of the server's node.

alias is the alias of the server the client is connected with.

sproc is the server's slump procname.

willingness is the server's willingness to accept new clients (0 - 100).

count is the number of connected clients.

For example:

```
(0 secs) server CMD40120 on anthill (domsrvr) willingness 98
2 connected clients (use pdm_listconn -c -s to list clients by server)
  vbop client 141.202.211.34|vbop2 connected 02/19/1999 10:53:16
  vbop client 141.202.211.34|vbop-0x40120000:anthill:0 connected 02/19/1999 10:44:17
```

-s -c

Lists connections by server, including client detail for each server. The utility displays several lines for each server:

```
(n secs) server alias (node|sproc) willingness willingness
count connected clients:
client_type node|cproc connected time
```

where:

n is the number of seconds it took for the server to respond.

node is the IP address of the server's node.

alias is the alias of the server the client is connected with.

sproc is the server's slump procname.

willingness is the server's willingness to accept new clients (0 - 100).

count is the number of connected clients.

client_type is "vbop," "animator daemon," or "web engine."

node is the IP address of the client's node.

cproc is the client's slump procname.

time is the connection time.

For example:

(0 secs) server CMD40120 on anthill (domsrvr) willingness 98

2 connected clients:

vbop client 141.202.211.34|vbop2 connected 02/19/1999 10:53:16

vbop client 141.202.211.34|vbop-0x40120000:anthill:0 connected 02/19/1999 10:44:17

-t nn

Specifies a timeout interval in seconds. Because pdm_listconn receives information from an unknown number of clients and servers, it terminates when the specified timeout interval elapses after the last message received.

Default: 2

proc1

Specifies one or more slump procnames, separated by spaces. The utility displays client or server information from them, as appropriate.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

pdm_load--Add, Update, and Delete Database Records

Important! Using pdm_load can be destructive so always back up your database before you perform a pdm_load, and use pdm_userload unless instructed to use pdm_load.

pdm_load updates a CA SDM database using an input file you specify, up to a maximum of 112 attributes.

Whenever you upload tickets (such as requests or issues), your ticket number should include a unique prefix or suffix in its string. CA SDM views this number as a string of characters not as a sequential number, and thus cannot guarantee that it will assign a unique number to the uploaded tickets. As long as you assign a unique prefix or suffix using awk or another text processor, you can upload tickets without CA SDM writing over previously assigned numbers.

Syntax

This command has the following format:

```
pdm_load [-c] [-h] [-m] [-r] [-u] [-v] -f filename
```

The input file entries follow this format:

```
TABLE table_name
fieldname1 fieldname2 . . . fieldnameN
{ "value11", "value12", . . . "value1N" }
{ "value21", "value22", . . . "value2N" }
.
.
.
{ "valueN1", "valueN2", . . . "valueNN" }
```

table_name is the name of the table to be loaded, as listed in the CA SDM database schema file, which is located in \$NX_ROOT/site/schema.sch (UNIX) or *installation-directory*\site\schema.sch (Windows). \$NX_ROOT or *installation-directory* is the directory where you installed CA SDM.

-f *filename*

Specifies an input ASCII file.

-c

Checks the input file against the database and reports on the updates that would be made, but does not make the updates.

-m

Specifies mass update. Specify when you are using pdm_load to add or delete a large number of records. This option suppresses all client notifications of updates and sends a cache refresh message for a table when pdm_load finishes processing the table.

-r

Removes database records that match input records.

Important! Make a backup copy of the database before running `pdm_load` with this option. After old database records are removed, you must restore the CA SDM database with this backup copy if you want to recover any deleted records.

-u

Updates existing records, but does not insert new records into the database.

More information:

[pdm_backup--Write Database to ASCII File](#) (see page 958)

[pdm_replace--Replace a Database Table](#) (see page 988)

[pdm_restore--Restore a Database](#) (see page 990)

[pdm_userload--Add, Update, and Delete Database Records](#) (see page 998)

pdm_logfile--Change stdlog Cutover Size

`pdm_logfile` lets you change your `stdlog.x` cutover size. The cutover can occur after a specified number of bytes are written. On UNIX, this value is reset with each `pdm_init`. On Windows, the settings are retained with each `pdm_halt` and `pdm_init`.

Syntax

This command has the following format:

```
pdm_logfile [-L|-h]
```

or

```
pdm_logfile [-g -h] [-b bytes]
```

Example

To change your `stdlog.x` files to cutover at 500,000 bytes, issue the following command:

```
pdm_logfile -f STD -b 500000
```

-L

Creates a listing of current cutovers.

-q

Runs `pdm_logfile` in quiet mode.

-b bytes

Specifies the number of bytes written before cutover occurs.

Restrictions

You can run `pdm_load` while CA SDM is active, but performance can become very slow. It is best to run `pdm_load` when no one is using CA SDM.

Important! On UNIX, the `LIBPATH` must be set before running several CA SDM utilities. Use `pdm_task` to set the `LIBPATH` before running a utility. For example, input "`pdm_task pdm_clean_attachments ...`".

pdm_log4j_config Utility--Modify the log4j properties File

The `pdm_log4j_config.pl` utility lets you configure the log4j properties file of CA SDM, web components, PDM_RPC, Support Automation, Rest, and CMDB Visualizer. Execute the utility batch script that is based on your environment. For Windows, execute `pdm_log4j_config` from the command line. For UNIX, execute the `pdm_log4j_config.sh` file.

This command has the following format:

```
pdm_log4j_config -f <component> -d
pdm_log4j_config -h
pdm_log4j_config -f <component> [-a | -n <name>] [-l <log level>] [-l <max # of log files>] [-s <max size of log files>]
[-t <log level threshold>]
```

-f

Specifies the log4j configuration of CA SDM or the component of CA SDM that you want to change. Enter one of the following values:

SDM_WEB, SDM_RPC, REST, SA, or Viz.

Note: Use the mandatory option along with the other options.

-d

Displays the current log4j.properties configuration.

-h

Displays help for the utility.

-a

Completes all changes to log4j.properties globally.

-n

Specifies that you only want to modify a specific class or package name.

Specify a specific class name, such as bop_logging, or a complete package name, such as com.ca.ServicePlus.

-l

Specifies the log level that you want to set.

Note: Specify the -a or -n option.

-i

Specifies the max file number index that you want to set.

Note: Specify the -a or -n option.

-s

Specifies the max file size that you want to set.

Note: Specify the -a or -n option.

Important! Change the appender in the log4j.properties file of Visualizer to Rolling File Appender before you execute the command with this parameter. If you do not change the appender, MaxFileSize generates logs in the same file.

-t

Specifies the log level threshold.

Note: Specify the -a or -n option.

Utility Usage Examples

The following list provides examples of using the pdm_log4j_config utility:

- Modify the log verbosity level for all loggers configured in the properties file by using the -l and -a variables.

For example, set all of the loggers configured in Support Automation to a level of DEBUG:

```
pdm_log4j_config -f SA -a -l DEBUG
```

- Modify the log verbosity level for a specific logger class or package name in the CA SDM log4j properties file by using the -l and -n variables.

For example, set the logger for the pdm_rpc package to DEBUG using one of the following code samples:

```
pdm_log4j_config -f SDM_RPC -n pdm_rpc -l DEBUG
```

```
pdm_log4j_config -f SDM_RPC -n com.ca.ServicePlus.pdm_rpc -l DEBUG
```


- Modify the maximum number of log files to create for all the appenders (MaxBackupIndex property) by using the -i and -a variables in the log4j properties file of REST.

For example, set the maximum number of files for all appenders to 9.

```
pdm_log4j_config -f REST -a -i 9
```

- Modify the maximum number of log files configured in the CA SDM log4j properties file to create for an appender for a specific class or set of classes (MaxBackupIndex property) by using the -i and -n variables.

For example, set the maximum number of files for bop_logging to 7.

```
pdm_log4j_config -f SDM_WEB -n bop_logging -i 7
```

- Modify the maximum size of each log file configured in the REST log4j properties file to create for any appender (MaxFileSize property) by using the -s and -a variables.

For example, set the maximum size of files for all appenders to 9 MB.

```
pdm_log4j_config -f REST -a -s 9MB
```

- Modify the maximum size of each log file configured in the CA SDM log4j properties file to create for an appender for a specific class or set of classes (MaxFileSize property) by using the -s and -n variables.

For example, set the maximum size of files for bop_logging to 7 MB.

```
pdm_log4j_config -f SDM_WEB -n bop_logging -s 7MB
```

- Modify the log level threshold for all the appenders configured in the Support Automation log4j properties file (Threshold property) by using the -t and -a variables.

For example, set the log level threshold to DEBUG.

```
pdm_log4j_config -f SA -a -t DEBUG
```

Note: The -t parameter log level threshold overrides the -l parameter log level. If you modify the log level and the threshold level, the DEBUG logs from the servlet do not appear in the file.

- Modify the log level threshold for an appender for a specific class or set of classes (Threshold property) configured in the CA SDM log4j properties file by using the -t and -n variables.

