CA Enterprise Log Manager

Administration Guide

r12.5

CA Product References

This document references the following CA products:

- CA Access Control
- CA Audit
- CA ACF2™
- CA Directory
- CA Embedded Entitlements Manager (CA EEM)
- CA Enterprise Log Manager
- CA Identity Manager
- CA IT Process Automation Manager (CA IT PAM)
- CA NSM
- CA Security Command Center (CA SCC)
- CA Service Desk
- CA SiteMinder®
- CA Spectrum®
- CA Top Secret®

Contact CA

Contact CA Support

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

- 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

Provide Feedback

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

If you would like to provide feedback about CA product documentation, complete our short customer survey, which is available on the CA Support website at http://ca.com/docs.

Documentation Changes

Updates have been made to the following areas since the last release of this documentation: Services and CA Adapters, Subscription, Filters and Profiles, Queries and Reports, Event Correlation and Incident Management, and Log Storage

Affected topics include the following:

- Services and CA Adapters:
  - Service Configurations—New topics covering the Correlation Service, Incident Service, and other new services.
- Subscription:
  - Subscription improvement topics—New and updated topics to cover the subscription download dashboard and other new and changed features.
  - On-Demand Updates—This topic updated to reflect changes in the on-demand update process.
- Filters and Profiles:
  - Global and Local Filters—This existing topic and other topics updated to reflect that filters can now be set against either the Event Database or the Incident Database.
- Queries and Reports:
  - How to Create a Query—This existing topic and related topics changed to reflect that queries can now be directed to either the Event Database or the Incident Database.
  - Approaches to Maintaining Keyed Lists—This existing topic and related topics changed to reflect user interface changes to keyed list configuration.
- Event Correlation and Incident Management—This new chapter contains conceptual material on event correlation rules and how they create incidents, as well as information about how to manage incidents.
  - Correlation Rules Tasks —This new topic and related topics contain information about correlation rules, designing and applying incident notifications, and examples of using predefined correlation rules.
  - Incident Management Tasks —This new topic and related topics contain information about incidents created by event correlation, including viewing incident details.
- Log Storage:
  - Data Integrity Checks—This new topic and related topics contain information about the CA Enterprise Log Manager data integrity system. The information includes how to schedule automatic checks, make checks on demand, and work with database signature keys.

More information:

Services Tasks (see page 120)
Edit Global Configurations (see page 121)
Service Configurations (see page 126)
View Global Subscription Status (see page 186)
About On-Demand Updates (see page 193)
Global and Local Filters (see page 206)
Create a Global Filter (see page 215)
Create a Local Filter (see page 217)
About Queries and Reports (see page 220)
How to Create a Query (see page 244)
Approaches to Maintaining Keyed Lists (see page 279)

Contents

Chapter 5: Services and CA Adapters 119

Chapter 6: Log Storage 151

Chapter 12: Suppression and Summarization 425

Chapter 13: Mapping and Parsing 447

Chapter 14: Integrations and Connectors 487

Chapter 1: Introduction

This section contains the following topics:

About this Guide

This *CA Enterprise Log Manager Administration Guide* addresses tasks performed after the Administrator installs CA Enterprise Log Manager and performs initial server configuration. Some of these tasks are performed to accommodate infrequent changes in the system; others are routine tasks performed on a scheduled basis; still others are ongoing monitoring tasks.

This guide is intended for all audiences, including the following:

- Administrators who maintain the configuration of the product and manage log storage and subscription updates
- Analysts who use reports to monitor the environment, create custom reports, and schedule the generation of alerts
- Auditors who schedule reports, use queries and reports to verify compliance with standards, and annotate reports

This guide includes a glossary of terms and an index. A summary of the contents follows:

Section	Describes how to
User Accounts	Configure user accounts with predefined roles and self-administer user accounts
Policies	Plan custom roles and associated policies by leveraging predefined roles and policies
Custom Roles and Policies	Restrict user access with custom roles, custom policies, and access filters
Services and CA Adapters	Configure the event log store, report server, subscription service, and certain event adapters
Log Storage	Configure auto-archive and restore archived databases.
Subscription	Maintain the subscription configuration, apply updates, and restore a subscription backup

Section	Describes how to
Filters and Profiles	Limit the data displayed in one report or query or in all reports and queries with filters. Limit the tag list, query list, and report list with profiles.
Queries and Reports	Create, edit, and import or export queries and reports to view current and recent event logs.
Action Alerts	Create an action alert to notify users or SNMP trap destinations or to run an IT PAM process when specified events occur
Scheduled Reports	Schedule and maintain report jobs; view and annotate generated reports
Suppression and Summarization	Create and use summarization and suppression rules to reduce server load and prevent collection or processing of unwanted events
Mapping and Parsing	Create and use mapping and parsing rules to refine raw events in various formats into standardized, CEG-compatible values and also create event forwarding rules
Integrations and Connectors	Create product integrations, which when deployed as connectors, let you refine and transmit events from a single event source to the CA Enterprise Log Manager server
Agents	Plan agent usage, prepare for agent installation, configure agents and agent groups, and apply subscription updates to agents
Custom Certificates	Implement custom certificates to replace the predefined certificates.
Accessibility Features	Use accessibility controls
Accessing Collected Events with ODBC/JDBC	Configure custom reports with a third-party reporting utility or retrieve selected log information with third-party products

Note: For details on operating system support or system requirements, see the *Release Notes*. For a tutorial on getting a single-system up and running so you can view results of queries on collected syslog and Windows events, see the *Overview Guide.* For step-by-step procedures on installing CA Enterprise Log Manager and performing initial configuration, see the *Implementation Guide.* For details on installing agents, see the *Agent Installation Guide*. For help on using any CA Enterprise Log Manager page, see the online help.

Chapter 2: User Accounts

This section contains the following topics:

Self-Administration Tasks

Users with access to CA Enterprise Log Manager can change their own passwords and unlock a locked user account if the configured user store is the default, the CA Enterprise Log Manager user store.

When the Administrator creates a new user account, a new password is assigned. The user changes that password during the first login session to a new password that conforms to the password policies for whether a password matching the username is permitted, minimum and maximum length, maximum number of repeating characters, and minimum number of numeric characters. It is the user's responsibility to change passwords within the frequency range specified by the password policies related to minimum and maximum password age.

Individual users administer their own accounts in the following ways:

- Change passwords in a manner consistent with password policies
- Unlock user accounts that have been locked, if permitted by the related password policy

More information:

Unlock a User Account (see page 22)
Change Your Password (see page 22)

Unlock a User Account

You can unlock a locked user account regardless of your role, if permitted by the password policy. When your account becomes locked, another user must unlock it so you can have access to the privileges granted to your role.

Locks and unlocks are controlled by the following two password policies:

- Lock user account after <n> failed logins
- Allow users to unlock passwords

User accounts can become locked if the password policy is set to lock user accounts after a certain number of failed logins and the user logs in with invalid credentials a number of times that exceeds the specified threshold.

Any user can unlock the account of another user if the password policy to allow users to unlock passwords is set. You need the user's password to unlock that user account.

To unlock a user account

1. Click the Administration tab and the User and Access Management subtab.
2. Click Unlock User on the left pane.
3. Enter the user name and password, and then click Unlock.

   The user account is unlocked.

Change Your Password

You can change your own password, regardless of your role. If the password policy for maximum password age is set, you should change your password with a frequency consistent with that policy.

Be sure to change your password as soon as possible in the following cases:

- If you give someone your password for the purpose of unlocking your account
- If you forget your password and the Administrator resets it for you

To change your password

1. Click the Administration tab and the User and Access Management subtab.
2. Click Change Password on the left pane.
3. Enter your old password.
4. Enter your new password twice.
5. Click OK.

Tasks Associated with Roles

Administrators assign roles to users based on the tasks they are to perform. You can assign users the predefined roles of Auditor, Analyst, and Administrator or to custom roles you create. To evaluate the impact of using predefined roles, review the tasks associated with each role.

More information:

Auditor Tasks (see page 24)
Analyst Tasks (see page 25)
Administrator Tasks (see page 26)

Auditor Tasks

Internal Auditors can perform tasks such as the following:

- Search for and select a query and view query results
- For a selected query result row, run an event/alert output process that is configured in CA IT PAM
- View current reports
- Schedule reports
- View the scheduled report job list
- View the generated report list
- View and annotate generated reports
- Set filters and profiles

You can assign the low-privileged role of Auditor when you create user accounts for third-party personal. For example, when a scheduled alert runs an event/alert output process at the query level, the alert sends a URL to CA Enterprise Log Manager that is appended to the description. For the third-party personnel to be able to browse to CA Enterprise Log Manager, they need user accounts.

Note: Analysts and Administrators can perform all Auditor tasks and their role-specific tasks.

External Auditors who are given temporary access to the CA Enterprise Log Manager for the period of the site audit can verify compliance to standards in areas such as the following:

- Verify that logs are collected from the expected sources.
- Verify that procedures are in place for preventing data loss. For example, verify that data is backed up frequently enough that nothing is missed.
- Verified that logs are being regularly reviewed to detect security breaches.
- Verify that logs are properly stored in a secure archive.
- Verify that the age of the archived data meets the standards for log retention.
- Verify that the content of the logs includes content mandated for retention.

More information:

How to Schedule a Report Job (see page 405)
View a Generated Report (see page 403)
Annotate a Generated Report (see page 404)
View a Query (see page 223)
View a Report (see page 224)
Example: Run an Event/Alert Output Process with Selected Query Results (see page 328)
Example: Send an Alert that Runs an IT PAM Process Per Query (see page 338)

Analyst Tasks

System analysts monitor the log collection network and then gather and distribute report data.

Administrators assign the *Analyst role* to users who are responsible for the following tasks:

- Create and edit custom reports and queries

  A report is a graphical or tabular display of event log data that is generated by executing predefined or custom queries with filters. The data can be from hot, warm, and defrosted databases in the event log store of the selected server and, if requested, its federated servers.

- Schedule action alerts

  An action alert is a scheduled query job, which can be used to detect policy violations, usage trends, logon patterns, and other information that can require near-term attention. Alert data can be viewed in the UI or through an RSS Feed. You can send a scheduled alert to email recipients, an SNMP trap destination, or a CA IT PAM event/alert output process. You can run the process once per row or once per query.

- Create tags

  A *tag* is a term or key phrase that is used to identify queries or reports that belong to the same category. To add a new report to a scheduled job configured to select reports with a specific tag, you add the common tag to the new report. A tag can also be a key phrase associated with a query, thus describing the query content and enabling key phrase-based classification and search.

- View RSS (Rich Site Summary) feeds

  An *RSS event* is an event generated by CA Enterprise Log Manager to convey an Action Alert to third-party products and users. The event is a summary of each Action Alert result and a link to the result file. The duration for a given RSS feed item is configurable.

Analysts can take the following approach as they become familiar with working with CA Enterprise Log Manager:

1. Examine the available predefined reports. (Auditors can also do this.)
2. Design custom reports, create tags for them, schedule, view, and annotate.
3. Schedule reports of interest for regular generation. (Auditors can also do this.)
4. Review generated reports and enter annotations. (Auditors can also do this.)
5. Identify criteria for sending an alert, the format to use, and the recipient. Then, schedule the alert to be generated when the criteria are met.

More information:

How to Create a Report (see page 261)
Tag Tasks (see page 222)

Administrator Tasks

Users assigned the role of Administrator have unlimited access to functionality available from all CA Enterprise Log Manager tabs. Only users assigned the role of Administrator have full access to the Administration tab. From the Administration tab, Administrators configure and maintain all aspects of log collection, all services, and all user access.

More information:

Log Collection Configuration and Customization (see page 27)
Services Configuration and Monitoring (see page 29)
User and Access Management (see page 31)

Log Collection Configuration and Customization

Only users with the role of Administrator can configure and maintain features related to log collection. Administrators perform log collection tasks from the Administration tab, Log Collection subtab.

Administrators use the Log Collection Explorer to configure connectors on agents, which is required for log collection. They also apply subscription updates to agents, when applicable.

Working with the event refinement library is optional. The out-of-the-box functionality, which is regularly updated, is designed to meet the needs of most customers.

Administrator tasks involving log collection include the following:

- Configure and manage the installed agents from a CA Enterprise Log Manager server dedicated to collection.
- Query the Archive Catalog on a CA Enterprise Log Manager reporting server.

  The *archive catalog* is the record of all databases that have ever been on the CA Enterprise Log Manager server. The archive catalog includes recently created databases, databases that have been backed up and moved, and databases that have been deleted before they were backed up, if any.
- Configure CA adapters used by CA Audit.
- Manage the event refinement library.
  - Work with predefined integrations and create new integrations from scratch or based on a copy of a predefined integration.

    Integration is the means by which unclassified events are processed into refined events so that they can be displayed in queries and reports.
  - Create a syslog listener from the predefined listener.
  - Create new parsing files from scratch or based on a copy of a selected predefined file.

    A message parsing file (XMP), associated with a specific event source type, applies parsing rules that break down the raw event into name/value pairs.
  - Create new mapping files from scratch or based on a copy of a selected predefined file.

    Data mapping (DM) files are XML files that use the CA Common Event Grammar (CEG) to transform events from the source format into one that can be stored for reporting and analysis in the Event Log Store.
  - Create summarization rules from scratch or based on a copy of a selected predefined rule.

    Summarization rules are rules that combine certain native events of a common type into one refined event.

- Create suppression rules from scratch or based on a copy of a selected predefined rule.

  Suppression rules prevent certain refined events from appearing in your reports.

- Create event forwarding rules.

  Event forwarding rules specify that selected events are forwarded to third-party products, such as those that correlate events, after being saved in the event log store.

- Create profiles.

  A profile specifies the set of data filters and tags that appear for selection. Data filters limit the data displayed in query or report; tag filters limit the tags displayed in the query tag list and in the report tag list.

More information:

How to Create an Agent Group (see page 544)
How to Configure Agent Management (see page 546)
How to Apply Subscription Updates (see page 552)

Services Configuration and Monitoring

Only users with the role of Administrator can configure and maintain the services accessible from the Administration tab, Services subtab. Configure all services soon after installing CA Enterprise Log Manager.

Administrator tasks involving services include the following:

- Configure global services including the following:
  - Update interval
  - Session timeout
  - Whether authentication is required to view alerts posted to RSS feeds
  - Tags to hide
  - Default profile
- Configure the event log store service.
  - At the global level, configure services that apply to all CA Enterprise Log Manager servers.
  - At the local level, configure auto-archive.

    The event log store on the collection CA Enterprise Log Manager server houses a hot database of new logs. The hot database is compressed into a warm database when it reaches the configured maximum rows.

    - If you configure auto-archive between the collection server and the reporting server, the warm database is copied to the reporting server and then deleted from the collection server.
    - If you configure auto-archive between the reporting server and a remote server that is not a CA Enterprise Log Manager server, warm databases are copied to the remote server on a daily or hourly basis, depending on your auto-archive settings. By default, the reporting server retains the databases until they reach the set disc space or time limits. They are automatically deleted from the reporting server only if you select the Remote ELM Server check box.
    - If you do not configure auto-archive from the reporting server to a remote storage server, manually create a backup of warm databases and move them to long-term storage. The warm database is stored in the event log store of a reporting server (or management/reporting server in a two-server deployment) for the number of days configured as max archive days unless the available space drops below the configured percentage set by archive disk space. In this case, warm databases are deleted, beginning with the oldest.
- Configure the report server service.

The report server service handles reports and alerts, including retention policies, the format for printed/emailed reports, and keyed values for reports and alerts. Additionally, it handles integration settings for CA IT PAM processes such as event/alert output and dynamic values and for SNMP trap destinations for alerts.

- Configure subscription updates.

  Subscription updates refer to the binary and non-binary files that are made available by CA Subscription Server to CA Enterprise Log Manager servers, the CA EEM component on the management server, and agents.

- Manage the federation of CA Enterprise Log Manager servers.

  At the management server, you can set a query to extend to federation children and peers. CA Enterprise Log Manager servers can be federated for two purposes:

  - The collection server auto archives each hot database into a warm database and sends it to the related reporting server. Create a federation between the collection server and the reporting server. When you query the management server with federation selected, you can get results not only from the local warm databases but also the hot database of the collection server.
  - Multiple reporting servers can be created to distribute the storage of warm databases among multiple event log stores. Reporting servers can be federated in a mesh with each collection server federated as a child to its reporting server. A federated query from any of the meshed reporting servers returns data from both its meshed (peer) servers and from all of their child servers.

    Note: If you create a restore point CA Enterprise Log Manager for the purpose of restoring archived databases from long-term storage, it is a good practice to leave such a server out of the federation.

- Monitor and manage the system status.

More Information:

Automating Backup and Restore (see page 155)
Query the Archive Catalog (see page 165)
Restore Auto-Archived Files (see page 167)
Restore–Script for Restoring Archived Databases (see page 168)
ODBC Server Considerations (see page 133)
Create a Diagnostics File for Support (see page 147)
Reboot a Host Server (see page 147)
Restart the ELM Services (see page 148)
Review Service Status and Version (see page 148)
Review System Status Self Monitoring Events (see page 149)

User and Access Management

Only users with the role of Administrator can configure and maintain user accounts, policies, and other application objects accessible from the Administration tab, User and Access Management subtab. To log on to CA Enterprise Log Manager, users must have a user account configured with a role and credentials for logging in. Predefined roles and policies enable Administrators to set up user access by defining user accounts. Creating custom roles and policies is optional.

Administrator tasks involving users and access include the following:

- Define new global users (if the user store is the default, the CA Enterprise Log Manager user store).

  When you add a new user, you create a global user. Details such as name, location, and telephone number are considered global because they can be shared. A *global user* is the user account information that excludes application-specific details.

- Retrieve referenced users (if the user store is a referenced user store).

  Global user details are stored in the configured user store, which can be an external directory.

- Assign predefined or custom application groups (roles) to new or referenced users.

  Application details are always stored in the repository of the management server. They are the details loaded in read-only format when you configure an external user store.

- Edit, delete, and view user accounts.

- Create custom application groups (roles) and associated policies.

  Creating user roles begins with defining a new application user group and then creating a policy defining the actions are permitted on specified resources. A user role can be a predefined application user group or a user-defined application group. Custom user roles are needed when the predefined application groups (Administrator, Analyst, and Auditor) are not sufficiently fine-grained to reflect work assignments. Custom user roles require custom access policies and modification of predefined policies to include the new role.

- Edit, delete, and view application groups and associated policies.

- Edit the CALM application access policy.

  The CALM Application Access policy is an access control list type of scoping policy that defines who can access the CA Enterprise Log Manager. By default, the [Group] Administrator, [Group] Analyst and [Group] Auditor are granted access.

- Create, edit, delete, and view access policies.

  An access policy is a rule that grants or denies an identity (user or user group) access rights to an application resource.

- Configure, edit, delete, and view access filters.

An access filter is a filter that the Administrator can set to control what event data non-Administrator users or groups can view. For example, an access filter can restrict the data specified identities can view in a report. Access filters are automatically converted into obligation policies.

More information:

Create a Global Group (see page 33)
Create a Global User (see page 34)
Assign a Role to a Global User (see page 35)
Back Up All Access Policies (see page 51)
Restore Access Policies (see page 55)
Configuring Custom User Roles and Access Policies (see page 73)
Add an Identity to an Existing Policy (see page 77)
Create a CALM Access Policy (see page 78)
Create a Dynamic User Group Policy (see page 87)
Create a Policy Based on an Existing Policy (see page 84)
Create a Scoping Policy (see page 81)
Create an Access Filter (see page 88)
Create an Application User Group (Role) (see page 75)
Grant a Custom Role Access to CA Enterprise Log Manager (see page 76)
Test a New Policy (see page 86)
Create a Calendar (see page 90)

How to Configure Accounts with Out-of-the-Box Settings

If you are setting up a temporary test environment, you can set up user and access management very quickly if you use out-of-the-box settings for User Accounts and configure only required fields. To complete minimal configuration with predefined settings, create user accounts for CA Enterprise Log Manager users as follows:

- If using the default user store, create an account with a user name, assign a predefined application group (Administrator, Analyst, Auditor), and assign a temporary password.
- If referencing an external user store, search for the global user by name, assign a predefined role (Administrator, Analyst, Auditor), and assign a temporary password.

More information:

Create a Global User (see page 34)
Assign a Role to a Global User (see page 35)
How to Manage Referenced User Accounts (see page 36)

Create a Global Group

The ability to create a global group depends on the configuration of the user store. Consider the following:

- If using the default user store, creating global groups is an optional task.
- If referencing an external user store, global groups and user accounts are automatically loaded into the default user store. You can, optionally, create custom policies for these global groups, but you cannot create new global groups.
- If the referenced user store is CA SiteMinder, you can use the global groups defined in this CA product as is or you can create new global groups from existing group memberships.

To create a global group

1. Click the Administration tab and then click the User and Access Management subtab.
2. Click Groups on the left pane.

   The Search Groups and User Groups panes appear.
3. Click the New Global Group button next the Global Groups folder.

   The New Global user Group pane appears.
4. Enter a name and, optionally, a description.
5. If this global group is to contain other global groups, do the following:
   a. Enter search criteria to display a group and click Search.
   b. Move the group you want to include to the Selected Global User Groups list.
   c. Repeat until the list contains all of the groups you want to select.
6. Click Save.

   A confirmation appears.

More information:

User Role Planning (see page 71)

Create a Global User

You can create new users only if the user store is configured as the CA Enterprise Log Manager user store, the default. Only Administrators can create new user accounts.

If referencing an external user store, user accounts are automatically loaded into the default user store as read-only records. If you need to create a new user, you must do so in the external user store. The new record is automatically loaded.

To use the CA Enterprise Log Manager product, a user must have a global user account. The account must be active at the time of login. Accounts can become inactive if suspended by the Administrator, locked due to violation of a password policy, or disabled due to the enabled account time having elapsed.

To create a new global user account

1. Click the Administration tab and the User and Access Management subtab.
2. Click the Users button.
3. Verify that the account you plan to create does not exist. Select Global Users and click Go. If the name does not appear in the results, proceed.
4. Click the New User button to the left of the Users tree.

   The New User page appears.
5. Enter the name of the user in the Name entry field.
6. (Optional) Assign an application user group.
   a. Click Add Application User Details.
   b. Select one or more available user groups and click the move button to move the selection to the Selected User Groups box.

      Note: If you do not do this now, you can edit the account of a global user later to assign an application user group.
7. Enter the General information for Global User Details.
8. (Optional) Assign a global user group.

9. Complete Authentication information:
    a. To set a threshold for the number of incorrect logins to accept before locking the account, enter a number for Incorrect Login Count. Configuring a count of 0 means there is no limit.
    b. Accept the cleared check box for Override Password Policy unless you want to permit this user to have passwords that do not conform to the password policy.
    c. Repeat your entry in the Confirm Password box.
    d. Select the Change Password at Next Login to permit the user to change the password.
    e. Leave Suspended clear when creating a new account.
    f. Enter a new password for New Password and Confirm Password.
    g. If this user is to have access only temporarily, enter a date range for enabling and disabling the user account.
    h. To defer the enabling of the user account to a later date, enter the date to enable the account.
10. Click Save.
11. Click Close.

Assign a Role to a Global User

You can search for an existing user account and assign the application user group for the role you want the individual to perform. If you reference an external user store, the search returns global records loaded from that user store. If your configured user store is the CA Enterprise Log Manager user store, the search returns records created for users in CA Enterprise Log Manager.

Only Administrators can edit user accounts.

To assign a role, or application user group, to an existing user

1. Click the Administration tab and the User and Access Management subtab.
2. Click Users on the left pane.

    The Search Users and Users panes appear.
3. Select Global Users, enter search criteria, and click Go.

    If the search is for loaded user accounts, the Users pane shows the path and the path labels reflect the referenced external directory.

    Important! Always enter criteria when searching to avoid displaying all entries in an external user store.

4. Select a Global User that has no membership in a CA Enterprise Log Manager application group.

   The User page displays with the folder name, global user details, and, if applicable, global group membership.

5. Click Add Application User Details.

   The "CAELM" User Details pane expands.

6. Select the desired group from Available User Groups and click the right arrow.

   The selected group appears in the Selected User Groups box.

7. Click Save.

8. Verify the addition.

   a. On the Search Users pane, click Application User Details and click Go.

   b. Verify that the name of the new application user appears in the displayed results.

9. Click Close.

How to Manage Referenced User Accounts

You can use global user account information when you reference an external user store. Although you cannot update the user record in the external user store from CA Enterprise Log Manager, you can assign application-level details.

Consider the following approaches to managing access for users with accounts stored in an external user store.

- Add a predefined application user group, or role, to the user account.
- Add the global group to the predefined policies that provide the access you want the user to have.
- Create custom roles and associated policies and add the custom role to the user account.

User Activation Guidelines

Consider the following guidelines when using account activation features:

- When you create a number of user accounts at once, you can use Enable Date to specify a future date on which to activate all or selected accounts. This lets you stage your roll-out of access rights to coordinate with any training you plan to provide your new users.
- When you create temporary accounts for external auditors, you can use the Enable Date and Disable Date settings to specify a given time interval.
- When you encounter suspicious behavior by any user, you can immediately mark the corresponding account as suspended and, thereby, prevent that user from successfully logging in to any CA Enterprise Log Manager server.
- When a user leaves the company, you can delete the entire user record, mark it suspended, or enter an expiration date to put it into disabled status.

More information:

Create a Global User (see page 34)
Edit a User Account (see page 37)
Delete a User Account (see page 40)

Edit a User Account

Only Administrators can create and edit user accounts. You can search for a user and display the selected user account information for any of the following reasons:

- To assign a CA Enterprise Log Manager role, that is, application group membership, to a global user, where account information was loaded from a referenced user store
- To update global user details for an account in the local user store
- To suspend the user account
- To reset the password for a user account either because the password was forgotten or the account is locked and the password policy does not allow users to unlock user accounts
- To disable a user account or reset the duration for the enabled account

Important! Make no entry in the Incorrect Login Count field in the Authentication area. The value displayed in this field is updated by the system.

To edit a user account

1. Click the Administration tab and the User and Access Management subtab.
2. Click Users on the left pane.

   The Search Users pane appears.
3. Specify search criteria on the Search Users pane in one of the following ways:

   - To add application details for global user, select Global Users, enter search criteria, and click Go.
   - To edit the account of a user with an existing CA Enterprise Log Manager role, select Application User Details, enter search criteria, and click Go.

   Note: For search criteria, use the operator LIKE when you specify a wildcard as the value and use the operator EQUAL when you specify the complete string. Examples follow:

   - Group Membership LIKE Aud*
   - Group Membership EQUALS Auditor

   The names of users meeting the search criteria appear in the Users pane.
4. Click the user name of the account to edit.

   The selected account appears in the right pane.
5. To add a role, click Add Application User Details, select the appropriate role from Available User Groups, and move it to Selected User Groups.
6. To update global user details, replace existing details with the new details in the Global User Details section.

   Note: You can update details only if the using the default user store.

7. To update authentication configuration, do any of the following:
   - Select Override Password Policy to exclude this user from checks performed by all password policies.
   - Select Suspended to prevent this user from logging in to any CA Enterprise Log Manager server.
   - Clear Suspended to activate this account so the user can log in.
   - If your password policy is set to disallow users to unlock passwords and this user has a locked password, select Reset Password, enter the new password twice, and select Change Password at Next Login.

     Note: The Incorrect Login Count field is automatically incremented for a failed login attempt and reset to 0 with a successful login attempt. A user account becomes locked if the incremented value reaches or exceeds the password policy value set for lock user account after the specified number of failed logins.
   - Set a duration for when this account is to be enabled by clicking Enable Date and setting a start date and then clicking Disable Date and setting an end date. Users have access from the beginning of day on the start date to the end of the day on the end date. To allow access for one day, specify the same date as start and end.
8. Click Save.

   Updates to the user account are saved and in force.

Reset a User Password

You can reset the password for users that forget their password. If a user gets locked out for exceeding the configured number of attempted logins that fail because of a forgotten password, you can reset the password and then the user can unlock the account, if allowed by the corresponding password policy.

To reset a user password

1. Click the Administration tab and the User and Access Management subtab.
2. Click the Users button.
3. Search for the user account to edit.
   a. Select Application User Details.
   b. Enter the user name in the Value field, where Attribute is set to User Name and Operator is set to LIKE.
   c. Click Go.

4. Click the user name under the Users tree.

   The selected User account details appear.

5. In the Authentication pane, select Reset Password.

   The New Password and Confirm Password fields appear.

6. Enter the new password in the New Password and Confirm Password fields.

7. Click Save and then click Close.

Delete a User Account

You can delete any global user account that was created in CA Enterprise Log Manager.

You can inactivate a user account without deleting it in either of the following ways:

- You can set a disable date to disable an account as of the specified date.
- You can suspend an account so the associated user cannot access the CA Enterprise Log Manager interface.

To delete a global user

1. Click the Administration tab, the User and Access management subtab, and the Users button.

   The Search Users and Users panes appear.

2. Select either Global Users or Application User Details, specify search criteria, and click Go.

3. Select the user to delete from the list of existing users.

   The record for the selected user appears in the right pane.

4. Click Delete.

   A confirmation to delete this user appears.

5. Click OK.

   The confirmation message: Global User deleted successfully appears.

   Note: If you click Go again in the Search Users pane, the displayed list does not contain the name of the deleted user.

Chapter 3: Policies

Creating custom roles requires the editing of predefined policies and the creation of custom policies. Before you begin these tasks, it is helpful to examine the predefined policies associated with each of the predefined roles. It is a good practice to back up predefined access policies before doing any editing.

This section contains the following topics:

Introduction to Policies

An *access policy* is a rule that grants or denies an identity (user or user group) access rights to an application resource or a global resource. CA Enterprise Log Manager determines whether policies apply to the particular user by matching identities, resources, resource classes, and evaluating the filters. That is, a policy states the actions that are granted or denied to specific identities on specific resources. Policies that deny access to a given resource have precedence over policies that grant access to the same resource.

CA Enterprise Log Manager supports the following types of access policies:

- CALM Access policies
- Delegation policies
- Dynamic User Group policies (an alternative approach to custom application groups)
- Obligation policies (created automatically when you create an access filter)
- Scoping policies

CA Enterprise Log Manager is installed with predefined CALM access policies and scoping policies for three CA Enterprise Log Manager application user groups: Administrator, Analyst, and Auditor. These policies are sufficient if you plan to assign only the out-of-the-box application user groups to users performing the various roles.

Important! We recommend that you take a backup of the predefined policies that are provided with CA Enterprise Log Manager. If a CALM Access policy is inadvertently deleted, users cannot access CA Enterprise Log Manager until that policy is restored from a backup.

Predefined Access Policies

If you use the out-of-the-box features, where you assign a predefined application group (Administrator, Analyst, or Auditor) as the role for each user, you do not need to create any access policies. All required policies are predefined and are ready for use.

More information:

Examine Policies for All Users (see page 42)
Examine Policies for Auditors (see page 45)
Examine Policies for Analysts (see page 47)
Examine Policies for Administrators (see page 49)
Resources and Actions (see page 66)

Examine Policies for All Users

You can examine policies for all users. Edit the CALM Application Access policy to define custom roles. All custom roles must be added as identities to this policy.

To examine policies for all users

1. Click the Administration tab and then click the User and Access Management subtab.
2. Click Access Policies in the left pane.
3. Display the CALM Application Access policy as follows:
   a. Select Show policies matching name.
   b. Enter CALM*.
   c. Click Go.
4. Examine the CALM Application Access policy.

   This policy grants read and write access to the listed resources for all members of the default application user groups (Administrator, Analyst, and Auditor) and to others who use the CA Enterprise Log Manager API:

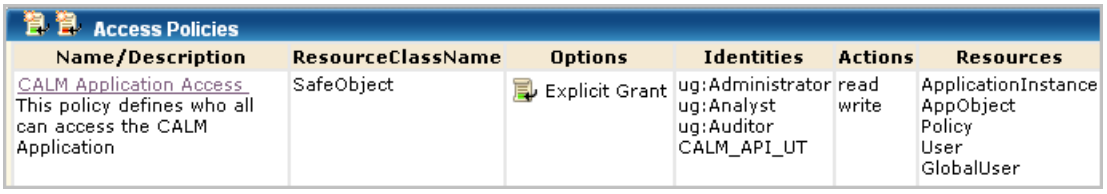

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
CALM Application Access This policy defines who all can access the CALM Application	SafeObject	Explicit Grant	ug:Administrator ug:Analyst ug:Auditor CALM_API_UT	read write	ApplicationInstance AppObject Policy User GlobalUser

The listed resources are as follows:

- The ApplicationInstance resource is CAELM, which refers to the CA Enterprise Log Manager product.
- Policy refers to access policies.
- User is any user added to a CA Enterprise Log Manager application user group.
- GlobalUser is any user defined in the user store within CA Enterprise Log Manager or referred to from CA Enterprise Log Manager.
- AppObject with the pozFolder value to the Profiles folder refers to profiles.
- AppObject with the pozFolder value to the flex folder refers to the dynamic time range XML used to populate the time ranges drop-down list in the Result Conditions step of query-based wizards.

The filter for CALM application access specifies the action limitations on each resource.

Filters

```
WHERE  ( req:resource     ==  val:ApplicationInstance
  AND    req:action       {}  val:read                    )
   OR  ( req:resource     ==  val:Policy
  AND    req:action       {}  val:read                    )
   OR  ( req:resource     ==  val:User
  AND    req:action       {}  val:read,write
  AND    name:cn          ==  req:identity                )
   OR  ( req:resource     ==  val:GlobalUser
  AND    req:action       {}  val:read
  AND    name:cn          ==  req:identity                )
   OR  ( req:resource     ==  val:AppObject
  AND    req:action       ==  val:read
  AND    name:pozFolder *--* val:/CALM_Configuration/Content/Profiles )
   OR  ( req:action       ==  val:read
  AND    req:resource     ==  val:AppObject
  AND    name:pozFolder *--* val:/CALM_Configuration/flex  )
```

5. Search for policies for all users as follows:
   a. Click Access Policies in the left pane.
   b. Select Show policies matching identity. Clear other selections.
   c. Enter [All Identities] in the Add identity field.
   d. Click Add.
   e. Click Go.

   Four policies appear, including the CEG Policy and the Default Data Access Policy. (If you do not explicitly enter [All Identities], many additional policies display.)

6. Examine the Default Data Access Policy.

   The predefined Default Data Access policy on the CALM resource class grants all users access to CA Enterprise Log Manager data to the extent specified in an access filter. An access filter is translated into an obligation policy with the FulfillOnGrant Action to dataaccess/CALM/Data.

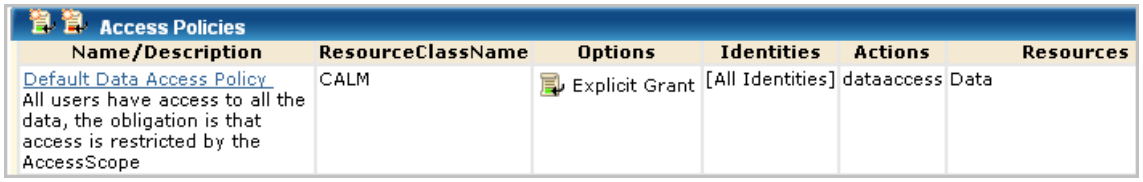

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Default Data Access Policy All users have access to all the data, the obligation is that access is restricted by the AccessScope	CALM	Explicit Grant	[All Identities]	dataaccess	Data

7. Examine the scoping policy, CEG Policy.

   The predefined CEG Policy grants all users with CALM Application Access the ability to view Common Event Grammar fields. Therefore the CEG fields appear in drop-down lists for simple and advanced filters for all users, because all users can set global and local filters for the queries they run. Users with rights to create and edit queries can set the filters for the queries they create and edit. This policy also helps ensure that all users can view the Global Configuration settings.