For example, set the log level threshold to WARN.

```
pdm_log4j_config -f SDM_WEB -n bop_logging -t WARN
```

- Execute the following command to view the current logger and appender configuration of the REST log4j properties file:

```
pdm_log4j_config -f REST -d
```

Important! Use -l, -i, -s, and -t variables together with one of the -a or -n options. Do not use both the options. The -f option is mandatory. The -h and -d options are mutually exclusive to any other option.

Modify the Log File Refresh Interval Manually

Administrators can modify how often CA SDM monitors the log4j.properties file for changes. By default, the refresh interval is set to 60 seconds. CA SDM components including SDM Servlets, PDM_RPC, Support Automation, CMDB Visualizer, and REST use log4j for logging.

Follow these steps:

1. Open the following directory on the CA SDM server:
NX_ROOT
2. Open the NX.env file for editing.
3. Modify the NX_LOG4J_REFRESH_INTERVAL variable with a value in milliseconds.
Note: If you enter a negative or nonnumeric value, the value defaults to 60 seconds.
4. Save the NX.env file.

Modify the jsrvr.log Appender

By default, servlets such as PDMContextListener, pdmweb, UploadServlet, and pdm_report log INFO level messages to the jsrvr.log file. You change the threshold level of the jsrvr.log appender to log any messages under the INFO level.

Follow these steps:

1. Modify the level in the log4j.properties file to the following threshold:

```
log4j.appender.jsrvrlog.Threshold=debug
```

2. Modify the log level of UploadServlet:

```
log4j.logger.com.ca.ServicePlus.uploadservlet=debug, jsrvrlog
```

3. Open the jsrvr.log file.
4. Confirm that the DEBUG log messages of UploadServlet appear.

Note: If you modify the log level without modifying the threshold level, the DEBUG logs from the servlet do not appear in the file. Not all servlets have explicit loggers attached. For example, the log4j.properties file does include pdmweb, BOServlet, pdm_export, pdm_report, and pdm_cache, which are part of the pdmweb servlet. To see DEBUG logs from these servlets, modify the pdmweb log level.

Modify the jstd.log Appender

All logs from nonwebapp applications dump into the jstd.log file separately. You can display logs for any of these applications, such as pdm_rpc by changing the log level of that specific application.

Follow these steps:

1. Modify the following log level:

```
log4j.logger.com.ca.ServicePlus.pdm_rpc=debug
```

2. Open log4j.properties and confirm that the log entries appear.

pdm_proctor_init--Start Proctor on Secondary Servers

Applies to UNIX only

Use pdm_proctor_init to start the proctor on secondary servers. All secondary servers should be started prior to starting the daemons from the primary server. After all daemons have been stopped from the primary server, use pdm_halt on the secondary server to stop this proctor.

Note: Do not use pdm_proctor_init on the primary server.

Syntax

This command has the following format:

```
pdm_proctor_init
```

pdm_replace--Replace a Database Table

pdm_replace deletes a table in a CA SDM database and replaces it with a table from a temporary file you specify with the -f option; the data from the input file is the only data that is in that table after running pdm_replace. Back up your table before running pdm_replace.

Note: As part of its processing, pdm_replace first shuts down the daemons (UNIX) or services (Windows).

pdm_replace accepts a text file as input, which is the same file format used by pdm_userload. You can create an input file for pdm_replace using pdm_extract; however, you cannot use the output of pdm_backup as input to pdm_replace.

Important! Be sure to name your input file with a name different from the table name you are attempting to replace. For example, if you are replacing a table named ca_contacts and you name the input file ca_contacts.dat, after you execute the pdm_replace command to point to the input file (ca_contacts.dat), it deletes the file after execution because it has the same name as the table.

Restrictions

- pdm_replace can be run only on a CA SDM server.
- Only the privileged user or root can run pdm_replace.
- Do not run pdm_replace when users are logged in to CA SDM.

Syntax

This command has the following format:

```
pdm_replace [-v] -f filename
```

-v

Specifies verbose mode.

-f filename

Specifies an ASCII file with the following format:

```
TABLE table_name  
  
fieldname1 fieldname2 . . . fieldnameN  
  
{ "value11", "value12", . . . "value1N" }  
  
{ "value21", "value22", . . . "value2N" }  
  
.  
  
.  
  
.  
  
{ "valueN1", "valueN2", . . . "valueNN" }
```

This format is the same file format used by `pdm_userload`. You can create an input file for `pdm_replace` using `pdm_extract`; however, you cannot use the output of `pdm_backup` as input to `pdm_replace`.

More information:

[pdm_extract--Extract Data from Database](#) (see page 969)

[pdm_userload--Add, Update, and Delete Database Records](#) (see page 998)

pdm_rest_util--Manage the CA SDM RESTful Web Services Application

CA SDM uses this utility automatically. You can run it manually if you require the utility, such as after an unexpected error occurs. This REST web services utility deploys the REST services web application. This utility deploys the REST application to the dedicated REST Tomcat instance. A batch file in the `NX_ROOT\bin` directory (`pdm_rest_util.bat` for Windows or `./pdm_rest_util.sh` for UNIX) lets you invoke the utility.

This command has the following formats and options:

```
pdm_rest_util -h | [-deploy] | [-undeploy]
```

-h

Prints command-line help.

-deploy

Generates, compiles, and deploys all Majic factories.

-undeploy

Undeploys REST Web Services on the local server.

Undeploy the REST Web Services Application

You can undeploy the REST Web Services application with the `pdm_rest_util` utility. For example, you want to perform CA SDM maintenance and prefer to undeploy REST during this operation.

Follow these steps:

1. Open a command prompt.
2. Execute the following command:

```
pdm_rest_util -undeploy
```

The REST Web Services application is undeployed.

Important! If you undeploy the REST application with this utility, REST can redeploy automatically when CA SDM restarts. We recommend that you use the CA SDM configuration to disable REST indefinitely.

pdm_restore--Restore a Database

`pdm_restore` stops CA SDM and then deletes all records from a CA SDM database and replaces them with records from a file you specify with the `-f` option. The data from the input file is the only data that will be in the CA SDM database after running `pdm_restore`.

The input file must be created using `pdm_extract` or `pdm_backup`, or otherwise formatted for `pdm_restore`. `pdm_backup` can back up non-database data, and `pdm_restore` can restore this data also. `pdm_backup` and `pdm_restore` are not recommended when other backup and restoration tools are available.

Note: As part of its processing, `pdm_restore` first shuts down the daemons (UNIX) or services (Windows).

Syntax

This command has the following format:

```
pdm_restore [-d] [-g] [-n] [-w] [-v] -f filename
```

Restrictions

`pdm_restore` can be run only on a CA SDM server. Only the privileged user or root can run `pdm_restore`.

Important! Use `pdm_restore` only with a full database backup created by `pdm_backup`, because your current database is deleted and replaced by the backup file. Do not run `pdm_restore` when users are logged in to CA SDM.

-d

Specifies that only database data is restored.

-g

Specifies that only non-database data be restored. Only windows (forms) and other non-database data are restored.

-n

Specifies that `NX.env` is restored if restoring non-database data. By default, `NX.env` is not restored. We recommend that the `NX.env` file not be restored unless the restore is to the same server the backup came from. Restoring an incorrect `NX.env` can affect unintended databases.

-w

Specifies that `web.cfg` is restored if restoring non-database data. By default, `web.cfg` is not restored.

-v

Specifies verbose mode.

-f filename

Specifies an input file that contains a full backup created by `pdm_backup`.

Important! On UNIX, the `LIBPATH` must be set before running several CA SDM utilities. Use `pdm_task` to set the `LIBPATH` before running a utility. For example, input "`pdm_task pdm_clean_attachments ...`".

More information:

[Database Restore](#) (see page 405)

[pdm_backup--Write Database to ASCII File](#) (see page 958)

[pdm_userload--Add, Update, and Delete Database Records](#) (see page 998)

pdm_status--Show Status of Daemons or Processes

pdm_status shows the status of all CA SDM daemons (UNIX) or processes (Windows) on the system from which the command is executed.

Output is displayed in this format:

DAEMON		STATUS	HOST	PID	SLUMP	CONNECT TIME
Agent anthill		Running	anthill	455	Tue Feb 17 17:55:12	
Ddict_rd	(ddictrd)	Completed	anthill			
Data Dictionary	(ddictbuild)	Completed	anthill			
...						
User Validation	(boplgln)	Running	anthill	456	Tue Feb 17 17:55:21	

Syntax

This command has the following format:

pdm_status

pdm_task--Set Environment Variables

Applies to UNIX only

The pdm_task utility sets environment variables for commands that do not have wrappers. For example, the pdm_task command must precede the report command on the same command line **only** when the command is invoked through a script or the command line. If you issue the report command from a menu, you do not need to include pdm_task, because all environment variables are set by the application.

Note: Reports cannot be generated from the command line on a client.

Syntax

This command has the following format:

pdm_task *command*

command

Specifies a command that does not have a wrapper and, therefore, does not automatically set environment variables. See [report--Generate Reports](#) (see page 1005) for more information about issuing pdm_task with a command.

pdm_text_cmd--Text API Command Line Interface

Use the `pdm_text_cmd` command for the Text API, which you can use to create and update various objects such as requests, change orders, issues, assets, and contacts.