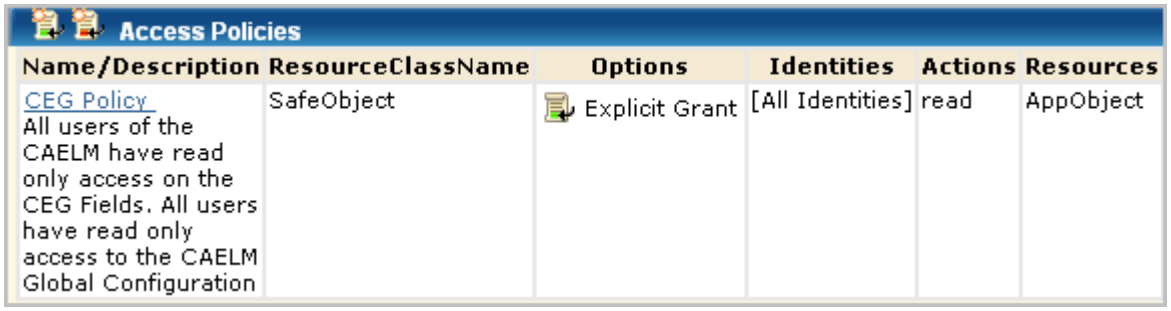

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
CEG Policy All users of the CAELM have read only access on the CEG Fields. All users have read only access to the CAELM Global Configuration	SafeObject	Explicit Grant	[All Identities]	read	AppObject

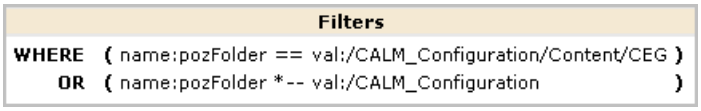

Filters	
WHERE	(name:pozFolder == val:/CALM_Configuration/Content/CEG)
OR	(name:pozFolder *-- val:/CALM_Configuration)

Examine Policies for Auditors

You can examine the predefined policies for Auditors to see how they limit application access to resources required to perform the following tasks.

- Schedule and annotate reports
- View reports

To examine predefined policies for Auditors

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies in the left pane.
3. Search for policies for Auditors as follows:
   a. Select Show policies matching identity.
   b. Enter ug:Auditor in the Add identity field.
   c. Click Add.
   d. Click Go.

   All policies for [All Identities] and ug:Auditor appear.
4. Examine the Auditor Schedule-Annotate Rights policy.

   All CALM access policies define the actions that can be performed against application-specific resources. This policy grants users assigned the application user group, Auditor, the ability to schedule and annotate reports.

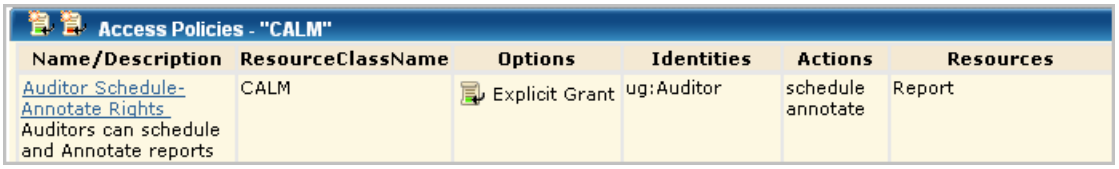
Access Policies - "CALM"

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Auditor Schedule-Annotate Rights Auditors can schedule and Annotate reports	CALM	Explicit Grant	ug:Auditor	schedule annotate	Report

Compare this policy with the Analyst Create-Schedule-Annotate policy and the Administrator Create policy.

5. Examine the Analyst Auditor Report Server Access Policy.

   This scoping policy gives Auditors the ability to set the report destination to any Report Server and to create a federated report, which requires access any Event Log Store. The resource listed in the policy is AppObject, where the application objects are the Report Servers and Event Log Stores.

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Analyst Auditor Report Server Access Policy Analyst ,Auditor can Schedule Reports and Alerts against all available Report Servers	SafeObject	Explicit Grant	ug:Analyst ug:Auditor ug:Administrator CALM_API_UT	read	AppObject

Filters	
WHERE	(name:pozFolder *--* val:CALM_Configuration/Modules/calmReporter)
OR	(name:pozFolder *--* val:CALM_Configuration/Modules/logDepot)

Note: For a given CALM Access policy, that is, policy for the CALM Resource Class, there is typically a related scoping policy for the SafeObject resource class.

6. Examine the Auditor View Report policy.

   This scoping policy grants users read access to reports. The resource listed in the policy is AppObject.

Scoping Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Auditor View Report Policy Auditor can view all the Reports	SafeObject	Explicit Grant	ug:Auditor	read	AppObject

   AppObject is limited to a specific application resource with a filter that grants the right to view reports. The path is an EEM folder path that stores the content of all reports.

Filters	
WHERE	(name:pozFolder *--* val:/CALM_Configuration/Content/Reports)

Examine Policies for Analysts

You can examine the predefined policies for Analysts to see how they limit application access to resources required to perform the following tasks:

- Schedule and annotate reports (Auditor tasks)
- View reports (Auditor task)
- Create reports and tags
- Create and schedule alerts (queries)
- Edit reports, alerts, and tags

To examine predefined policies for Analysts

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies in the left pane.
3. Search for policies for Analysts as follows:
   a. Clear the checkmark for Show policies matching name.
   b. Select Show policies matching identity.
   c. Enter ug:Analyst in the Add identity field.
   d. Click Add.
   e. Click Go.
4. All policies for ug:Analyst appear, including [All Identities] that includes this user group.
5. Examine the Analyst Create-Schedule-Annotate policy.

   This CALM access policy defines the actions that can be performed against application-specific resources. The policy grants users assigned the CA Enterprise Log Manager application user group, Analyst, the ability to create, schedule, and annotate reports, create and schedule action alerts, and create tags. (Auditors can only schedule and annotate reports.)

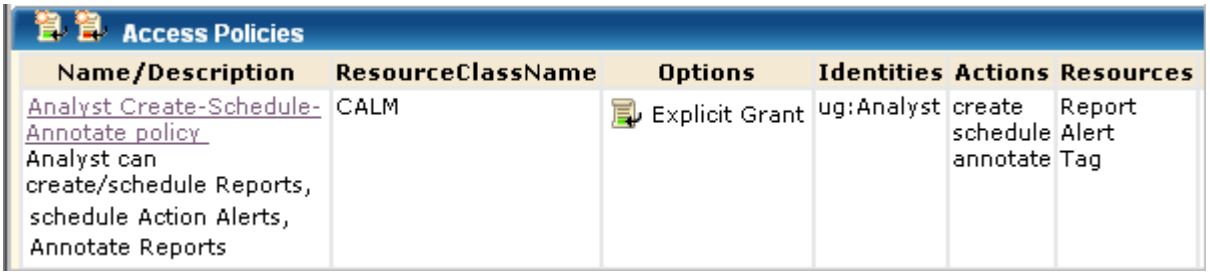
Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Analyst Create-Schedule-Annotate policy Analyst can create/schedule Reports, schedule Action Alerts, Annotate Reports	CALM	Explicit Grant	ug:Analyst	create schedule annotate	Report Alert Tag

6. Examine the Analyst Auditor Report Server Access Policy.

   This scoping policy grants Analysts schedule rights for any Report Server. The resource listed in the policy is AppObject.

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Analyst Auditor Report Server Access Policy Analyst ,Auditor can Schedule Reports and Alerts against all available Report Servers	SafeObject	Explicit Grant	ug:Analyst ug:Auditor ug:Administrator CALM_API_UT	read	AppObject

   AppObject is limited to specific resources with filters.

Filters
WHERE (name:pozFolder *--* val:CALM_Configuration/Modules/calmReporter) **OR** (name:pozFolder *--* val:CALM_Configuration/Modules/logDepot)

   - The filter ending in calmReporter grants read access to all Report Servers. When a report is scheduled, the destination Report Servers from which the generated report can be viewed is specified.
   - The filter ending in logDepot grants access to all Event Log Stores. When a report is defined as federated, queries are run against the data in all eligible Event Log Stores. Eligibility depends on the position in the hierarchy of the server on which the report is initiated, in hierarchical federations.

7. Examine the Analyst Report View-Edit policy.

   This scoping policy grants users assigned the Analyst role the ability to view, edit, or delete any report. The resource specified in the policy is AppObject.

Scoping Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Analyst Report View-Edit Policy Analyst can View/Edit any Report	SafeObject	Explicit Grant	ug:Analyst	read write	AppObject

   AppObject is limited to reports by the following filter, which grants the right to view generated reports saved in the EEM Folder /CALM_Configuration/Content/Reports.

Filters
WHERE (name:pozFolder *--* val:/CALM_Configuration/Content/Reports)

   Note: The ability to edit reports granted by this policy is extended by the CEG policy, which grants the right to add filters to reports using CEG columns.

Examine Policies for Administrators

Administrators assign the Administrator role to users who are to have full access to the CA Enterprise Log Manager application and all of its features. You can examine the predefined policies for Administrators to see how to grant access to those users who are to perform the following tasks:

- Create EventGrouping, that is, create suppression and summarization rules using common event grammar.
- Create Integration, that is, create data mapping and message parsing files using common event grammar.
- Create EventForwarding, that is, create rules to forward events to third-party systems.
- Run Database, that is, query with Archive Catalog Query for the names of databases that have been backed up and moved to an external archive solution.
- View or edit policies.
- View or edit user-defined calendars.
- View or edit any application object. Application objects are report templates, query templates, scheduled report jobs, alert jobs, profiles, service configurations, data mapping (DM) files, message parsing (XMP) files, suppression and summarization rules, and event forwarding rules.
- Create filters with the iPoz attribute of AppObject.
- View the folders listed under Administration, User and Access Management, EEM Folders and edit any user-defined data in these folders.
- View or edit details for any application user, application user group, or global user.
- Any Analyst or Auditor tasks.

To examine predefined policies for Administrators

1. Click the Administration tab and then click the User and Access Management subtab.
2. Click Access Policies in the left pane.
3. Search for policies for Administrators as follows:
   a. Select Show policies matching identity.
   b. Enter ug:Administrator in the Add identity field.
   c. Click Add.
   d. Click Go.

   All policies for [All Identities] and ug:Administrator appear.

4. Examine the CALM access policy, Administrator Create Policy.

   This policy defines the actions that can be performed against application-specific resources. The policy grants users assigned the application user group, Administrator, the ability to perform the specified actions as they apply to the specified resources.

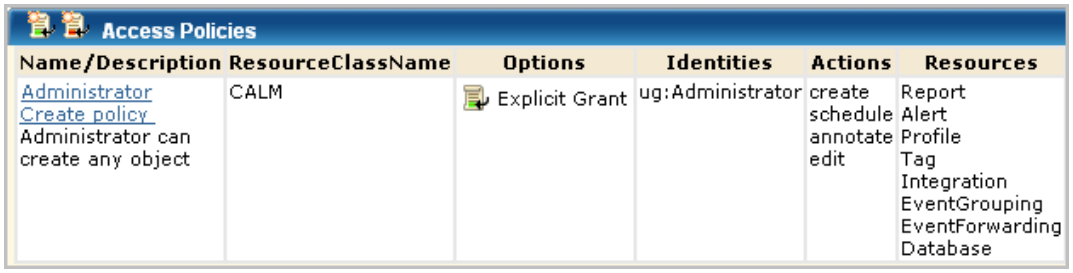

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Administrator Create policy Administrator can create any object	CALM	Explicit Grant	ug:Administrator	create schedule annotate edit	Report Alert Profile Tag Integration EventGrouping EventForwarding Database

5. Examine the CALM access policy, Admin Agent Manager Policy.

   The policy grants Administrators the right to create agent groups, edit all agent groups, configure connectors, and create integrations. The policy lets Administrators edit the Agent Authentication Key for the application instance of the CA Enterprise Log Manager server to which the agent transfers collected events. By default, the configured Agent Authentication Key applies to all CA Enterprise Log Manager servers across application instances, but can be set to be unique to the application instance.

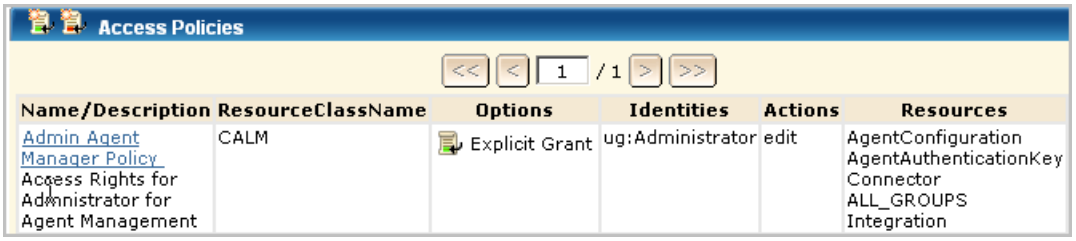

Access Policies

<< < 1 / 1 > >>

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Admin Agent Manager Policy Access Rights for Administrator for Agent Management	CALM	Explicit Grant	ug:Administrator	edit	AgentConfiguration AgentAuthenticationKey Connector ALL_GROUPS Integration

6. Examine the scoping policy, Administrator Default Policy.

   This policy grants Administrators the right to view, edit, or delete the listed resources. The listed resources are not specific to CA Enterprise Log Manager, and AppObject. AppObject refers to application-specific objects, which are resources listed in the CALM Administrator Create policy and in the CALM Admin Agent Manager policy.

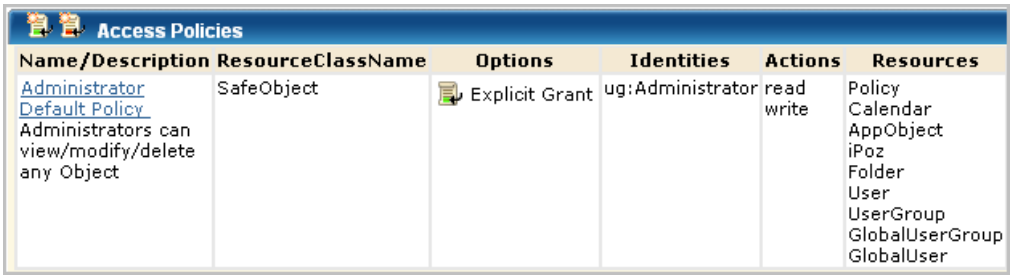

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
Administrator Default Policy Administrators can view/modify/delete any Object	SafeObject	Explicit Grant	ug:Administrator	read write	Policy Calendar AppObject iPoz Folder User UserGroup GlobalUserGroup GlobalUser

Access Policies for Registered Products

When a product is registered with CA Enterprise Log Manager, a new certificate is generated and certain access policies are updated to allow read only access to all tags, queries, and reports. Specifically, the certificate name that is used to authenticate the registered product is added as the Identity cert name to the following policies:

- CALM Application Access
- Analyst Auditor Report Server Access Policy
- Analyst Report View-Edit policy

The addition of the certificate name to the policies lets users of any CA product, third-party product, or CA customer get a list of queries and reports by tag. These users can display the lists within their own user interface and retrieve the refined event data they need.

Back Up All Access Policies

Exporting predefined access policies is a recommended way of preserving a backup in the event an access policy is inadvertently deleted or corrupted.

Important! Since corruption to policies can occur during a system or CA EEM service restart, it is important to have a current backup to restore. In addition, you should back up CA EEM periodically, for example, after an installation of a new CA Enterprise Log Manager and after creating custom policies.

You can export all of the policies for each type of access policy. When you export policies, an XML file is generated for each policy of the selected type. The XML files are zipped into a zip file named CAELM[1].xml.gz that contains the CAELM[1].xml document. You save the exported zip file to a directory of your choice.

Before you can restore your saved backup file, you need to copy them to the following directory of the CA Enterprise Log Manager with the internal user store: /opt/CA/LogManager/EEM. You can do this copy after a save to your local directory or wait and copy them only if a restore is needed.

The format in which policies are exported depends on the number of objects being exported.

- *filename*.tar.gz is used if there is a huge number of exported objects
- *filename*.xml.gz is used if there is a small to moderate number of exported objects

It is a good practice to rename *filename* (CAELM[n]), in a meaningful way when you do the export. For example, export the files from the three policy folders containing predefined policies as CAELM_CalmAccessPolicies, CAELM_EventPolicies, and CAELM_ScopingPolicies.

Note: The same extensions, xml.gz or tar.gz, must be maintained.

You can extract the XML file containing the access policy definition from the zip file and use it as input to the safex utility, used to restore the access policy.

To back up all access policies

1. Click the Administration tab and then click the User and Access Management subtab.
2. Back up the predefined CALM access policies as follows:
   a. Click the Access Policies button.
   b. Click CALM.

      That policy table, Access Policies - "CALM" appears
   c. Click the Export button.
   d. The File Download dialog appears with options to open or save.
   e. (Optional) Click Open to open the zip file, CAELM[1].xml.gz. Double click CAELM[1].xml to examine the file in XML format.
   f. Click Save to save the file.

      The Save As dialog appears.
   g. Select the target folder for save in, change the file name if desired, and click Save.

      If you do not change the file name, the zip file is saved as CAELM[1].xml.gz.
   h. Click Close.

      The Download Complete dialog closes. The policy list remains displayed in the left pane.

3. Back up the predefined Event Policies as follows:
   a. Click Event Policies.

      That policy table, Event Policies appears

   b. Click the Export button.

   c. The File Download dialog appears with options to open or save.

   d. Click Save to save the file.

      A message appears asking whether you want to replace the existing CAELM[1].xml.gz file.

   e. Click No.

   f. Enter a unique name in the file name field and click Save. For example, edit the entry to CAELM[2].xml.gz or enter a name for the policy type such as CAELM_EventPolicies.

   g. Click Close.

      The Download Complete dialog closes. The policy list remains displayed in the left pane.

4. Back up the predefined Scoping Policies as follows:
   a. Click Scoping Policies.

      That policy table, Scoping Policies appears

   b. Click the Export button. You may need to scroll horizontally to view the button in the top right corner.

   c. The File Download dialog appears with options to open or save.

   d. Click Save to save the file.

      A message appears asking whether you want to replace the existing CAELM[1].xml.gz file.

   e. Click No.

   f. Enter a unique name in the file name field and click Save. For example, edit the entry to CAELM[3].xml.gz or enter a name for the policy type such as CAELM_ScopingPolicies.

   g. Click Close.

      The Download Complete dialog closes. The policy list remains displayed in the left pane.

5. Click Close.

   The Access Policies list closes.

Example --CAELM[1].xml for CALM Access Policies

Following is an entry for one policy in the CAELM[1].xml file.