Important! You cannot use a single or double quote as the parameter of the `pdm_text_nxd` or `bop_cmd` commands.

Syntax

This command has the following format:

```
pdm_text_cmd -t table [-u from_userid -p from_persid] [-o operation] [-f input file] [-T timeout] [-h]
```

-t *table*

(Required) Specifies the table to process. The *table* name can be one of the following values (not case sensitive):

- Asset
- Contact
- Change
- Issue
- Request

Note: See the [OPTIONS] section of the `text_api.cfg` file for a complete list of valid table names.

-u *from_userid* | **-p** *from_persid*

(One option required) Identifies the contact for this operation:

-u *from_userid*

Identifies the contact using the User ID value.

-p *from_persid*

Identifies the contact using the unique object identifier for the contact record. *from_persid* must be of the form **cnt:xxxx**. *xxxx* is the persistent ID of the object.

Note: The value that you specify with this option is appended to the end of the input for the `pdm_text_cmd` command using the appropriate keyword, `%FROM_USERID` or `%FROM_PERSID`.

-o operation

Specifies the operation to perform. The *operation* must be one of the following values (not case sensitive):

- NEW—Creates an object. This value is the default if no operation is specified.
- UPDATE | UPD—Creates an object if not found or updates an existing object if found.
- UPDATE_ONLY | UPDO—Updates object if found; otherwise, does nothing.

Both UPDATE and UPDATE_ONLY require the %SEARCH keyword in the command input. You can perform only one operation transaction with each invocation of pdm_text_cmd.

-f input_file

Specifies the full path of the file to process, which is a text file containing valid Text API commands. If you omit this parameter, commands are used from STDIN. The Text API uses the following basic format for input:

%keyword=value

You can issue multiple commands within the input by separating the command request by at least five percent signs (%%%%%).

Note: For more information about valid keywords and about formatting input to the Text API, see the file text_api.cfg.

-T timeout

Specifies the number of seconds to wait for a response from the server before timing out. The default is 30 seconds.

Note: pdm_text_cmd shows the text-based replies received back from the Text API, which include success or error messages, and the original text sent using the API for processing. pdm_text_cmd returns zero if the command completes successfully without warnings or errors or one if the command completes successfully, but with warnings. Any other return value indicates that an error occurred.

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

Note: When passing the parameters from command prompt, use Ctrl+Z in Windows and Ctrl+D in POSIX.

(new related group 1)

[Using the Text API](#) (see page 423)

Input Examples

pdm_text_cmd is the command line interface for the Text API that you can use to create and update various objects such as requests, change orders, issues, assets, and contacts.

Example: Use an Input File to Create an Issue

The following example demonstrates how to use a pdm_text_cmd input file to create an issue:

```
%DESCRIPTION=This is my Test.  
%PRIORITY=3
```

To process this file, assuming its full path is c:\input.txt, issue the following command:

```
pdm_text_cmd -t Issue -u user01 -f c:\input.txt
```

Example: Use an Input File to Update an Issue

The following example demonstrates an input file to update issue 123 to a priority of 2:

```
%SEARCH=ISSUE_ID  
%ISSUE_ID=123  
%PRIORITY=2
```

To process this file, assuming its full path is c:\update.txt, issue the following command:

```
pdm_text_cmd -t Issue -o UPDATE_ONLY -u user01 -f c:\update.txt
```

Example: Use an Input File to Create Multiple Requests

The following example demonstrates creating multiple requests with one input file. This command can be helpful creating test data on a test system.

```
%DESCRIPTION=This is Test 1.  
%PRIORITY=3  
%%  
%DESCRIPTION=This is Test 2.  
%PRIORITY=2  
%%  
%DESCRIPTION=This is Test 3.  
%PRIORITY=None
```

To process this file, assuming its full path is c:\testdata.txt, issue the following command:

```
pdm_text_cmd -t Request -u user01 -f c:\testdata.txt
```

pdm_uconv--Convert Local Charset to UTF-8

The pdm_uconv utility assists you in converting data from previous releases of CA SDM or integrations with other CA Technologies products. The most common usage of this utility is to convert from a local charset to UTF-8 and from UTF-8 to a local charset.

Syntax

This command has the following format:

```
pdm_uconv -h [-V] [-s] [-l] [--list-code  
| --default-code | -L] [--cannon] [-x] [--to-callback | -c] [--from-callback | -i] [--fallback | --no-fallback] [-b] [-f] [--t]  
[--add-signature] [--remove-signature] [-o] [file ...]
```

-h

Opens the help menu.

-V

Prints the program version.

-s

Uses silent operation and suppresses messages.

-v

Displays the progress information of the utility.

-l

Lists all available encoding. The following are valid:

--list-code

Lists only the given encoding.

--default-code

Lists only the default encoding.

-L

Lists all available transliterators.

--cannon

Prints the list in cnvtrs.txt(5) format.

-x

Runs the progress through transliteration.

--to-callback *callback*

Uses callback on destination encoding.

-c

Omits invalid characters from the output.

--from-callback *callback*

Uses callback on original encoding.

-i

Ignores invalid sequences in the input.

--callback *callback*

Uses callback on both encoding.s.

-b

Specifies the block size.

Default: 4096

--fallback

Uses fallback mapping.

--no-fallback

Does not use fallback mapping.

-f

Sets the original encoding.

-t

Sets the destination encoding.

--add-signature

Adds U+FEFF Unicode signature character (BOM).

--remove-signature

Removes U+FEFF Unicode signature character (BOM)

-o

Writes output to file.

Examples:

From local charset to UTF-8

```
pdm_uconv -t utf-8 inputfile.txt > outputfile.txt
```

From specific charset (iso-2022-jp) to UTF-8

```
pdm_uconv -f iso-2022-jp -t utf-8 inputfile.txt > outputfile.txt
```

From UTF-8 to local charset

```
pdm_uconv -f utf-8 inputfile.txt > outputfile.txt
```

From UTF-8 to specific charset

```
pdm_uconv -f utf-8 -t iso-2022-jp inputfile.txt > outputfile.txt
```

The `pdm_uconv` utility has the following are valid callbacks:

- substitute
- skip
- stop
- escape
- escape-icu
- escape-java
- escape-c
- escape-xml
- escape-xml-hex
- escape-xml-dec
- escape-unicode

pdm_userload--Add, Update, and Delete Database Records

The `pdm_userload` utility updates a CA SDM database using an input file you specify.

Important! You should always backup your database before you perform a `pdm_userload`.

Whenever you upload tickets (such as issues or requests), your ticket number should include a unique prefix or suffix in its string. CA SDM views this number as a string of characters not as a sequential number, and thus, cannot guarantee that it will assign a unique number to the uploaded tickets. As long as you assign a unique prefix or suffix using `awk` or another text processor, however, you can upload tickets without CA SDM writing over previously assigned numbers.

Syntax

This command has the following format:

```
pdm_userload [-a] [-c] [-h] [-r] [-v] [-u] [-m] -f filename
```

Input File Format

The input file entries follow this format:

TABLE *table_name*

fieldname1 fieldname2 . . . fieldnameN

{ "value11", "value12", . . . "value1N" }

{ "value21", "value22", . . . "value2N" }

.

.

.

{ "valueN1", "valueN2", . . . "valueNN" }

table_name is the name of the table to be loaded, as listed in the CA SDM database schema file, which is located in \$NX_ROOT/site/schema.sch (UNIX) or *installation-directory*\site\schema.sch (Windows), where \$NX_ROOT or *installation-directory* is the directory where you installed CA SDM.

-f filename

Specifies an input ASCII file.

-a

Updates all existing records, regardless of whether more than one existing record matches a single input record. Without this option, input records that match more than one existing record are rejected.

Important! Use this option carefully.

-c

Checks the input file against the database and reports on the updates that would be made, but does not make the updates.

-r

Removes database records that match input records. The -a option can be used with the -r option.

Note: Make a backup copy of the database before running pdm_userload with this option. Once old database records are removed, you must restore the CA SDM database with this backup copy if you wish to recover any deleted records.

-v

Specifies verbose mode.

-u

Updates existing records, but does not insert new records into the database.

-m

Means mass update. Specify when you are using pdm_userload to add or delete a large number of records. This option suppresses all client notifications of updates and sends a cache refresh message for a table when pdm_userload finishes processing the table.

-x

Uses locale sensitive numeric input formats.

-t

Specifies the name or UUID of tenant to associate all loaded data with the specified tenant. This argument is valid only when multi-tenancy is installed.

Pdm_userload supports new arguments on the TABLE statement, "Truncate" and "NoNewID". These arguments are specified in an optional parenthesized option after the table name. For example:

```
TABLE Call_Req (TRUNCATE, NONEWID)
```

Truncate

Causes pdm_userload to issue a database specific TRUNCATE command for the table prior to loading any data. In addition, it forces pdm_userload logic to use insert-only logic regardless of command line arguments, as all records are new.

NoNewID

Causes pdm_userload to use the id value from its input control file for new rows in the table, rather than generating a new ID for inserted data (the default logic of the pdm_userload -i option).