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
- <Safex>
    <Attach label="CAELM" />
    <Add />
  - <AddOrModify>
    - <Policy folder="/" name="Auditor Schedule-Annotate Rights">
        <Description>Auditors can schedule and Annotate reports</Description>
        <ResourceClassName>CALM</ResourceClassName>
        <PolicyType>policy</PolicyType>
        <Disabled>False</Disabled>
        <ExplicitDeny>False</ExplicitDeny>
        <PreDeployment>False</PreDeployment>
        <RegexCompare>False</RegexCompare>
        <Resource>Report</Resource>
        <Action>schedule</Action>
        <Action>annotate</Action>
        <Identity>ug:Auditor</Identity>
        <Attribute name="CreateTimestamp">20080926053329</Attribute>
      </Policy>
```

More information:

Manually Backing Up Archived Databases (see page 170)

Restore Access Policies

You can restore an access policy that has been deleted or changed in a way that causes problems. If an access policy is accidentally deleted or corrupted, users referenced as Identities in that policy cannot access CA Enterprise Log Manager until that policy is redefined or restored.

Restoring access policies requires running the safex utility for policies.

Use one of the two following procedures, depending on whether your export created a backup file with the xml.gz extension or the tar.gz extension.

To restore access policies from a backup named filename.xml.gz

1. Copy your saved backup files to following directory of the management CA Enterprise Log Manager, typically the first server installed.

   /opt/CA/LogManager/EEM

2. Run the following command to retrieve the XML file:

   gunzip *filename*.xml.gz

   This creates *filename*.xml.

3. (Optional) If you want to restore only one of the policies in the group that you backed up, do the following:

   a. Open the XML file.

   b. For the policies you do not want to restore, delete the XML lines beginning and ending with the following tags:
   <Policy folder="/ name=*policyname*> and </Policy>

   c. Save the file.

4. Execute the following command, where *eemserverhostname* refers to the host name of the management CA Enterprise Log Manager.

   ./safex –h *eemserverhostname* –u EiamAdmin –p *password* –f *filename*.xml

   When the CA Enterprise Log Manager server is in FIPS mode, be sure to include the -fips option.

   The policy or policies defined in filename.xml being restored are added to the appropriate policy type and put into effect.

To restore access policies from a backup named filename.tar.gz

1. Copy your saved backup files to following directory of the management CA Enterprise Log Manager, typically the first server installed.

   /opt/CA/LogManager/EEM

2. Run the following command to retrieve the XML file.

   gunzip *filename*.tar.gz

   This creates *filename*.tar.

3. Run the following command:

   tar –xvf *filename*.tar

   This creates *filename*.xml.

4. (Optional) If you want to restore only one of the policies in the group that you backed up, do the following:

   a. Open the XML file.

   b. For the policies you do not want to restore, delete the XML lines beginning and ending with the following tags:
   <Policy folder="/ name=*policyname*> and </Policy>

   c. Save the file.

5. Execute the following command, where *eemserverhostname* refers to the host name of the management CA Enterprise Log Manager.

   ./safex –h *eemserverhostname* –u EiamAdmin –p *password* –f *filename*.xml

To recreate the CALM Access Policy if you have no backup

If you have no backup, you can recreate the CALM Application Access policy.

1. Recreate the CALM Application Access policy. See "Predefined Policies."
2. Define the filters as shown in the following illustration. The partial paths are:
   - /CALM_Configuration/Content/Profiles
   - /CALM_Configuration/flex

Logic	(	Left type/value	Operator	Right type/value	)
NONE	(	request resource	STRING EQUAL ==	value ApplicationInstance	
AND		request action	STRING WITHINSET {}	value read	)
OR	(	request resource	STRING EQUAL ==	value Policy	
AND		request action	STRING WITHINSET {}	value read	)
OR	(	request resource	STRING EQUAL ==	value User	
AND		request action	STRING WITHINSET {}	value read,write	
AND		named attribute cn	STRING EQUAL ==	request identity	)
OR	(	request resource	STRING EQUAL ==	value GlobalUser	
AND		request action	STRING WITHINSET {}	value read	
AND		named attribute cn	STRING EQUAL ==	request identity	)
OR	(	request resource	STRING EQUAL ==	value AppObject	
AND		request action	STRING EQUAL ==	value read	
AND		named attribute pozFolder	STRING CONTAINS *--*	value ration/Content/Profiles	)
OR	(	request action	STRING EQUAL ==	value read	
AND		request resource	STRING EQUAL ==	value AppObject	
AND		named attribute pozFolder	STRING CONTAINS *--*	value ALM_Configuration/flex	)

The presence of this policy enables any Administrator to log in and create the other policies.

Chapter 4: Custom Roles and Policies

This section contains the following topics:

Guidelines for Creating Policies

All CALM access policies and scoping policies state the actions that are granted or denied to specific identities on specific resources. Policies for the CALM resource class grant or deny specified identities the ability to perform actions on application resources, also known as CALM resources. Policies for the SafeObject resource AppObject grant or deny specified identities write and read actions on an application-level resource, as indicated in the filters. Other policies for the SafeObject resource class grant or deny specified identities write and read actions on global resources.

The type of policy or policies you create depends on the resource to which you want to limit access. A summary of the policy requirements by resource follows:

- Resources requiring a CALM policy and scoping policies for AppObject
  - Event Forwarding
  - Event Grouping
  - Integration (non-agent)
  - Profile
  - Report

- Resources requiring only a CALM policy
  - AgentAuthenticationKey
  - AgentConfiguration
  - Alert
  - ALL_GROUPS
  - Connector
  - Database
  - Integration (agent-related)
  - Tag
- Resources requiring only scoping policies for the global resource
  - Calendar
  - Folder
  - GlobalUser
  - GlobalUserGroup
  - iPoz
  - Policy
  - User
  - UserGroup

The following guidelines highlight the differences in the approaches for creating policies, where differences are based on the resources you want to control.

To control access to EventForwarding, EventGrouping, Integration, Profile, and Report

The following approach applies only to policies on the CALM resources, EventGrouping, Integration, Profile, and Report. These application resources require a CALM policy and two scoping policies.

1. Create a CALM policy for one or more application resources such as Report or Integration. Specify one or more application-specific actions that are valid for the specified resources such as create, schedule, or annotate. Add the Identities to which the actions are granted or denied.
2. Create a companion scoping policy on the AppObject resource with both read and write actions. Specify the write action to let the identity edit or delete the resource, but not create it. Specify the read action to let the identity display or view the resource. Create a filter that ties back the AppObject resource to the related application resource. Specify in the filter the EEM folder path that stores the content for the specified resource or is a module for which access is required for the related application resource. Add the same Identities to this policy that you added to the related CALM policy.
3. Create a second companion scoping policy on the AppObject resource with the read action. Specify the read action to let the identity display or view the resource. Create a filter that ties back the AppObject resource to the related application resource. Specify in the filter the EEM folder path that stores the content for the specified resource or is a module for which access is required for the related application resource. Add lower-privileged users or user groups as Identities to this policy.

To control access to Alert, Database, Tag, and agent-related resources

The following approach applies to application resources that require only a CALM policy to grant or restrict access.

- Create a CALM access policy for a resource such as Connector or Tag. Specify the edit action to let the identity create, edit, and delete the resource and take any other valid action. Add the Identities to which this action is granted or denied.

  Note: Access to agent-related resources makes available the buttons from the Agent Explorer folder or its subfolders on the Log Collection subtab of the Administration tab. Access to the Alert resource lets the Identity access the Alerts tab. Access to the Tag resource lets the identity create a Tag for a custom query or report. Access to Database lets the Identity run an archive query.

To control access to global resources used in the CAELM application

The following approach applies to global resources, which require only a scoping policy to limit access.

1. Create a scoping policy for one or more global resources such as User or Policy. Specify the write action to let the identity create, edit, or delete the resource. Add the Identities to which this action is granted or denied.
2. Create a scoping policy for one or more global resources such as User or Policy. Specify the read action to let the identity view the global resource. Add the Identities to which this action is granted or denied.

   Note: Global resources are available with buttons on the User and Access Management subtab of the Administration tab.

More information:

CALM Access Policy Types (see page 62)
Resources and Actions (see page 66)
CALM Resources and EEM Folders (see page 68)
Global Resources and CA EEM Functionality (see page 71)
Create a CALM Access Policy (see page 78)

CALM Access Policy Types

When you create an access policy for CALM or a scoping policy, you select one of the following three types:

- Access policy
- Access control list
- Identity access control list

This choice impacts the level of detail for access policy configuration, where access policy is the broadest.

Note: The examples shown here are access policies for the CALM resource class and therefore, include actions and resources specific to CA Enterprise Log Manager.

An access policy specifies actions that are valid for any of the selected resources that are granted to all selected identities. When you create a generic policy for CA Enterprise Log Manager, you add resources belonging to the CALM resource class, and then you select actions from the displayed list. The actions you select apply to the selected resources for which they are valid. In this example, the policy allows each selected action to be performed on all the selected resources for which create is valid.

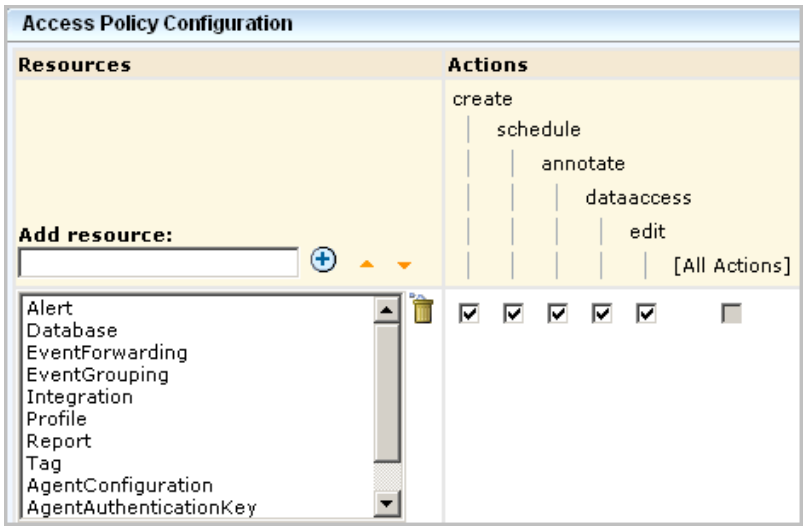

An access control list specifies actions valid for each resource separately for the selected identities. When you create a resource-centric policy, you specify what actions are permitted on each resource. You do not need to select actions for a given resource simply because they are valid. For example, you can allow create for reports but not allow create for alerts, even though create is a valid action for alerts. The access control list is the most fine-grained policy when it is implemented for one identity at a time.

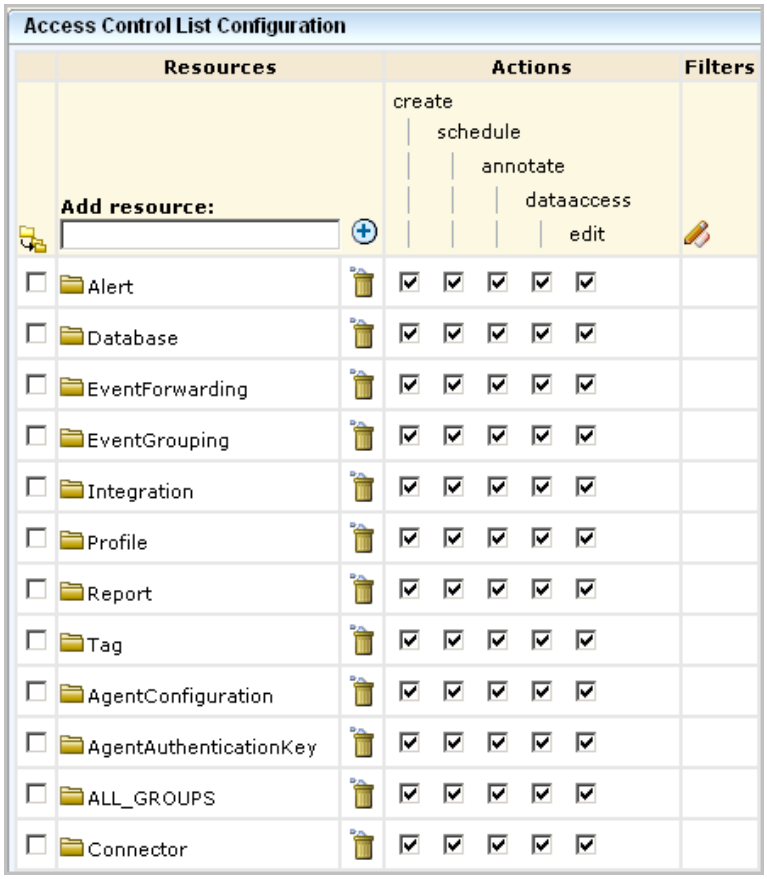

An identity access control list specifies the actions granted to selected identities for all applicable selected resources. When you create an identity-centric policy, specify which identities can perform which actions (create, schedule, annotate, edit) on all the listed resources to which each action applies. If you want to restrict the Auditor from scheduling alerts, you would leave schedule blank. However, leaving schedule blank would also restrict the Auditor from scheduling reports.

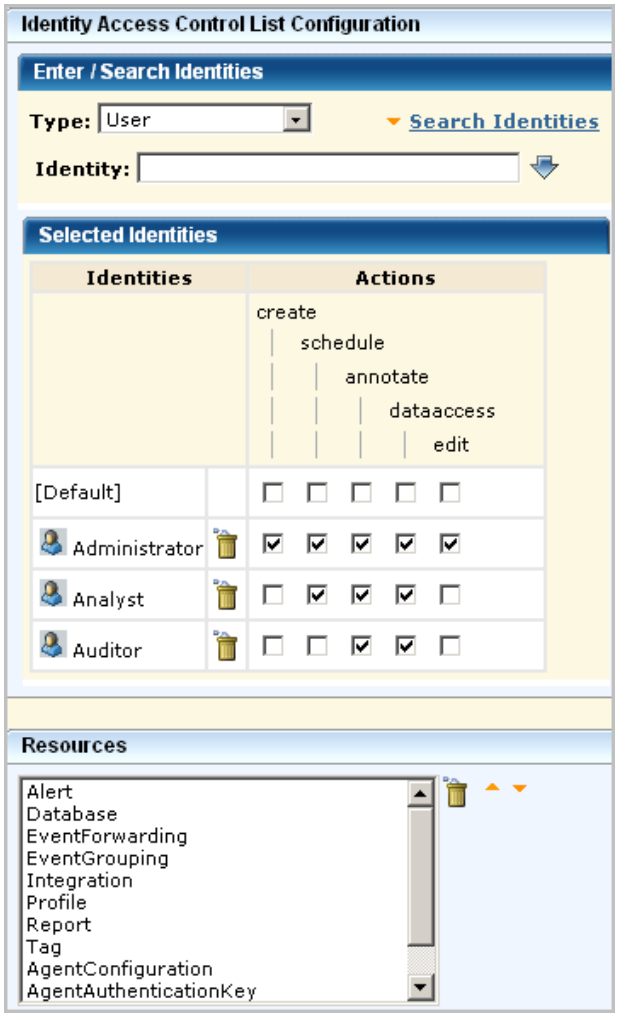

Resources and Actions

When creating policies, configure an access policy for which an access filter is needed. An access filter is a filter that the Administrator can set to control what event data non-Administrator users or groups can view. For example, an access filter can restrict the data that appears on reports viewed by the specified users or groups. Access filters are automatically converted into EEM Obligation Policies. Access filters are often expressed in terms of the relative paths for the objects to which user access is limited. You can view these relative paths in the EEM Folders area of the interface.

Typically, policies that authorize actions such as create and schedule are defined with the CALM resource class and CALM resources such as reports, tags, DM and MP files, and suppression and summarization rules. Policies that authorize the read and write actions are defined with the SafeObject resource class and the AppObject resource. The Edit action is the only valid action for agent-related resources in the CALM resource class.

More specifically, actions that can be authorized for objects belonging to the CALM resource class follow:

Action	Resource	Description
Annotate	Report	Record comments on reports
Create	EventForwarding	Create rules to forward specific events to specific third-party applications.
Create	EventGrouping	Create suppression and summarization rules using common event grammar
Create	Integration	Create data mapping and message parsing files using common event grammar
Create	Profile	Create profiles
Create	Report	Create reports and queries
Create	Tag	Create tags for reports and queries
Dataaccess	Data	Access the CALM event data, which can be limited by data access filters.
Edit	AgentConfiguration	Create agent groups. Configure installed agents with sources for collection and destination for processing
Edit	AgentAuthenticationKey	Create and edit the agent authentication key that is specified during agent installation

Action	Resource	Description
Edit	ALL_GROUPS	Edit all available agent groups **Note:** Access can be restricted to a particular agent group by specifying the Agent Group name as the resource
Edit	Connector	Configure connectors
Edit	Database	Determine the logs that exist that match the archive catalog query criteria and recatalog the database
Edit	Integration	Edit integration details
Schedule	Alert	Schedule action alerts
Schedule	Report	Schedule reports and queries

The actions that allow users to view or edit an object belonging to the SafeObject resource class follow:

Action	Resource	Description
Read	AppObject	View report templates, query templates, tags, scheduled report jobs, alert jobs, service configurations, data mapping (DM) files, message parsing (XMP) files, suppression and summarization rules, and event forwarding rules.
Read	Calendar	View the calendars listed under Administration, User and Access Management, Calendars
Read	Folder	View the folders listed under Administration, User and Access Management, EEM Folders
Read	GlobalUser	View information displayed for users listed when you query for Global Users under Administration, User and Access Management, Users
Read	iPoz	View the user store setting under Administration, User and Access Management, User Store View the password policy settings under Administration, User and Access Management, Password Policies
Read	Policy	View the policies listed under Administration, User and Access Management, Access Policies
Read	User	View User details when you query for Application User Details under Administration, User and Access Management, Users
Read	UserGroup	View the application group membership for users listed when you query for Application User Details under Administration, User and Access Management, Users

Action	Resource	Description
Write	AppObject	Edit or delete report templates, query templates, tags, scheduled report jobs, alert jobs, service configurations, data mapping (DM) files, message parsing (XMP) files, suppression and summarization rules, and event forwarding rules.
Write	Calendar	Edit user-defined calendars
Write	Folder	Edit user-defined data added to the EEM Folders structure
Write	GlobalUser	Edit global user details
Write	iPoz	Configure user store and password policies
Write	Policy	Edit user-defined and predefined policies
Write	User	Edit application user details
Write	UserGroup	Create, edit, or delete an application user group

CALM Resources and EEM Folders

For every custom CALM policy involving EventForwarding, EventGrouping, Integration, Profile, or Report that you create from scratch, you create a scoping policy on AppObject. The scoping policy has read/write access that filters on the EEM paths for each CALM resource listed in the corresponding CALM policy. The same user groups that are Identities for the CALM policy are assigned as Identities to this policy. To complete the policy set, create an additional scoping policy for read only, assign an Identity that can only view the resource, and enter a filter with an EEM folder path.

Note: Whether a CALM policy requires a supporting scoping policy depends on the resource that the CALM policy uses. For example, the Database, Tag, and Alert resources are purely CALM resources and no scoping policies are required for them. Scoping policies are not required for agent-related resources either.

You can view EEM folders from the Administration tab, User and Access Management subtab. When you select a folder such as Suppression, that path is displayed. See the following example:

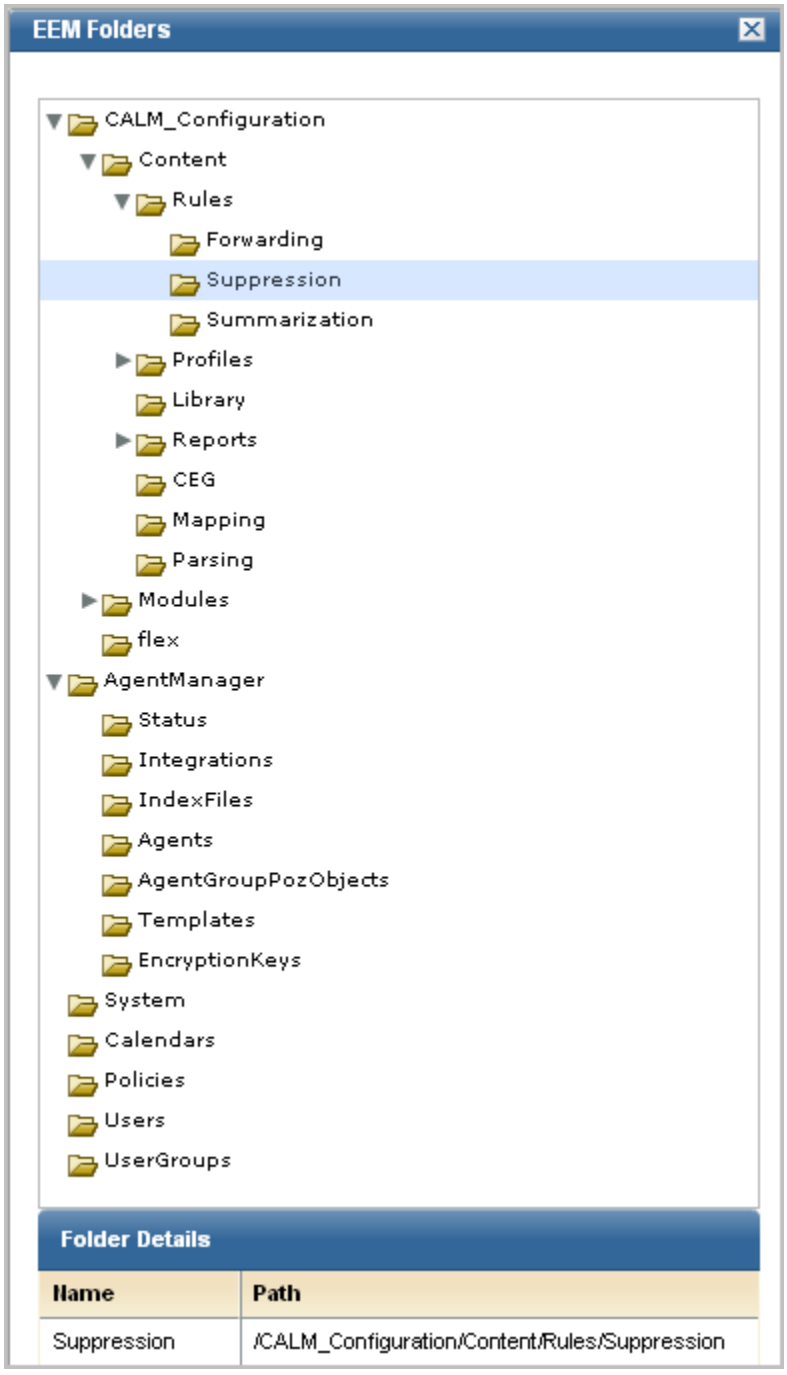

You specify the EEM folder path as the value in an expression that begins with pozFolder CONTAINS, as shown in the Filters section of a Policy definition. An example follows:

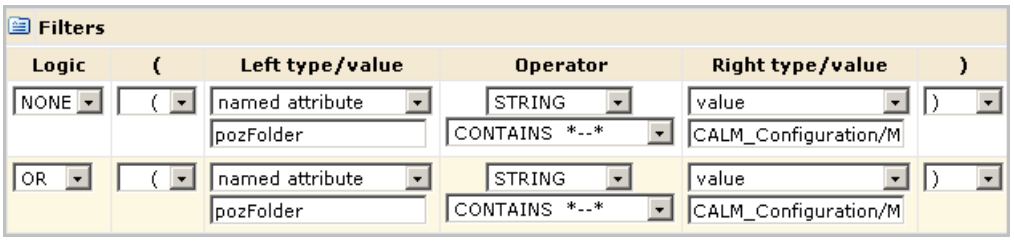

The following tables provide guidelines for the value specified in the filter of a scoping policy that is related to a CALM policy that grants, or denies, access to specific CALM resources.

Note: There is not a one-to-one correspondence between the CALM resources and the folders.

When creating a scoping policy that grants access to the content of this CALM resource	Add a filter specifying this EEM Folder path
EventForwarding	pozFolder CONTAINS /CALM_Configuration/Content/Rules/Forwarding
EventGrouping	pozFolder CONTAINS /CALM_Configuration/Content/Rules/Summarization pozFolder CONTAINS /CALM_Configuration/Content/Rules/Suppression
Integration (Server)	pozFolder CONTAINS /CALM_Configuration/Content/Mapping pozFolder CONTAINS /CALM_Configuration/Content/Parsing
Profile	pozFolder CONTAINS /CALM_Configuration/Content/Profiles
Report	pozFolder CONTAINS /CALM_Configuration/Content/CEG pozFolder CONTAINS /CALM_Configuration/Content/Reports

When creating a scoping policy that requires access to this CALM module	Add a filter specifying this EEM Folder path
Agent Manager	pozFolder CONTAINS /CALM_Configuration/Modules/AgentManager
Event Log Store	pozFolder CONTAINS /CALM_Configuration/Modules/logDepot
Report Server	pozFolder CONTAINS /CALM_Configuration/Modules/calmReporter
Subscription Module	pozFolder CONTAINS /CALM_Configuration/Modules/Subscription

Global Resources and CA EEM Functionality

You can create a scoping policy that is similar in intent to a CALM policy, except the resources are global rather than product-specific. Global resources are those resources that are used across multiple CA products. You can create policies that grant or deny access to specific global resources, accessed by buttons on the Administration tab, User and Access Management subtab.

Use the following table as a guide when creating a scoping policy that grants or denies specified Identities the ability to read and write where the resource specified is a global resource.

Task	Action	Global Resource
Show, create, edit, or delete a global user, a global user group, and an application user group (role); add an application group (role) to a global user or create a global user with a role.	read, write	User UserGroup GlobalUser GlobalUserGroup
Create, edit, copy, export, disable, test, view or delete a policy; add a calendar to a policy	read, write	Policy Calendar
Create, edit, copy, view, or delete an access filter; view EEM Folders	read, write	Policy
Create a calendar	read, write	Calendar
Configure the user store; create, edit, or view password policies	read, write	iPoz

When creating a filter for a global resource, refer to the filter for the CALM Application Access policy as an example. One of the things the filter does is to specify what actions go with what resource. If you click Edit on a predefined policy, you can examine the source for an example of how to enter the logic.

User Role Planning

If the predefined application user groups, Administrator, Analyst, and Auditor, are insufficient for your needs, you can create custom roles with new application user groups. For example, to assign a small group of individuals to manage user accounts, where these individuals have no access to unrelated functionality in the CA Enterprise Log Manager, you could define a UserAccountAdministrator role, create a scoping policy for that role, add that role to the CALM Application Access Policy, and assign that role to the users who are to manage user accounts.

User planning for CA Enterprise Log Manager involves the following steps:

- Determining the number of users needed to administer, analyze, and audit CA Enterprise Log Manager
- Identifying the users to grant CA Enterprise Log Manager access

If you are considering creating custom roles with associated access policies, consider the following approach:

- Identify the role to assign each CA Enterprise Log Manager user
- Identify the type of access to CA Enterprise Log Manager resources is required for each role

You can also consider the following alternatives to user-defined roles (application groups):

- Configure dynamic user group policies that create dynamic user groups.
- Create global groups and treat them as application groups. That is, assign them to users and assign them to policies as Identities.

  This approach may be useful if policies are created for the purpose of restricting access by geographical location and you want the same users to have the same level of rights on multiple CA products. For example, a global group for Location-A_Admin could be assigned to users that were to administer several CA products at Location-A. Policies for each CA product could be created that grant administrative rights to the servers on which that product was installed in Location-A.

More information:

Create a Global Group (see page 33)

Configuring Custom User Roles and Access Policies

A *user role* can be a predefined application user group or a user-defined application group. Custom user roles are needed when the predefined application groups (Administrator, Analyst, and Auditor) are not sufficiently fine-grained to reflect work assignments. Custom user roles require custom access policies and modification of predefined policies to include the new role.

Administrators can create user roles and corresponding policies as follows:

1. For each role assumed by users of CA Enterprise Log Manager:
   - Identify the resources to which access must be granted.
   - Identify the actions you want to permit on each resource.
   - Identify the identities, or individuals, to whom this role applies.

     Note: Identities can be other application groups designed to make up a super group.
2. If a predefined application group is too broad for your needs, create a new application group and assign this application group to the individuals you identified. It is good practice to name a user-defined application group with a term that describes the role the assigned users are to perform.
3. Add the new application group to the CALM Application Access policy, which is an Access Control List type.
4. If the new role needs to be able to take an action on one or more resources, such as create, do the following:
   a. Configure a CALM policy that allows the new application group to create or take other valid actions the identified CA Enterprise Log Manager resources.
   b. Configure a scoping policy that grants the new application group read and write access to the AppObject resource and specify a filter that states where the identified resource is stored in the EEM folders. For each filter, enter the named attribute, pozFolder, CONTAINS and a value, where the value is the EEM Folder path beginning with /CALM_Configuration.
5. If the new role only needs to view a specific CA Enterprise Log Manager resource, configure a scoping policy that permits read access to AppObject and specify a filter where the named attribute, pozFolder, CONTAINS a value, where the value is the EEM Folder path beginning with /CALM_Configuration where that resource is stored.
6. Test the policies.
7. Assign the new role to user accounts.

Administrators can also create restrict user access with access filters. If a particular kind of restricted access applies to only one individual, you can omit assigning that person an application group, or role. To limit the access of a user:

1. Create a user but assign no role.
2. Give the user access to the CA Enterprise Log Manager application by adding the user to the CALM access policy.
3. Create a scoping policy that grants read or write access to the SafeObject, AppObject and specify a filter where the named attribute pozFolder is equal to the value of the EEM folder for the resource. For example, if the resource is reports, set the named attribute calmTag equal to the value of a report tag.
4. Create a custom access filter.

Administrators can customize user access to the CA Enterprise Log Manager resources. Consider the following examples:

- Create roles to assign specific administration responsibilities to different groups of administrators. For example, create a role such as UserAccountAdministrator. Create a policy that grants users with this role access to only the functionality needed to maintain users and groups. Such a policy must define read and write access to the GlobalUser resource as well as to the User and UserGroup resources.
- Create roles to distribute responsibilities of analysts to the various types of reports and queries based on tags. For example, create roles such as SystemAccessAnalyst and PCIAnalyst and assign analysts to just one of the restricted analyst roles. Then create policies that grant access to a subset of these resources based on tag. For example, create a policy that grants the SystemAccessAnalyst role access to reports and queries that have the System Access tag and another that grants the PCIAnalyst role access to reports and queries that have the PCI tag. Create other roles and policies based on other tags. Policies that restrict access in this way do so with access filters.

Administrators can create server-based policies using either of the following approaches:

- Restrict data

  You can restrict access to specific logs by creating a data access filter, setting the filter for receiver_name field, and specifying a value such as systemstatus or syslog.

- Restrict configuration

  You can restrict access to a particular CA Enterprise Log Manager server by creating a policy on the SafeObject resource class with AppObject as the selected resource. That is, to restrict access just to the report server configuration on a particular host, define a filter such as the following:

  pozFolder contains /CALM_Configuration/Modules/calmReporter/LogServer01

More information:

Sample Policies for Custom Integrations (see page 116)
Sample Policies for Suppression and Summarization Rules (see page 117)
Create an Access Filter (see page 88)
Restricting Data Access for a User: Win-Admin Scenario (see page 98)
Restricting Access for a Role: PCI-Analyst Scenario (see page 110)

Create an Application User Group (Role)

You can create a new application user group to support the roles you need. Once you create a new application user group, you must create access policies for that group.

One case where new access policies are not needed for a new group is when that group is given memberships to existing groups. Consider the scenario where you need one role for individuals who are dedicated to creating data mapping and message parsing files, another role for individuals dedicated to creating suppression and summarization rules, and a third role for those who can perform either of these two tasks. You might define one application user group called AdminDMMP with a policy that grants create access to the Integration resource and another group called AdminSS with a policy that grants create access to the EventGrouping resource. You could then create a third group called AdminDMMPSS with memberships to the AdminDMMP group and the AdminSS group. This third group would automatically inherit the policies from the two membership groups.

Rather than creating new application groups or roles, you can expand the roles of the predefined Analyst and Auditor roles. For example, if you want Analysts to be able to create suppression and summarization rules and you want Auditors to be able to view these rules, you could create a CALM policy that grants the ability to create summarization and suppression rules and a scoping policy that grants the ability to view or edit custom rules and assign the Analyst role to those policies. You could then create a scoping policy that grants users the ability to view suppression and summarization rules and assign the Auditor group to that policy.

Only Administrators can create new roles.

To create a new application user group (role)

1. Click the Administration tab and the User and Access Management subtab.
2. Click Groups.
3. Click the New Application Group button to the left of the Application Groups folder in the User Groups list.
4. Provide the group name and description.

5. If this new user group is to have access you have already defined for two or more user-defined application groups, select those application groups for membership. Otherwise, make no selection.

   Note: If this new group is composed of existing groups, existing policies for each of the component groups will apply to this group. No additional policies are required.

6. Click Save.

7. Click Close.

More information:

Step 2: Create the PCI-Analyst Role (see page 112)
Sample Policies for Suppression and Summarization Rules (see page 117)

Grant a Custom Role Access to CA Enterprise Log Manager

When you create an application user group, or role, be sure to add it to the predefined CALM Application Access policy. Only identities that are explicitly added to this policy can access CA Enterprise Log Manager. Identities can be individual users or members of a user group.

If you encounter a situation where users assigned to a new user group cannot log into the CA Enterprise Log Manager, verify that the Identities of the CALM Application Access policy include this group.

To grant CA Enterprise Log Manager access to a user-defined application user group

1. Select the Administration tab, click User and Access Management, and then click Access Policies on the left pane.

2. Click Scoping Policies and select CALM Application Access.

3. Under Identities, search for the new application group as follows:

   a. For Type, select Application Group.

   b. Click Search Identities.

   c. Leave Name as the attribute and LIKE as the operator. Click Search.

      The name of the new application group appears in the displayed list of identities.

   d. Select the name of the new application group and click the Move button to move the group name to the Selected Identities box.

4. Click Save.

Add an Identity to an Existing Policy

When you create a new application user group, you can add the new group to existing policies, if applicable. When you create a user that has no role but has access limited with an access filter, you can add such a user to existing policies.

Important! When working with the installed access policies, take special care not to delete them as they are not locked or protected.

If a predefined access policy is accidentally deleted, users will be unable to access the CA Enterprise Log Manager server until it is restored. You can restore policies using the safex utility.

To add an identity to an existing policy

1. Select the Administration tab, click User and Access Management, and then click Access Policies on the left pane.
2. Click the policy type, and then select the policy that applies to the new application user group. View the Identities pane.
3. For Type, select Application Group.
4. Click Search Identities.
5. Leave Name as the attribute and LIKE as the operator. Click Search.

   The name of the new application group appears in the displayed list of identities.
6. Select the name of the new application group and click the move button to move the group name to the Selected Identities box.
7. Click Save.

More information:

Step 4: Add PCI-Analyst to Existing Policies (see page 113)

Create a CALM Access Policy

You can create a CALM access policy to grant (or deny) one or more valid actions on one or more CALM resources.

The following CALM resources are application-specific; that is, they are used only by the CA Enterprise Log Manager product:

- Alert
- AgentConfiguration
- AgentAuthenticationKey
- ALL_GROUPS
- Connector
- Data
- Database
- EventGrouping
- Integration
- Profile
- Report
- Tag

To create a new CALM policy from scratch

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies.
3. Click the New access policy button to the left of the CALM folder.
4. Enter a meaningful name for the policy and, optionally, a short description.
5. If this policy is temporary, select the Calendar with the date range to which it applies.
6. Accept CALM as the resource class name.

7. Select Type in the General panel according to the following criteria:
   - Select access policy to grant or deny to all selected identities the ability to perform all selected actions on all of the selected resources to which they apply.
   - Select access control list to grant or deny to all selected identities the ability to perform only the selected actions on a selected resource.

     Note: It is not possible to save filters for multiple resources. The workaround is to create separate policies, one for each resource/filters combination.
   - Select identity access control list to grant or deny to each selected identity the ability to perform selected actions on all the selected resources to which they apply.
8. Use the Identities area to select the users or groups to which this policy applies as follows:
   a. Select Application Group for Type or one of the other options, click Search Identities, and click Search.
   b. Select identities from those available and click the Move button to move them to the Selected Identities box.
9. If the policy type is access policy, complete the access policy configuration as follows:
   a. Enter a CALM resource in the Add resource field and click Add.
   b. Select each Action that the selected identities are to be able to perform on any selected Resource, where valid actions include the following: annotate, create, dataaccess, edit, and schedule. You cannot grant the ability to perform a given action on one resource and not another where it is valid.
10. If the policy type is access control list, complete access control list configuration as follows:
    a. Enter a CALM resource in the Add resource field and click Add.
    b. Select each Action that the selected identities are to be able to perform on this Resource, where valid actions include one or more of the following: annotate, create, dataaccess, edit, and schedule.
    c. Repeat the last two steps for each resource to be addressed by this policy.

       With this type, you can grant the ability to perform an action such as create on one resource but not on another.
11. If the policy type is identity access control list, complete the identity access control list configuration as follows:
    a. For each identity selected, select the Actions to be granted or denied on all resources for which they are valid.
    b. For each resource to be added, enter a CALM resource name in the Add resource field and click Add:

12. Review the check boxes at the top and select any that apply:
    - Select Explicit Deny to change the policy from one that grants access to one that denies access
    - Select Disabled to inactivate this policy temporarily, if new.
    - Select Pre-Deployment and then select Assign Labels and add the labels if using this policy for testing purposes and you want to categorize the policies with custom labels.
13. Click Save and then click Close on the left pane.

Create a Scoping Policy

You can create a scoping policy on any global resource. Actions on scoping policies are limited to read and write.

- The following global resources are used by many CA products (applications):
  - Calendar
  - GlobalUser
  - GlobalUserGroup
  - iPoz
  - Policy
  - User
  - UserGroup
  - AppObject
- The global resource, AppObject, lets you create scoping policies on application-specific resources and modules. You do this by adding a filter that designates the relevant EEM folder where the application-specific content or module is stored.
  - EEM content folders you can use in filters with the AppObject resource include the following:
    - EventGrouping
    - Integration (Server)
    - Profile
    - Report
  - CA Enterprise Log Manager module folders you can use in filters with the AppObject resource include the following:
    - Event Log Store
    - Report Server
    - Subscription

You can create a policy from scratch if no policy exists from which you can leverage the settings. If you are creating a scoping policy associated with a CALM policy you have created, specify the same identities as those in the related CALM policy.

Only Administrators can create, edit, delete, and view access policies.

To create a new explicit grant scoping policy

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies.

3. Click the New Scoping Policy button to the left of the Scoping Policies folder.
4. Create a meaningful name for the policy. For example, include the role or roles to which it applies and the tasks that are scoped. View the names of the predefined policies for examples of how this standard can be used.
5. Enter a short description that more fully describes what the more cryptic name implies.
6. Typically, you will accept SafeObject as the resource class name.
7. Select Type in the General panel according to the following criteria:
   - Select access policy to grant or deny to all selected identities the ability to perform all selected actions on all of the selected resources to which they apply.
   - Select access control list to grant or deny to all selected identities the ability to perform only the selected actions on a selected resource.

     Note: It is not possible to save filters for multiple resources. The workaround is to create separate policies, one for each resource/filters combination.
   - Select identity access control list to grant or deny to each selected identity the ability to perform selected actions on all the selected resources to which they apply.
8. If the policy type is access policy or access control list, use the Identities area to select the users or groups to which this policy applies.
   a. Select Application Group for Type, click Search Identities, and click Search.
   b. Select identities from those available and click the Move button to move them to the Selected Identities box.

9. If the policy type is access policy, all actions are selected for all resources by default. To customize this, complete access policy configuration as follows:

   a. Select a resource from the Add resource drop-down list and click Add.

      - Select AppObject if the resources to which read or write access is to be configured are CA Enterprise Log Manager-specific resources.
      - Select User and GlobalUser for access to Users buttons on the Administration tab, User and Access Management subtab.
      - Select UserGroup and GlobalUserGroup for access to Groups buttons on the Administration tab, User and Access Management subtab.
      - Select Policy for access to the Access Policies, EEM Folders, and Test Policies buttons on the Administration tab, User and Access Management subtab.
      - Select Calendar for access to the Calendars button on the Administration tab, User and Access Management subtab.
      - Select iPoz for access to the Password Policy and User Store buttons on the Administration tab, User and Access Management subtab.

   b. Select read to grant/deny view access; select write to grant/deny edit access. If you select neither, all actions are selected.

      Note: To grant/deny create access, you must define a CALM access policy and select CA Enterprise Log Manager resources individually.

   c. Add a generic filter that applies to the selected resources, if needed.

10. If the policy type is access control list, complete access control list configuration as follows:

    a. Select a resource from the Add resource drop-down list and click the Add (+) button.

    b. Select read, write, or both for Actions.

    c. Click the Edit Filters button to open the filter form. Create a filter for the associated resource by selecting or entering values for the Left type/value, Operator type/value, and Right type/value.

    d. If the filter includes a resource name as a value, select the check box labeled Treat resource names as regular expressions. Otherwise, leave this check box cleared.

    Important! Define one policy for each resource/filters combination.

11. If the policy type is identity access control list, complete the identity access control list configuration as follows:
    a. For Type, select one of the displayed options. For example, select Application Group, click the Search Identities link, and click the Search button to display the members of the type you selected.
    b. Select the identities and click the move button to populate the Selected Identities pane.
    c. For each identity selected, specify read or write or both.

       The identity-specific actions apply to all the selected resources. That is, a given identity can view, view and edit, or just edit all of the selected resources.
    d. Add the resources to which the identity-specific actions are to be granted or denied.
12. Review the check boxes and select any that apply:
    - Select Explicit Deny to change the policy from one that grants access to one that denies access
    - Select Disabled to inactivate this policy temporarily, if new.
    - Select Pre-Deployment and then select Assign Labels and add the labels if using this policy for testing purposes and you want to categorize the policies with custom labels.
13. Click Save and then click Close on the left pane.

More information:

Step 3: Create Win-Admin System Access Policy (see page 101)

Create a Policy Based on an Existing Policy

You can create a new access policy by copying an existing access policy and modifying the copy. This procedure can save you the time it takes to manually duplicate the specifications of an existing policy that requires only minor modifications to satisfy your current requirements.

Only Administrators can create, edit, delete, or view access policies.

To create a policy based on an existing policy

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies.
3. Select either CALM or Scoping Policies, depending on the type of policy you want to use as a template.

4. Click the name link to open the policy to copy.
5. Click Save As.

   The Explorer user Prompt dialog appears.
6. Enter the name for the new policy to be based on the open policy and click OK.
7. Make the needed modifications.

   For example, replace the copied Identity with the name of the role (user-defined application user group) to which this policy applies. Consider modifying the actions permitted on the copied resources. Consider clicking Filters and specifying an additional filter for the new role.
8. Click Save, and then click Close.
9. Verify the new policy definition.
   a. Re-select the policy type to display the view of all policies.
   b. Compare the new policy with the original policy and verify that all planned changes are reflected in the new policy.
   c. Click Close.
10. Test the policy.

More information:

Step 5: Create a Policy Based on Analyst Report View-Edit Policy (see page 113)

Test a New Policy

You can test whether a new policy is syntactically correct with the Test Policies feature. The Test Policies feature lets you run ad-hoc queries against the access policies you define. You can consider a permission as a request: "Can {identity} perform {action} against the resource of type {resource class} and of name {resource} [with the following attributes}] [at the {specified time}]?" A result of ALLOW means that the identity you entered can perform the specified action on the specified resource with the specified attributes at the specified time.

Before you begin, have your policy at hand.

To test a policy

1. Click the Administration tab and the User and Access Management subtab.
2. Click Test Policies.

   The Permission Check Parameters page appears.
3. If the policy you plan to check was one where you selected Pre-Deployment and added labels, then check the check box that indicates you want to include pre-deployment policies and add the associated labels.
4. Complete the entry fields. If your policy includes filters, specify the filters in the order that they appear in the policy.
5. Click Run Permission Check.
6. Observe the result and proceed in one of the following ways:
   - If the result is ALLOW, log on to CA Enterprise Log Manager as a user specified as an identity in this new policy and test the effectiveness, scope, and coverage of the before putting it into production use.
   - If the result is DENY, verify your entries on the query. If they are correct, return to the policy or and make the needed correction there.

Create a Dynamic User Group Policy

A *dynamic user group* is composed of global users that share one or more common attributes. A dynamic user group is created through a special dynamic user group policy where the resource name is the dynamic user group name and membership is based on a set of filters configured on user and group attributes.

You can create a dynamic group composed of Users, Application Groups, Global Groups, or Dynamic Groups. For example, you can create a dynamic group of Global Groups or Application Groups based on Name, Description, or Group Membership. Or, you can create a dynamic group of Users with different roles based on a common attribute in their global user profile, for example:

- Job title
- Department or office
- City, state, or country

Only Administrator can create Dynamic User Group Policies.

To create a dynamic user group policy

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies.
3. Click New Dynamic Group Policy.

   The New Dynamic Group Policy page appears.
4. For Name, enter a group name that indicates what this group of users has in common. Optionally, enter a description.
5. Select a policy type. The default is Access Policy.
6. Select Identities as follows:
   a. For Type, select User, Application Group, Global Group, or Dynamic Group and click Search Identities.
   b. For Attribute, Operator, and Value, enter the expression that sets the criteria for membership in this group and click Search.

      For example, if you selected User, you could enter Job Title Like Manager and click Search to find all of the users who have the job title of Manager.
   c. Select from the displayed identities those who are to be members of this dynamic group and click the Move arrow to move your selections to the Selected Identities box.

7. For Actions, select belong.
8. In the Add resource field, enter the value you entered in the Name field and click the Add button. This indicates that the selected identities belong to the dynamic group resource you just created.
9. Optionally, add more filters.
10. Click Save.
11. Click the Dynamic User Group Policies link and verify the new dynamic user group you created. For example:

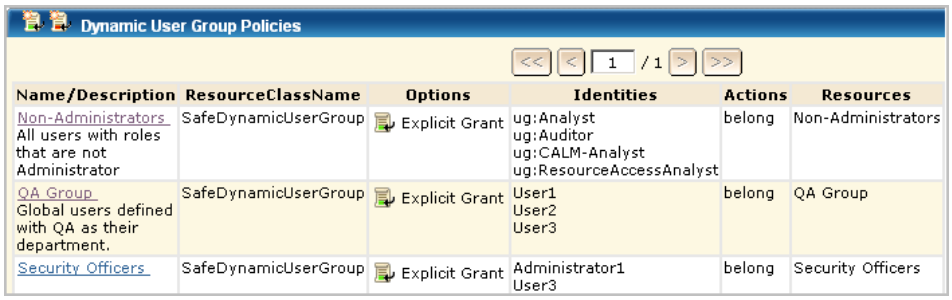

Create an Access Filter

You can create an access filter to restrict access to log data meeting the filter criteria. By default all CA Enterprise Log Manager application users have query access to event log data stored in the event log stores of the active CA Enterprise Log Manager server, peer servers in a meshed federation or descendant servers in a hierarchical federation.

You can restrict access to the event log store of one or more specific CA Enterprise Log Manager servers by creating a data access filter. You can apply an access filter to an individual or a group.

To create an access filter for a user-defined role

1. Click the Administration tab and the User and Access Management subtab.
2. Click New Access Filter.

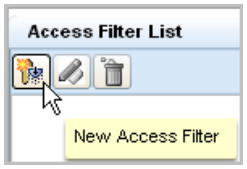

The Access Filter Design wizard appears.

3. For Details, enter the name and description for the filter.

4. Click Identities. Select an Identity type, click the search button to show available identities, and use the shuttle control to select the ones to which this access filter applies.

   For example, select the application group you created for this purpose.

5. Set the access filters.

   a. Click Access Filters.

   b. Click the New Event Filter button.

   

   c. Add one or more expressions that define the access filter.

   d. Click Save and Close.

   The Access Filter you created appears.

6. Click Close.

More information:

Step 4: Create Win-Admin Data Access Filter (see page 104)
Create an Application User Group (Role) (see page 75)

Maintaining User Accounts and Access Policies

As an Administrator, you can perform the following maintenance tasks on user accounts and access policies:

- Lock a user account so the user cannot log onto CA Enterprise Log Manager
- Unlock user accounts that have been locked, if the password policy does not permit any user to unlock a locked user account.
- Add new user accounts
- Edit existing user accounts
- Lock or delete user accounts that belong to individuals who no longer need access to CA Enterprise Log Manager
- Edit existing access policies
- Delete access policies that are no longer needed
- Create, edit, or delete delegation policies
- Create, edit, or delete access filters with their corresponding auto-generated obligation policies
- Create a super role from existing roles with limited access
- Add a new custom role and corresponding access policies

Create a Calendar

You can create a new calendar to help restrict user access during certain time periods. Calendars work as part of access policies. When you define a calendar, you can include or exclude time blocks in hours, days of the week, or dates.

To create a calendar

1. Click the Administration tab, then click User and Access Management, then click the Calendars button.

   The Calendars page appears.

2. Click the New Calendar icon at the top left of the calendar list.

   The New Calendar details pane appears.

3. Enter a name that specifies the target policy, and provide a description of the intended use.
4. Use the calendar icons to set start and end dates for the calendar.
5. Click Add Include Time Block or Add Exclude Time Block to create exception periods within the main effective period of the calendar.
6. Click Save, and then click Close.

More information:

Add a Calendar to a Policy (see page 91)

Add a Calendar to a Policy

When creating a policy, you can select an existing calendar that specifies when the specified identities can perform the selected actions on the specified resources. A calendar can define start and end dates and time blocks in hours or days of the week.

To add a calendar to a policy

1. Click the Administration tab and the User and Access Management subtab.
2. Open the policy to which this calendar applies
   a. Click Access Policies
   b. Select the policy type.
   c. Select the policy.
3. Open the Calendar drop-down list and select the calendar you created for this policy.

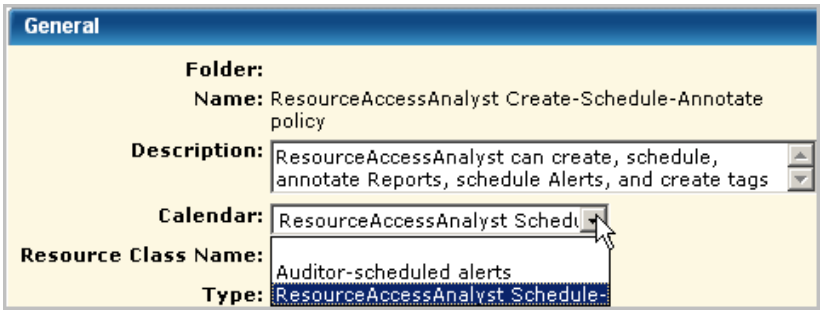

4. Click Save to save the addition of the calendar to an existing policy.

More information:

Create a Calendar (see page 90)

Example: Limit Access to Work Days

You can restrict the time of day or days of the week a given user group can access CA Enterprise Log Manager by creating a calendar for the times to grant access, creating a custom role, creating a new policy based on the policy providing CA Enterprise Log Manager access, and assigning the calendar and custom role to this policy.

Example--Limit External Auditors' Access to CA Enterprise Log Manager to Weekdays

To limit certain group's access to CA Enterprise Log Manager to business days, create a calendar for weekdays and add it to the policies that grant auditors specific access.

For example, if you wanted to limit External Auditor's access to CA Enterprise Log Manager to business hours, create a calendar that specifies weekdays, Monday through Friday, 9 a.m. to 5 p.m., for all months of the year.

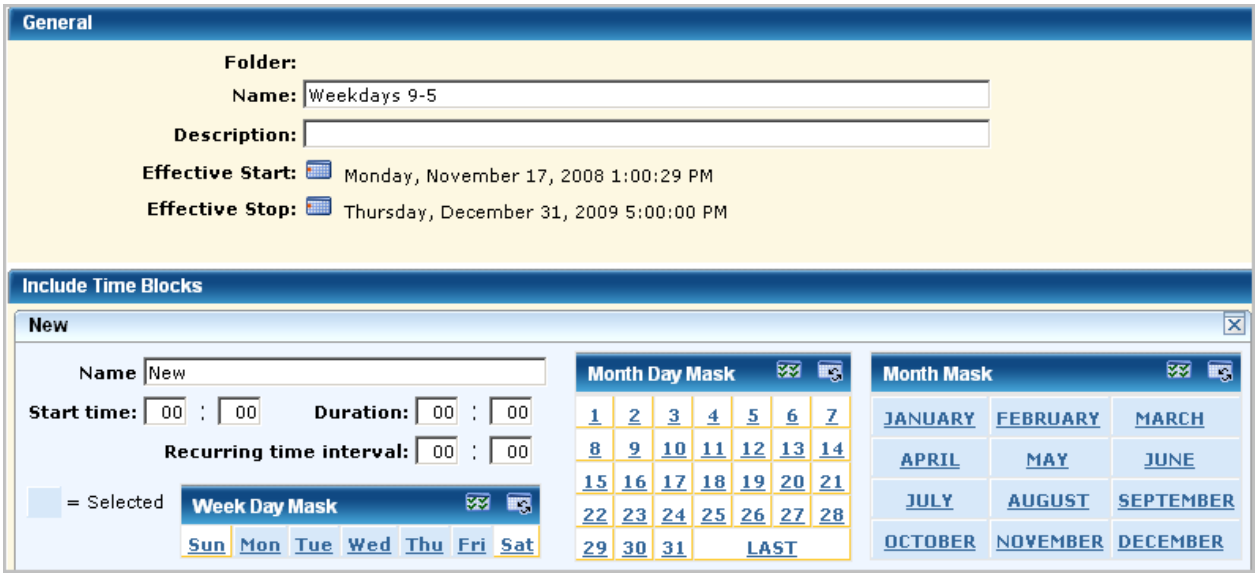

Create a role for external auditors.

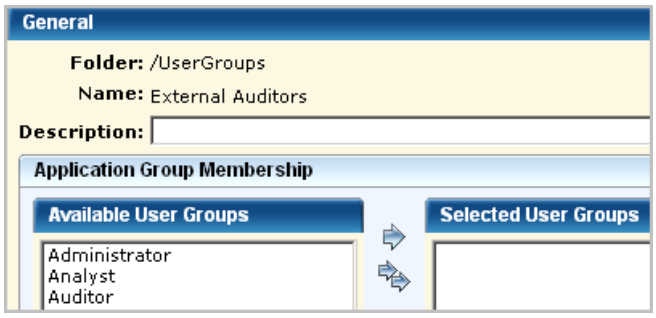

Open the CALM Application Access scoping policy and save it as ExternalAuditors-CALM Application Access, select the Weekdays 9-5 calendar and select the user group External Auditors as the identity.

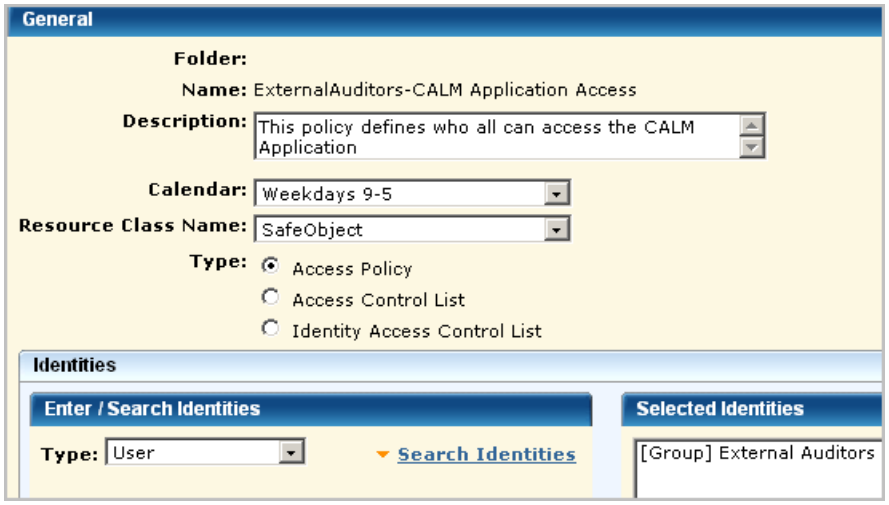

Important! Use the Calendar feature only with policies that grant access. Do not use it with policies that deny access.

More information:

Create a Calendar (see page 90)
Create an Application User Group (Role) (see page 75)
Create a Policy Based on an Existing Policy (see page 84)
Add a Calendar to a Policy (see page 91)

Export Access Policies

You can export all policies of a selected type at any time, both predefined policies and custom policies. Exporting policies is a good way to keep a current backup.

An export creates an XML file for each selected policy, where all XML files are zipped into a file named CAELM[1].xml.gz.

To export access policies

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies.
3. Select the type of access policy to export and click Export.

   The File Download dialog appears.
4. Click Save and save the file with a unique name.
5. Click Close.

More information:

Back Up All Access Policies (see page 51)

Delete a Custom Policy

You can delete a custom policy for any of the following reasons:

- You saved it under a different name and plan to make no other changes, so you need to delete the duplicate.
- You no longer have active memberships in the Identities defined to the policy, so the policy is no longer in use.

Important! Take care not to delete a predefined policy. Should this occur, you can restore it if you exported a backup.

To delete a custom policy

1. Click the Administration tab and the User and Access Management subtab.
2. Click Access Policies.
3. Select either CALM or Scoping Policies, depending on the type of policy you want to delete.
4. Click the name of the policy to be deleted.
5. Click Delete.
6. Click OK to confirm the deletion.

Delete an Access Filter and Obligation Policy

You can delete an access filter and the obligation policy generated by the filter to remove the data access limitation.

Do not delete the obligation policy generated by the filter from Access Policies.

To delete an access filter and its obligation policy

1. Click the Administration tab, click User and Access Management.

   The Access Filter List appears at the top of the left pane.

2. Select the filter to delete and click the Delete Access Filter button.

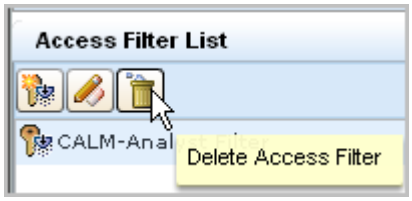

   The Delete Access Filter Confirmation warning message appears.

3. Click Yes to remove the selected access filter and the associated obligation policy.

Example: Allow a Non-Administrator to Manage Archives

Suppose you want to allow a non-Administrator group to manage auto-archiving. You could create a group called ArchiveAdministrator, a CALM policy that allows the edit action on the resource database. This allows read access on the archive catalog of databases for querying, write access on the archive catalog for ReCatalog, and the ability to use LMArchive utility for manual archiving or the restore-ca-elm shell script for restoring auto-archived databases.

To allow specified non-Administrators to handle archiving

1. Create a role called ArchiveAdministrator.
   a. Select the Administration tab and then the User and Access Management subtab.
   b. Select Groups.
   c. Click New Application Group.
   d. Enter ArchiveAdministrator as the name.
   e. Click Save.

      The Application Group, or role, ArchiveAdministrator is created.

   f. Click Close.

2. Create a CALM policy to allow edit access to the database resource.
   a. Click Access Policies.
   b. Click New Access Policy to create a new CALM policy.
   c. Type ArchiveAdministrator policy in the Name field.
   d. Type ArchiveAdministrator can run the LMArchive utility and the restore-ca-elm shell script for the description.
   e. For Identities, select Application Group as the Type, click Search Identities, and then click Search.
   f. Select ArchiveAdministrator and then click the move arrow.
   g. Type Database under Add resource and click Add.
   h. Select edit as the Action.

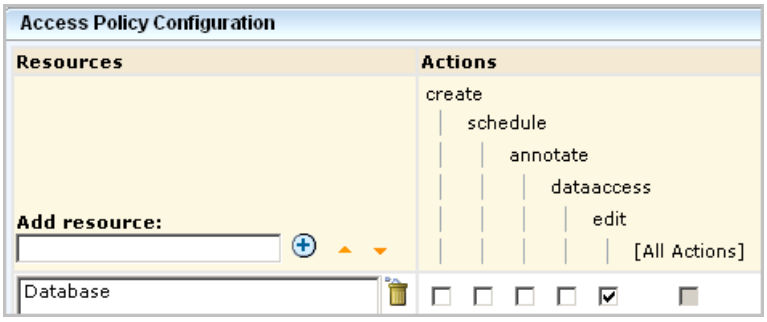

   i. Click Save. Click Close

3. Test the policy and verify that the result is ALLOW.

Result	Policy	Identity	Resource Class	Resource	Action
ALLOW	ArchiveAdministrator policy	ug:ArchiveAdministrator	CALM	Database	edit

4. Grant the ArchiveAdministrator role the ability to log on to CA Enterprise Log Manager.
   a. Click CALM under Access Policies.
   b. Select CALM Application Access.
   c. Under Identities, search for the Application Group ArchiveAdministrator, and move it to Selected Identities.

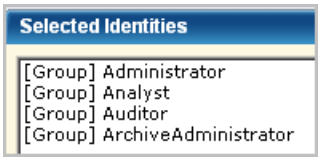

d. Click Save. Click Close. Click Close.

The User and Access Management tab appears with the buttons in the left pane.

5. Assign the ArchiveAdministrator role to one or more users.

a. Click Users.

b. Enter the name of a person to whom you want to assign this role as the Value under Search Users and click Go.

The selected user name appears under the Users folder.

c. Select the link for the selected user.

d. Click Add Application User Details.

e. Move Archive Administrator to the Selected User Groups list.

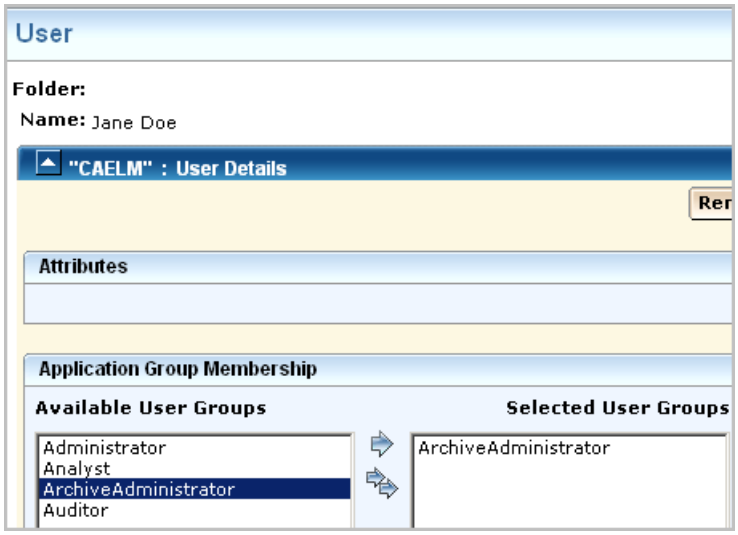

f. Click Save. Click Close.

g. Repeat for each user to whom you want to assign this role.

h. Click Close.

6. (Optional) Review the results from CA Enterprise Log Manager.

a. Click Log Out to log out as the Administrator.

b. Log in as a user to whom you assigned the role ArchiveAdministrator.

c. Click the Administration tab, Log Collection subtab.

d. Select Archive Catalog Query.

e. Observe that you use the Query and ReCatalog buttons.

7. (Optional) Run the restore-ca-elm restore script with the credentials of the user defined with the ArchiveAdministrator role to verify the policy works as expected.

More information:

Restore Auto-Archived Files (see page 167)

Restricting Data Access for a User: Win-Admin Scenario

You can limit reports users can view to those with a specified tag. You can limit the data users can view on those reports to data generated from specified event sources. Limiting access to reports with a given tag is done with an access policy. Limiting data access to events returned to a particular CA Enterprise Log Manager server is done with an access filter. With an access filter defined, a role assignment is optional. That is, you can create a new user, assign no role, and limit data access for that user with an access filter.

Consider the scenario for ABC company with four data centers in the U.S. The Administrator wants to give the Windows Administrator in the Houston region read access to Windows events processed by the domain controller in the Houston area. Windows events processed by the CA Enterprise Log Manager server installed on the Houston domain controller are sent from sources where the host names begin with the string, ABC-HOU-WDC.

This example walks you through creating a user called Win-Admin and ensuring that this user can only view reports that have a System Access tag and that the data on these reports is limited to events from event sources with host names that begin with the known naming convention.

The example provides details for each of the following steps:

1. Create the new user, Win-Admin.
2. Give Win-Admin basic access to CA Enterprise Log Manager. Add this identity to the CALM Application Access policy.
3. Restrict access to reports for Win-Admin to those tagged as System Access. Create a scoping policy with read access to AppObject with filters that specify the EEM folder where Reports are stored and specify the calmTag is equal to System Access. Test the policy.
4. Limit the data Win-Admin can view to that generated by the domain controller in Win-Admin's region. Create an access filter, named Win-Admin Data Access, that limits the query and report data Win-Admin can view to Windows events from event sources with a hostname that begins with ABC-HOU-WDC.
5. Log onto CA Enterprise Log Manager as the Win-Admin user and evaluate the access provided by the policies.
6. If the access is too limited for the user to perform intended tasks, extend the access with additional policies.

Step 1: Create the Win-Admin User

You can create a user without a role (application group) if you specify data access with an access filter.

Step 1 of the complete process for restricting data access in this way is to create the user.

You create a user only if the global user account is not available for import from an external directory. When creating such an account, do not add application user details. In this example scenario, Win-Admin is the user name.

If you search for users, the new name appears in the list

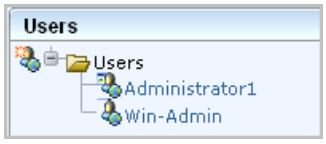

More information:

Create a Global User (see page 34)

Step 2: Add Win-Admin to the CALM Application Access Policy

The second step in restricting data access of a user named Win-Admin is to grant this identity access to the CA Enterprise Log Manager application.

Add the new user to the CALM Application Access policy. The procedure is the same as granting CA Enterprise Log Manager access to a new role except when you search identities, you specify Type as User.

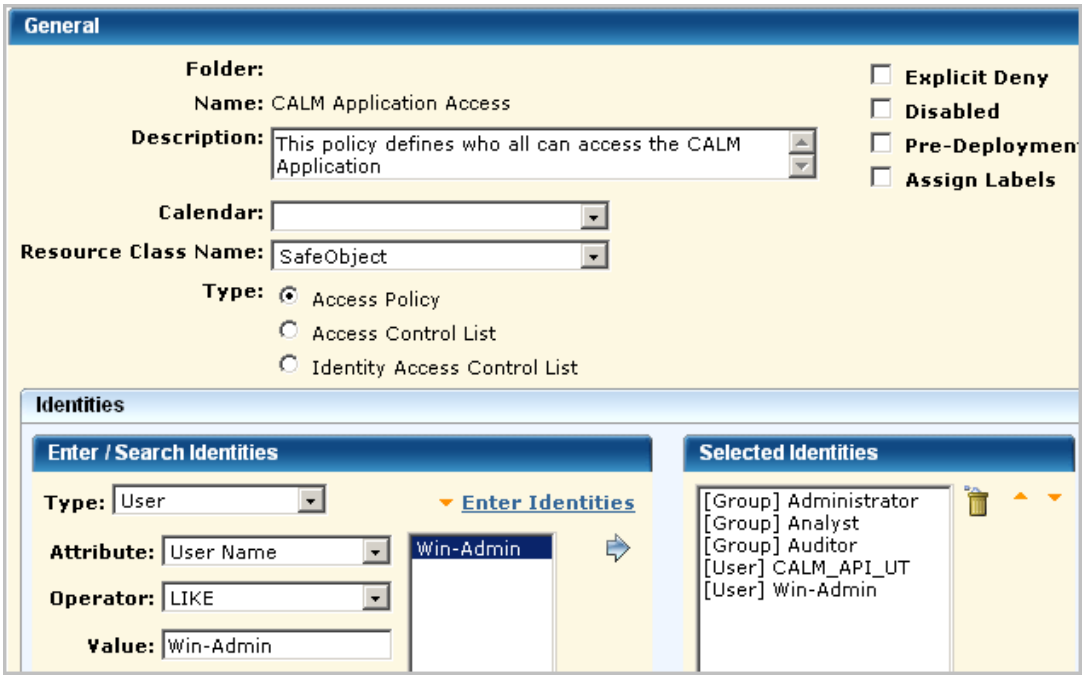

More information:

Grant a Custom Role Access to CA Enterprise Log Manager (see page 76)

Step 3: Create Win-Admin System Access Policy

Step 2 grants access to log on to the CA Enterprise Log Manager application.

Step 3 restricts, or scopes, access to the CA Enterprise Log Manager application after logon. At the broadest level, you can grant read only access or both read and write access to the specified identities.

Selection of policy type determines the granularity at which you can specify permitted actions.

- Access Policies permit selected actions on applicable selected resources.
- Access Control List Policies let you specify what actions are permitted on each added resource.
- Identity Access Control List policies let you specify what actions on applicable resources are permitted by each identity.

You can permit limited access to a resource by creating a filter that specifies the EEM folder for that resource and then specifying restrictions on the folder.

This example demonstrates how to restrict access to read access in general with further restrictions to a specific feature. Specifically, Step 3 restricts the Win-Admin user to viewing system access reports. The following example shows how to create a scoping policy called Win-Admin System Access that grants read access to the SafeObject, AppObject and specifies filters that restrict report access to those with the System Access tag. It also demonstrates how to test the policy, and after verification how to remove the pre-deployment setting.

The General area of a scoping policy designed to specify application access to read only or both read and write specifies SafeObject as the resource class name. The following example shows the policy name of Win-Admin System Access. It is a good practice to select Pre-deployment for a new policy until you have tested it and are satisfied it is ready for use in a production environment.

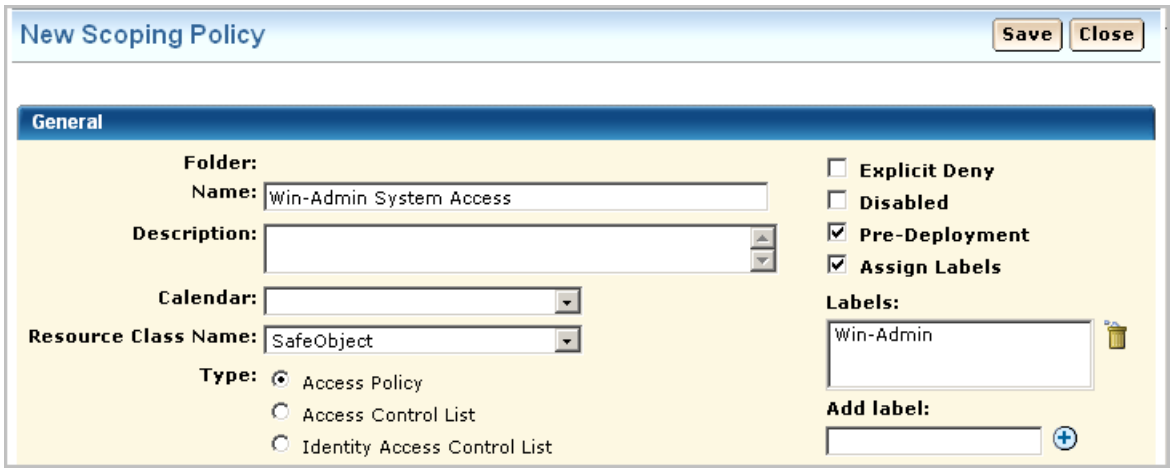

You can grant access to either users or groups. In this example, access is granted to the new user Win-Admin.

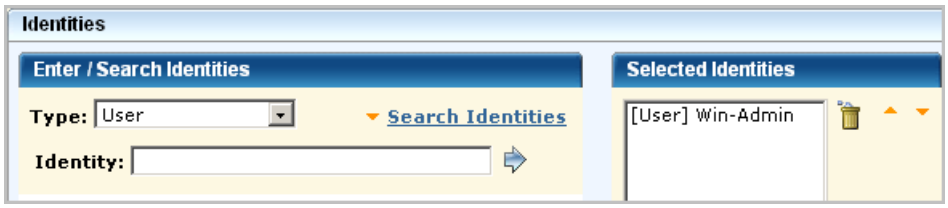

The "highest-level" policy created for CA Enterprise Log Manager is the CALM Access Policy, where CAELM is the application instance. This scoping policy is to specify that the read action is allowed on the application objects, AppObject, which refers to all the application features.

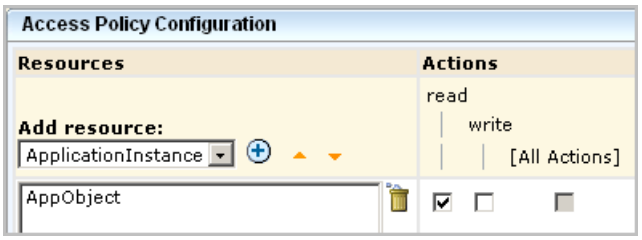

You can further limit the specified action allowed on all objects by specifying filters. Filters are often specified in pairs, where the first filter specifies the CA EEM folder where data related to a given feature is stored and the second filter specifies a restriction on objects in that location. The first filter in the following example limits CA EEM folder access to the folder where the reports resource is stored. Specifically, it specifies that the pozFolder contains /CALM_Configuration/Content/Reports. The second filter limits access to reports with the tag System Access by specifying calmTag is equal to System Access.

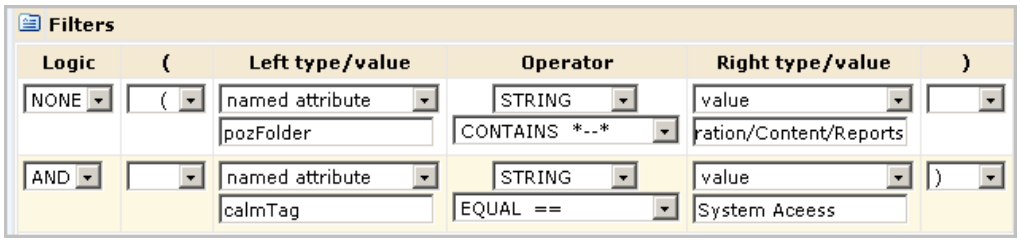

After saving a policy, you can search for it to review. You can search for policies by name, identity or resource. You can enter a partial value. You can also enter multiple criteria. Examples for this scenario follow.

Searching by the full name displays the one policy you need.

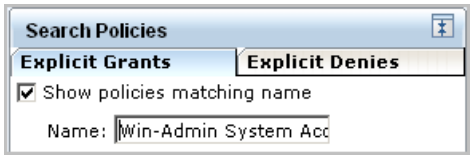

Searching by identity only displays all policies that apply to this identity including those that apply to all identities.

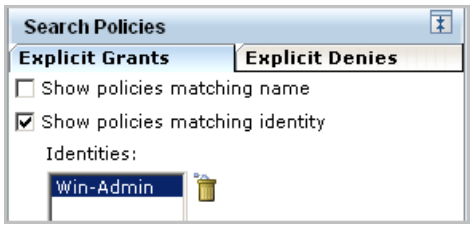

Searching by resource only where the resource is AppObject displays all system supplied and custom policies that grant read or read and write access to any identity.

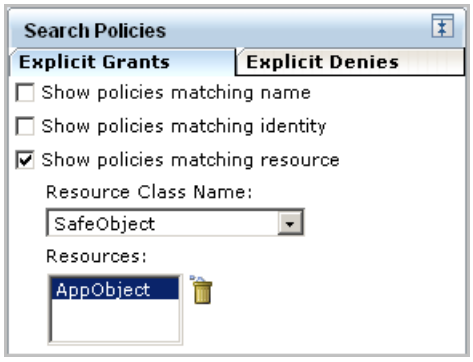

When the custom policy you search for displays on the policy table, examine the values, including the filters. If you notice anything that requires correction, you can click the name link to redisplay the policy for editing.

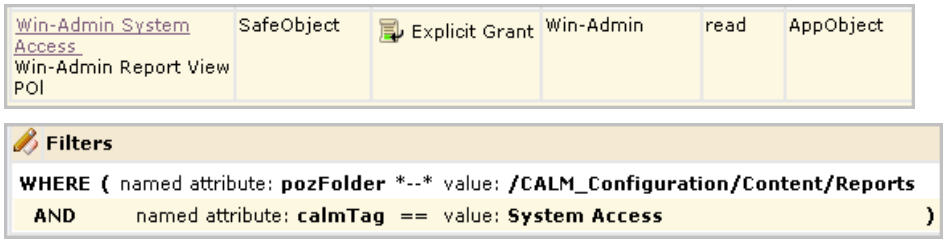

It is a good practice to test each new policy. Be sure to enter the attribute/value pairs in the order you entered the filters, with the higher-level attribute first.

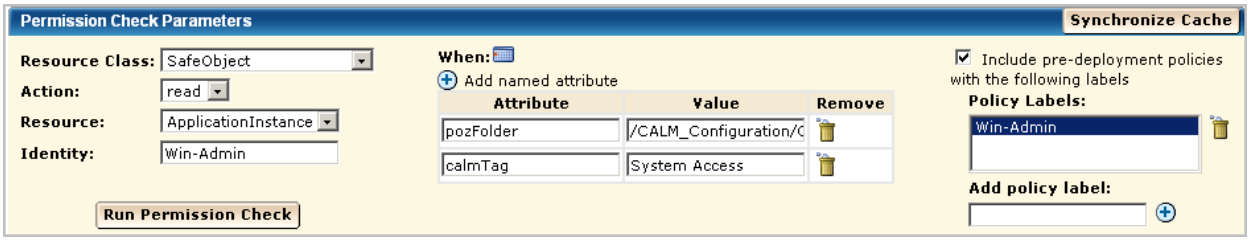

Verify that the result is ALLOW.

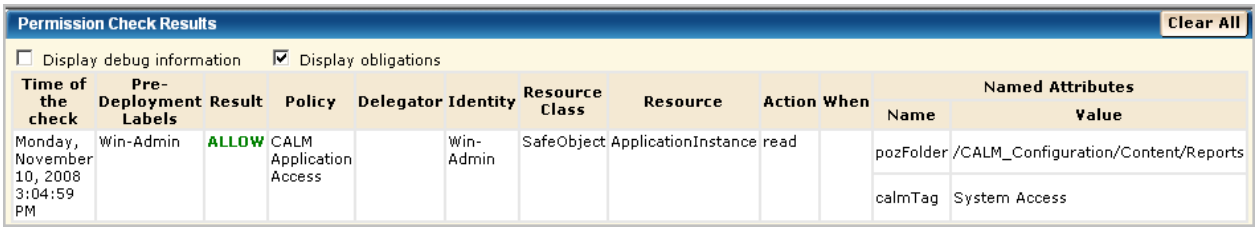

After verifying the result is ALLOW, clear the Pre-Deployment setting on the policy. Otherwise, you will not be able to log in as Win-Admin to evaluate what this user can do.

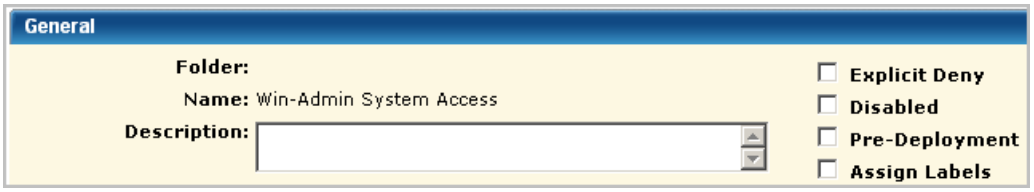

More information:

Test a New Policy (see page 86)

Step 4: Create Win-Admin Data Access Filter

Step 3 restricts the Win-Admin user to viewing system access reports. This level of access allows the Win-Admin user to view system access reports across all four regions of the ABC company.

Step 4 creates an access filter to limit the data the Win-Admin user can view to the system access reports for the Houston domain controller.

Creating a data access filter begins with specifying its name. The name used in this scenario is Win-Admin Data Access.

You specify the identities to which the access filter applies in the Identities area. A filter can apply to users or groups. In this scenario, this access filter applies only to the Win-Admin user.

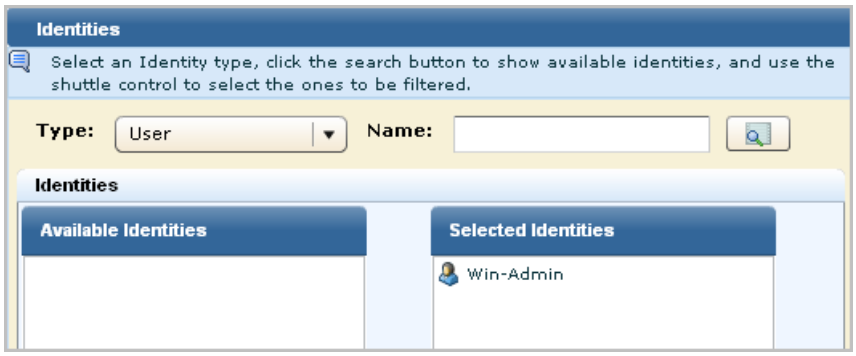

For Access Filters, you define each condition in terms of the value for a CEG column. Values following the LIKE operator can contain either of the following wildcard characters:

- _ (underscore character) - represents any single character
- % (percent sign) - represents a string containing any number of characters

The first filter for this scenario takes advantage of the fact that all Windows events are prefixed with NT-. To limit the data to Windows events, you can specify that the event_logname CEG column must have data that includes the string NT-%. To further limit Windows events to those from a specific domain controller, this example specifies that event_source_hostname must have data that includes a string using local conventions. The value ABC-HOU-WDC% is based on a naming convention of a hyphenated name composed of abbreviations for the company, the region, and the prefix for the domain controller type.

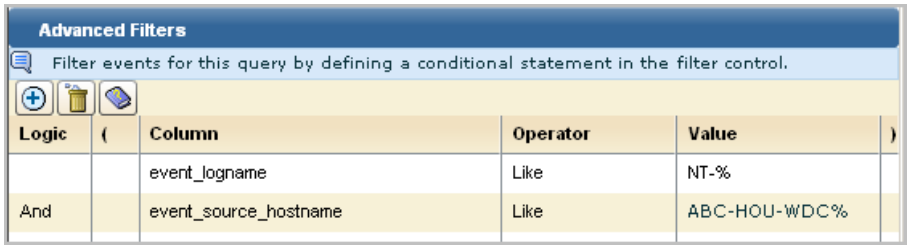

Note: In the absence of event sources with a standardized naming convention, you can create a keyed values list with the desired event_source_hostnames and use the keyed values list name as the value.

When there are only two filters and the logic is AND, parentheses are not required. If you enter a complex expression, such as the following, parentheses are required.

(event_logname like NT-%
And event_source_hostname=ABC-%)
Or (event_logname like CALM-%
And event_source_hostname=XYZ-%)

When you save a data access filter, its name is displayed in the access filter list.

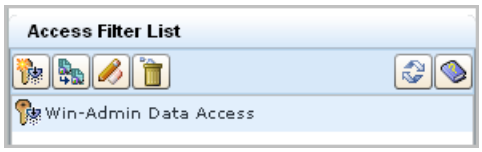

A search for policies matching the Win-Admin user name displays the three policies for All Identities plus another three: the CALM Application access policy where Win-Admin was added, the Win-Admin System Access policy created from scratch, and the data policy that is added automatically when you define an access filter. The data policy is listed first in the following. You can also view it under Obligation policies. You never create Obligation policies directly with CA Enterprise Log Manager.

Access Policies

CALM Application Access This policy defines who all can access the CALM Application	SafeObject	Explicit Grant	ug:Administrator ug:Analyst ug:Auditor Win-Admin	read write	ApplicationInstance Policy User GlobalUser
9e0855ba-calmrhbuildtest0249187347-73694758-1f **Labels:** DataPolicy	SafeObligation	Explicit Grant	Win-Admin	FulfillOnGrant	dataaccess/CALM/Data
Win-Admin System Access **Labels:** Win-Admin	SafeObject	Explicit Grant Pre-Deployment	Win-Admin	read	AppObject

More information:

Create an Access Filter (see page 88)

Step 5: Log on as Win-Admin User

Before you create policies for a given user or application user group, log on as that user or group member and determine what you can and cannot do. First, verify that the restrictions you expect to be in place are working. Second, verify that you can perform the tasks you expect such users to do.

For this scenario, you expect to be able to view only reports or action alerts that are tagged with System Access. In the example, the only available query tag filter is System Access. Therefore, the expectation is verified.

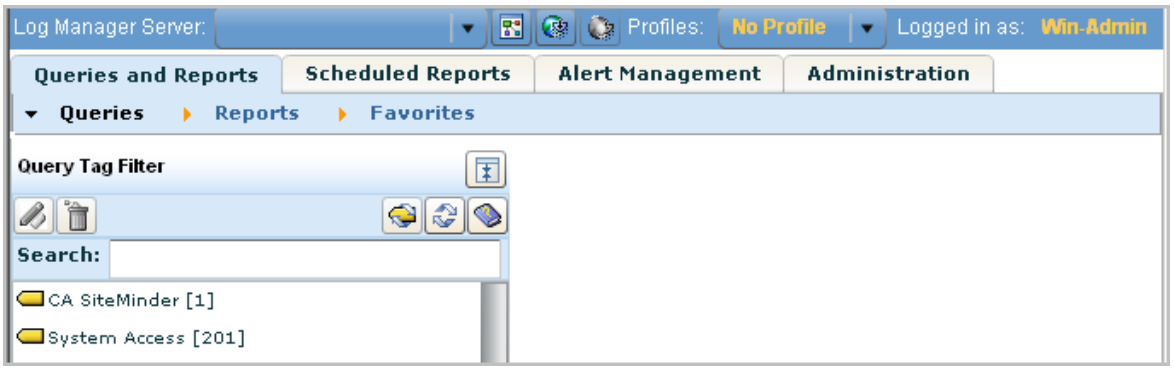

A quick way to test an access filter is to use the Prompts function. However, this function is not available to the Win-Admin user. All the prompt queries have the tag "Event Viewer". Access to Prompt Filters can be granted with the policy filter calmTag=Event Viewer.

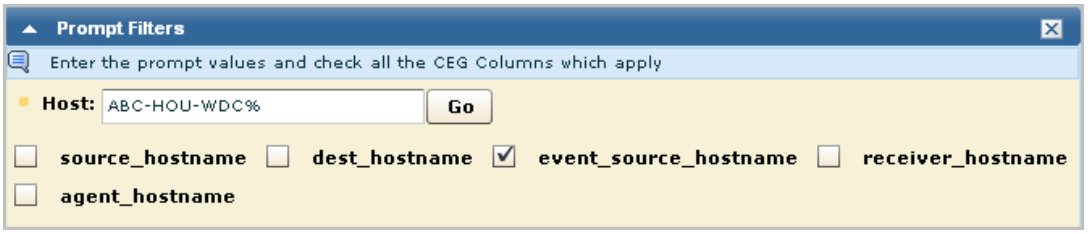

The best way to test an access filter is to review the data displayed on a report. Consider the following access filter. The event_logname CEG column begins with NT- and the CEG column event_source_hostname begins with ABC-HOU-WDC, an abbreviation for the ABC company, Houston location, Windows Domain Controller.

event_logname Like NT-% AND event_source_hostname Like ABC-HOU-WDC%

The following example shows a report viewed by the user to whom this access filter applies. Notice the data in the Log Name column begins with NT- and the data in the Source column begins with ABC-HOU-WDC.

System Access - All Events

Severity	Date Time	Source	User	Action	Log Name	Class
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Wed Jun 4 2008 08:01:0(	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Tue Jun 3 2008 08:01:00	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Tue Jun 3 2008 08:01:00	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Tue Jun 3 2008 08:01:00	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Tue Jun 3 2008 08:01:00	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Tue Jun 3 2008 08:01:00	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Tue Jun 3 2008 08:01:00	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management
Information	Tue Jun 3 2008 08:01:00	ABC-HOU-WDC05	user01	Account Creation	NT-Security	Account Management

Step 6: Extend Granted Actions

The policies and access filter defined in steps 2, 3, and 4 of this example enable the Win-Admin user to view System Access reports, with limits on the data. With this access alone, the Win-Admin user cannot schedule a report, schedule an alert, or annotate a report. To do these things, add Win-Admin to the Analyst Auditor Report Server Access Policy and the Analyst Create-Schedule-Annotate policy. Example of these policies with Win-Admin added follow:

Analyst Auditor Report Server Access Policy Analyst ,Auditor can Schedule Reports and Alerts against all available Report Servers	SafeObject	Explicit Grant	ug:Analyst ug:Auditor Win-Admin	read	AppObject
Analyst Create-Schedule-Annotate policy Analyst can create/schedule Reports, schedule Action Alerts,Annotate Reports	CALM	Explicit Grant	ug:Analyst Win-Admin	create schedule annotate	Report Alert Tag

For Win-Admin to be able to create a report, this user needs write access added to the Win-Admin System Access policy. This requires opening the Win-Admin System Access policy for editing and adding write to the permitted actions.

For Win-Admin to be able to use prompts, the filter for Win-Admin System Access can be modified such that the attribute calmTag equals either System Access or Event Viewer.

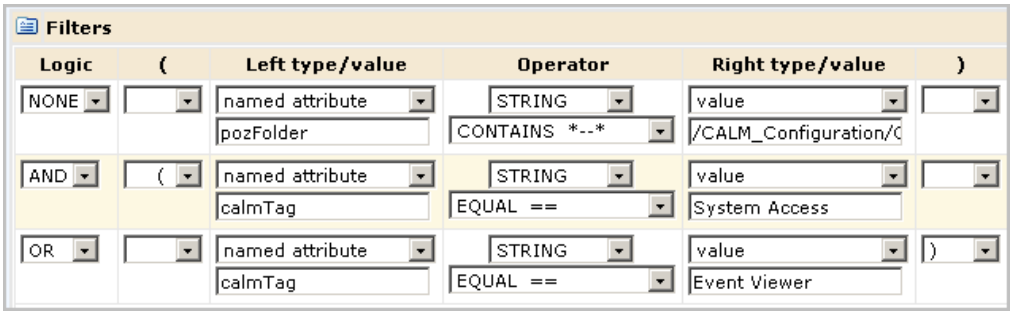

Restricting Access for a Role: PCI-Analyst Scenario

You can create a role that is similar to a predefined role and quickly create policies modeled on predefined policies. The user-defined role may be similar to the predefined role in that it provides the same access to the same resource types, but different in that it limits access based on a filter not present in the predefined role. There may be several policies to which this predefined role has been added as an identity. If the configuration of any policy is such that it applies to your new role, you just add the new role to the existing policy. If the configuration is such that you need to change the type, the resources, the actions, or the filters, you can create a new policy from a copy of the existing one.

This example walks you through creating a role for an analyst that is to work only with reports that have a PCI tag. The associated policy for this role is created from a copy of an existing policy for all Analysts.

The process involves the following procedures:

1. Plan the policies you need. Begin by identifying the existing policies to leverage for the new role.
2. Create the new application user group (role), PCI-Analyst.
3. Give the PCI-Analyst basic access to CA Enterprise Log Manager. Add this identity to the CALM Application Access policy.
4. Give the PCI-Analyst the same access to report servers and create report ability that Analysts have. Add the PCI-Analyst identity to the identified policies.
5. Restrict report access to those reports tagged with the PCI calmTag. Use the policy that grants the ability to view and edit reports as a template to modify.
6. Assign the PCI-Analyst role to a test user for evaluation.
7. Log in as the test user and evaluate access.

   If the access allowed by the role and policies is what you expect, assign the role to all individuals who are to analyze PCI reports.

Step 1: Plan the Role and Policies to Create

Suppose you want to create a role similar to Analysts, but restrict access to PCI-related reports and queries. Plan a name for the role that describes its function, for example, PCI-Analyst.

Before you begin creating new roles, or application user groups, consider the policies that are required to support the new role. It is a good practice to identify the existing policies that are candidates for use as templates. Under Identities, look for the role that is similar to the one you are planning.

In this example scenario, that role is ug:Analyst. Under Search Policies, check Show policies matching identity, enter the identity, **ug:Analyst**, and click Go. The policies displayed include those for All Identities and those where ug:Analyst is explicitly named under Identities.

Access Policies

Name/Description	ResourceClassName	Options	Identities	Actions	Resources
CALM Application Access This policy defines who all can access the CALM Application	SafeObject	Explicit Grant	ug:Administrator ug:Analyst ug:Auditor	read write	ApplicationInstance Policy User GlobalUser
Analyst Auditor Report Server Access Policy Analyst ,Auditor can Schedule Reports and Alerts against all available Report Servers	SafeObject	Explicit Grant	ug:Analyst ug:Auditor	read	AppObject
Analyst Create-Schedule-Annotate policy Analyst can create/schedule Reports, schedule Action Alerts,Annotate Reports	CALM	Explicit Grant	ug:Analyst	create schedule annotate	Report Alert Tag
Analyst Report View-Edit Policy Analyst can View/Edit any Report	SafeObject	Explicit Grant	ug:Analyst	read write	AppObject

The policy names that include this role follow:

- Under Scoping, CALM Application Access
- Under Scoping, Analyst Auditor Report Server Access Policy
- Under Scoping, Analyst Report View-Edit Policy
- Under CALM, Analyst Create-Schedule-Annotate policy

For each of the candidate policies, examine the definition and determine which of the following actions to take:

- Add the new role as an identity to which this policy applies. This is the best choice if the policy applies to the new role without change.

  This action is appropriate for the following policies in this example:

  - CALM Application Access, which defines all identities who can access CA Enterprise Log Manager.
  - Analyst Auditor Report Server Access Policy, which defined all identities who can schedule reports and alerts against all available report servers. This role requires no limitation on report servers.
  - Analyst Create-Schedule-Annotate policy

- Save the policy as a new name and modify its definition.

  This action is appropriate for the following policy in this example, where the new copy would include only the new identity and would have an additional filter to limit the read/write access to reports tagged PCI:

  - Analyst Report View-Edit Policy

Step 2: Create the PCI-Analyst Role

You can create a custom role to represent any task that multiple users perform with the CA Enterprise Log Manager application. A role is the same thing as an application user group.

Step 1 of the process of restricting access for a role is to create the role.

When creating a custom role that is not a superset of an existing role, do not make a selection from Available User Groups.

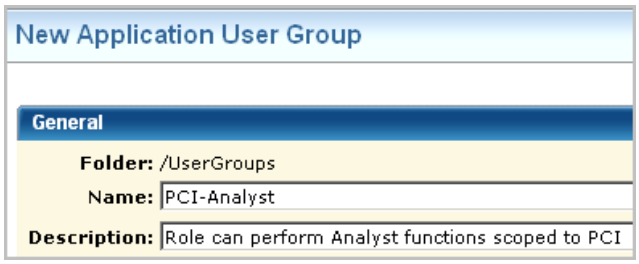

More information:

Create an Application User Group (Role) (see page 75)

Step 3: Add PCI-Analyst to the CALM Application Access Policy

After creating any new role, the next step is to give this role basic logon ability to the CA Enterprise Log Manager application. By default, only the predefined roles have logon access. Add this application group to the CALM Application Access policy.

More information:

Grant a Custom Role Access to CA Enterprise Log Manager (see page 76)

Step 4: Add PCI-Analyst to Existing Policies

Once you identify the policies that apply to an application user group of which the new role is a subset, you add the new role to the current list of Identities.

For this scenario, add the PCI-Analyst role to the following existing policies:

- Analyst Auditor Report Server Access Policy
- Analyst Create-Schedule-Annotate policy

More information:

Add an Identity to an Existing Policy (see page 77)

Step 5: Create a Policy Based on Analyst Report View-Edit Policy

When you create a policy based on an existing policy, you copy the existing policy and save it to a new name. Then, you rename it, edit the description to fit the new role, and replace the existing identities with your new identity. When the policy you are using as a template provides access that is too broad for your new role, you create filters to limit that access.

For the PCI-Analyst scenario, you copy the Analyst Report View-Edit policy, save it to a new name, open the new policy, replace the identity with the PCI-Analyst group, and add a filter to limit report access to those reports tagged with the PCI calmTag.

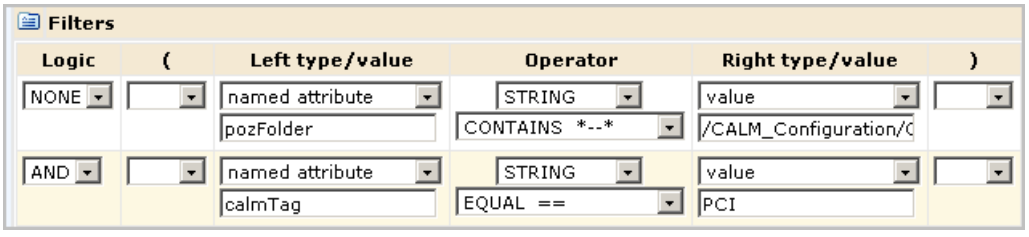

It is a good practice to test a policy based on an existing policy just as you would a policy you created from scratch. When you test a policy with a filter, be sure to enter the filter exactly as it is entered in the policy. When you enter a group name for identity, be sure to prefix it with ug:, for example, ug:PCI-Analyst.

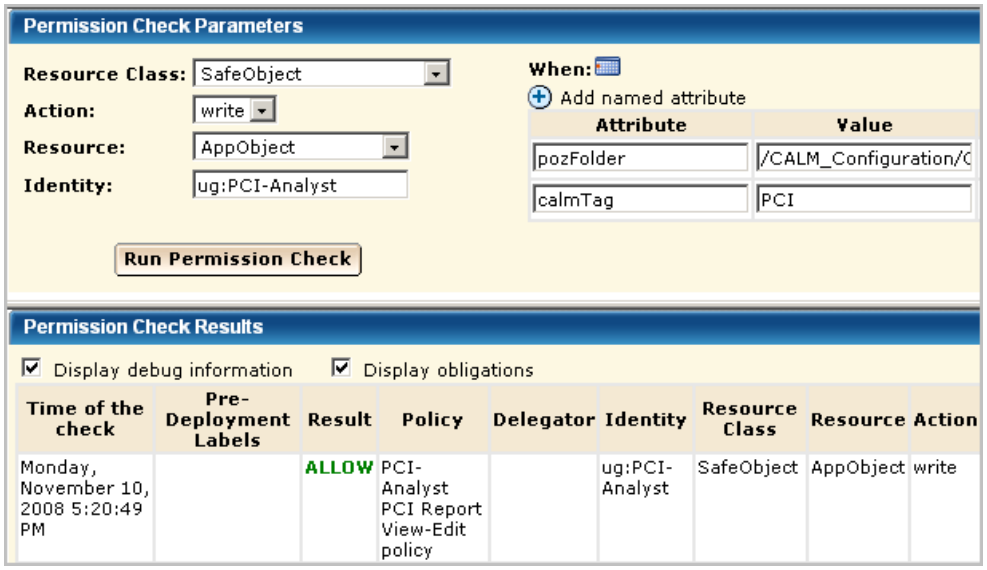

More information:

Create a Policy Based on an Existing Policy (see page 84)
Test a New Policy (see page 86)

Step 6: Assign the PCI-Analyst Role to a User

After creating a new role and its supporting policies, it is a good practice to log on as a user with just this role assigned to evaluate whether the access provided is what is needed. Once verified, the new role can be added to the accounts of all users who are to perform the tasks for which the role was designed.

You can create a temporary user account for the purpose of testing a new role and then delete that account when testing is complete. Or, you can create a user called Test-User and replace the role assignment at each reuse.

More information:

Assign a Role to a Global User (see page 35)

Step 7: Log in as a PCI-Analyst and Evaluate Access

Verify that the policies are sufficient to limit access to reports and alerts tagged as PCI. Assign the PCI-Analyst role to a user and log into the CA Enterprise Log Manager as that new user.

View report tags. Verify that the reports you can view are only those with the PCI tag.

Schedule a report. Verify that the reports you can schedule are only those with the PCI tag.

Create a report. Verify that the only available tag for the new report is PCI.

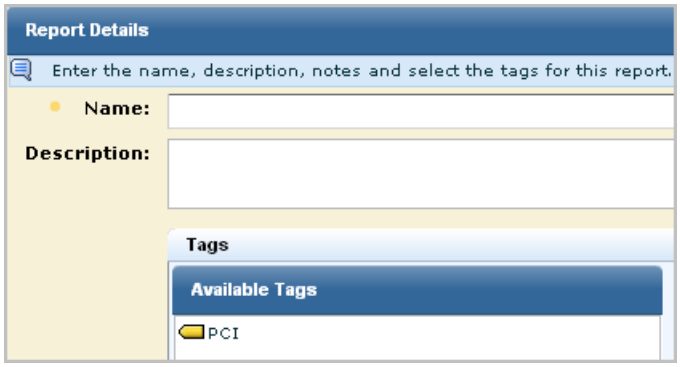

Sample Policies for Custom Integrations

You can give non-Administrators the ability to create custom integrations by creating one custom role, one CALM policy, and one scoping policy. You can give other non-Administrators the ability to view custom integrations by creating an additional custom role with an associated scoping policy. You add both custom roles to the CALM Application Access policy and assign users to these roles.

The following example procedure shows you how to do this:

1. Create an application user group called Create-DM-XMP-Files.
2. Create an application user group called View-DM-XMP-Files.
3. Grant Create-DM-XMP-Files and View-DM-XMP-Files access to the CA Enterprise Log Manager product.

Name/Description	ResourceClassName	Identities	Actions	Resources
CALM Application Access This policy defines who all can access the CALM Application	SafeObject	ug:Administrator ug:Analyst ug:Auditor ug:Create-DM-XMP-Files ug:View-DM-XMP-Files	read write	ApplicationInstance Policy User GlobalUser

4. Create a CALM policy that grants Create-DM-XMP-Files the ability to create data mapping files and message parsing files using common event grammar while logged on to CA Enterprise Log Manager.

Name/Description	ResourceClassName	Identities	Actions	Resources
Integration-Create policy Can create data mapping and message parsing files using common event grammar.	CALM	ug:Create-DM-XMP-Files	create	Integration

5. Create a scoping policy that grants Create-DM-XMP-Files the ability to edit and view the custom DM files and XMP file saved to the EEM folder /CALM_Configuration/Content/Mapping or /CALM_Configuration/Content/Parsing using common event grammar.

Name/Description	ResourceClassName	Identities	Actions	Resources
Edit-DM-XMP-Files with CEG Policy Can edit data mapping files and message parsing files using common event grammar saved to the EEM folders /CALM_Configuration/Content/...	SafeObject	ug:Create-DM-XMP-Files	read write	AppObject

Filters	
WHERE	name:pozFolder *--* val:/CALM_Configuration/Content/Mapping
OR	name:pozFolder *--* val:/CALM_Configuration/Content/Parsing

6. Create a scoping policy that grants View-DM-XMP-Files the ability to view the custom DM files and XMP file saved to the EEM folder /CALM_Configuration/Content/Mapping or /CALM_Configuration/Content/Parsing.

   Note: The CEG policy grants all Identities rights to view the Common Event Grammar.

Name/Description	ResourceClassName	Identities	Actions	Resources
View-DM-XMP-Files Can view data mapping files and message parsing files.	SafeObject	ug:View-DM-XMP-Files	read	AppObject

Filters	
WHERE	(name:pozFolder *--* val:/CALM_Configuration/Content/Mapping
OR	name:pozFolder *--* val:/CALM_Configuration/Content/Parsing)

7. Test the policies.
8. Assign users to both Create-DM-XMP-Files and View-DM-XMP-Files.

Sample Policies for Suppression and Summarization Rules

You can authorize non-Administrators to create custom suppression rules and custom summarization rules by creating one custom role, one CALM policy and one scoping policy. You can give other non-Administrators the ability to view custom suppression rules and custom summarization rules by creating an additional custom role with an associated scoping policy. You add both custom roles to the CALM Application Access policy and assign users to these roles.

The following example procedure shows you how to do this.

1. Create an application user group called Create-SUP-SUM-Rules.
2. Create an application user group called View-SUP-SUM-Rules.
3. Grant both roles access to the CA Enterprise Log Manager product.

Name/Description	ResourceClassName	Identities	Actions	Resources
CALM Application Access This policy defines who all can access the CALM Application	SafeObject	ug:Administrator ug:Analyst ug:Auditor ug:Create-Sup-Sum-Rules ug:View-Sup-Sum-Rules	read write	ApplicationInstance Policy User GlobalUser

4. Create a CALM policy that grants Create-SUP-SUM-Rules users the ability to create summarization and suppression rules or import them while logged on to CA Enterprise Log Manager.

Name/Description	ResourceClassName	Identities	Actions	Resources
EventGrouping-Create policy Can create custom summarization and suppression rules or import rules.	CALM	ug:Create-Sup-Sum-Rules	create	EventGrouping

5. Create a scoping policy that grants Create-SUP-SUM-Rules users the ability to view or edit custom summarization or suppression rules that have been saved to the EEM folder, /CALM_Configuration/Content/Rules/Suppression or /CALM_Configuration/Content/Rules/Summarization.

Name/Description	ResourceClassName	Identities	Actions	Resources
View-Edit-SUP-SUM-Rules Can view or edit suppression and summarization rules that have been saved to Content/Rules/Summarization or Content/Rules Suppression folders.	SafeObject	ug:Create-Sup-Sum-Rules	read write	AppObject

Filters			
WHERE	name:pozFolder	*--*	val:/CALM_Configuration/Content/Rules/Summarization
OR	val:pozFolder	*--*	val:/CALM_Configuration/Content/Rules/Suppression

6. Create a scoping policy that grants View-SUP-SUM-Rules users the ability to view custom summarization or suppression rules.

Name/Description	ResourceClassName	Identities	Actions	Resources
View-SUP-SUM-Rules Can view suppression and summarization rules that have been saved to Content/Rules/Summarization or Content/Rules Suppression folders.	SafeObject	ug:View-Sup-Sum-Rules	read	AppObject

Filters			
WHERE	name:pozFolder	*--*	val:/CALM_Configuration/Content/Rules/Summarization
OR	val:pozFolder	*--*	val:/CALM_Configuration/Content/Rules/Suppression

7. Test the policies
8. Assign users to the new roles. For example, external auditors may want to be able to view your summarization and suppression rules. To permit this, you could assign a role similar to View-Sup-Sum-Rules to such users.

An alternative to creating two new roles explicitly for the create/edit/view task and the view-only task is to expand the roles of the predefined Analyst and Auditor roles. For example, you could eliminate steps 1, 2, 3, and 8 of the previous procedure and instead assign the Analyst as the identity to EventGrouping Create policy and View-Edit-SUP-SUM-Rules and assign the user group Auditor as the identity to View-SUM-SUP-Rules.

More information:

Create an Application User Group (Role) (see page 75)
Grant a Custom Role Access to CA Enterprise Log Manager (see page 76)
Test a New Policy (see page 86)
Assign a Role to a Global User (see page 35)

Chapter 5: Services and CA Adapters

This section contains the following topics:

Services Tasks

You can set global configurations that apply to all CA Enterprise Log Manager servers. You can view and edit two types of individual service configurations: A global service configuration applies to all the instances of a single service in your environment, and a local service configuration only to a selected individual service host.

Note: Global configurations are distinct from global *service* configurations: the first controls the behavior of all CA Enterprise Log Manager servers, and the second that of a chosen service. For example, you can set the update interval for all services (global configuration), or report retention policies for all report servers (global service configuration).

You can also view self-monitoring events from the service configuration areas.

Available services include:

- Agent Manager
- Alerting Service
- Correlation Service
- Event Log Store
- Incident Service
- ODBC Server
- Query Service
- Report Server
- Rule Test Service
- Subscription Service
- System Status

You can display some services by service name, or by host. You can use the System Status service to gather information about, and to control, an individual CA Enterprise Log Manager server.

More information:

Edit a Global Service Configuration (see page 124)
Edit a Local Service Configuration (see page 125)

Delete a Service Host

If you uninstall a CA Enterprise Log Manager server, you must delete the host configuration from the management server repository. The removal of this reference will keep the server up to date with the list of its registered CA Enterprise Log Manager servers.

To delete a service host

1. Click the Administration tab, and then click the Services subtab.

   The Service List appears.

2. Click Host in the Show Services By dialog at the top of the list.

   An expandable tree list of service hosts appears.

3. Select the host you want to delete, and click Delete.

   The host is removed from the list.

Important! No warning appears when deleting a host. Clicking Delete immediately removes the host, so you must be sure you want to delete the host.

Edit Global Configurations

You can set global configurations for all services. If you attempt to save values outside the allowed range, CA Enterprise Log Manager defaults to the minimum or maximum as appropriate. Several of the settings are interdependent.

To edit global settings

1. Click the Administration tab and the Services subtab.

   The Service List appears.

2. Click Global Configuration in the Service List.

   The Global Service Configuration details pane opens.

3. Change any of the following configuration settings:

   Update Interval

   Specifies the frequency, in seconds, at which server components apply configuration updates.

   Minimum: 30

   Maximum: 86400

Session Timeout

Specifies the maximum length of an inactive session. If auto-refresh is enabled, a session never times out.

Minimum: 10

Maximum: 600

Allow Auto Refresh

Lets users auto-refresh reports or queries. This setting lets administrators globally disable auto-refresh.

Auto Refresh Frequency

Specifies the interval, in minutes, at which the report views refresh. This setting depends on the selection of Allow Auto Refresh.

Minimum: 1

Maximum: 60

Enable Auto Refresh

Sets auto-refresh in all sessions. Auto-refresh is not enabled, by default.

Viewing Action Alerts Requires Authentication

Prevents Auditors or third-party products from viewing Action Alert RSS feeds. This setting is enabled by default.

Default Report

Specifies the default report.

Enable Default Report Launch

Displays the default report when you click the Reports subtab. This setting is enabled by default.

4. Change any of the following report or query tag settings:

Hide Report Tags

Prevents specified tags from appearing in any tag list. Hiding report tags streamlines the view of the available reports.

Hide Query Tags

Lets you hide chosen tags. Hidden tags do not appear in the main query list or the action alert scheduling query list. Hiding query tags customizes the view of the available queries.

5. Change any of the following Dashboard settings:

   Enable Default Dashboard Launch

   Displays the default dashboard when you click the Queries and Reports tab. This setting is enabled by default.

   Default Dashboard

6. Change any of the following Profiles settings:

   Enable Default Profile

   Lets you set the default profile.

   Default Profile

   Specifies the default profile.

   Hide Profiles

   Lets you hide chosen profiles. When the interface refreshes or the Update Interval expires, the hidden profiles do not appear. Hiding profiles customizes the view of the available profiles.

   Note: Click Reset to restore the last saved values. You can reset a single change or multiple changes until you save changes. After you save changes, reset your changes individually.

7. Click Save.

Edit a Global Service Configuration

You can edit global service configurations, which are settings that apply to all instances of a given service in your environment. A global service configuration does *not* override any local service setting that differs from the global setting.

The maximum and minimum configuration values are detailed in the specific service sections. If you attempt to save values outside the allowed range, CA Enterprise Log Manager defaults to the minimum or maximum as appropriate.

To edit a global service configuration

1. Click the Administration tab, and then click the Services subtab.

   The Service List appears.

2. Select the service whose configuration you want to edit.

   The Global Service Configuration display opens in the details pane.

3. Make the configuration changes you want.

   Note: You can click Reset to restore the entry fields to the last saved value. You can reset a single change or multiple changes up to the point you click Save. Once you have saved changes you must reset your changes individually.

4. Click Save when you are finished making changes.

   Any configuration changes you are applied to all hosts of the selected service, unless they have different local settings.

Edit a Local Service Configuration

You can view or edit local service configurations by service or by host server. Local service configurations let you control services or settings that may not apply, or be required, for your entire environment, overriding global settings only for specific hosts. For example, you may want a specific CA Enterprise Log Manager server to retain action alerts longer than others. You control this using a local configuration.

The maximum and minimum configuration values are detailed in the specific service sections. If you attempt to save values outside the allowed range, CA Enterprise Log Manager defaults to the minimum or maximum as appropriate.

To edit a local service configuration

1. Click the Administration tab, and then click the Services subtab.

   The Service List appears.

2. Click the arrow beside the service whose configuration you want to edit.

   The service display expands, showing individual service hosts.

3. Click the service host you want.

   The service configuration you select opens in the details pane.

4. Make the configuration changes you want. Every entry field, menu, or control in the local configuration displays a local/global configuration button which can be toggled to one of two states.

   Global configuration: 

   Local configuration: 

   Clicking the button changes it from the global to the local setting, and makes its associated value available for use. The value must remain set for local configuration for the setting to take effect: If it is set for global configuration, the global setting for that listener is in effect.

   Note: Clicking Reset shows the most-recently saved configuration values for all the available configurations. You can reset a single change or multiple changes up to the point you click Save. Once you have saved changes you must reset your changes individually.

5. Click Save when you are finished making changes.

   Any changes you make are applied to the selected service host only.

Service Configurations

This section includes details and service guidelines to review when making configuration changes in the following CA Enterprise Log Manager services:

- Alerting Service - controls delivery settings for action alerts, including SMTP servers, CA IT PAM and SNMP trap settings.
- Correlation Service - controls correlation rules and event routing for incident creation.
- Event Log Store - stores all refined and recorded raw events.
- Incident Service - controls the creation and storage of incidents resulting from event correlation.
- ODBC Server - provides access to the CA Enterprise Log Manager event log store from an external application such as BusinessObjects Crystal Reports.
- Report Server - controls distribution, formatting, and the retention of reports and alerts.
- Subscription Service - routes content and configuration updates to the management server and binary updates to subscription clients.

More information:

Subscription (see page 185)

Alerting Service Considerations

The Alerting Service controls the delivery of action alerts. You can perform the following tasks from the alerting service configuration area:

- Set the mail server, admin email, and SMTP port and authentication information for alert delivery in the Email Settings area.
- Configure CA IT PAM integration for action alerts.
- Set up SNMP traps for action alert delivery.

More information:

Configure Integration with an SNMP Trap Destination (see page 128)

Configure CA IT PAM Integration for Event/Alert Output

You can configure CA IT PAM integration to leverage either or both of the following types of CA IT PAM processes:

- Event/alert output process--a process that invokes processing on a third-party system
- Dynamic values process--a process that accepts an input key and returns current values for that key as a comma-separated values (*.csv) file

The following procedure addresses both the common settings and settings specific to event/alert output. Refer to the details you recorded as you configure IT PAM integration for event/alert output.

To configure IT PAM integration for the event/alert output process

1. Click the Administration tab and the Services subtab.
2. Click Alerting Service

   The Global Service Configuration: Alerting Service dialog appears.
3. Scroll to the IT PAM area.
4. Enter the fully qualified host name of the server on which CA IT PAM is installed, accept the default port number, 8080, and enter valid login credentials for CA IT PAM.
5. If you have imported the sample EventAlertOutput.xml for use, accept the default entry for Event/Alert Output Process. If not, replace this entry with your custom event/alert output process name preceded by its path.

   Note: You can view the Name and Path of the process under Folders in the ITPAM Client.
6. If you have imported the sample EventAlertOutput.xml for use, define the default values for ReportedBy, Severity, Priority, and EndUser as follows:
   a. Select a parameter and click Add Default Value.

      The Add Value dialog appears.
   b. Enter the default value and click OK.
7. If you specified a custom event/alert output process, delete the displayed parameters and add your own. Then define the default value for each.
8. Click Save.

   The following message appears: Confirmation: Configuration changes saved successfully.

Configure Integration with an SNMP Trap Destination

Configure SNMP integration as part of the Global Service Configuration for Report Server. The configuration is the IP address and port of one SNMP trap destination.

You can configure SNMP integration either before or after preparing the destination product to receive and interpret SNMP traps from CA Enterprise Log Manager.

When you create an alert destined for an SNMP trap recipient, you can specify one or more destinations. This configuration serves as the default. This default applies to all servers listed under Report Server.

To configure SNMP integration

1. Click the Administration tab and the Services subtab.
2. Click Alerting Service

   The Global Service Configuration: Alerting Service dialog appears.
3. Scroll to the SNMP Configuration area.
4. Enter the IP address or host name of the destination server for the SNMP traps.
5. Accept the default port, 162, or change it.
6. Click Save.

Correlation Service Considerations

The Correlation Service controls the rules applied on the correlation server. When you apply a rule it becomes active.

You can associate notification destinations with rules, and enable or disable rules from the Correlation Service Configuration page. You can choose which CA Enterprise Log Manager servers route events to the selected correlation server, or set an Event Limit.

Event Limit

Defines how many events are retained per incident when accumulation is enabled. The Event Limit helps prevent undue traffic caused by correlation in periods of high activity. When this limit is reached, additional events are lost. For example, if your limit is set to 100, a single rule can accumulate up to 100 recorded events, including the initial qualifying event or events. Accumulation continues until the event limit is reached, or more usually, gap or limit values reset the rule.

You can also remove applied correlation rules, making them inactive.

To remove rules from the applied list

1. Highlight the row for the correlation rule you want to remove.
2. (Optional) You can control-click or control-shift to highlight multiple rows.
3. Click Remove.

   The highlighted rules are removed from the active list.
4. Click Save to confirm the configuration. If you have not saved, you can click Reset to restore removed rules to the list.

   Note: This procedure only removes correlation rules from the active list. They are not removed from the rule library.

Event Log Store Considerations

The event log store uses a federated event log store system, with each host server maintaining its own local event log store and the ability to contact other event log stores in your environment. When you query a server for event information, it can search its own local event log store as well as all others connected through the federation. This arrangement allows for flexible storage and archiving of event data.

The event log store archive settings let you specify how often data is archived and where it is stored. Both hot (active) event log stores and warm (archived) event log information are queried. Event information in cold storage (remote) is not queried.

You can configure the following event log store and archiving settings:

Maximum Rows

Sets the maximum number of events your event log store's hot database can contain. When the event count reaches this value, the event log compresses all event information in the hot database and moves it to the warm database.

Minimum: 50000

Maximum: 100000000

Max Archive Days

Sets the number of days archived files are retained in the archive before being deleted.

Minimum: 1

Maximum: 28000

Archive Disk Space

Defines the percentage of remaining disk space which triggers automatic deletion of the oldest archive files. For example, the default value is 10. When the available event log store space falls below 5 percent, the event log removes the oldest archive files to make more room.

Minimum: 10

Maximum: 90

Export Policy

Defines the number of hours a file restored from an outside backup source to the archive (defrosted) will be retained in the event log store before being deleted.

Minimum: 0

Maximum: 168

Summarization/Suppression Rules

Controls which of the available summarization or suppression rules are applied to received events. New summarization or suppression rules must be applied by an administrator before they begin refining events.

Forwarding Rules

Controls which of the available event forwarding rules are applied to received events.

Federation Children

Controls which of the available event log stores are set as children of the current server. This lets you set up separate federation "trees", controlling query access levels. This setting is only available as a local setting.

Logging settings control how individual CA Enterprise Log Manager modules record internal messages. They are only available as local settings. Logging settings are usually used for troubleshooting purposes. It is not normally necessary to change these settings, and you should have a good understanding of log files and logging before doing so.

Log Level

Defines the type and level of detail recorded in the logging file. The drop-down list is arranged in order of detail, with the first choice providing least detail, and the last providing most detail.

Apply to all loggers

Controls whether the Log Level setting overrides all log settings from the log's properties file. This setting only applies when the Log Level setting is lower (showing more detail) than the default setting.

Auto Archive Settings enable and control scheduled database archiving jobs, which move warm databases to a remote server.

Note: Before you move scheduled database jobs from one CA Enterprise Log Manager server to another, or to a remote server, you must configure non-interactive authentication between the servers. See the Configuring Non-interactive Authentication section of the *CA Enterprise Log Manager Implementation Guide* for more information.

You can set the following auto archive values:

Enabled

Sets an auto archive job to run. The auto archive uses the scp utility as controlled by the other settings.

Backup Type

Controls the backup type: A full archive that copies all database information, or an incremental archive that copies all databases that have not yet been backed up.

Default: Incremental

Frequency

Specifies whether the archive job runs daily or hourly. A daily job runs at the time you set using the Start Time clock. An hourly job runs every hour on the hour.

Start Time

Sets the time a daily archive job runs, in whole hours, based on the server's local time. The value is a 24-hour clock.

Limits: 0-23, where 0 means midnight and 23 means 11:00 p.m.

EEM User

Specifies the user who can perform an archive query, recatalog the archive database, run the LMArchive utility, and run the restore-ca-elm shell script to restore archive databases for examination. This user must be assigned the predefined role of Administrator or a custom role associated with a custom policy that permits the edit action on the Database resource.

Default: Log Manager administrator user

EEM Password

Specifies the password for the user who has the rights defined in the EEM user field.

Remote Server

Specifies the hostname or IP Address of the remote server to which the auto archive job copies the database information.

Remote User

Specifies the username that the scp utility uses to connect to the remote server.

Default: caelmservice

Remote Location

Specifies the archive file destination on the remote server.

Default: /opt/CA/LogManager

Remote ELM Server

Specifies whether the remote server is a management server or not. If it is, then the auto archive job will delete the databases from the local machine when the transfer is complete and notify the remote machine to recatalog itself.

Correlation Event Reception Span

Controls how wide a time variance is tolerated for the creation of incidents. The two values allow you to set a value after the current <CALM >server time (future) and before the current CA Enterprise Log Manager server time (past). If an event falls outside that window it is not forwarded for correlation.

More information:

Log Storage (see page 151)
Apply a Suppression or Summarization Rule (see page 437)

Incident Service Considerations

You can control the way in which the incident service stores events and creates incidents for a selected CA Enterprise Log Manager server. You can set the following values:

Expiration Time

Specifies how long in days the service retains incidents in the incident database. If the value is 0, events are never deleted. Expired incidents are not displayed.

Incident Generation Limit values

Specifies how often a single correlation rule can create incidents, allowing you to reduce unwanted multiple incidents. For the purposes of incident generation limits, different versions of a rule are considered separate rules. So if you have applied multiple versions of a rule in your environment, they are limited separately. Limit values include:

Enabled

Indicates whether incident generation limits are applied.

Count

Sets a threshold for the number of incidents generated by a single rule. This value works with the Time value, if that value is above 0. After these numbers are reached, the incident service applies the Blocked Time limit. So if you set Count to 3, and the Time to 10, the limit applies after a single rule generates more than 3 incidents in 10 seconds.

Time

Sets a threshold, in seconds, for the number of incidents generated by a single rule. This value works with the Count value, if that value is above 0. After these numbers are reached, the incident service applies the Blocked Time limit. So if you set Count to 3, and the Time to 10, the limit applies after a single rule generates more than 3 incidents in 10 seconds.

Blocked Time

Specifies an interval in seconds, when a rule is blocked from creating further incidents. When this limit is reached, the rule creates no incidents until the time expires.

ODBC Server Considerations

You can install an ODBC client or a JDBC client to access the CA Enterprise Log Manager event log store from an external application like SAP BusinessObjects Crystal Reports.

You can perform the following tasks from this configuration area:

- Enable or disable ODBC and JDBC access to the event log store.
- Set the service port used for communications between the ODBC or JDBC client and CA Enterprise Log Manager server.
- Specify whether communications between ODBC or JDBC client and CA Enterprise Log Manager server are encrypted.

The field descriptions are as follows:

Enable Service

Indicates whether the ODBC and JDBC clients can access data in the event log store. Select this check box to enable external access to events. Clear the check box to disable external access.

The ODBC service is not currently FIPS-compatible. Clear this check box to prevent ODBC and JDBC access if you intend to run in FIPS mode. This prevents non-compliant access to event data. If you intend to disable the ODBC and JDBC service for FIPS mode operations, ensure that you set this value for *each* server in a federation.

Server Listening Port

Specifies the port number used by the ODBC or JDBC services. The default value is 17002. The CA Enterprise Log Manager server refuses connection attempts when a different value is specified in the Windows Data Source or the JDBC URL string.

Encrypted (SSL)

Indicates whether to use encryption for communications between the ODBC client and the CA Enterprise Log Manager server. The CA Enterprise Log Manager server refuses connection attempts when the corresponding value in the Windows Data Source or JDBC URL does not match this setting.

Session Timeout (minutes)

Specifies the number of minutes to keep an idle session open before it is closed automatically.

Log Level

Defines the type and level of detail recorded in the logging file. The drop-down list is arranged in order of detail, with the first choice providing least detail.

Apply to all loggers

Controls whether the Log Level setting overrides all log settings from the properties file of the log. This setting only applies when the Log Level setting is lower (showing more detail) than the default setting.

Report Server Considerations

The Report Server controls the administration of automatically delivered reports, and their appearance in PDF format, and Action Alert and report retention. You can perform the following tasks from the report server configuration area:

- Control the company name and logo, fonts, and other PDF reports settings in the Report Configurations area.
- Set the total Actions Alerts retained, and number of days they are retained in the Alert Retention area:

  Maximum Action Alerts

  Defines the maximum number of action alerts the reporting server retains for review.

  Minimum: 50

  Maximum: 1000

  Action Alerts Retention

  Defines the number of days action alerts are retained, up to the maximum number.

  Minimum: 1

  Maximum: 30

- Set the retention policy for each scheduled report recurrence type in the Report Retention area.
- Set whether or how often the retention utility searches for reports to delete automatically based on those policies. For example, if the report retention utility runs daily, it deletes reports daily that are older than the specified maximum age.

Rule Test Service Considerations

The rule test service controls how CA Enterprise Log Manager tests correlation rules. You can set the following rule test values.

Event Limit

Defines how many events are retained per incident when accumulation is enabled. The Event Limit helps prevent undue traffic caused by correlation in periods of high activity. When this limit is reached, additional events are lost. For example, if your limit is set to 100, a single rule can accumulate up to 100 recorded events, including the initial qualifying event or events. Accumulation continues until the event limit is reached, or more usually, gap or limit values reset the rule.

Maximum Concurrent Rule Tests

Defines the number of rule tests that can be run simultaneously on a single CA Enterprise Log Manager server.

Subscription Considerations

A Proxy/Client server system delivers subscription updates. The first server you install is set as your Default Subscription Proxy server, which contacts the CA Subscription Server periodically to check for updates. Subsequent installations are configured as clients of that proxy server, contacting it periodically for updates. If they fail to make contact, a self-monitoring event is logged.

The default system reduces network traffic by eliminating the need for each server to contact the CA Subscription Server directly, but is fully configurable. You can add proxy servers as needed.