Restrictions

You can run pdm_userload while CA SDM is active, but performance can become very slow. It is best to run pdm_userload when no one is using CA SDM.

More information:

[pdm_backup--Write Database to ASCII File](#) (see page 958)

[pdm_replace--Replace a Database Table](#) (see page 988)

[pdm_restore--Restore a Database](#) (see page 990)

pdm_webstat--Return Web Usage Statistics

Use pdm_webstat to return CA SDM session and user statistics for one or more web engine processes. The pdm_webstat command shows cumulative sessions, most sessions at a time, and current active sessions. It can also provide information about the individual users.

Syntax

This command has the following format:

```
pdm_webstat [-r] [-d | -D] [-i] [-t timeout] [-p webengine process] [-n] [-h]
```

-r

Specifies raw text mode, without titles and other formatting. Within the output for a single web engine process there are no line breaks; however, the output for each web engine process always starts on a new line.

Raw text mode displays data in exactly the same order as when you use pdm_webstat without the -r option. Use raw text mode if you want to use the resulting data in a spreadsheet or other type of report. For example, the following syntax:

```
pdm_webstat -r
```

shows the following output:

```
10/11/2005 10:31:49 web:local:0 12 4 2
10/11/2005 10:31:49 web:local:1 9 2 2
```

-d

Specifies detailed output that displays user sessions. The -d option lists all current sessions in the form of *userid@IPaddress*. If a session is displayed without a user ID, it usually means that the session has not logged on yet. For example, the following syntax:

```
pdm_webstat -d
```

shows the following output:

```
PDM_Webstat: Invoked at 10/11/2005 10:27:31
=====
Report from Webengine: web:local:0
=====
Cumulative sessions so far = 12
Most sessions at a time   = 4
Currently active sessions = 2
    @192.168.1.16
    usery@192.168.1.20
=====
Report from Webengine: web:local:1
=====
Cumulative sessions so far = 7
Most sessions at a time   = 2
Currently active sessions = 2
    SvcPlus@192.168.1.14
    userx@192.168.1.8
```

-D

Provides more detailed output for debugging purposes. This option should be specified only when specifically requested by CA support. It adds internal information about each session to the detailed output. For example, the following syntax:

```
pdm_webstat -D
```

shows the following output:

```
PDM_Webstat: Invoked at 10/11/2007 10:28:10
=====
Report from Webengine: web:local:0
=====
Cumulative sessions so far = 12
Most sessions at a time   = 4
Currently active sessions = 2
  @192.168.1.16      SessionStat  1
  userx@192.168.1.20 SessionStat  5
=====
Report from Webengine: web:local:1
=====
Cumulative sessions so far = 7
Most sessions at a time   = 2
Currently active sessions = 2
  SvcPlus@192.168.1.14 SessionStat  7
  userx@192.168.1.8   SessionStat 13
```

-i

Specifies an interval in seconds between successive reports. When the -i argument is specified, pdm_webstat runs continuously. The pdm_webstat command outputs its report in the format requested by other arguments, waits for the interval specified, and outputs the report again. With -i, pdm_webstat terminates only when explicitly cancelled, normally with Ctrl+C. For example, the following syntax:

```
pdm_webstat -i 5 -p web:local -r
```

shows the following output:

```
09/21/2007 16:27:25 web:local 14 10 6
09/21/2007 16:27:30 web:local 17 10 9
09/21/2007 16:27:35 web:local 18 10 10
09/21/2007 16:27:41 web:local 21 13 13
```

-t timeout

Specifies a time-out value in seconds. This parameter causes pdm_webstat to wait for a number of seconds for a response before terminating. The default is 30 seconds.

-p webengine_process

Specifies the process name of the web engine for which you want to report. By default, all web engine processes are reported. The process name (also called slump-name) is the same name that would be viewed in an slstat output and always starts with "web:".

-n

Suppresses the normal log entry for each reported web engine. By default, a log entry summarizing each web engine is created.

Note: By default, if you do not specify any parameters, pdm_webstat displays summary data for all running processes.

pdm_webstat returns zero if the command completes successfully, and a nonzero value if errors occur (for example, a time-out or no active web engine processes).

Example: pdm_webstat Output When No Parameters Are Specified

The following example shows the output of pdm_webstat when it is executed with no parameters:

```
PDM_Webstat: Invoked at 10/11/2007 10:26:42
```

```
=====
```

```
Report from Webengine: web:local:0
```

```
=====
```

```
Cumulative sessions so far = 12
```

```
Most sessions at a time   = 4
```

```
Currently active sessions = 2
```

```
=====
```

```
Report from Webengine: web:local:1
```

```
=====
```

```
Cumulative sessions so far = 7
```

```
Most sessions at a time   = 2
```

```
Currently active sessions = 2
```

Important! On UNIX, the LIBPATH must be set before running several CA SDM utilities. Use *pdm_task* to set the LIBPATH before running a utility. For example, input "pdm_task pdm_clean_attachments ...".

report--Generate Reports

Applies to UNIX only

The report program lets you generate a report from the command line on the server. To issue the report command at the command line or in a script, you must include `pdm_task`. The `pdm_task` command sets up environment variables for commands that do not have a wrapper. Enter `pdm_task` with the report command on the same command line *only* when the report is invoked through a script or the command line. If you issue the report command from a menu, you do not need to include `pdm_task`, because all environment variables are set by the application.

Syntax

This command has the following format:

```
pdm_task report [-h] [-e] [-f] [-F ffstring] [-p pagelength] filename [command line arguments]
```

-e

Echoes compiled script (for debugging purposes).

-f

Uses form feeds between pages.

-F ffstring

Sets the optional form-feed string.

-p pagelength

Sets the page length. The default page length is 66.

filename

The report template. If you are not running the report command from the directory in which the template file is located, include the complete file path name. The command sends the output as standard output (stdout).

command line arguments

Specifies parameters received by the report template. If the report is designed to accept command line arguments, you must enter a command line argument for each parameter in the report template. If the argument is empty, enter the null string.

For example, the following command supplies the command line arguments Smith, Jane, and L. The report template requires these three parameters to generate the Affected Contacts Report.

For example, enter:

```
pdm_task report /opt/CAisd/samples/sdk/reports/affected.rpt
```

In the next example, Jane Smith does not have a middle initial:

```
pdm_task report /opt/CAisd/samples/sdk/reports/affected.rpt Smith Jane "
```

rpt_srv--Generate Reports

Valid for Windows only

The report program lets you generate a report from the command line on the server. To issue the report command at the command line or in a script, you must include `pdm_task`. The `pdm_task` command sets up environment variables for commands that do not have a wrapper. Enter `pdm_task` with the report command on the same command line only when the report is invoked through a script or the command line. If you issue the report command from a menu, you do not need to include `pdm_task`, because all environment variables are set by the application.

Syntax

This command has the following format:

```
rpt_srv -m [-h] [-e] [-f] [-F ffstring] [-p pagelength] [-C] [-B] filename [ command line arguments]
```

-m

Signifies that the report is manually being run from the command line.

-e

Echoes compiled script (for debugging purposes).

-f

Uses form feeds between pages.

-F ffstring

Sets the optional form-feed string.

-p pagelength

Sets the page length. The default page length is 66.

-C

Changes encoding from UTF-8 to another charset. The default output is UTF-8.

Example: To convert the output to JIS, you would run "-C iso-2022-jp"

Example: To encode to the operating system's native charset, use "DEFAULT" or "NATIVE".

-B

Suppresses the Byte Order Mark if the variable NX_ADD_UTF8_BYTE_ORDER_MARK is set.

The NX_ADD_UTF8_BYTE_ORDER_MARK option is a signature into a file. It allows editors that support UTF-8 to maintain the UTF-8 integrity of the file.

Note: This is only needed for non-ASCII data. If this is not installed, the default behavior omits the Byte Order Mark (BOM). If installed, set it to "1" or "Yes".

filename

Specifies the report template. If you are not running the report command from the directory in which the template file is located, include the complete file path name. The command sends the output as standard output (stdout).

command line arguments

Specifies parameters received by the report template. If the report is designed to accept command line arguments, you must enter a command line argument for each parameter in the report template. If the argument is empty, enter the null string.

For example, the following command supplies the command line arguments Smith, Jane, and L. The report template requires these three parameters to generate the Affected Contacts Report.

For example, enter the following command:

```
rpt_srv -m c:\reports\affected.rpt Smith Jane L
```

In the next example, Jane Smith does not have a middle initial:

```
rpt_srv -m c:\reports\affected.rpt Smith Jane "
```

uniconv--Start UNIX CA NSM Event Converter Daemon

Applies to UNIX only

When CA SDM is integrated with CA NSM, you can use uniconv to send generic event data to filter daemons in CA SDM. uniconv is used in a message action in CA NSM Event Management.

Syntax

This command has the following format:

```
uniconv -h &opnode -e &'text' [-n &
nodeid] [-u &userid] [-d &datem] [-t &
time]
```

Restrictions

uniconv must be run from \$NX_ROOT/bin on UNIX. Your site must be integrated with CA NSM to use this utility.

-h &opnode

Specifies the node name of the machine on which you are executing uniconv (required).