You can also reduce internet traffic still further by creating offline proxy servers, which store update information locally and provide it to clients when contacted. Support any offline proxy servers by manually copying everything in the download path of the online proxy to the download path of the offline proxy. Offline proxies must be configured in environments where there are CA Enterprise Log Manager servers that cannot access the Internet or an internet-connected server.

When configuring the Subscription Service, consider the following information about certain settings and their interactions:

Default Subscription Proxy

Defines the default proxy server for the Subscription Service. The default subscription proxy must have internet access. If no other subscription proxies are defined, this server gets subscription updates from the CA Subscription server, downloads binary updates to all clients, and distributes content updates to the CA Enterprise Log Manager user store. If other proxies are defined, clients contact this server for updates when no subscription proxy list is configured or when the configured list is exhausted. The default value is the first server installed in your environment. This value is only available as a global setting.

Public Key

Defines the key used to test and verify the signature used to sign the updates. When a public-private key pair is updated, the proxy downloads the update to the public key value, and the proxy updates the public key. This value is only available as a global setting.

Important! Never manually update this value.

Subscription Proxy

Controls whether the local server is a subscription proxy. If the subscription proxy check box is cleared, the server is a subscription client.

Update Now

Starts an on demand update cycle immediately for the selected server. You can perform an on demand update for only one server at a time; this option is not available globally. Update a subscription proxy server before you update its subscription client.

Online Subscription Proxy

Controls whether the local server is an online subscription proxy. An online subscription proxy uses its internet access to get subscription updates from the CA Subscription server and distribute them to the CA Enterprise Log Manager environment. To designate a server as an online subscription proxy, select both the Subscription Proxy check box and the Online Subscription Proxy option. This value is only available as a local setting.

Offline Subscription Proxy

Controls whether the local server is an offline subscription proxy. An offline subscription proxy is a server that gets subscription updates through a manual directory copy (using scp) from an online subscription proxy. Offline subscription proxies do not need internet access. To designate a server as an offline subscription proxy, select both the Subscription Proxy check box and the Offline Subscription Proxy option. This value is only available as a local setting.

Note: For important information about configuring offline subscription, see the *CA Enterprise Log Manager Administration Guide.*

RSS Feed URL

Defines the URL of the CA Subscription server. Online subscription proxies use this URL to access the CA Subscription server and download subscription updates.

Modules Available for Download

Lets you select from the modules available for download the modules that apply to your CA Enterprise Log Manager environment. Click Browse to display this dialog; the modules you select appear in the Modules list.

Modules selected are downloaded from the CA Subscription Server during subscription updates. Modules for download can be selected at the global level; other configured subscription proxies download these modules by default during update. Modules for download can also be selected at the local level for individual proxy and client servers. Doing so overrides global settings so that only the selected modules download to the given server. Modules selected for clients are used to update corresponding modules installed on the client. You can select a module to download for a client that is not selected for its proxy. The proxy retrieves it for the client, but does not install it on itself.

Note: If not populated, set the RSS Feed URL. This setting lets the system read the RSS Feed and, at the next update interval, display the list of available modules to download.

Modules Selected for Download

Displays the modules selected in the RSS Feed Browser dialog. The default subscription proxy and all other online proxies download these modules from the CA Subscription Server during the update process. The modules listed can be modules chosen for download at the global level, or can reflect modules selected for a given server at the local level.

HTTP Proxy Server

Controls whether this server contacts the CA Subscription Server through an HTTP Proxy for updates, rather than directly.

Proxy Address to Use

Specifies the full IP Address of the HTTP Proxy.

Port

Specifies the Port number used to contact the HTTP Proxy.

HTTP Proxy User ID

Specifies the user ID used to contact the HTTP Proxy.

HTTP Proxy Password

Specifies the password used to contact the HTTP Proxy.

Schedule

Specifies the start time and frequency for CA Enterprise Log Manager servers to request subscription updates. Online subscription proxy servers (including the default proxy server) contact the CA Subscription Server, and proxy clients contact their proxy servers, according to this schedule. The schedule can be set globally for all CA Enterprise Log Manager servers; it can also be overridden locally for a given server.

Subscription Proxy for Client Updates

Lets you set which proxies are contacted, in a round robin fashion, for product and operating system updates by all clients or the selected client. You can use the up/down arrows to control the order in which the client contacts the subscription proxies. The client downloads updates from the first proxy it successfully reaches. If none of the configured proxies are available, the client contacts the default subscription proxy.

Subscription Proxy(s) for Content Updates

Lets you select which proxies are used to distribute content updates to the user store. You can select offline proxies or online proxies. This value is only available as a global setting.

Note: Consider selecting more than one server to act as a subscription proxy for content updates, for redundancy.

More information:

Free Disk Space for Updates (see page 195)
About Subscription Public Keys (see page 196)

System Status Service

You can use the System Status service to gather information about, and to control, a CA Enterprise Log Manager server. You display system status only for individual CA Enterprise Log Manager servers. All settings and options apply at the local level.

The System Status service offers the following tabs:

- Administration - control services and host servers, and create a support diagnostics file
- Status - review the status and version of system services and processes
- Self Monitoring Events - review events related to system status and component status

More information

System Status Tasks (see page 146)
Create a Diagnostics File for Support (see page 147)
Reboot a Host Server (see page 147)
Restart the ELM Services (see page 148)
Review Service Status and Version (see page 148)
Review System Status Self Monitoring Events (see page 149)

CA Adapters Configuration Tasks

Local listeners receive and collect native events from certain types of sources using various types of CA adapters.

You can view and edit two types of individual adapter configurations.

- A global configuration applies to all the instances of a single adapter in your environment, such as all SAPI collector instances.
- A local configuration only to a selected individual adapter host, such as a single SAPI collector.

You can also view self-monitoring events for each adapter service or adapter host from the individual adapter's global or local configuration areas.

More information:

Edit a Global Adapter Configuration (see page 140)
Edit a Local Adapter Configuration (see page 141)
View Adapter Self-Monitoring Events (see page 142)
View Adapter Status (see page 143)
SAPI Service Considerations (see page 144)
iTechnology Event Service Considerations (see page 145)

Edit a Global Adapter Configuration

You can edit global adapter configurations, which are settings that apply to all instances of a given CA Adapter in your environment. For example, you could make configuration changes that apply to all SAPI collectors running in your environment. A global adapter configuration does *not* override any local adapter settings that differ from the global setting.

To edit a global adapter configuration

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the CA Adapters folder.

   The folder expands, displaying subfolders for each adapter.

3. Select the folder for the adapter whose configuration you want to edit.

   The Global Service Configuration display opens in the details pane.

4. Make the configuration changes you want.

   Note: Clicking Reset restores the configuration values to the most-recently saved states. You can reset a single change or multiple changes up to the point you click Save. Once you have saved changes you must reset your changes individually.

5. Click Save when you are finished making changes.

   Any configuration changes you make will be applied to all hosts of the selected adapter, unless they have differing local settings.

Edit a Local Adapter Configuration

You can view or edit local adapter configurations. Local adapter configurations allow you to control settings that may not apply, or be required, for your entire environment. They override global settings only for specific adapter hosts. For example, you may want a specific SAPI adapter host to listen on a different port. You can set this behavior using a local configuration.

To edit a local adapter configuration

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the CA Adapters folder.

   The folder expands, displaying subfolders for each adapter.

3. Select the folder for the adapter whose configuration you want to edit.

   The service display expands, showing adapter hosts.

4. Click the adapter host you want.

   The host configuration you select opens in the details pane.

5. Make the configuration changes you want. Every value entry field, menu, or control in the local configuration displays a local/global configuration button which can be toggled to one of two states.

   Global configuration: 

   Local configuration: 

   Clicking the button changes it from the global to the local setting, and makes its associated entry field available for use. The entry field must remain set for local configuration for the setting to take effect: If it is set for global configuration, the global setting for that adapter is in effect.

   Note: Clicking Reset shows the most-recently saved configuration values for all the available configurations. You can reset a single change or multiple changes up to the point you click Save. Once you have saved changes you must reset your changes individually.

6. Click Save when you are finished making changes.

   Any changes you make are applied to the selected adapter host only.

View Adapter Self-Monitoring Events

You can monitor adapter service activity and troubleshoot problems by viewing self-monitoring events for each adapter service host. You can see pre-screened events from the individual adapter's global or local configuration areas.

To view adapter self-monitoring events

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the CA Adapters folder.

   The folder expands, displaying subfolders for each adapter service.

3. Select the folder for a adapter service to view self-monitoring events for that service, or expand the folders and choose a adapter host to view only self-monitoring events for that individual adapter host.

   The adapter configuration appears in the details pane

4. Click the Self-Monitoring Events tab.

   An event viewer window showing appropriately-filtered events appears. For example, if you select the iTechnology Event Plugin folder in Step 3, you see self-monitoring events for all instances of the iTechnology Event Plugin. If you select a specific host from the iTechnology Event Plugin folder, you see only events relating to that specific iTechnology host.

   Note: Your federation structure controls which events are visible. If no federation is set up then you will only see local events, regardless of which host you select.

More information:

View Adapter Status (see page 143)
Edit a Global Adapter Configuration (see page 140)
Edit a Local Adapter Configuration (see page 141)

View Adapter Status

You can view the current status of certain CA adapter services, including start time, running state, and event delivery information and statistics. You cannot view the status of the iTechnology Event plug-in service.

To view an adapter status

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the CA Adapters folder.

   The folder expands, displaying subfolders for each adapter service.

3. Select the folder for the adapter whose status you want to view.

   The service display expands, showing individual adapter hosts.

4. Click the adapter host you want.

   The host configuration you select opens in the details pane.

5. Click the Status tab.

   The Status information appears.

   Note: Status information appears only in the local configuration panel.

SAPI Service Considerations

CA Enterprise Log Manager uses two instances of a CA Audit Submit Application Programming Interface (SAPI) service, one installed as the SAPI Collector, the other as the SAPI Router. The SAPI services are generally used to receive events from existing CA Audit clients and integrated products. You can configure the SAPI adapters using the following settings:

Enable Listener

Activates the selected service. This setting is enabled by default.

SAPI Port

Sets a specific port number for the selected service, if it is not registered with the portmapper. The default value, 0, allows the service to use a randomly-determined port, if the Register check box is selected.

Note: The port number must be different for the SAPI Collector and Router. If the same ports are set for both services, the second one set will not function/.

Register

Controls whether the service registers with the system portmapper. If you select Register and enter 0 in the SAPI Port field, a random port is selected each time the service starts. This is the default setting for both fields. If Register is not selected, you must specify a SAPI port.

Encryption Key

Defines the encryption key, if you use a non-standard encryption key in your CA Audit environment, which the SAPI adapter uses to read incoming SAPI events.

Event Ordering

Ensures that events are sent to the event log store in the same order in which they are received. If event ordering is disabled, the order may be changed if some events are parsed and sent onward more quickly than others. Enabling event ordering may affect performance by increasing the size of the event queue.

Event Throttling

Defines the maximum number of events in the event processing queue, allowing control of processing resources. Entering 0 in this field means that no throttling occurs. Events that exceed this threshold will be delayed at the source.

Thread Count Per Queue

Defines the number of processing threads for each protocol. Using many processing threads will speed up processing if event ordering is disabled. If event ordering is enabled, the thread count will have no effect. Using many threads may have performance implications.

Cipher and Data Mapping

- The Cipher shuttle control determines which of the available ciphers the service uses to decrypt incoming messages.

- The Data Mapping file shuttle control determines which of the available DM files the service uses for event mapping.

Logging settings control how individual CA Enterprise Log Manager modules record internal messages. They are only available as local settings. Logging settings are usually used for troubleshooting purposes. It is not normally necessary to change these settings, and you should have a good understanding of log files and logging before doing so.

Log Level

Defines the type and level of detail recorded in the logging file. The drop-down list is arranged in order of detail, with the first choice providing least detail, and the last providing most detail.

Apply to all loggers

Controls whether the Log Level setting overrides all log settings from the log's properties file. This setting only applies when the Log Level setting is lower (showing more detail) than the default setting.

iTechnology Event Service Considerations

The iTechnology service controls events sent through the iGateway daemon. You can configure the service by setting which of the available data mapping (DM) files the service uses for event mapping, using the DM file shuttle control.

The event plug-in service is preconfigured to include most of the major data mapping files.

Logging settings control how individual CA Enterprise Log Manager modules record internal messages. They are only available as local settings. Logging settings are usually used for troubleshooting purposes. It is not normally necessary to change these settings, and you should have a good understanding of log files and logging before doing so.

Log Level

Defines the type and level of detail recorded in the logging file. The drop-down list is arranged in order of detail, with the first choice providing least detail, and the last providing most detail.

Apply to all loggers

Controls whether the Log Level setting overrides all log settings from the log's properties file. This setting only applies when the Log Level setting is lower (showing more detail) than the default setting.

System Status Tasks

You can do the following things from within the System Status service:

- Review the status and version of system services.
- Review self-monitoring events related to system components and usage.
- Create a Support diagnostics file.
- Restart the ELM services.
- Reboot the host server on which a CA Enterprise Log Manager server is running.
- Enable FIPS and non-FIPS mode operation.

More information:

Create a Diagnostics File for Support (see page 147)
Reboot a Host Server (see page 147)
Restart the ELM Services (see page 148)
Review Service Status and Version (see page 148)
Review System Status Self Monitoring Events (see page 149)

Create a Diagnostics File for Support

You can review the status and version for services running on a selected CA Enterprise Log Manager server. Clicking Support Diagnostics executes the LmDiag.sh script provided with CA Enterprise Log Manager.

This utility packages system information and log files into a compressed .tar file for transmission to CA Support personnel. You can transfer this file using FTP or another file transfer method.

Note: Some of the information in the resulting file can be sensitive, for example, IP addresses, system configurations, hardware logs, and process logs. Use a secure method for storing and transporting this file.

To create a diagnostic file

1. Click the Administration tab and then click the Services subtab.
2. Expand the System Status entry.
3. Select a specific CA Enterprise Log Manager server.

   The System Status Service Configuration displays the Administration tab.
4. Click Support Diagnostics.
5. Select a file location for the download of the generated diagnostics file.

   The utility creates the file and downloads it to the specified location. The utility closes automatically when the file is copied.

Reboot a Host Server

You can review the status and version for services running on a selected CA Enterprise Log Manager server.

Important! Use this feature only when necessary, or when directed to do so by CA Support. Rebooting a CA Enterprise Log Manager server causes it to stop receiving, parsing, and storing event logs until the reboot is complete. If you reboot the management server, the managed CA Enterprise Log Manager sessions on other, associated servers must log out and log back in again.

To reboot a host server

1. Click the Administration tab, and then click the Services subtab.
2. Expand the System Status entry.
3. Select a specific CA Enterprise Log Manager server.

   The System Status Service Configuration displays the Administration tab.
4. Click Reboot Host.

Restart the ELM Services

You can restart the ELM services running on a selected CA Enterprise Log Manager server.

Important! Use this feature only when necessary, or when directed to do so by CA Support. Restarting the ELM services causes the affected CA Enterprise Log Manager server to stop receiving, parsing, and storing event logs until the restart is complete. If you restart the management server, the current session and all other CA Enterprise Log Manager sessions on other servers must log out and log back in again.

To restart the ELM services

1. Click the Administration tab and then click the Services subtab.
2. Expand the System Status entry.
3. Select a specific CA Enterprise Log Manager server.

   The System Status Service Configuration displays the Administration tab.
4. Click Restart ELM Services.

Review Service Status and Version

You can review the status and version for services running on a selected CA Enterprise Log Manager server.

To review status

1. Click the Administration tab and then click the Services subtab.
2. Expand the System Status entry.
3. Select a specific CA Enterprise Log Manager server.
4. Click the Status tab.

Review System Status Self Monitoring Events

You can review the status and version for services running on a selected CA Enterprise Log Manager server. The status messages include events related to processor and disk space usage, CPU load averages, memory use, hardware access and usage, and other events.

To review self monitoring events

1. Click the Administration tab and then click the Services subtab.
2. Expand the System Status entry.
3. Select a specific CA Enterprise Log Manager server.
4. Click the Self Monitoring events tab.

Chapter 6: Log Storage

This section contains the following topics:

About Log Storage

You can manage two aspects of log storage through CA Enterprise Log Manager:

- Backing up the databases of log files in the archive directory of each reporting server to an archive directory you create on a remote storage server. The remote storage server is an interim location for holding archived databases until they can be moved off-site.
- Restoring the databases of log files from the archive directory on a remote storage server to the original reporting server or a CA Enterprise Log Manager you have dedicated as a restore point server. Once restored, you can examine the contents with queries and reports.

You can manage backups of event log databases in one of two ways:

- (Preferred) Configure CA Enterprise Log Manager to use auto-archive to move warm databases from a CA Enterprise Log Manager reporting server to a remote storage server on a scheduled basis. The auto-archiving process notifies the reporting server that the databases have been backed up.

  Note: See "About Auto Archive" in the *CA Enterprise Log Manager Implementation Guide.*

- Back up the databases on the CA Enterprise Log Manager server manually and copy them to an on-site storage location. Use LMArchive utility to notify the CA Enterprise Log Manager server to mark these databases as backed up.

  Moving backed up files to an off-site location is a task you perform outside of CA Enterprise Log Manager, as is the move back to the network, when needed for restoration.

You can query the archive catalog to identify database files to restore. You can restore databases on an as needed basis in one of two ways:

- You can restore them to the original reporting server with either of the following methods:
  - If you configure non-interactive authentication between the remote storage server and the original reporting server, run the restore-ca-elm.sh script to restore the archived databases to the original reporting server.

    Once the files are restored, query and report on them for the length of time in days configured for the life of warm files.
  - If you backed up the archive databases manually, copy the files back to the same archive directory and then notify this CA Enterprise Log Manager of the restoration. You use an option of the LMArchive command-line utility to inform CA Enterprise Log Manager that the databases are restored.

    Once the files are restored, query and report on them for the length of time in hours configured for the life of defrosted files.
- You can restore archived databases to a restore point server dedicated to examining restored event logs, with either of the following methods.
  - If you configure non-interactive authentication from the remote storage server to the CA Enterprise Log Manager restore point, you can run the restore-ca-elm.sh script to restore archived databases to the restore point.
  - If you have not configured non-interactive authentication, manually copy the archive databases from the remote storage server to the archive directory of the restore point server. Then, notify this CA Enterprise Log Manager of the restoration with a Recatalog from the Archive Catalog Query in the Log Collection Explorer.

    This notification results in the rebuilding of the catalog, which makes the database files available for querying and reporting. This availability is contingent on the age in days configured for warm files before deletion being set to a value that exceeds the age of the restored files. Therefore, it is important that the maximum age for warm files is set appropriately on any dedicated restore point.

More information:

Configuring Non-Interactive Authentication for Restore (see page 160)
Restore Auto-Archived Files (see page 167)
Manually Backing Up Archived Databases (see page 170)
Configure Max Archive Days for Restored Archives (see page 181)
Manually Restoring Archives to the Original Event Log Store (see page 174)
Manually Restoring Archives to a New Event Log Store (see page 179)
Restore–Script for Restoring Archived Databases (see page 168)
LMArchive–Backup/Restore Tracking (see page 183)

Event Log Database States

When you configure auto-archive across three servers (collection, reporting, and remote storage), all event log databases progress through three states: hot, warm, and cold. With this architecture, a hot database of uncompressed logs exists only on the collection server. The reporting server holds compressed warm databases; the remote storage server holds only cold databases. When a cold database is restored with the restore shell script, it is restored in a warm state. When it is restored manually with LMArchive utility, it is restored in a defrosted state.

The following four event log storage states describe uncompressed, compressed, backed up and moved, and restored databases, respectively:

Hot

A *hot database state* is the state of the single uncompressed database in the event log store of a collection server where newly processed events are inserted. You can configure the maximum number of new records to store in a hot database (Maximum Rows) before compressing it into a warm database. You can schedule auto-archiving to move warm databases from the collection server to the configured reporting server hourly. (A hot database also exists on the reporting server for inserting self-monitoring events.)

Warm

The *warm database state* is the state of databases retained in the event log store of the reporting server. If you configure daily auto-archiving between the reporting server and a remote storage server, warm databases are retained until they are moved to the remote storage server; then they are auto-deleted from the reporting server. If you do not configure auto-archiving between the reporting server and a remote storage server, warm databases can remain on the reporting server until their age in days reaches the configured value for Max Archive Days or when the configured Archive Disk Space threshold is reached, whichever comes first. When one of these thresholds is met, the database is deleted and its state changed to cold. Without auto-archiving, you must manually back up warm databases with a third party tool before they are deleted and then run LMArchive utility to notify CA Enterprise Log Manager of the names of the databases you backed up and moved. The warm state is also applied when a recatalog is performed with the restore-ca-elm.sh script or the Recatalog button after restoring a cold database.

Cold

A *cold database state* applies to a database on the remote storage server. A record of a cold database is created on the reporting server when the database is auto-archived to the remote management server and deleted from the reporting server. If handled manually, a record of the cold database is created when the LMArchive utility is run with the -notify arch option. You can query the archive catalog of a reporting server to identify cold databases to restore.

Defrosted

A *defrosted database state* is the state applied to a physical cold database that has been restored to the archive directory after the Administrator runs the LMArchive utility with the -notify rest option to notify CA Enterprise Log Manager that it has been restored. Defrosted databases are retained for the number of hours configured for the Export Policy.

You can query databases in every state. A normal query returns event data from the hot and warm databases on the reporting server and defrosted databases, if they exist. A federated query returns event data from all servers in the federation, including the federated collection servers that include hot databases. An archive query returns a list of databases that no longer exist on the reporting server, that is, databases in a cold state. The physical databases represented by an archive query can exist on the remote storage server used for on-site storage or in off-site storage.

More information

Automating Backup and Restore (see page 155)
Manually Backing Up Archived Databases (see page 170)
Manually Restoring Archives to the Original Event Log Store (see page 174)
Manually Restoring Archives to a New Event Log Store (see page 179)

Automating Backup and Restore

A backup process ensures that no data is lost due to the deletion of aged databases. The preferred method of backing up archived databases is to use auto archive. Auto archive is a scheduled, automated transfer of archived databases between pairs of servers. Auto archiving between a source server and a destination server requires non-interactive authentication. Non-interactive authentication uses RSA public-key authentication without a passphrase. You can configure non-interactive authentication and auto archiving:

- From each collection server to its reporting server.
- From each reporting server to its remote storage server.

Note: For details, see the *Implementation Guide*.

A restore process moves archived databases from remote storage to a CA Enterprise Log Manager server for investigation. The preferred method of restoring archived databases is to use the restore-ca-elm.sh script. This restore utility automates the transfer of archived databases. Like the auto archive process, the restore-ca-elm.sh script uses non-interactive authentication. You can configure non-interactive authentication and run the restore script:

- From the remote storage server to each original reporting server.
- From the remote storage server to a single restore point server.

You configure auto archiving to take place on a regular schedule. You invoke restore on an as needed basis.

More information:

Configuring Non-Interactive Authentication for Restore (see page 160)
Restore Auto-Archived Files (see page 167)

Data Integrity Checks

You can check archived or recataloged data for tampering if you are logged in as a user with Administrator rights. Checking allows you to secure your archived data, and meet regulatory requirements. CA Enterprise Log Manager uses digital signatures to validate the databases. If the database is corrupted or if its signature is missing or corrupted, the data integrity check considers the database tampered.

You can set up data integrity checks in the following ways:

- Automatically when you restore and recatalog data
- At scheduled times on selected servers
- On demand whenever you want

You can view the results of any of these checks from the Data Integrity interface. Any tampered databases are quarantined, and appear in the Quarantined Databases list.

More information:

Enable Automatic Integrity Check (see page 156)
Schedule a Data Integrity Check (see page 157)
Check Data Integrity on Demand (see page 157)
Sign Quarantined Databases (see page 158)

Enable Automatic Integrity Check

You can set a data integrity check to occur automatically whenever you restore or recatalog data.

To enable an automatic data integrity check

1. Click the Administration tab, and expand the Services subtab.
2. Select the Event Log Store node to enable global automatic checks, or expand the Event Log Store node and select a CA Enterprise Log Manager server to enable a local automatic check.
3. Select Validate Integrity on Recatalog and Restore.

Schedule a Data Integrity Check

You can schedule daily data integrity checks to occur at set times and on selected CA Enterprise Log Manager servers. Any tampered databases detected by a scheduled integrity check are automatically quarantined.

To schedule a data integrity check

1. Click the Administration tab, and expand the Services subtab.
2. Select the Event Log Store node to schedule a global check, or expand the Event Log Store node and select a CA Enterprise Log Manager server to set a local check.
3. Select Enabled.
4. (Optional) Select Federated to run the scheduled check on any federated servers visible from the selected server.
5. Set the Daily Start Time.

Check Data Integrity on Demand

You can run a data integrity check at any time on a selected CA Enterprise Log Manager server.

To run a data integrity check on demand

1. Click the Administration tab, and click the Archival Management subtab.
2. Expand the Data Integrity Folder, and select the CA Enterprise Log Manager server where you want to run a check.
3. Select a time range for your check.
4. (Optional) Select Federated to run the check on any federated servers visible to the selected server.
5. Click Validate Now.

   A list of the checked databases appears. Any tampered databases detected are displayed with a red icon. They also appear in the Quarantined Databases list.

Sign Quarantined Databases

You can regenerate the digital signature on a quarantined database, making it available for queries.

To regenerate a quarantined database signature

1. Click the Administration tab, and click the Archival Management subtab.
2. Expand the Data Integrity Folder, and select the CA Enterprise Log Manager server where the quarantined databases are located.
3. Click the Quarantined Databases tab in the right pane.
4. Select the database for which you want to regenerate a signature.
5. Click Generate Signature.

   A confirmation message appears

Rotate Keys

You can rotate the registration keys used to secure archived databases for improved security.

Registration keys use a public/private encryption key combination to secure the database files. When you rotate keys, the the old public key is retained, so that CA Enterprise Log Manager can verify files which used the old private key.

To rotate registration keys

1. Click the Administration tab, and expand the Services subtab.
2. Select the Data Integrity Folder, and click Rotate Keys

   A confirmation message appears.

Import Keys

You can import public registration keys used to secure archived databases from an outside source. Importing allows you to retain keys used by other older CA Enterprise Log Manager servers, or by previous servers. For example, if you maintain a backup list of public keys, you could import to be able to verify old database signatures on a newly built CA Enterprise Log Manager server.

To import registration keys

1. Click the Administration tab, and expand the Services subtab.
2. Select the Data Integrity Folder, and click Import Keys.

   An import file dialog appears.
3. Browse for the XML key file you want to import and click OK.

   A confirmation message appears.

   Note: You can only import keys in XML format.

Export Keys

You can export the public registration keys used to secure archived databases. Exporting allows you to back up the keys for later import to other CA Enterprise Log Manager servers.

To export registration keys

1. Click the Administration tab, and expand the Services subtab.
2. Select the Data Integrity Folder, and click Export Keys.

   An export dialog appears.
3. Select the location you want and click Save.

Configuring Non-Interactive Authentication for Restore

After you configure non-interactive ssh authentication between the remote storage server and the destination server, you can use the restore-ca-elm shell script to restore the archived databases on demand. For restore, the remote storage server is the source and the reporting CA Enterprise Log Manager or restore point CA Enterprise Log Manager is the destination.

The processes are slightly different, depending on whether the destination is a reporting server or a dedicated restore point.

- If you use a dedicated restore point, you set up non-interactive authentication once and then use it for every restore. The procedure sets up the .ssh directory on the restore point with the required ownership and sets permissions on the key file.
- If you restore archived databases from the remote storage server to multiple reporting servers, you set up non-interactive authentication between each server pair. You create the key pair once, but you copy the same public key of the key pair to each destination reporting server. For example, copy the public key as authorized_keys_RSS from the remote storage server to each reporting server. On each reporting server, you concatenate the authorized_keys_RSS file to the existing authorized_keys file. The existing file contains the public keys copied from each collection server.

Both processes assume that you previously prepared the remote storage server to act as the destination server for auto archive, which requires non-interactive authentication. If preparation has not been done, see "Create a Directory Structure with Ownerships on the Remote Storage Server" in the *Implementation Guide* for guidance.

More information:

Example: Configure Authentication From Remote Storage to a Restore Point (see page 161)
Example: Configure Authentication From a Storage Server to a Reporting Server (see page 163)

Example: Configure Authentication From Remote Storage to a Restore Point

Dedicating a CA Enterprise Log Manager server to use as a restore point makes setting up non-interactive authentication easy. Once you set up authentication between the remote store server and the restore point, you can use the restore-ca-elm.sh script for every restore without any additional steps for authentication.

The process for configuring non-interactive authentication from a storage server to a restore point CA Enterprise Log Manager involves the following procedures:

1. From the remote storage server, generate the RSA public/private key pair. Copy the public key as authorized_keys to the /tmp directory on the restore point.
2. From the restore point, create the .ssh directory in /opt/CA/LogManager and set ownership to caelmservice. Copy authorized_keys from the /tmp directory to the .ssh directory. Change ownership and set permissions on authorized_keys.
3. Validate successful non-interactive authentication between the remote storage server and the restore point.

More information:

Generate Keys and Copy the Public Key to the Restore Point (see page 161)
Prepare the Public Key File for Use (see page 162)

Generate Keys and Copy the Public Key to the Restore Point

From the remote storage server, generate an RSA key pair as the caelmservice user. Then, copy the public key file id_rsa.pub as authorized_keys, to the /tmp directory on the restore point CA Enterprise Log Manager. A restore point is a server dedicated to investigating restored data.

It is assumed that the /opt/CA/LogManager/.ssh directory structure exists on the storage server with the ownership set to caelmservice user and group. It contains authorized_keys copied from reporting servers. When you generate the key pair, you save id_rsa.pub to the /opt/CA/LogManager/ssh directory.

To generate the RSA public/private key pair for remote storage to restore point server authentication

1. Log on to the remote server used for storage through ssh as the caelmadmin user.
2. Switch users to the root account.

   su –
3. Switch users to the caelmservice account.

   su – caelmservice
4. Generate an RSA key pair as the caelmservice user.

   ssh-keygen –t rsa

5. Press Enter to accept the default when each of the following prompts appears:
   - Enter file in which to save the key (/opt/CA/LogManager/.ssh/id_rsa):
   - Enter passphrase (empty for no passphrase):
   - Enter same passphrase again:
6. Change directories to /opt/CA/LogManager.
7. Change the permissions of the .ssh directory using the following command:

   chmod 755 .ssh

8. Navigate to .ssh, where id_rsa.pub key is saved.

   cd .ssh

9. Copy the public key as authorized_keys to the /tmp directory on the restore point server.

   scp id_rsa.pub caelmadmin@<restore_point>:/tmp/authorized_keys

Prepare the Public Key File for Use

You create the .ssh directory on the restore point server and set ownership to caelmservice. Then, you copy authorized_keys from the /tmp directory to the .ssh directory. Last, you set ownership and permissions on the public key file.

To prepare the public key on the restore point server for non-interactive authentication

1. Log into the restore point CA Enterprise Log Manager server through ssh as caelmadmin.
2. Switch users to root.
3. Change directories to the CA Enterprise Log Manager directory.

   cd /opt/CA/LogManager

4. Create the .ssh directory:

   mkdir .ssh

5. Change the ownership of .ssh to the caelmservice user and group:

   chown caelmservice:caelmservice .ssh

6. Change directories to /opt/CA/LogManager/.ssh.
7. Copy the authorized_keys file from /tmp to .ssh:

   cp /tmp/authorized_keys .

8. Change ownership of the authorized_keys file to caelmservice:

   chown caelmservice:caelmservice authorized_keys

9. Change permissions on the authorized_keys file:

   chmod 755 authorized_keys

Example: Configure Authentication From a Storage Server to a Reporting Server

You can restore archived databases from a remote storage server back to their original reporting server, that is, the server from which they were auto archived. The advantage of this method is that you do not have to recatalog the CA Enterprise Log Manager archive database. The databases of log files you are restoring are already known to the reporting server. If you have multiple reporting servers, you configure non-interactive authentication between the remote storage server and each reporting server. The authorized_keys file exists in the .ssh directory on the reporting server. This authorized_keys file has the public keys of each key pair generated on a collection server that auto archives to this reporting server. Therefore, you create an authorized keys file with a suffix and then concatenate that file to the original authorized_keys.

The process for configuring non-interactive authentication from a remote storage server to a reporting CA Enterprise Log Manager involves the following procedures:

1. From the remote storage server:
   a. Configure the RSA public/private key pair for remote storage to reporting server authentication.
   b. Copy the public key as authorized_keys_RSS from the storage server to the /tmp directory on the reporting server.
2. From the reporting server:
   a. Copy the current authorized_keys from .ssh to /tmp.
   b. Concatenate authorized_keys_RSS in the /tmp directory to the authorized_keys file.
   c. Copy the appended authorized_keys file back to the .ssh directory.
3. From the remote storage server, validate successful non-interactive authentication between servers.
4. Repeat these steps for each remote storage server to reporting server combination.

More information:

Generate Keys and Copy the Public Key to a Reporting Server (see page 164)
Update the Existing Public Key File (see page 165)

Generate Keys and Copy the Public Key to a Reporting Server

From the remote storage server, generate an RSA key pair as the caelmservice user and then copy the public key as authorized_keys_RSS to the /tmp directory on a reporting CA Enterprise Log Manager server. The reporting server typically has an authorized_keys file in the .ssh directory that contains a concatenation of public keys from various collection servers. Send the key with a unique name so that it can be appended to the existing authorized_keys file.

To generate the RSA public/private key pair and copy the public key from the remote storage to a reporting server

1. Log on to the remote storage server through ssh as the caelmadmin user.
2. Switch users to root.
3. Switch users to the caelmservice account.

   su – caelmservice
4. Generate an RSA key pair as the caelmservice user.

   ssh-keygen –t rsa
5. Press Enter to accept the default when each of the following prompts appears:
   - Enter file in which to save the key (/opt/CA/LogManager/.ssh/id_rsa):
   - Enter passphrase (empty for no passphrase):
   - Enter same passphrase again:
6. Change the permissions of the .ssh directory using the following command:

   chmod 755 .ssh
7. Navigate to the .ssh directory.
8. Copy id_rsa.pub as authorized_keys_RSS to the /tmp directory on the reporting server.

   scp id_rsa.pub caelmadmin@<reporting_server>:/tmp/authorized_keys_RSS

Update the Existing Public Key File

You copied the public key, authorized_keys_RSS, to the /tmp directory on the reporting server. Now you prepare the existing public key file for use. Preparation involves appending the authorized_keys_RSS to authorized_keys. The correct ownership and permissions are already set on the existing authorized_keys file.

To append authorized_keys_RSS to authorized_keys and copy it to the correct location

1. Log into the reporting CA Enterprise Log Manager server through ssh as caelmadmin.
2. Switch users to root.
3. Change directories to the /tmp directory containing authorized_keys_RSS.
4. Copy the existing authorized_keys from .ssh to the current directory, /tmp.

   ```
   cp /opt/CA/LogManager/.ssh/authorized_keys .
   ```

5. Add the contents of the public key from the remote storage server to the authorized_keys file that contains public keys from collection servers.

   ```
   cat authorized_keys_RSS >> authorized_keys
   ```

6. Change directories to /opt/CA/LogManager/.ssh.
7. Copy the authorized_keys file from /tmp to.ssh, the current directory:

   ```
   cp /tmp/authorized_keys .
   ```

Query the Archive Catalog

You can create queries to search the local archive catalog for cold (remote-stored) databases, using quick or advanced filters. The query results can help you identify the backed up database files that must be restored to conduct an investigation.

To query the archive catalog

1. Click the Administration tab, and then click the Archival Management subtab.

   The Archival Explorer folder list appears.
2. Click the Archival Query and Recatalog folder.

   The Archive Query dialog appears in the details pane.
3. Select or enter the time period for your query.

4. Click Add Filter, select a column, and enter the column search value. You can add multiple filters.
5. Select Exclude to query for all logs *except* those logs with the value you enter.

   Note: If you create a filter specifying a column that is *not* in the catalog, CA Enterprise Log Manager returns all the databases in the specified time range, rather than an empty set. You are not required to know all the cataloged columns to create a useful archive query.
6. (Optional) Click the Advanced Filters tab to add advanced filters. Include the event information if the column bears the appropriate relation to the value you enter. Select a column, select an Operator, and then select or enter a value. Operator descriptions follow:

   Relational Operators

   Equal to, Not Equal to, Less than, Greater than, Less than or equal to, Greater than or equal to.

   Like

   Includes the event information if the column contains a pattern matching your entry of text with the wildcard character, %. L% includes values beginning with L. %L% includes values containing L, excluding an L that is either the first or last character.

   Not like

   Includes the event information if the column does not contain the pattern you specify.

   In set

   Includes the event information if the column contains one or more of the values in the quote-delineated set you enter. Multiple values in the set must be separated with commas.

   Not in set

   Includes the event information if the column contains one or more of the values in the quote-delineated set you enter. Multiple values in the set must be separated with commas.

   Matches

   Includes any event information that matches one or more of the characters that you enter, allowing you to search for key words.

Keyed

Includes any event information that is set as a key value during Report Server configuration. Use key values to set business relevance or other organizational groups.

Not Keyed

Includes any event information that is not set as a key value during Report Server configuration. Use key values to set business relevance or other organizational groups.

7. Click Query.

   The query results appear. The files containing records matching your query criteria are displayed with the full relative path, relative to $IGW_LOC. Examples follow:

   ../../LogManager/data/archive/<databaseFilename>

   <remoteHostname>-.../../LogManager/data/archive/<databaseFilename>

Restore Auto-Archived Files

If you copy archived files from external storage to the remote server configured for auto-archiving, you can restore them with the restore-ca-elm.sh script. This alternative is preferred over the manual use of LMArchive utility.

To restore auto-archived files

1. Use your caelmadmin credentials to log on to the CA Enterprise Log Manager server with the event log store where you want to restore the databases.
2. At the command prompt, switch users to root, that is:

   su - root

3. Change directories to /opt/CA/SharedComponents/iTechnology with the following shortcut:

   cd $IGW_LOC

4. At the command prompt, switch users to the caelmservice account.

   su - caelmservice

5. Run the following command, where *userid* and *pwd* are the credentials of a CA Enterprise Log Manager user account with the Administrator role.

   restore-ca-elm.sh -euser *userid* -epasswd *pwd* -rhost *hostname* -ruser *userid* -rlocation *path* -files file1,file2,file3...

   Note: To allow a non-Administrator to run the restore-ca-elm shell script, create a custom role and custom policy. Then, users to whom you assign this custom role can specify their credentials for userid and pwd.

More information:

Restore–Script for Restoring Archived Databases (see page 168)
Example: Allow a Non-Administrator to Manage Archives (see page 95)
Automating Backup and Restore (see page 155)

Restore–Script for Restoring Archived Databases

You cannot query or report on data that resides in a cold database on a remote storage server. To query and report on such data, that data must reside in a warm state on a CA Enterprise Log Manager server. The restore shell script, restore-ca-elm.sh, is a command-line utility that moves a specified cold database and its digital signature to a specified CA Enterprise Log Manager server and restores it to a warm state. You can use the restore utility to move a database back to the original reporting server or to a dedicated restore point. Configuring non-interactive authentication is a prerequisite to running the restore script.

You run the restore script from the CA Enterprise Log Manager server to which you want to restore the files. The remote host you identify in the command refers to the remote storage server. Cold databases reside in the archive directory of the remote storage server.`

Requirements for restoring database files to either the original reporting server or a restore point server follow:

- The RSA key file ownership has been set on the remote server.
- Permission to the /opt/CA/LogManager folder has been granted to caelmservice on the remote server.

If restoring files to a restore point server, take the following additional actions:

1. Copy the RSA key from the remote storage server to the restore point server
2. Set the RSA key file ownership on the restore point server.

The command has the following format:

restore-ca-elm.sh -euser *userid* -epasswd *pwd* -rhost *hostname* -ruser *userid* -rlocation *path* -files *file1,file2,file3...*

-euser *username*

Specifies the user name of a CA Enterprise Log Manager user account with the Administrator role.

-epasswd *pwd*

Specifies the CA Enterprise Log Manager password associated with the user name.

-rhost *host*

Specifies the hostname or IP address of the remote host where cold database files reside in the archive directory. The remote host is not a CA Enterprise Log Manager server.

-ruser remote *user*

Specifies the user account with permissions to the /opt/CA/LogManager path and ownership of the .ssh folder containing the authorized keys file. Typically, this user account is the caelmservice user account.

-rlocation *path*

Specifies the path to the database files on the remote storage server. If the remote storage server is a UNIX server, the path is /opt/CA/LogManager/data/archive.

files *file1,file2,file3...*

Specifies a comma-separated list, without spaces, of the database files to restore.

Example: Restore Shell Script

The following example command is run from the CA Enterprise Log Manager to which the archived database files are to be restored. It is run by a user with account credentials of Administrator1, calm_r12. The remote server to which the archived databases have been moved from off-site storage is named NY-Storage-Svr. This remote server has been configured with a caelmservice account that has ownership of the .ssh folder where the RSA public key has been copied. This account also has full privileges on the directory structure /opt/CA/LogManager. This command specifies that the files to be restored are in the data/archive directory path of the NY-Storage-Svr server and identifies the database file to be restored as NY-Storage-Svr_20081206192014.db.cerod.

restore-ca-elm.sh -euser Administrator1 -epasswd calm_r12 -rhost NY-Storage-Svr -ruser caelmservice -rlocation /opt/CA/LogManager/data/archive -files NY-Storage-Svr_20081206192014.db.cerod

More information:

Restore Auto-Archived Files (see page 167)
Configuring Non-Interactive Authentication for Restore (see page 160)
Example: Configure Authentication From Remote Storage to a Restore Point (see page 161)
Example: Configure Authentication From a Storage Server to a Reporting Server (see page 163)

Manually Backing Up Archived Databases

CA Enterprise Log Manager creates a new archived database automatically each time data is moved from hot to warm storage, as controlled by your settings. Although it is recommended that you configure auto-archive to move warm databases to a remote server, you can use your own tools to perform backups of the archive databases and then run the LMArchive utility to notify the system that you performed the backup.

We recommend that you back up your warm databases daily, using either the automated method or the manual method described here. This is important since warm-storage archive files are automatically deleted after the time you specify or when disk space drops to the percentage you specify.

Backing up warm databases manually involves the following steps:

1. Identify the warm databases that are not yet backed up.
2. Perform the backups.
3. Record the backups.

More information

Identify Databases Not Backed Up (see page 170)
Perform the Backups (see page 172)
Record the Backups (see page 172)

Identify Databases Not Backed Up

You can view a list of archived databases that are not yet marked as backed up with the LMArchive utility. The ability to get reliable results assumes that someone has run this utility with the *-notify arch* option each time archived database are backed up.

Important! To avoid confusion, keep current with notifying CA Enterprise Log Manager of completed backups.

To display the names of all current archived database files not marked as backed up

1. Use your **caelmadmin** credentials to log on to the CA Enterprise Log Manager server with the event log store that contains the databases that need to be backed up for archiving.
2. At the command prompt, switch users to root, that is:

   su - root
3. Change directories to /opt/CA/SharedComponents/iTechnology with the following shortcut:

   cd $IGW_LOC
4. Execute the following command, where *username* and *pwd* are the credentials of a CA Enterprise Log Manager user account with the Administrator role.

   LMArchive -euser *username* -epassword *pwd* –list inc

Example: Display All Current CA Enterprise Log Manager Archived Files Not Marked as Backed Up

The following command issued by an Administrator requests the list of all warm databases that are not marked as backed up.

```
LMArchive -euser Administrator1 -epassword calmr12 -list inc
```

A list of archive files not marked as backed up appear, in a format similar to the following:

```
CAELM Archived Files (not backed up):
    calm04_20091206192014.db.cerod
    calm04_20091206192014.db.sig
```

Perform the Backups

If you have not configured auto-archive from a CA Enterprise Log Manager reporting server to a remote storage server that is not a CA Enterprise Log Manager server, you must manually back up the archived databases and move them to a safe storage location, such as a separate disk or server.

Important! Be sure to back up the databases and move them before they are *deleted* from the CA Enterprise Log Manager reporting server.

Warm databases are automatically deleted when the configured value for Max Archive Days is reached or the percentage of disk space falls below the configured value for Archive Disk Space. To prevent data loss from deleted files, perform the backups regularly.

To back up warm databases manually

1. Use your caelmadmin credentials to log on to the CA Enterprise Log Manager reporting server with the event log store that contains the target databases.
2. Switch users to root, that is:

   ```
   su - root
   ```

3. Change directories to /opt/CA/LogManager/data/archive.
4. Back up the warm databases with the backup tool of your choice and move them to an on-site interim storage server or to an off-site location for long-term storage, according to your site procedures.

Record the Backups

Each time you perform a backup of one or more archived database, be sure to record that fact in the CA Enterprise Log Manager where you took the backup.

Note: Failure to record each backup causes incorrect data to be reported when using the LMArchive utility to list backed up databases.

To record the backups of specific archived databases

1. Use your caelmadmin credentials to log on to the CA Enterprise Log Manager server with the event log store that contains the databases that you backed up.
2. At the command prompt, switch users to root, that is:

   su - root
3. Change directories to /opt/CA/SharedComponents/iTechnology with the following shortcut:

   cd $IGW_LOC
4. Execute the following command, where *username* and *pwd* are the credentials of a CA Enterprise Log Manager user account with the Administrator role.

   LMArchive -euser *username* -epassword *pwd* –notify arch –files *file1*,*file2*,*file3*...

Example: Notify CA Enterprise Log Manager that Certain Files Have Been Backed Up

The following command issued by the Administrator named Administrator1notifies the CA Enterprise Log Manager event log store that the warm database, calm04_20091206192014.db.cerod, has been backed up. Backed up databases can be manually moved to external storage for long-term retention.

LMArchive -euser Administrator1 -epassword calmr12 -notify arch -files calm04_20091206192014.db.cerod

The archive file notification appears, in a format similar to the following:

Archive notification sent for file calm04_20091206192014.db.cerod...

Manually Restoring Archives to the Original Event Log Store

You may occasionally need to restore cold database files for querying or reporting to the archive directory on a CA Enterprise Log Manager server. This may be required to investigate a security breach or for an annual or semi-annual compliance audit. The procedures you use depend on the following two factors:

- Whether you used auto-archive to back up the files you now want to restore
- Whether you are restoring the files to the original reporting server or to a different server, such as a dedicated restore point server

  If restoring to a different server, see "Restoring Archives to a New Event Log Store."

When you are restoring files to the original reporting server, use the following procedures:

1. Prepare to restore archived databases by identifying the files to restore and determining the archive directory.
2. Move the databases from external storage to the archive directory on either the remote server location you configured for auto-archive or to the original reporting server.
3. If you moved the archived files to the remote storage server configured for auto-archiving, log on to the reporting CA Enterprise Log Manager and restore the auto-archived files from the remote storage server with the restore-ca-elm.sh script.
4. If you moved the archived files to the archive directory on their original reporting CA Enterprise Log Manager server, restore the manually archived files with LMArchive utility.
5. Verify that the defrosted database can be queried by running a query with the end date set to the date of the restored database and examining the query results.

More information

Prepare to Restore Archived Databases (see page 175)
Move Archived Databases to an Archive Directory (see page 177)
Restore Auto-Archived Files (see page 167)
Restore Manually Archived Files (see page 178)
Verify Restoration (see page 179)

Prepare to Restore Archived Databases

Before you restore archived databases, you need to know the following:

- The names of the files to restore
- The archive directory path into which you copy the files you retrieve from off-site storage. The path is always /opt/CA/LogManager/data/archive

You can query the archive catalog through the CA Enterprise Log Manager Administration tab, Log Collection Explorer, where you can specify simple or advanced filters. Or you can use the command line utility as described here.

If you already have the needed information at hand, skip this procedure.

To prepare to restore archived databases

1. Use your caelmadmin credentials to log on to the CA Enterprise Log Manager server with the event log store where you want to restore the databases.
2. At the command prompt, switch users to root, that is:

   ```
   su - root
   ```

3. Change directories to /opt/CA/SharedComponents/iTechnology with the following shortcut:

   ```
   cd $IGW_LOC
   ```

4. Identify the databases that you want to restore from a list of files that includes those that have been backed up and moved to external storage. To view the list of all archived files in this archive catalog, execute the following command, where userid and pwd are the credentials of a CA Enterprise Log Manager user account with the Administrator role.

   ```
   LMArchive -euser userid -epassword pwd -list all
   ```

   The list of all archived files appears.

5. (Optional) If restoring from manual backups, determine the location of the archive directory to which the identified cold archive files are to be copied. Execute the following command, where userid and pwd are the credentials of a CA Enterprise Log Manager user account with the Administrator role.

   ```
   LMArchive -euser userid -epassword pwd -list loc
   ```

   The archive directory appears.

Example: Display All Current CA Enterprise Log Manager Archive Files

The following command issued by the CA Enterprise Log Manager Administrator, Administrator1, requests a list of all databases located in the event log store's archive directory.