-e &'text'

Specifies the full text of the message (required).

-n &nodeid

Specifies the node name from which the message originated (required).

-u &userid

Specifies the login ID of the person who originated the message.

-d &datem

Specifies the system date in *mm/dd/yy* format.

-t &'time'

Specifies the system time.

Appendix D: Form Groups

This section contains the following topics:

[Customer Forms Group](#) (see page 1009)

[Employee Forms Group](#) (see page 1010)

[Analyst Forms Group](#) (see page 1012)

Customer Forms Group

The following web forms are included in the Customer forms groups:

- about.html
- bin_form_np.html
- chg_lr.html
- cr_lr.html
- detail_iss.html
- detail_issalg.html
- detail_KD.html
- generic.html
- home.html
- iss_lr.html
- issue_status_change.html
- list_iss.html
- list_isscat.html
- list_KD.html
- menu_frames.html
- std_body.html
- std_body_site.html
- std_footer.html
- std_footer_site.html
- std_head.html
- std_header.html
- std_head_site.html

Employee Forms Group

The following web forms are included in the Employee forms groups:

- about.html
- bin_form_np.html
- buttons.html
- change_status_change.html
- chg_lr.html
- cr_lr.html
- detail_alg.html
- detail_chg.html
- detail_chgalg.html
- detail_cr.html
- detail_in.html
- detail_KD.html
- generic.html
- home.html
- iss_lr.html
- list_chg.html
- list_chgcat.html
- list_cr.html
- list_in.html
- list_KD.html
- list_pcat.html
- list_pcat_cr.html
- list_pcat_in.html
- menu_frames.html
- request_status_change.html
- show_error.html
- std_body.html
- std_body_site.html
- std_footer.html
- std_footer_site.html

- std_head.html
- std_header.html
- std_head_site.html

Analyst Forms Group

The following web forms are included in the Analyst forms groups:

A

- about.html
- acctyp_role_tab.html
- acctyp_web_auth_tab.html
- acctyp_wsp_tab.html
- admin_empty.html
- admin_main_role.html
- admin_tab_dflt.html
- admin_tree.html
- attmnt_content_tab.html
- attmnt_fields.html
- attmnt_permissions_tab.html
- attmnt_upload_popup.html
- att_mgs_event.html
- att_stype_event.html
- aty_chg_ntfr_tab.html
- aty_chg_svy_tab.html
- aty_cr_ntfr_tab.html
- aty_cr_svy_tab.html
- aty_iss_ntfr_tab.html
- aty_iss_svy_tab.html
- aty_kdComment_ntfr_tab.html
- aty_kd_ntfr_tab.html
- aty_mgs_ntfr_tab.html

B

- bhvtpl_todo_tab.html
- bhvtpl_trans_info_tab.html
- bin_form_np.html

C

- cancel.html
- cancel_empty.html
- category_content_tab.html
- category_permissions_tab.html
- chgcat_auto_assignment_tab.html
- chgcat_prptpl_tab.html
- chgcat_wftpl_tab.html
- chg_accumulate.html
- chg_causedreq_tab.html
- chg_close_all_child.html
- chg_expedite.html
- chg_lr.html
- chg_relchg_tab.html
- chg_relreq_tab.html
- cia_bmhier_tab.html
- cia_export_bmhier.html
- cia_export_nr.html
- cia_nr_tab.html
- cia_pwd_tab.html
- cia_sync_stat.html
- cnote_tracker.html
- cnt_addr_tab.html
- cnt_auto_assignment_tab.html
- cnt_env_tab.html
- cnt_grp_tab.html
- cnt_mem_tab.html
- cnt_notif_tab.html
- cnt_org_tab.html
- cnt_rem_tab.html
- cnt_role_tab.html
- cr_attach_chg.html
- cr_close_all_child.html
- cr_detach_chg.html

- cr_lr.html
- cr_relreq_tab.html

D

- dcon_constraint_tab.html
- dcon_sql_tab.html
- detail.template
- detail_acctyp.html
- detail_act_type_assoc.html
- detail_ADMIN_TREE.html
- detail_alg.html
- detail_arcpur_rule.html
- detail_asset.html
- detail_atev.html
- detail_atomic_cond.html
- detail_atmnt_edit.html
- detail_atmnt_error.html
- detail_atmnt_folder.html
- detail_atmnt_ro.html
- detail_attr_alias.html
- detail_aty.html
- detail_audlog.html
- detail_bhvtpl.html
- detail_bmcls.html
- detail_bmhier.html
- detail_bmrep.html
- detail_BU_TRANS.html
- detail_ca_cmpny.html
- detail_chg.html
- detail_chgalg.html
- detail_chgcat.html
- detail_chgstat.html
- detail_chgtype.html

- detail_CI_ACTIONS.html
- detail_CI_ACTIONS_ALTERNATE.html
- detail_CI_DOC_TEMPLATES.html
- detail_CI_STATUSES.html
- detail_CI_WF_TEMPLATES.html
- detail_cmth.html
- detail_cnote.html
- detail_cnt.html
- detail_cost_cntr.html
- detail_country.html
- detail_cr.html
- detail_crs.html
- detail_crsq.html
- detail_cr_prptpl.html
- detail_ctab.html
- detail_ctimer.html
- detail_ctp.html
- detail_dcon.html
- detail_dept.html
- detail_dmn.html
- detail_doc_rep.html
- detail_DOC_VERSIONS.html
- detail_EBR_ACRONYMS.html
- detail_EBR_LOG.html
- detail_EBR_NOISE_WORDS.html
- detail_EBR_SYNONYMS_ADM.html
- detail_event_log.html
- detail_evt.html
- detail_fmgrp.html
- detail_grc.html
- detail_g_cnt.html
- detail_g_loc.html
- detail_g_org.html

- detail_g_prod.html
- detail_g_qname.html
- detail_g_srvrs.html
- detail_g_tblmap.html
- detail_g_tblrule.html
- detail_help_set.html
- detail_hier_edit.html
- detail_hier_ro.html
- detail_ical_event_template.html
- detail_imp.html
- detail_in.html
- detail_iss.html
- detail_issalg.html
- detail_isscat.html
- detail_issstat.html
- detail_iss_wf.html
- detail_kc.html
- detail_KCAT.html
- detail_KD.html
- detail_KD_FILE.html
- detail_KD_QA.html
- detail_KD_SAVE_AS.html
- detail_KD_TASK.html
- detail_KD_TASK_cancel_rework.html
- detail_KD_TASK_retire.html
- detail_KD_template.html
- detail_KEIT_TEMPLATES.html
- detail_KT_ACT_CONTENT.html
- detail_KT_BLC.html
- detail_KT_FILE_TYPE.html
- detail_KT_FLG_TYPE.html
- detail_ldap.html
- detail_ldap_group.html

- detail_loc.html
- detail_LONG_TEXTS.html
- detail_lr_ro.html
- detail_macro.html
- detail_macro_type.html
- detail_menu_bar.html
- detail_menu_tree_name.html
- detail_mfrmod.html
- detail_mgs.html
- detail_mgsalg.html
- detail_mgsstat.html
- detail_NOTIFICATION.html
- detail_no_contract_sdsc.html
- detail_nr.html
- detail_nrf.html
- detail_nr_com.html
- detail_ntfl.html
- detail_ntfm.html
- detail_ntfr.html
- detail_options.html
- detail_org.html
- detail_O_COMMENTS.html
- detail_O_EVENTS.html
- detail_pcat.html
- detail_perscnt.html
- detail_position.html
- detail_pr.html
- detail_prefs.html
- detail_pri.html
- detail_prod.html
- detail_projex.html
- detail_prptpl.html
- detail_prpval.html

- detail_prpval_rule.html
- detail_prp_edit.html
- detail_QUERY_POLICY.html
- detail_rc.html
- detail_response.html
- detail_role.html
- detail_role_go_form.html
- detail_rptmeth.html
- detail_rrf.html
- detail_rss.html
- detail_sapolicy.html
- detail_saprobtyp.html
- detail_sdsc.html
- detail_sdsc_map.html
- detail_seq.html
- detail_sev.html
- detail_site.html
- detail_slatpl.html
- detail_srvr_aliases.html
- detail_srvr_zones.html
- detail_state.html
- detail_svc_contract.html
- detail_svy_atpl.html
- detail_svy_qtpl.html
- detail_svy_tpl.html
- detail_tab.html
- detail_tenant.html
- detail_tenant_group.html
- detail_tskstat.html
- detail_tskty.html
- detail_tspan.html
- detail_typecnt.html
- detail_tz.html

- detail_urg.html
- detail_USP_PREFERENCES.html
- detail_usp_servers.html
- detail_vpt.html
- detail_web_form.html
- detail_wf.html
- detail_wftpl.html
- detail_wrkshft.html
- dmn_dcon_tab.html
- edit_prop_dyn.html
- ed_image_pane.html
- evt_action_info.html
- evt_config_info.html

G

- generic.html
- get_comment.html
- g_profile_browser.html
- g_profile_browser2.html
- g_profile_browser3.html
- g_profile_browser_frameset.html
- g_profile_jump.html
- g_profile_scratchpad.html

H

- hierload_admin_tree.html
- hierload_KCAT.html
- hiersel_admin_tree.html
- hiersel_KCAT.html
- hourglass.html
- html_editor_create_change_order.html
- html_editor_create_ticket.html
- html_editor_frames.html
- html_editor_insert_image.html

- [html_editor_insert_link.htmlpl](#)
 - [html_editor_insert_table.htmlpl](#)
 - [html_editor_tabs.htmlpl](#)
 - [html_editor_toolbar.htmlpl](#)
- I
- [insert_iss_wf.htmlpl](#)
 - [insert_wf.htmlpl](#)
 - [in_relreq_tab.htmlpl](#)
 - [isscat_auto_assignment_tab.htmlpl](#)
 - [isscat_prptpl_tab.htmlpl](#)
 - [isscat_wftpl_tab.htmlpl](#)
 - [issue_status_change.htmlpl](#)
 - [iss_accumulate.htmlpl](#)
 - [iss_close_all_child.htmlpl](#)
 - [iss_custfld_tab.htmlpl](#)
 - [iss_expedite.htmlpl](#)
 - [iss_lr.htmlpl](#)
 - [iss_reliss_tab.htmlpl](#)
 - [iss_resol_tab.htmlpl](#)
- K
- [kd_action_forward.htmlpl](#)
 - [kd_action_publish.htmlpl](#)
 - [kd_action_reject.htmlpl](#)
 - [kd_action_unpublish.htmlpl](#)
 - [kd_action_unretire.htmlpl](#)
 - [kd_attachments_tab.htmlpl](#)
 - [kd_attributes_tab.htmlpl](#)
 - [kd_categories_tab.htmlpl](#)
 - [kd_content_tab.htmlpl](#)
 - [kd_file_prop_tab.htmlpl](#)
 - [kd_permissions_tab.htmlpl](#)
 - [kd_qa_attributes_tab.htmlpl](#)

- kd_qa_content_tab.html
- keit_tmpl_export_fields_tab.html
- keit_tmpl_export_filter_tab.html
- keit_tmpl_import_settings_tab.html
- keit_tmpl_name_tab.html
- kt_admin_attachments.html
- kt_admin_automated_policies.html
- kt_admin_document_settings.html
- kt_admin_faq_options.html
- kt_admin_general_settings.html
- kt_admin_integration.html
- kt_admin_knowledge.html
- kt_admin_parse_settings.html
- kt_admin_report_card.html
- kt_admin_search_config_cr.html
- kt_admin_search_config_iss.html
- kt_admin_search_options.html
- kt_admin_survey_settings.html
- kt_admin_workflow_settings.html
- kt_architect.html
- kt_architect2.html
- kt_architect3.html
- kt_architect_delete_KCAT.html
- kt_architect_delete_KD.html
- kt_architect_frameset.html
- kt_architect_init.html
- kt_architect_javascript.html
- kt_architect_KCATs.html
- kt_architect_KCAT_path.html
- kt_architect_KDs.html
- kt_dtbuilder.html
- kt_dtbuilder2.html
- kt_dtbuilder3.html

- kt_dtbuilder_frameset.html
- kt_dtbuilder_node.html
- kt_dtbuilder_prompt_window.html
- kt_dtbuilder_save_dialog_window.html
- kt_dtbuilder_save_tree_form.html
- kt_dtbuilder_tree.html
- kt_email_document.html
- kt_faq_tree.html
- kt_main.html
- kt_main2.html
- kt_main3.html
- kt_main_role.html
- kt_permissions.html

L

- list.template
- list_acctyp.html
- list_act_type_assoc.html
- list_alg.html
- list_all_fmgrp.html
- list_all_lr.html
- list_architect_KDs.html
- list_architect_KDs_Pref.html
- list_arcpur_hist.html
- list_arcpur_rule.html
- list_atev.html
- list_atomic_cond.html
- list_atmnt.html
- list_attr_alias.html
- list_aty.html
- list_audlog.html
- list_bmcls.html
- list_bmhier.html

- list_bmrep.html
- list_bm_task.html
- list_bool.html
- list_ca_cmpny.html
- list_ca_logical_asset.html
- list_chg.html
- list_chgalg.html
- list_chgcat.html
- list_chgsched.html
- list_chgsched_config.html
- list_chgstat.html
- list_chgtype.html
- list_CI_ACTIONS.html
- list_CI_ACTIONS_ALTERNATE.html
- list_CI_DOC_TEMPLATES.html
- list_CI_STATUSES.html
- list_CI_WF_TEMPLATES.html
- list_cmth.html
- list_cnote.html
- list_cnt.html
- list_cost_cntr.html
- list_country.html
- list_cr.html
- list_crs.html
- list_crsq.html
- list_crs_cr.html
- list_crs_in.html
- list_crs_pr.html
- list_crt.html
- list_cr_kt.html
- list_ctab.html
- list_ctimer.html
- list_ctp.html

- list_dcon.html
- list_dept.html
- list_dmn.html
- list_DOC_VERSIONS.html
- list_EBR_ACRONYMS.html
- list_EBR_LOG.html
- list_EBR_NOISE_WORDS.html
- list_EBR_SYNONYMS_ADM.html
- list_event_log.html
- list_evt.html
- list_evtdly.html
- list_grc.html
- list_grpmem.html
- list_g_chg_queue.html
- list_g_cnt.html
- list_g_cr_queue.html
- list_g_iss_queue.html
- list_g_loc.html
- list_g_org.html
- list_g_prod.html
- list_g_qname.html
- list_g_srvrs.html
- list_g_tblmap.html
- list_g_tblrule.html
- list_g_tenant.html
- list_help_item.html
- list_help_set.html
- list_ical_event_template.html
- list_imp.html
- list_in.html
- list_iss.html
- list_issalg.html
- list_isscat.html

- list_issstat.html
- list_iss_kt.html
- list_iss_wf.html
- list_kc.html
- list_KCAT_LINKED.html
- list_KCAT_QA.html
- list_KCAT_tree.html
- list_KD.html
- list_kdsched.html
- list_kdsched_config.html
- list_KD_ATTMNT.html
- list_kd_CI_DOC_LINKS.html
- list_KD_FILE.html
- list_kd_INDEX_DOC_LINKS.html
- list_KD_QA.html
- list_KEIT_export_transactions.html
- list_KEIT_IMPORT_PACKAGES.html
- list_KEIT_import_transactions.html
- list_KEIT_TEMPLATES.html
- list_KT_ACT_CONTENT.html
- list_KT_BLC.html
- list_KT_FILE_TYPE.html
- list_KT_FLG_TYPE.html
- list_KT_FREE_TEXT.html
- list_KT_LIFE_CYCLE_REP.html
- list_ldap.html
- list_ldap_group.html
- list_loc.html
- list_LONG_TEXTS.html
- list_lr.html
- list_macro.html
- list_macro_type.html
- list_menu_bar.html

- list_menu_tree_name.html
- list_mfrmod.html
- list_mgs.html
- list_mgsalg.html
- list_mgsstat.html
- list_NOTIFICATION.html
- list_no_contract_sdsc.html
- list_nr.html
- list_nrf.html
- list_nr_com.html
- list_ntfl.html
- list_ntfm.html
- list_ntfr.html
- list_OA_TABLES.html
- list_options.html
- list_org.html
- list_O_EVENTS.html
- list_pcat.html
- list_pcat_cr.html
- list_pcat_in.html
- list_pcat_pr.html
- list_perscnt.html
- list_position.html
- list_pr.html
- list_pri.html
- list_prod.html
- list_prod_list.html
- list_prpval.html
- list_prpval_rule.html
- list_QUERY_POLICY.html
- list_QUERY_POLICY_ACTIONS.html
- list_rc.html
- list_rel_cat.html

- list_response.html
- list_role.html
- list_role_tab.html
- list_rptmeth.html
- list_rrf.html
- list_rss.html
- list_sapolicy.html
- list_saprobtyp.html
- list_sdsc.html
- list_sdsc_map.html
- list_seq.html
- list_sev.html
- list_showgrp.html
- list_site.html
- list_srvr_aliases.html
- list_srvr_zones.html
- list_state.html
- list_svc_contract.html
- list_svy_atpl.html
- list_svy_qtpl.html
- list_svy_tpl.html
- list_tab.html
- list_tenant.html
- list_tenant_group.html
- list_tskstat.html
- list_tskty.html
- list_tspan.html
- list_typecnt.html
- list_tz.html
- list_urg.html
- list_usp_servers.html
- list_vpt.html
- list_web_form.html

- list_wf.html
- list_wrkshft.html
- load_properties.html
- load_wait.html
- loc_address_tab.html
- loc_auto_assignment_tab.html
- log_reader.html
- log_reader_banner.html
- log_reader_fs.html
- log_sol_4itil.html

M

- macro_atomic_cond_tab.html
- macro_cnt_tab.html
- macro_ctp_tab.html
- macro_ntfl_tab.html
- macro_rrf_tab.html
- mactyp_exescript_tab.html
- mactyp_valscript_tab.html
- mapped_contracts_tab.html
- menubar.template
- menubar_admin.html
- menubar_architect.html
- menubar_chg_sched.html
- menubar_dtbuilder.html
- menubar_html_editor.html
- menubar_kt.html
- menubar_no.html
- menubar_sd.html
- menubar_sd_chg_manager.html
- menubar_sd_cust_mgr.html
- menubar_sd_cust_rep.html
- menubar_sd_hd_manager.html

- menubar_sd_inc_manager.html
- menubar_sd_know_analyst.html
- menubar_sd_know_manager.html
- menubar_sd_l1_analyst.html
- menubar_sd_l2_analyst.html
- menubar_sd_prb_manager.html
- menubar_sd_vendor_analyst.html
- menu_frames.html
- menu_tree_editor.html
- menu_tree_editor2.html
- menu_tree_editor3.html
- mgs_cnt_tab.html
- mgs_ctp_tab.html
- mgs_ini_tab.html
- mgs_ntfl_tab.html
- mgs_rem_tab.html
- multiframe.template
- multiframe_reports_admin.html
- multiframe_reports_chg_manager.html
- multiframe_reports_cust_mgr.html
- multiframe_reports_inc_mgr.html
- multiframe_reports_know_analyst.html
- multiframe_reports_know_mgr.html
- multiframe_reports_prb_mgr.html
- multiframe_reports_sd_mgr.html

N

- new_lr.html
- nf.html
- nosession.html
- nr_bm_tab.html
- nr_chg_tab.html
- nr_contact_tab.html

- nr_inc_tab.html
- nr_inv_tab.html
- nr_iss_tab.html
- nr_loc_tab.html
- nr_log_tab.html
- nr_org_tab.html
- nr_prb_tab.html
- nr_projex_tab.html
- nr_rel_tab.html
- nr_reqitil_tab.html
- nr_req_tab.html
- nr_serv_tab.html
- ntfl_ntfr_tab.html
- ntfr_aty_tab.html
- ntfr_cnt_tab.html
- ntfr_ctp_tab.html
- ntfr_ntfl_tab.html

O

- order_status_change.html
- org_address_tab.html
- org_env_tab.html

P

- pcat_auto_assignment_tab.html
- pcat_prptpl_tab.html
- pcat_wftpl_tab.html
- power_user_tips.html
- profile_browser.html
- profile_browser2.html
- profile_browser3.html
- profile_browser_frameset.html
- profile_envcnt.html
- profile_envorg.html

- profile_histcnt_chg.html
- profile_histcnt_cr.html
- profile_histcnt_in.html
- profile_histcnt_iss.html
- profile_histcnt_pr.html
- profile_historg_chg.html
- profile_historg_cr.html
- profile_historg_in.html
- profile_historg_iss.html
- profile_historg_pr.html
- profile_infocnt.html
- profile_infoorg.html
- profile_menu.html
- profile_qtemplate.html
- pr_attinc_tab.html
- pr_relreq_tab.html

R

- reports.html
- reports.html.tpl
- request_status_change.html
- role_auth_tab.html
- role_fnacc_tab.html
- role_goform_tab.html
- role_kt_ct_tab.html
- role_kt_docs_tab.html
- role_webform_tab.html
- role_web_interface_tab.html

S

- sapolicy_ac_tab.html
- sapolicy_pt_tab.html
- saprobtyp_dh_tab.html
- saprobtyp_rd_tab.html

- scoreboard.html
- scratchpad.html
- screen_reader_usage.html
- sdsc_chg_slatpl_tab.html
- sdsc_chg_wf_slatpl_tab.html
- sdsc_cr_slatpl_tab.html
- sdsc_iss_slatpl_tab.html
- sdsc_iss_wf_slatpl_tab.html
- sdsc_map_cnt_tab.html
- sdsc_map_grp_tab.html
- sdsc_map_nr_tab.html
- sdsc_map_pri_tab.html
- sdsc_map_urg_tab.html
- sd_kt_admin.html
- sd_main.html
- sd_main_role.html
- search_child_KCATs_filter.html
- show_error.html
- show_main_detail.html
- std_body.html
- std_body_site.html
- std_footer.html
- std_footer_site.html
- std_head.html
- std_head_site.html
- suggest_knowledge_isscat.html
- suggest_knowledge_list_isscat.html
- suggest_knowledge_list_pcat.html
- suggest_knowledge_pcat.html
- suggest_knowledge_search_options.html
- T
- tab_detail.template

- tenant_address_tab.html
- tenant_groups_tab.html
- tenant_group_members_tab.html
- tskty_tskstat_tab.html

U

- update_lrel_bmrep.html
- update_lrel_chg.html
- update_lrel_chgcat.html
- update_lrel_cnt.html
- update_lrel_cr.html
- update_lrel_ctp.html
- update_lrel_goform.html
- update_lrel_help_content.html
- update_lrel_in.html
- update_lrel_iss.html
- update_lrel_isscat.html
- update_lrel_loc.html
- update_lrel_macro.html
- update_lrel_nr.html
- update_lrel_ntfl.html
- update_lrel_ntfr.html
- update_lrel_org.html
- update_lrel_pcat.html
- update_lrel_pr.html
- update_lrel_role.html
- update_lrel_tab.html
- update_lrel_tenant.html
- update_lrel_tenant_group.html
- update_lrel_tskstat.html
- update_lrel_webform.html
- update_lrel_wrkshft.html
- upd_chg_sched.html

- upload_file.html
- upload_success.html
- usq_update.html
- usq_update_control.html
- usq_update_fin.html
- usq_update_select.html
- usq_update_tree.html

V

- v30_date_helper.html

W

- wfdef.html
- wftpl_auto_assignment_tab.html
- wftpl_bhvtpl_tab.html
- working.html
- workitems.html
- wrkshft_auto_assignment_tab.html
- wrkshft_schedule_tab.html
- wspmain.html

X

- xfer_esc_chg.html
- xfer_esc_cr.html
- xfer_esc_iss.html
- xx_attmnt_tab.html
- xx_candp_tab.html
- xx_nr_tab.html
- xx_prop_tab.html
- xx_solnalg_tab.html
- xx_stype_tab.html
- xx_template_tab.html
- xx_wf_tab.html

Chapter 22: Troubleshooting Performance Problems

Identifying Performance Problems in CA SDM

Customers that experience performance problems in CA SDM can gather information about their environment before contacting CA Support Online to help identify and resolve CA SDM problems. Use the CA SDM Diagnostic Report Tool and Interval Logging Script to collect diagnostic information about your CA SDM environment. For example, understand the computer configuration to determine the Operating System type, version, and available system resources.

Important! Install and configure the CA SDM server before running the diagnostic tool. Execute the diagnostic utility on all the primary and secondary CA SDM server installations. The tool does *not* change or modify the CA SDM server.

Note: We recommend that you contact CA Support Online before you use the diagnostic tool.

The following process describes how to identify performance problems:

1. [Define the Performance Problem](#) (see page 1036).
2. Gather the following site-specific environmental data:
 - [Execute the CA SDM Diagnostic Report tool](#) (see page 1036).
 - Windows report
 - UNIX report
 - Collected diagnostic report
 - [Gather database server environment details](#) (see page 1042).
 - Gather network topology and topology information or other products that you integrate with CA SDM.

For example, locate information about the products from available PDF files or diagrams.
3. [Execute the Interval Logging Script](#) (see page 1041).
4. [Review general tuning recommendations](#) (see page 1043).

Define the Performance Problem

Define your performance problem by collecting system information initially. Review the list of questions to determine if they are relevant for your issue. The following example questions can help you define the problem.

- What did users experience to disrupt their tasks?
- When did users first discover the problem?
- Has your environment changed recently, such as hardware upgrades?
- What functionality of the product experiences issues?
- How many users does the issue impact, and what type of users?
- What users are not impacted?
- What is the geographic location of the users that see the problem?
- What is the Access Level of the impacted users?
- Do you host CA SDM on a VMWare ESX server or other virtualized environment?
- What other software are you running on the ESX server or host machine?
- What are the specifications of this environment?
- How many CPUs do you have on each computer?
- How much memory did you configure for each computer?

Use the CA Diagnostic Report Tool

You can use the diagnostic tool on Windows, Oracle Solaris, AIX, and Linux operating system. After you execute `supp_diag`, the tool detects the type of operating system to determine what commands to utilize for data collection from the CA SDM installation.

The tool creates a .CAZ file on the Windows OS and .tar.gz on UNIX in the `$NX_ROOT\diag\rpt` directory that you upload with your issue at <http://support.ca.com>.

Windows

Set `$NX_ROOT` as your CA SDM installation root directory. The default for `$NX_ROOT` is the `C:\Program Files\CA\Service Desk\` on Windows. You can change the default when you want to change the default directory during the installation process.

Important! The diagnostic tool requires `pslist.exe` to run correctly on Windows. Download `pslist` from the Microsoft support website. Install `pslist` on the CA SDM computer and add its directory path variable to the system path variable before you use the diagnostic tool.

UNIX/Linux

Set \$NX_ROOT as your CA SDM installation root directory. The default for \$NX_ROOT is — /opt/CAisd/ on a Unix or Linux operating system. You can change the default when you want to change the default directory during the installation.

Follow these steps:

1. Complete one of the following actions:
 - For Windows, verify that the system path variable includes \$NX_ROOT\bin.
 - For UNIX, verify that your \$PATH includes \$NX_ROOT/bin
2. Execute supp_diag.cmd on Windows, or execute supp_diag.sh on UNIX.

The diagnostic tool can take five to 10 minutes to complete.
3. If the data collection process does not complete, view the \$NX_ROOT\diag\<host_name>_supp_diag.log log file to determine the errors when collecting diagnostic information.

Note: If you want to cancel the background batch job, use CTRL-C to cancel the batch file. Some processes still run in the background, such as MSINFO32.exe. If you have any questions about using this diagnostic tool, contact the CA Support Online.
4. The script directory structure displays the location of the script files, diagnostic zip files, and log files:
 - The \$NX_ROOT\diag\bin directory contains script files.
 - The \$NX_ROOT\diag\rpt directory contains the diagnostics zip file (in .caz format on Windows and .tar.gz format on UNIX systems).
 - The \$NX_ROOT\diag\misc_logs directory contains the log files that can be automatically included in the zip file.
5. Complete the appropriate steps to unzip the gathered files that are based on your operating system:

Windows

- Open a command prompt.
- Cd to \$NX_ROOT\diag\rpt or any directory where the .CAZ file is located.
- Execute the following command.

```
$NX_ROOT\diag\bin\cazipxp -u <package_name>.CAZ
```

UNIX

- Open a command prompt
- Cd to \$NX_ROOT/diag/rpt or any directory where the .tar or .tar.gz file is located.
- Uncompress and untar the file:

```
gunzip -d <package_name>.tar.gz  
tar -xvf <package_name>.tar
```

Windows Report

The following list describes the Windows report files that are created and included in the CAZ\tar file package.

Ca.olf

Specifies the CA licensing information from ca_lic directory.

Lic98.log

Specifies the log file that is related to CA licensing from ca_lic directory.

Lic98version.log

Specifies the log file that is related to CA licensing from ca_lic directory.

Licdebug.log log

Specifies the file that is related to CA licensing from ca_lic directory.

Drwatsoninfo.txt

Specifies the Dr Watson configuration of the computer.

<host name>_env.txt

Specifies the environment variables that are set on the computer.

<host name>_slstat.txt

Specifies the output of slstat command.

<host name>_pdm_status.txt

Specifies the output of slstat command.

<host name>_dir_listing.txt

Specifies the Service Desk install directory listing.

<host name>_pslist.txt

Specifies the process listing when the pslist Microsoft tool is installed.

<host name>_MSINFO32.NFO

Specifies the MSINFO output gathering system information.

<host name>_SYSTEMINFO.TXT

Specifies the system information.

<host name>_appevents.csv

Specifies the application event logs created in the past seven days.

<host name>_sysevents.csv

Specifies the application event logs created in the past seven days.

<host name>_hostinfo.txt

Specifies the computer information.

<host name>_prodinstallinfo.txt

Specifies CA products installation information.

<host name>_caprod_registry.txt

Specifies the Registry information of installed CA products.

<host name>_softfeatures.txt

Specifies the list of software features that are installed for Service Desk.

<host name>_ipconfig.txt

Specifies the IP configuration information.

<host name>_supp_diag.log

Specifies the log created for running the supp_diag tool.

UNIX Report

The following list describes the UNIX report files that are created and included in the CAZ\tar file.

Ca.olf

Specifies the CA licensing information from ca_lic directory.

Lic98.log

Specifies the log file that is related to CA licensing from ca_lic directory.

<host name>_env.txt

Specifies the environment variables that are set on the computer.

<host name>_slstat.txt

Specifies the output of slstat command.

<host name>_pdm_status.txt

Specifies the output of slstat command.

<host name>_dir_listing.txt

Specifies the Service Desk install directory listing.

<host name>_pslist.txt

Specifies the process listing when the pslist microsoft tool is installed.

<host name>_uname.txt

Specifies the output of uname operating system command.

<host name>_diskinfo.txt

Specifies the output of df operating system command.

<host name>_freemem.txt

Specifies the output of memory information.

<host name>_supp_diag.log

Specifies the log created for running the supp_diag tool.

<host name>_prtconf.txt

Specifies the output of prtconf operating system command on Solaris and AIX computers.

<host name>_solrev.txt

Specifies the OS version and patches Information on Solaris computers.

<host name>_netconf.txt

Specifies the IP configuration Information on AIX computers.

Collected Diagnostic Report

The default CA SDM documentation file\directory list includes the following reports in the CAZ\tar file. Add additional files\directories that you include in the CAZ\tar file by placing it in the \$NX_ROOT\diag\misc_logs directory.

- \$NX_ROOT/GENLEVEL or \$NX_ROOT/.GENLEVEL
- \$NX_ROOT/<COMPUTERNAME>.his
- \$NX_ROOT/NX.env
- \$NX_ROOT/NX.env.last
- \$NX_ROOT/log\
- \$NX_ROOT/pdmconf\

- \$NX_ROOT/pdmconf\version
- \$NX_ROOT/bopcfg\www*.cfg
- \$NX_ROOT/site\mods\
- \$NX_ROOT/site\ddict.sch
- \$NX_ROOT/site\eh\
- \$NX_ROOT/bopcfg\www\CATALINA_BASE\logs\
- \$NX_ROOT/bopcfg\www\CATALINA_BASE\webapps\CAisd\WEB-INF\web.xml
- \$NX_ROOT/bopcfg\www\CATALINA_BASE\webapps*.xml

Execute the Interval Logging Script

Interval logging determines resource usage on your server. Confirm whether any processes consume high amounts of CPU, Memory, or Disk I/O on your server. The Interval Logging script (batch file on Windows Server or Shell Script on Unix/Linux) starts automatically when the Service Desk Manager services start. The batch script collects CA SDM and OS statistics every 3 minutes.

For Windows, the batch file generates the CPUStats_Hostname.txt, pslistx_hostname.out, pslistm_hostname.out, netstat_hostname.out files for both Primary and Secondary servers in the \$NX_ROOT/log folder. For only Primary servers, the batch file generates dbstats_hostname.out and status_hostname.out.

For UNIX/Linux, the script creates a file named interval_log_<hostname>.out in the \$NX_ROOT/log directory. The script stores all backup files in the Interval_logging folder with the timestamp in the \$NX_ROOT/log directory.

Note: If your UNIX system contains the TOP tool, add that tool to the Internal Logging command sets. For CA Support Online, take note of the TOP output OS statistics for the top 50 processes that take up the most of your system resources.

For Windows 2008, this utility requires the pslist.exe tool (part of pstools.zip) to be installed and added to the windows OS %PATH% environment variable on each of the CA SDM servers.

Complete the following steps to download PS Tools Suite:

1. Download PS Tools Suite from Microsoft.
2. Extract pstools.zip to any directory of your choice and add the directory to the Windows %PATH% environment variable.
3. Execute pslist.exe once manually from the command prompt as the Local System user and accept the license agreement. To execute pslist.exe as the Local System user, run the following command:

```
psexec.exe -s -i pslist.exe
```

Note: You can start the CA SDM Windows service under a different user account other than the Local System account. For that service, execute pslist.exe under that account and accept the license agreement. If pslist.exe is added after installing & configuring CA SDM, restart the CA SDM services after accepting the license agreement of pslist.

Gather Database Server Environment Details

Gather details about your database server. This information can help identify performance problems in CA SDM.

Follow these steps:

1. Determine the location of your database server, such as local or remote.
2. Determine the DBMS version, Operating System version, and patch level.
3. Complete the appropriate steps for your database type:
 - Execute the following queries for SQL Server and note the results:

```
select @@version  
SELECT SERVERPROPERTY('productversion'), SERVERPROPERTY ('productlevel')
```
 - Execute the following query for Oracle:

```
select * from v$version where banner like 'Oracle%';
```
4. Confirm the version of the database client that you installed on the application server.
5. If available, gather additional information about Environmental Data, such as the Operating System, other databases, and so on.

Review General Tuning Recommendations

We recommend that you monitor the key performance indicators and resource consumption regularly. You can also help ensure that routine maintenance is applied to help identify small problems before they become serious.

If you suspect you see an issue, or if users complain about slow performance, open a CA Support Online issue and send the output from the Interval Logging Script, CA SDM Diagnostic Tool, CA SDM standard logs, and information about the problem. Send this information to CA Support from each CA SDM primary and secondary server.

The following list describes common signs of performance problems:

- Long messages in the CA SDM standard logs (stdlog.x files)
- Search for messages in the stdlogs that state "The following query took #### milliseconds...".
- Queuing on database agents in pdm_vdbinfo output
- Search for *Queued Requests(#)* in the pdm_vdbinfo output. You can execute this command at the OS prompt manually. The Interval Logging script also executes this command.
- Large number of connected users to each web engine. For example, more than 200 users.
Note: Execute pdm_webstat to display the number of concurrent users per web engine. You can execute this command at the OS prompt manually. The Interval Logging script also executes this command.
- User complaints