```
LMArchive -euser Administrator1 -epassword calmr12 -list all
```

A list of the current archive files appears, in a format similar to the following:

```
CAELM Archived Files:
    calm04_20091206191941.db.cerod
    calm04_20091206191958.db.cerod
    calm04_20091206192014.db.cerod
    calm04_20091206191941.db.sig
    calm04_20091206191958.db.sig
    calm04_20091206192014.db.sig
```

Example: Display the CA Enterprise Log Manager Archive Directory

The following command issued by a CA Enterprise Log Manager Administrator, Administrator1, requests the directory location of the archived databases:

```
LMArchive -euser Administrator1 -epassword calmr12 -list loc
```

The following is a typical response:

```
CAELM Archive Location (localhost) :
../../LogManager/data/archive
```

More information:

Query the Archive Catalog (see page 165)

Move Archived Databases to an Archive Directory

If you moved your archived files to an off-site location, use your site procedures to retrieve them and bring them back on-site.

Move the archived databases back to the archive directory of either the original CA Enterprise Log Manager server or a remote server configured for non-interactive authentication. The archive directory is /opt/ca/LogManager/data/archive.

To move an archived database from an external storage to your network

1. Move the database files to restore from external storage back to your network in one of the following ways:
   - If you use auto-archive to automatically move your archived files to the remote server, then copy them back to the archive directory on this remote server. (This remote server is already configured for non-interactive authentication with the CA Enterprise Log Manager server to which the archived databases are to be restored.)
   - If you do not use auto-archive, copy your archived files back to the archive directory on their original CA Enterprise Log Manager server.
2. Proceed in one of the following ways, depending on the location of the archived files.
   - If the archived files are on the remote server configured for auto-archiving, restore auto-archived files with the restore-ca-elm.sh script.
   - If the archived files are in the archive directory on their original CA Enterprise Log Manager server, notify CA Enterprise Log Manager that the archived files have been restored with LMArchive utility. Upon notification, the restored files are put into a defrosted state.

More information:

Restore Auto-Archived Files (see page 167)
Restore Manually Archived Files (see page 178)

Restore Manually Archived Files

After you have restored one or more databases from long-term storage to the archive directory, you must change ownership of the archive directory to the caelmservice user before notifying CA Enterprise Log Manager that the databases have been restored with the LMArchive utility. Archived files owned by root are not recognized by the LMArchive utility.

Running LMArchive with the -notify rest option changes the state of the archived database files from cold to defrosted so that they are available for querying and reporting.

Administrators configure the number of hours a defrosted archived database is retained before being automatically deleted from the archived directory with the Export Policy setting on the Event Log Store service configuration.

To restore manually archived database files

1. Use your **caelmadmin** credentials to log on to the CA Enterprise Log Manager server with the event log store that contains the restored databases.
2. At the command prompt, switch users to root, that is:

   su - root
3. Change directories to the /data directory. For example:

   cd /opt/CA/LogManager/data
4. Assign ownership of the archive directory (/opt/CA/LogManager/data/archive) to the caelmservice account.

   chown -R caelmservice:caelmservice archive

   Ownership of the archive files is changed to caelmservice, the internal operating system user, which is a non-login account.
5. Change directories to /opt/CA/SharedComponents/iTechnology with the following shortcut:

   cd $IGW_LOC
6. Execute the following command, where *username* and *pwd* are the credentials of a CA Enterprise Log Manager user account with the Administrator role.

   LMArchive -euser *username* -epassword *pwd* –notify rest –files *file1*,*file2*,*file3*

   A restoration confirmation appears. CA Enterprise Log Manager defrosts the specified files. Defrosted files are retained for the configured number of hours, where you can configure retention for up to seven days.

Note: You can now run queries and reports against the event data contained in the restored archive files.

Example: Notify CA Enterprise Log Manager that Certain Databases Have Been Restored

The following command issued by a CA Enterprise Log Manager user with an Administrator role notifies the CA Enterprise Log Manager event log store that the specified cold database, calm04_20091206192014.db, has been copied into the archive directory.

```
LMArchive -euser Administrator1 -epassword calmr12 -notify rest -files calm04_20091206192014.db.cerod
```

A restoration confirmation appears in a format similar to the following:

```
Restore notification sent for file calm04_20091206192014.db.cerod
```

More information:

LMArchive–Backup/Restore Tracking (see page 183)

Verify Restoration

You can quickly verify that the restored database is available for examination by running a quick query. Normal queries display data from restored databases in warm and defrosted states.

Consider the following process.

1. Copy a subscription query designed to display the type of event details that the restored database holds.
2. Advance to the query design wizard step where you set result conditions and enter a date range that corresponds to the newly defrosted database files.
3. Save the query.
4. Run the query.

Manually Restoring Archives to a New Event Log Store

You may occasionally need to restore cold stored files for querying or reporting, as for an annual or semi-annual compliance audit. If you designate one CA Enterprise Log Manager to act as a restore point for investigations on cold data, you must force a rebuilding of the catalog each time you restore a new database to this CA Enterprise Log Manager. A rebuilding of the catalog, or recatalog, is required only when restoring data to a different server than the one on which it was generated.

Important! Ensure the Max Archive Days setting for this server's event log store is adequate. Otherwise, restored files are immediately deleted.

A recatalog is automatically performed when iGateway is restarted, if needed. If databases were incompletely cataloged before iGateway was shut down, the recataloging process completes when iGateway is restarted. If one or more databases are added to the archive database directory while iGateway is down, the recatalog process is performed at the next startup of iGateway.

Restoring archived files from external storage to a different CA Enterprise Log Manager from where they were backed up involves the following steps:

1. Identifying the databases that you want to restore. For help, query the archive catalog with filters.
2. Moving the identified cold archive files from external storage to your network.
3. Copying the moved databases to the archive directory. To display the archive directory, run the LMArchive utility with the -list loc option.
4. Rebuilding the archive catalog (recatalog).

   Rebuilding the archive catalog to add a single database can take several hours. After waiting long enough for the recatalog process to complete, you can begin your investigation by running queries and reports on the event logs from the restored databases.
5. Verify restoration by issuing a query.

Note: If you dedicate a CA Enterprise Log Manager as restore point, be sure to exclude it from the federation.

More information:

Move Archived Databases to an Archive Directory (see page 177)
Configure Max Archive Days for Restored Archives (see page 181)
Add Restored Databases to the Catalog (see page 182)
Verify Restoration (see page 179)
Example: Allow a Non-Administrator to Manage Archives (see page 95)

Configure Max Archive Days for Restored Archives

When you configure the Event Log Store for a CA Enterprise Log Manager dedicated as a restore point, we recommend you override the global setting for Max Archive Days and set it to the maximum, 28000. If the number of days to store archived database files before deleting it is set to a lower value than the age of the restored database files, those files are deleted by the system immediately after being restored to a warm state.

Note: This procedure applies only to files restored to a new event log store.

To set the max archive days to accommodate the age of restored files

1. Click the Administration tab and the Services subtab.
2. In the Service List, expand the Event Log Store folder and select the CA Enterprise Log Manager that is the dedicated restore point.
3. Click the toggle button next to Max Archive Days to switch to local configuration and enable the entry field.
4. Set the field to a value in days that accommodates the oldest file you will restore. The maximum value is 28000.
5. Click Save.

Add Restored Databases to the Catalog

If you copy the restored database directly into the archive directory on a different server than the one where it was generated, rebuild the archive catalog to add the restored database.

Do *not* use recatalog in either of the following cases:

- If you use the restore-ca-elm.sh script to restore an archived database. The restore shell script does the recataloging for you.
- If you copy the restored database directly into the archive directory on the same server where it was generated and then notify CA Enterprise Log Manager that the database is restored with the LMArchive -notify rest option.

The recatalog process sets the restored database to a "warm" state, not a "defrosted" state as with the LMArchive -notify rest option. Therefore, it follows normal archiving rules, rather than the Export Policy set in the event log store configuration.

To rebuild the archive catalog to add the restored database

1. Click the Administration tab, and then the Archival Management subtab.

   The Archival Management folder list appears.

2. Click the Archival Query and Recatalog folder.

   Three buttons, including Recatalog, appear above tabs for Quick Filters and Advanced Filters.

3. Click Recatalog.

   A success confirmation message appears. The restored database is added to the catalog in a warm state.

More information:

Example: Allow a Non-Administrator to Manage Archives (see page 95)

LMArchive–Backup/Restore Tracking

The *LMArchive* is the command-line utility that tracks the backup and restoration of warm databases to the event log store on a CA Enterprise Log Manager server. Use LMArchive to query for the list of warm database files that are ready for archiving. After backing up the listed database and moving it to long-term (cold) storage, use LMArchive to create a record on the CA Enterprise Log Manager server that this database was backed up. After restoring a cold database to its original CA Enterprise Log Manager server, use LMArchive to notify CA Enterprise Log Manager, which in turn defrosts the cold database files to a defrosted state that can be queried.

The command has the following format:

LMArchive -euser *username* -epassword *pwd* {-list [loc|all|inc] | -notify [arch|rest] -files *file1*,*file2*,*file3*...}

-euser *username*

Specifies the user name of a CA Enterprise Log Manager user account with the Administrator role.

-epassword *pwd*

Specifies the CA Enterprise Log Manager password associated with the user name.

-list [loc | all | inc]

Requests a list of one of the following: the archive directory locations, the names of all warm and cold databases, the names of just the warm databases

loc

Requests the location of the archive directory.

all

Requests a list of all filenames located in the archive directory of the event log store.

inc

Requests an incremental list of filenames of the current warm database files that have not been archived. The request returns the names of files that have not been backed up, moved to external storage, and changed to a cold state. Files are set to a cold state upon notification of the move through this utility's notify command.

-notify [arch | rest]

Notifies the CA Enterprise Log Manager event log store that the specified files have been successfully backed up or restored.

arch

Notifies the CA Enterprise Log Manager event log store that the specified files have been successfully backed up.

rest

Notifies the CA Enterprise Log Manager event log store that the specified files have been successfully restored.

-files *file1,file2,file3...*

Specifies the names of the database files you have backed up or restored.

More information:

About Log Storage (see page 151)
Identify Databases Not Backed Up (see page 170)
Record the Backups (see page 172)
Prepare to Restore Archived Databases (see page 175)
Restore Manually Archived Files (see page 178)

Chapter 7: Subscription

This section contains the following topics:

Upgrading to CA Enterprise Log Manager Version 12.5 through Subscription

To upgrade CA Enterprise Log Manager to version 12.5, first upgrade to version 12.5 of the Log Manager product, then update all other CA Enterprise Log Manager modules, such as Content, Integration and Agent modules. You perform all upgrade tasks through Subscription.

Important! Upgrade the management CA Enterprise Log Manager server before you install any new CA Enterprise Log Manager servers in your network. Doing so allows the new servers to register properly.

To upgrade to CA Enterprise Log Manager version 12.5

1. Upgrade to Log Manager version 12.5.

   Note: The upgraded CA Enterprise Log Manager 12.5 user interface lists both Subscription Module and Subscription Service under the Administration tab, Services subtab. Subscription Module reflects the interface and functionality previous to the 12.5 update, and is present to help ensure proper communication between all CA Enterprise Log Manager servers during the upgrade to 12.5. Once you have upgraded the Log Manager product on a given CA Enterprise Log Manager server to version 12.5, use only Subscription Service to perform all further subscription tasks and configuration changes.

2. Update all other CA Enterprise Log Manager modules.

   Important! After you have performed Step 1, the upgraded CA Enterprise Log Manager 12.5 user interface lists both Subscription Module and Subscription Service. Use only Subscription Service, and not Subscription Module, to perform all further subscriptions tasks, including this step. Subscription Module is present only to help ensure proper communication between all CA Enterprise Log Manager servers during the upgrade to 12.5; do not use it to perform subscription functions post-upgrade.

3. If Agent or Connector modules were among the updates, install updated agents or connectors.

4. Reregister third-party and other CA products, like CA Access Control, that display CA Enterprise Log Manager reports in their native interfaces using the Open API calls.

   Completing this step updates the certificates that changed in this release. See the *CA Enterprise Log Manager API Programming Guide* for more information.

Note: For detailed information about performing these steps, see *Upgrading to CA Enterprise Log Manager Version 12.5 through Subscription* in the *CA Enterprise Log Manager 12.5 Release Notes*. Also, review the Release Notes for any Known Issues related to subscription upgrades.

More information:

Apply Subscription Updates to Agents and Connectors (see page 202)

View Global Subscription Status

The Subscription Service downloads and distributes selected subscription modules to your CA Enterprise Log Manager servers, either according to a schedule you configure or upon an on demand update. You can view the current subscription status of your CA Enterprise Log Manager servers through the Subscription Dashboard.

The Subscription Dashboard displays the progress of any updates a CA Enterprise Log Manager server is current downloading or installing. You can also see the state of any content updates currently in progress, and a list of all content updates previously installed.

To view the global status of subscriptions

1. Click the Administration tab and the Services subtab.
2. Click Subscription Service.

   The Global Service Configuration for the Subscription Service appears, with the Subscription Dashboard tab selected.

3. To view the status of content updates, examine the Content Subscription Status window.

   If a content update is currently in progress, a progress bar and status messages indicate the status of the update. Click the Refresh button to display the most recent progress of the update.

4. To view a list of all content updates installed to date, click Browse Catalog.

   The Content window appears. Click on a content type to display the name, date and version of all associated updates installed to date.

5. To view the subscription status of any CA Enterprise Log Manager server, examine the Servers window.

   If an update is in progress on any CA Enterprise Log Manager server, a progress bar appears in the Progress column for that server. Click the Refresh button to display the most recent progress of the update.

   The Servers window also displays information for each CA Enterprise Log Manager server, including the following:

   State

   Displays the subscription state of a given CA Enterprise Log Manager server. Subscription states are Idle, Waiting, Downloading, Installing, Completed, and Failed.

   Message

   Displays any message relating to the subscription status of a given CA Enterprise Log Manager server.

   Date

   Displays the date of the last update action taken by a given CA Enterprise Log Manager server.

View a Server's Subscription Status

The Subscription Service downloads and distributes selected subscription modules to your CA Enterprise Log Manager servers, either according to a schedule you configure or upon an on demand update. You can view the current subscription status of a specific CA Enterprise Log Manager server through the global Subscription Dashboard, or through the local State window of that server.

To view the subscription status of a CA Enterprise Log Manager server

1. Click the Administration tab and the Services subtab.
2. Click Subscription Service.

   The Global Service Configuration for the Subscription Service appears, with the Subscription Dashboard tab selected.
3. Do one of the following:

   - In the Servers window, examine the status information for the desired server. If an update is in progress, a progress bar appears in the Progress column for that server. Click the Refresh button to display the most recent progress. The Servers window also displays additional information for each CA Enterprise Log Manager server, including State, Message and Date.

   or

   - For more detailed information, in the Servers window, click the name of the desired server to open the local State window of that server. If an update is in progress, a progress bar and status messages appear in the Status window, indicating the progress of the update. Click the Refresh button to display the most recent progress. The Servers window also displays information for each CA Enterprise Log Manager server, including the following:

     State

     Displays the subscription state of a given CA Enterprise Log Manager server. Subscription states are Idle, Waiting, Downloading, Installing, Completed, and Failed.

     Message

     Displays any message relating to the subscription status of a given CA Enterprise Log Manager server.

     Date

     Displays the date of the last update action taken by a given CA Enterprise Log Manager server.

Edit the Global Subscription Configuration

During the implementation phase, you can configure global subscription settings for your environment. All CA Enterprise Log Manager servers inherit and use these global settings, unless you override a global setting by configuring local settings for an individual server.

By default, the first server installed is the Default Subscription Proxy and all servers installed after the first are configured as subscription clients. If no other proxy is configured or available, the default proxy downloads updates to subscription clients.

You can change the global settings at any time. All servers inherit any changes you make, unless a specific server is configured locally to override those settings.

Settings that can only be made and edited at the global level include:

- Default Subscription Proxy
- Public Key - Version
- Subscription Proxy(s) for Content Updates

To edit the global subscription configuration

1. Click the Administration tab and the Services subtab.
2. Click Subscription Service.

   The Global Service Configuration for the Subscription Service appears.
3. Click the Administration tab.
4. Examine the Global Services Configuration: Subscription Service settings in the right pane. Review the settings and make any required changes.

   Note: See online help for field-level details.
5. Click Save.

More information:

How to Apply Subscription Updates (see page 552)

Edit a Server's Local Subscription Configuration

During subscription implementation, you can configure global subscription settings, such as a subscription schedule and proxy list, for your environment. Individual CA Enterprise Log Manager servers inherit these global settings.

At any time, you can override global settings by configuring local settings for an individual server. Consider the subscription role of a server when planning to override global settings.

To edit a server's subscription configuration

1. Click the Administration tab and the Services subtab.
2. Expand Subscription Service and select the server to configure.

   The Subscription Service Configuration for the selected CA Enterprise Log Manager server appears.
3. Click the Administration tab.
4. If you want to change the subscription role of the server, do one of the following:
   - To designate it as a subscription client, clear the Subscription Proxy check box
   - To designate it as an online subscription proxy, select the Subscription Proxy check box, and select Online Subscription Proxy.
   - To designate it as an offline subscription proxy, select the Subscription Proxy check box, and select Offline Subscription Proxy.

   Note: Before changing a server's role from subscription proxy to subscription client, consider the configurations of the subscription clients that use this proxy. If you change a server from proxy to client, immediately remove the server from global and local Subscription Proxy(s) for Client Updates lists.

   Important! Before changing a server's role to offline subscription proxy, consider whether it is included in any proxy lists. In a mixed subscription environment, where you have configured both online and offline subscription proxies, do not include offline proxies in the proxy list for any online subscription client.
5. To override an inherited global setting, click the global/local button to switch to local service configuration for the selected field.

   Note: To restore the global setting, click the button again. The global setting is restored at the next update start time.
6. If the modules you want to download for this server are different from the settings inherited from the global settings, switch to local configuration, click Browse, and select the desired modules.

   Note: Be sure that the modules selected for a subscription proxy include at minimum all modules selected in the download lists of any clients that receive updates from this proxy.

7. If this server is to download subscription updates through an HTTP proxy server that is different from the inherited one, switch to local configuration and configure the desired HTTP proxy.
8. If this server is to download updates on a schedule different from the inherited one, switch to local configuration and edit the schedule.
9. If this server is to download updates from another CA Enterprise Log Manager subscription proxy switch to local configuration and add the desired CA Enterprise Log Manager proxy servers to the Subscription Proxy(s) for Client Updates list.

   This server contacts the specified proxy servers to download subscription updates, creating a tiered proxy structure.
10. Click Save.

Download & Select Modules for Offline Subscription

Offline subscription update files are available at the CA offline subscription FTP site, packaged in .zip files. As new modules become available, they appear on the FTP site. Periodically monitor the list of available modules to be sure that you have downloaded the most recent updates. You can also check the CA Support Site at http://ca.com/support to learn when new Log Manager service packs and versions are available.

Before you can select modules to download for offline subscription proxies, download the offline update file package from the CA FTP site and manually copy it to your offline proxies. You can then select which modules to download and install. Offline subscription clients, by contrast, automatically receive all updates that are manually installed on their offline proxy, regardless of any modules selected for the client at the local level.

To download and select offline subscription modules

1. On a system with Internet or FTP access, navigate to the FTP offline subscription site:

   ftp://ftp.ca.com/pub/elm/connectors/ftp/outgoing/pub/elm/ELM_Offline_Subscription

   The directory index displays a folder for each major CA Enterprise Log Manager release.
2. Download the appropriate .zip file for the update you want to perform.

   Note: The folder for the CA Enterprise Log Manager r12.5 release contains a subfolder as well as a .zip file. The subfolder contains modules for upgrading from any previous version to version r12.5. The .zip file contains the modules for performing routine, periodic updates to version r12.5. If you are using offline subscription to upgrade from a previous version to r12.5, see the Upgrading to CA Enterprise Log Manager topic in the *Release Notes*. If you are performing a routine update, select the .zip file.

3. Using physical media such as a disk, or using scp, manually copy the .zip file to the following file path on your offline proxies:

   /opt/CA/LogManager/data/subscription/offline.

4. Log in to a system in your CA Enterprise Log Manager environment.
5. Click the Administration tab and the Services subtab.
6. Expand Subscription Service and select an offline proxy server to configure.

   The Subscription Service Configuration for the selected CA Enterprise Log Manager server appears.

   Note: Offline subscription clients automatically receive all modules that are manually installed on their offline proxy. The contents of the proxy server control which updates the subscription client receives. Modules selected at the local level for an offline client have no effect.

7. Click the Administration tab.
8. In the File drop-down, select the offline update .zip file you copied to the server, and click Browse.

   The Modules Available for Download dialog appears.

9. Select the modules you want to download.
10. Click OK.

    The Modules Available for Download dialog closes, and the modules you selected appear in the Modules Selected for Download list.

11. Click Save.

    Offline subscription clients can now download these modules automatically according to the subscription schedule you set, or on demand when you begin a manual update.

12. (Optional) Click Update Now.

    The offline proxy server updates itself with the selected modules.

    Note: Though you can allow the offline proxy to update itself according to the subscription schedule you set, it is good practice to perform a manual update whenever you transfer new files. This practice ensures that the updates are available when offline subscription clients request them.

About On-Demand Updates

An on-demand update is different from a scheduled update in that it is performed immediately and updates only the selected server. You can invoke an on-demand update for only one CA Enterprise Log Manager server at a time. Before you perform an on demand update on a subscription client, first update its proxy server.

Typically, scheduled updates keep all of your servers updated with the latest binary files and your management server updated with the latest configuration and content files. However, it is sometimes appropriate to invoke an off-schedule update to a single server.

Consider using on-demand updates in the following cases:

- A subscription event failure or warning is reported for the management server, for example:

  Failed to connect to EEMServer

  Error installing content in EEM

  Select the management server and click Update Now. If the management server is configured as the proxy for content updates, the server downloads new updates from the CA Subscription Server. Then the server pushes the content updates to the content repository. If this server is configured as a proxy for clients, it makes available the binary updates to its clients on a first-come-first-serve basis.

  Note: Alternatively, you can wait until the next scheduled update session, when any incomplete processing resumes and completes.

- A subscription event failure message indicates that download is disrupted. If download is disrupted on only one proxy server, performing an on-demand update attempts the download again. An on-demand update is advisable if you notice the failure shortly after it occurs. If the start time between proxy and client is sufficient, an on-demand update of the proxy can complete before the scheduled update begins by clients who get updates from that proxy.
- You install a new CA Enterprise Log Manager, configure it as a proxy, and want to help ensure that the latest updates are applied before using it.
- You notice that a module needed by a client was not selected as a module to download. However, that module was selected as a module to download by the proxy. Running Update Now on the client installs the missing updates.
- You have downloaded an offline subscription package and copied it to an offline subscription proxy. Performing a manual update on the offline proxy whenever you transfer new files helps ensure that the updates are available when offline subscription clients request them.

Start an On-Demand Update

You can update servers on demand using the Update Now function. When you update two or more servers in a series, be sure to verify that processing completes on one before beginning processing on the next. Review self-monitoring events for confirmation.

If you recently changed any subscription configurations, be sure to wait for the update interval (300 seconds by default) to elapse before running an on-demand update. CA Enterprise Log Manager generates a self-monitoring event after the update completes.

Note: If a scheduled update is currently in progress, clicking Update Now has no effect. If the start time for a scheduled update occurs while update now processing is in progress, the scheduled process does not run. When the on-demand processing completes, the scheduled update cycle resumes.

Important! If the modules to download contain Content updates, update the content proxy before performing any other on-demand update.

To update a server on demand

1. Click the Administration tab and then click the Services subtab.
2. Expand Subscription Service in the Service List
3. Select the server that is the proxy for content updates and examine the modules selected for download. If they include Content updates, click Update Now.

   The server retrieves updates from the CA Subscription Server. As proxy for content updates, it pushes content updates to the content repository. As an online proxy, it downloads binary updates as well.
4. Select the server to update and examine its role.
5. If it is a subscription proxy for client updates, click Update Now.

   The server is updated with the modules selected for download.
6. If it is a client with an online proxy, do the following:
   a. Identify a proxy for this client from the selected items on the Subscription Proxy for Client list.
   b. Select the proxy for the client and click Update Now.
   c. Select the client and click Update Now.

   The server is updated with the modules selected for download.

7. If it is an offline subscription proxy, do the following:
   a. Download the offline update package from the CA offline subscription FTP site.
   b. Manually copy the updates to the download path of the offline proxy:

      .../data/subscription/offline

   c. Select the offline proxy, select the offline update .zip file, click Browse and select the desired modules.
   d. Click Update Now.

   The server is updated with the modules selected for download.
8. If it is a client with an offline proxy, do the following:
   a. Update the offline proxy as described in step 7.
   b. Select the client and click Update Now.

   The server is updated with all modules installed to its offline proxy.

   Note: Offline subscription clients automatically receive all modules that are manually installed on their offline proxy. Modules selected at the local level for an offline client have no effect.

Free Disk Space for Updates

You can help ensure successful subscription updates on CA Enterprise Log Manager servers by regularly cleaning up the disk. If available disk space is less than 5 GB when the subscription update process begins, the Subscription Service issues a self-monitoring event and suspends the download process.

To ensure sufficient disk space for subscription updates

1. Monitor available disk space on a regular basis. Alternatively, set up an action alert to notify you when available disk space drops below a specified amount.
2. Free disk space with a disk cleanup tool.
3. If you are warned of the need for cleanup with a self-monitoring event, free enough disk space to allow the download to succeed.

More information:

Example: Create an Action Alert for Low Disk Space (see page 388)

About Subscription Public Keys

The subscription proxy maintains a set of public keys corresponding to the private keys used by the CA Subscription server. The subscription proxy downloads subscription updates as a zip file that is digitally signed using a private key. The update identifies the public key to be used to check the signature of the update. Verifying the signature is how the subscription proxy ensures the update is from the CA Subscription server. There is only one public-private key pair used for a given subscription operation. A private key is used in signing the update; a public key is used for verifying the signature. The public key is stored at each CA Enterprise Log Manager server and can be updated.

CA Enterprise Log Manager stores the initial version of the public key in the Subscription Config file during installation. If a new private key is required, the associated public key is downloaded with the subscription update prior to the update cycle where the new key is needed.

Important! Do not manually update the Public Key field for subscription without explicit directions from CA Technical Support.

Self-Monitoring Events for Subscription

You can monitor the successes and failures of the subscription update processes that involve CA Enterprise Log Manager servers configured as the following:

- Default subscription proxy
- Additional online subscription proxies, if any
- Offline subscription proxies, if any
- Subscription clients

Note: The events described here do not track subscription updates for agents.

Successes are reported for the following:

- Startup and shutdown of CA Enterprise Log Manager servers, both subscription proxies and subscription clients.
- Successful download of a component from an online or offline subscription proxy to a subscription client
- Successful installation of a component by a subscription client

Failures and errors are reported for the following:

- Failure to download a component from an online or offline subscription proxy to a subscription client
- Failure to install a component by a subscription client
- Error conditions

Monitor Subscription Events

You can monitor the successes and failures of the subscription update processes by viewing self-monitoring events for subscription.

To monitor subscription processing events

1. Click the Administration tab and the Services subtab.
2. Click Subscription Service in the Service List.
3. Click the Self Monitoring Events tab for either the Subscription Service or for a host listed under Subscription Service.

4. Review the displayed event details.

   Some of the fields displayed for self monitoring events reflect the CEG standard way of referring to common actions across vendors. The Ideal Model is based on this hierarchy: event_category, event_class (within the category), event_action (within the class), event_result, and event_severity. The CA Severity shown in this report is the CA interpretation of the severity.

   CA Severity

   The severity of the event, where

   - Unknown - Unknown events, events not mapped to CEG, or unclassified events
   - Information - General system operations information, security related information, or a notice
   - Warning - Unusual changes, a normal but significant condition, failed operations, or degraded performance
   - Minor Impact - Minor impact to a system, a function, or security
   - Major Impact - Major impact to a system, a function, or security
   - Critical - Immediate action is required; there is a likely security breach.

   Date

   When the event occurred.

   Receiver

   The component that failed to successfully complete the current process. This could be either a subscription proxy or a subscription client. If it is a subscription proxy, it could be the default proxy, an online proxy, or an offline proxy.

   Receiver Host

   The host name of the CA Enterprise Log Manager server that is the Receiver.

   Category

   The category of the event. Configuration Management is the event_category for subscription service events.

   User

   The username or identity that initiated the action expressed in the event information. This is the source_username field in CEG.

   Action

   The action that caused the event to be generated. This is the event_action field in CEG.

Result

S for Success; F for Failure. The result of the event action is the event_result field in CEG. Other actions are Accept, Reject, Drop, and Unknown.

Result Description

The message text. If the Result column displays Failure, view the Result Description text.

5. View details in the Event Viewer. Here you can view changes such as those made to a configuration value such as update start time.

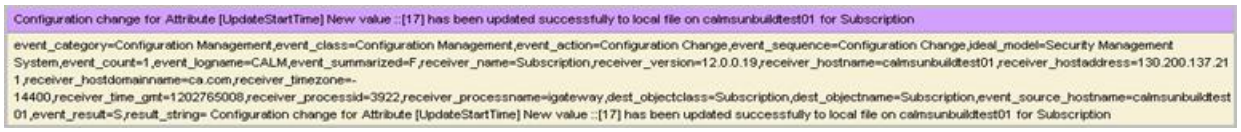

View Subscription Event Details

After you configure subscription, you can view the self-monitoring events. After a subscription update, you can verify that the update to each server completed successfully. As upgrades complete, look for the following self-monitoring event messages on each affected server:

- <component> is downloaded successfully to <proxyname> proxy and installed in EEM.
- <component> is downloaded successfully to <proxyname> proxy.
- <component> is installed successfully to <clientname> client.

You can also view subscription self-monitoring events for troubleshooting purposes.

To view subscription event details in the event viewer

1. Click the Administration tab and the Services subtab.
2. Click Subscription Service in the Service List.
3. Click the Self Monitoring Events tab for either the Subscription Service or for a host listed under Subscription Service.
4. Examine the Result Description column. For example, this column can display events such as "Subscription client has no modules selected for getting updates."

Result Description
Subscription Client - calmrhbuildtest01 is communicating with the proxy - calmrhbuildtest01 for getting the Subscription updates.
Subscription client has no modules selected for getting updates. Please select the modules for getting the updates from the Proxy.
Unable to connect to the specified RSSFeed URL. Please check the URL again.
Unable to connect to the specified RSSFeed URL. Please check the URL again.
No modules are selected for getting updates. Please select the modules for getting the updates from the Subscription server.

5. Double-click the result description for which you would like to review details.

   The event viewer opens.
6. Scroll down to the results section and review the text displayed for result_string.

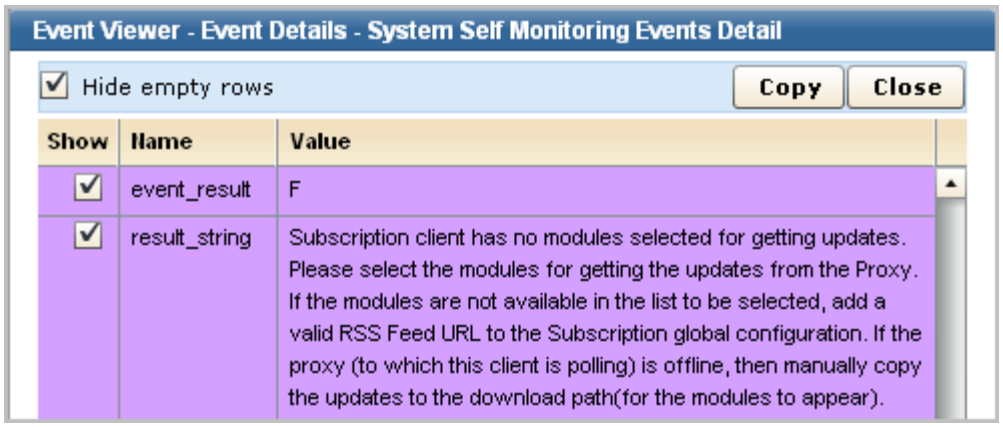
Event Viewer - Event Details - System Self Monitoring Events Detail
Hide empty rows
Copy
Close
Show
Name
Value
event_result
F
result_string
Subscription client has no modules selected for getting updates. Please select the modules for getting the updates from the Proxy. If the modules are not available in the list to be selected, add a valid RSS Feed URL to the Subscription global configuration. If the proxy (to which this client is polling) is offline, then manually copy the updates to the download path(for the modules to appear).

Apply Subscription Updates to Agents and Connectors

Periodic updates, service packs, and point releases are all delivered by subscription. Often the modules to download include Agents and Integrations. When these modules are downloaded to a subscription client that manages agents, you must apply these updates to the agents after verifying that the subscription client that manages the agents has been successfully updated. Updates to agents must be applied before updates to connectors.

To upgrade CA Enterprise Log Manager agents with subscription updates

1. If the upgrade includes the Agents module, update your agents by platform as follows:
   a. Click the Administration tab and click the Log Collection subtab.
   b. Determine whether to apply agent updates to all agents at once, a selected agent group, or an individual agent, depending on the level on which a single platform applies.
      - If all of your agents are installed on the same platform, select Agent Explorer and then click Subscription.
      - If your agent groups are composed of agents installed on the same platform, expand Agent Explorer, select an agent group, and then click Subscription.
      - Otherwise, expand Agent Explorer, expand an agent group, select an agent, and then click Subscription.

      The Subscription Wizard appears.
   c. If you selected Agent Explorer or an agent group, select Agent Updates, select the platform from the Platform drop-down list, click Search, and click the Version Selection Step.
   d. If you selected an agent, select Agent Updates and click the Version Selection step.
   e. Select the Update Version for each listed agent.
   f. Click Save and Close.
   g. Verify success. Click Status and Command. Click Configuration success. Note the version of the configuration that was applied.
   h. Repeat as needed to update all agents.

2. If the upgrade includes the Integrations module, update connectors for your agents as follows:
   a. Click the Administration tab and click the Log Collection subtab.
   b. Determine whether to apply connector updates to all connectors on all agents at once, connectors on agents in a selected agent group, connectors on an individual agent.
   c. Select Agent Explorer, an agent group, or an agent. Then, click Subscription.
   d. Select Connector Updates on the Updates Selection List.
   e. Optionally, select a value from one or more of the following drop-down lists to change the default, All: Agent Group, Platform, Integration. Click Search.
   f. Click the Version Selection step.
   g. Click Select all to select all members of the list or select each row corresponding to a connector you want to update. For each selected row, select the update version to apply.
   h. Click Save and Close.
3. Verify the updates. Run the Subscription Wizard again. Select the Version Selection step to view the current version and verify that it is the version you selected for update. Click Cancel.

More information:

How to Apply Subscription Updates (see page 552)
Open Updates List Wizard (see page 552)
Select Agents or Connectors for Update (see page 553)
Update Agent or Connector Integration Versions (see page 554)

Chapter 8: Filters and Profiles

This section contains the following topics:

Global and Local Filters

You can set or edit filters to refine the displayed event or incident information. You can access the global filter dialog from the main CA Enterprise Log Manager window. You can add local filters from within an individual query or report display, or from the Incidents area. You can also use the Global filter interface to set application-wide query settings.

Each type of filter has its own creation dialog that is launched by a unique button:

Global Filter 

Applies to all reports, queries you view in the *current* session only, and provides a way to view a wide variety of event types qualified in an identical way. The Global Filter button appears at the top of the main CA Enterprise Log Manager window beside the Log Manager Server menu. You can use a global filter to view all events received in the last week, or from a certain host, for example. You can also set global filters for Incidents, which are constructed in the same way, but apply only to incidents and their component event information.

Note: A global filter returning the last 6 hours of data is the default setting.

Local Filter 

Applies only to the current report, query or incident view. The Local Filter button appears at the top of the details pane in query or report displays, and at the top of the Incidents pane. When you view a new report, the local filter is not applied or saved, unless you save the report as a favorite with that filter set. Local filters let you narrow a current view, to see only one host in a multi-host report view, for example, without changing other report views.

About Simple Filters

Before your first use of the Query Design wizard or the Profile Design wizard, become familiar with the simple filter types.

Examples of Simple Filters

An example of each type of simple filter follows:

Filter type	Value	Description
Ideal Model	Antivirus	Displays only event data that products such as the following generate: ■ CA Anti-Virus ■ McAfee VirusScan ■ Symantec Antivirus Corporate Edition ■ TrendMicro OfficeScan
Event Category/ Event Class	System Access/ login activity	Displays only event data related to users logging in to a system.
Event Log Name	Cisco PIX Firewall	Displays only event data that Cisco PIX Firewall devices generate.

With the exception of Event Log Name, the filter types are based on the Common Event Grammar (CEG).

- To learn about product technology-based filters, see the Ideal Models List.
- To learn about product category/class/action-based filters, see the Event Categories List and the Event Classes List.

Find these lists in the online help under the "Common Event Grammar" section.

Set a Simple Filter

You can set simple filters to establish criteria for the event data you want displayed or reported. When set in the Query Design wizard, simple filters let you limit the event data returned by a query used in a report or alert. When set in the Profile Design wizard, simple filters limit the data displayed in the report or query results when the profile is applied.

1. Open the wizard.
2. Determine the type of simple filter to set:
   - technology-based.
   - category-based, category and class-based, or category, class, and action-based.
   - product-based.
3. To set a technology-based filter, click the Ideal Model is check box and then select a value from the Ideal Model drop-down list.
4. To set a filter based on a security event category, category and class, or category, class, and action, do the following:
   a. Click the Event Category is check box and then select a value from the associated drop-down list.
   b. (Optional). Click the Event Class is check box and then select a value from the drop-down list.
   c. (Optional). If you selected Event Class, click the Event Action is check box, and then select a value from the drop-down list.

   Note: You can also set this type filter under a technology-based filter.
5. To set a product-based filter, click the Event Log Name is check box and then select a value from the drop-down list.
6. Complete the wizard.

About Profile Filters

A profile is a set of filters. You can create a profile with tag filters, data filters, or a combination. The query tag filter limits the queries displayed for selection; the report tag filter limits the reports displayed for selection. The data filters limit the data displayed in a report or in query results. The profile filters apply to queries, reports, scheduled alerts, and scheduled reports.

You can select tag filters for reports and queries separately. Tag filters include, but are not limited to, the following:

- Standards-based tags, such as COBIT, FISMA,GLBA, HIPAA, NERC, PCI, SAS 70, SOX.

  Standards-based tag filters apply to report tags, not query tags.
- Security event category tags, such as content security, host security, network security, operational security, resource access, system access.
- Product tags, such as CA Access Control, CA Identity Manager, and CA SiteMinder.

You can select a simple data filter or you can create an advanced data filter. A brief description of each follows:

- Simple data filters can be based on any of the following:
  - A selected technology (system software, host application software and services, network application software and services).
  - A selected CEG event category, a selected CEG event category and class, or a selected CEG event category, class, and action.
  - A selected product.
- Advanced data filters are based on a user-defined SQL query composed of one or more WHERE clauses. The query selects a CEG column with a WHERE clause composed of that CEG column, a selected operator, and a specified value.

How to Create a Profile

You can create profiles, which allow users to narrow their CA Enterprise Log Manager views, according to your environmental needs. For example, you could create a CA Access Control profile that would show only reports, queries and events relevant to Access Control.

The process of creating a profile, using the profile wizard, has the following steps:

1. Opening the profile wizard
2. Naming the profile and entering description information
3. Identifying the information shown using simple and advanced filters
4. Selecting which queries and reports are displaying using tag filters

More information:

Add Profile Details (see page 210)
Create Data Filters (see page 211)
Create Tag Filters (see page 212)

Open the Profile Wizard

To create a new profile, or edit an existing one, you must open the profile wizard.

To open the profile wizard

1. Click the Administration tab, and then the Library subtab.

   The Library folder list appears.

2. Select the Profiles folder.

   The Profiles buttons appear in the details pane.

3. Click New Profile: 

   The Profile Wizard opens.

   When using the wizard:

   - Click Save to save the rule file without closing the wizard.
   - Click Save and Close to save the rule file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Add Profile Details

You must name a profile. You can also enter optional description information for reference.

To name a profile

1. Open the profile wizard.
2. Type a name for the new profile. The name may be up to 80 characters long, and may include special characters.
3. (Optional) Type description information.
4. Advance to the Data Filters step.

Create Data Filters

You filter the information shown by your profile using simple or advanced filters. Each profile must have at least one filter.

To set profile data filters

1. Open the profile wizard.
2. Enter the profile name, if not already specified, then advance to the Data Filters step.

   The filters dialog appears, displaying the Simple Filters Tab.
3. Create any simple filters you want, to search for stated CEG field values. For example, you could select the Event Log Name check box, and enter "CA Access Control" to search for CA Access Control events.
4. (Optional) Click the Advanced Filters tab.

   The advanced filters dialog appears.
5. Create advanced filters as needed.
6. Click the appropriate arrow to advance to the profile wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new profile appears in the list, otherwise the wizard step you choose appears.

More information:

Create a Simple Event Filter (see page 479)
Create an Advanced Event Filter (see page 482)
Using Advanced Filters (see page 408)

Create Tag Filters

You can create tag filters for your profile, controlling which queries or report category tags appear in the CA Enterprise Log Manager interface when a user applies the profile. For example, if you create a tag filter for CA SiteMinder, the CA Enterprise Log Manager interface displays only those reports and queries with the CA SiteMinder tag.

To create a tag filter

1. Open the profile wizard.
2. Enter the profile name, if not already specified, then advance to the Tag Filters step.

   The filters dialog appears, displaying the Report Tag Filters subtab.
3. Click New Event Filter.

   The first row of the tag filter table becomes active.
4. Click the Tag cell and select or type the query or report tag name you want to display. If you type, the display narrows available tag names as you type.
5. (Optional) Click New Event Filter again to add additional filters.

   The second row of the tag filter table becomes active, displaying AND in the logic column.
6. (Optional) Click the logic cell to select either an AND or OR operator.
7. (Optional) Click the Tag cell and select or type the tag name you want to display. If you type, the display narrows available tag names as you type.
8. (Optional) Click the open and closed parentheses cells and enter the number of parentheses you need.
9. (Optional) Click the Query Tag Filters subtab, and repeat steps 3 through 8 to create any query tag filters that you need.
10. Click Save when you have entered all the filter statements you want.

More information:

Create Data Filters (see page 211)

Import a Profile

You can import a profile, allowing you to move profiles from one environment to another. For example you could import a profile created in a test environment to your live environment.

To import a profile

1. Click the Administration tab, and then the Library subtab.

   The Library folder list appears.

2. Click the arrow beside the Profiles folder to expand it.

   The profiles buttons appear in the details pane.

3. Click Import Profile.

   The import file dialog appears.

4. Browse to find the file you want to import, and click OK.

   The profile wizard appears, displaying the details of the profile you selected.

5. Make any changes you want, and click Save and Close. If the imported profile shares a name with one already in your management database, you are prompted to change the name.

   The imported profile appears in the appropriate folder.

Export a Profile

You can export a profile. This lets you share profiles between environments. For example, you could export a profile created in a test environment to your live environment.

To export a profile

1. Click the Administration tab, and then the Library subtab.

   The Library folder list appears.

2. Click the arrow beside the Profiles folder to expand it.

   The profiles folders appear.

3. Click the folder which contains the profile you want to export.

   The folder expands, showing the individual files.

4. Select the profile you want to export, and then click Export Profile.

   An export location dialog appears.

5. Enter or browse to the location where you want to store the exported profile, and click Save.

   An export successful confirmation dialog appears.

6. Click OK.

   The profile is exported.

Set a Profile

You can select any available profile to apply to your environment, restricting the queries and reports available, depending on the terms of the profile. To set a profile, select the profile you want from the Profiles drop-down menu at the top of the main CA Enterprise Log Manager window.

Note: To set the selected profile as a default profile of your environment, click the Set as default profile option at the top of the main CA Enterprise Log Manager window. The selected profile is set as the default profile of the logged in user.

Create a Global Filter

You can create a global filter. Global filters let you view all queries and reports, or all incidents using the same qualifying factors. When you create a global filter you choose whether it applies to events or to incidents. A single global filter cannot apply to both. You can also use the Global filter interface to set application-wide query settings.

To create a global filter

1. Click the Global Filters button at the top of the main window.

   The Global Filters and Settings dialog appears.

2. Click the Events tab or the Incidents tab to select where you want the global filter applied.

3. Specify the time period you want your filter to search, using the Time Range drop-down list.

4. Select the Match check box to enter a specific value by which you want to filter all available raw events.

   Note: You can search for multiple values, phrases, or partial values in the raw events by using the specialized Match syntax.

5. Click Add Filter to specify event fields that you want to include in the filter.

   The Column drop-down menu and Value entry field appear.

6. Choose the event field you want to include in the filter, and type the value that the field must have to be displayed in the filtered reports. You can enter multiple event field names and values by clicking Add Filter again. Selecting the Exclude button includes every value *but* the one you enter for the chosen event field name.

   Note: If you create a global filter on a string-type field, it is added to the Quick filters list. If you create a filter on a numeric or time field, it is added to the Advanced filters list.

7. (Optional) Click the Advanced Filters tab to add additional complex qualifiers.

8. (Optional) Click the Settings tab to choose and global settings. These setting are applied to the whole application.

9. (Optional) Select Set as Default at the bottom of the dialog to retain the filter settings for any future sessions, as long as you are logged in as the same user.

10. Click Save.

    The Global Filters and Settings dialog closes, and the new filter is applied to reports.

More information:

Using Advanced Filters (see page 408)
Configure Global Query Settings (see page 216)

Configure Global Query Settings

Using the Global Filter dialog, you can set application-wide conditions that apply to all reports and queries in your environment. Global settings apply throughout the current session unless you set them as a default.

To configure global query settings

1. Click the Global Filters button at the top of the main window.

   The Global Filters and Settings dialog appears, displaying the Quick Filters tab.

2. Click the Settings tab

   The tab opens, displaying the following values.

   Local Time Zone

   Controls the time zone for all date/time fields in reports and queries. Your reports and queries adopt the time zone you select from the drop-down list rather than using the CA Enterprise Log Manager server time zone.

   Execute queries on federated data

   Allows the query to be applied to all available federated servers. This setting is enabled by default. Disabling this setting confines queries to only the event data stored in the local event log store. This lets you quickly check your local event log store when you know your target events are local.

   Enable auto refresh for queries

   Allows the display to automatically refresh at the set interval for each query.

3. (Optional) Select Set as Default at the bottom of the dialog to preserve the settings as defaults that are retained after the current session.

4. Make any changes you want, and click Save.

   The Global Filters and Settings dialog closes, and the new filter is applied.

Edit a Global Filter

You can edit an existing global filter.

To edit a global filter

1. Click the Global Filters button at the top of the main window.

   The Global Filters and Settings dialog appears, displaying the Quick Filters tab.

2. Change or add parameters as needed. You can remove an individual quick filter parameter by clicking the Delete icon beside it.

3. Click Save.

   The Global Filters and Settings dialog closes, and the edited filter is applied.

Remove a Global Filter

You can remove a global filter, returning all reports to their default state.

To remove a global filter, click Clear Global Filters at the top of the main CA Enterprise Log Manager window: 

Create a Local Filter

You can create a local filter to narrow the scope of the query or report you are viewing, or the incidents displayed in the Incidents area.

To create a local filter

1. Open the query or report you want to filter, or click the Incidents tab, and click the Local Filters button at the top of the pane.

   The Local Filters dialog appears, displaying the Quick Filters tab.

2. (Optional) Click the Incidents tab if you want to filter displayed incidents rather than events.

3. (Optional) Select the Match check box to type a specific value by which you want to search events or incidents.

   Note: You can search for multiple values, phrases, or partial values by using the specialized Match syntax.

4. Click Add Filter.

5. Choose the event field you want to include in the filter, and type the value that the field must have to be displayed in the filtered reports. You can enter multiple column values by clicking Add Filter again. Selecting the Exclude button includes every value *but* the one you enter for the chosen event field name.
6. (Optional) Click the Advanced Filters tab to add additional qualifiers.
7. Click Save.

   The filter is applied to the display. You can save the report view by setting it as a Favorite.

Edit a Local Filter

You can edit an existing local filter.

To edit a local filter

1. Click the Local Filters button at the top of the query or report pane, or at the top of the Incidents view.

   The Local Filters dialog appears, displaying the Quick Filters tab.
2. Change or add values as needed. You can remove individual filter settings by clicking the Delete icon beside each one, or remove a Match value by clearing the check box.
3. Click Save.

   The edited filter is applied to the display.

Remove a Local Filter

You can remove a local filter, returning a query, report, or incident view to its previous state.

To remove a local filter, click the Clear Local Filter button at the top of the query, report or incident display you are viewing: 

Chapter 9: Queries and Reports

This section contains the following topics:

About Queries and Reports

You can use queries in the following ways:

- You can run a query to view event or incident data in near-real time.
- You can select a predefined report to view results of multiple related queries.
- You can create a report composed of queries you select.
- You can use prompt queries to search for specific preselected information.
- You can schedule queries to run on recent data as action alerts that notify responsible parties through email. Action alerts are also added to an RSS feed that can be viewed using third-party readers.
- You can create your own queries for viewing, reporting, or creating action alerts.

There are two types of queries and reports:

- *Subscription* queries and reports are predefined by CA and come with the CA Enterprise Log Manager application at installation or are added with a subscription update.
- *User* queries and reports are those created by a user. You can create a query or report from scratch or you can create one based on a subscription query or report that you want to modify.

CA Enterprise Log Manager offers a comprehensive list of queries and reports by subscription. If you are assigned a role of Auditor, Analyst, or Administrator, you can view all Subscription queries and reports. In addition, you can take the following actions on any subscription query or report you are viewing:

- Refresh the displayed data
- Edit local filters to hide the data you do not want to view.
- Clear the local filters to re-display the unfiltered query or report.
- Add the displayed query or report to your list of favorites.
- Print the query
- Change the option to show the selected query or report
- Close the displayed query or report

Only users who are assigned a role of Analyst or Administrator can take the following actions:

- Create a new User query or report from scratch
- Copy a subscription query or report and use it as the basis for a User query or report.
- Edit a User query or report
- Export a User query or report
- Delete a User query or report
- Save changes to the selected User query or report
- Import a User query or report definition

Example of Queries and Related Report

Consider the query tag Firewall Activity by DMZ. Notice that it is associated with six separate queries on this topic.

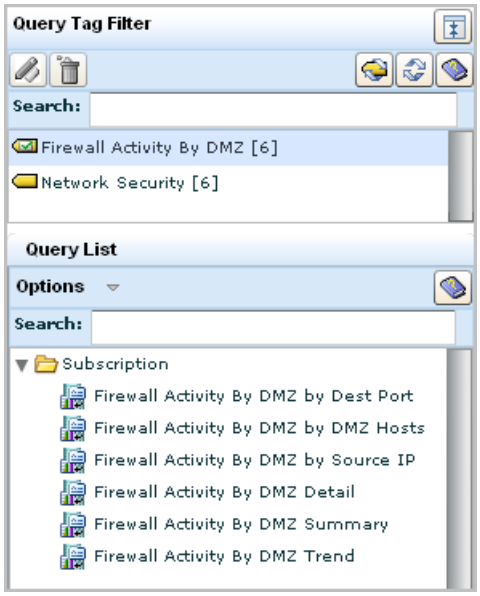

The queries you view on the query list are used in reports. From the Reports tab, you can display a report called Firewall Activity By DMZ.

The following illustration shows just the names. Notice that each name reflects one of the six queries in the report. Most reports include query results for summary, trend, and detail.

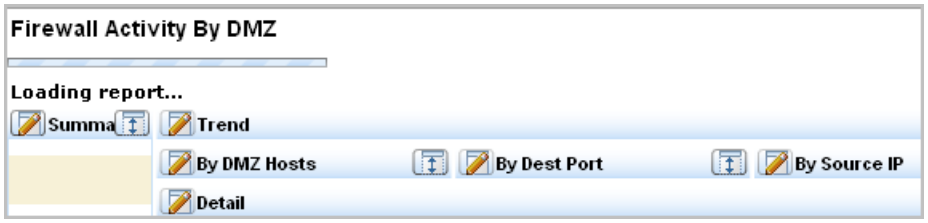

Tag Tasks

Tags let you attach your reports and queries to categories for easy reference, and provide an organizational framework for reporting on your environment. Category tags also allow simple division of labor by role or type of event.

You can use the pre-defined tags or create your own custom tags for reports or queries. For example, you can create a "Monthly" tag to add to any reports you want to schedule every month for easy reference and viewing. This lets you add or remove reports from the report jobs without editing the jobs themselves, by simply adding the Monthly tag to a new job, or removing it from an old one.

You can add custom tags for individual queries or reports as part of the query or report creation or editing process. Once you create a new tag, its title appears in the tag list, and you can select it to add it to other reports or queries.

You can rename or delete custom tags. You can remove custom tags from reports or queries that include them by editing the report or query.

View a Query

All users who are assigned a role of Auditor, Analyst, or Administrator can view all queries. Predefined queries are listed under the Subscription folder. When the first custom query is defined, a User folder is added to the query list to hold the custom query. After that, all custom queries are added to this User folder.

To view a query

1. Click the Queries and Reports tab, then the Queries subtab.

   The Query Tag Filter maximize button, the Query List, and Options menu and a Search text box appear in the left pane.

2. Select the query to view in any of the following ways:

   - Scroll through the query list and select a query to view
   - Enter a keyword in the Search box to show only the queries with the matching word in their names
   - Click the maximize button to display the query tag filter list. Either select a displayed tag or enter a keyword in the tag Search box to limit the tags displayed. Select a tag to show the related queries. Select the query to view.
   - If searching for a custom query, contract the Subscription folder, expand the User folder, and then scroll through the User folder list

   The selected query displays in the main pane of the page.

3. (Optional) Take any of the following actions:

   - Click Edit Local Filters to set filters to display only the data you want to view. To restore the original query display, click Clear the Local Filters.
   - Click Add to Favorites to add the displayed query or report to your list of favorites.
   - Click Refresh to refresh the data with that which was most recently added.
   - Click Print to print the query.

4. Click Close to close the displayed query.

View a Report

All users who are assigned a role of Auditor, Analyst, or Administrator can view all reports. Predefined reports are listed under the Subscription folder. When the first custom report is defined, a User folder is added to the report list to hold the custom report. After that, all custom reports are added to this User folder.

Selecting a report from the report list runs the queries that make up the report on log records currently residing in the internal event log stores. The report results, displayed on the right pane, are from the event log stores of the active CA Enterprise Log Manager server and its child servers.

To view a report

1. Click the Queries and Reports tab, then the Reports subtab.

   The Report Tag Filter maximize button, a Search entry field, the Report List, and the Options menu appear in the left pane.

2. From the Options menu, select Show Selected Report, if not already selected.

   This lets you display any selected report in the right pane.

3. Select the report to view in any of the following ways:

   - Scroll through the report list and select a report to view
   - Enter a keyword in the report Search entry field and select a report to view from the filtered list
   - Click the maximize button to display the report tag filter list. Either select a displayed tag or enter a keyword in the tag Search box to limit the tags displayed. Select a tag to show the related reports. Select the report to view.
   - If searching for a custom report, contract the Subscription folder, expand the User folder, and then scroll through the User folder list

   The selected report displays in the main pane of the page.

4. (Optional) Take any of the following actions:

   - Click Edit Local Filters to set filters to display only the data you want to view. To restore the original report display, click Clear the Local Filters.
   - Click Add to Favorites to add the displayed report or report to your list of favorites.
   - Click Refresh to refresh the data with that which was most recently added.
   - Click Print to print the report.

5. Click Close to close the displayed report.

Disable Show Selected Report

You can set your report list so that you can make changes without loading reports. Normally, selecting a report from the list displays it in the details window.

Disabling this default mode saves time by letting you select a report from the list and edit it immediately, without waiting for it to display. This is especially useful if you have multiple reports to edit and already know what changes you plan to make.

Since only users with the Administrator or Analyst mode can create or edit reports, only these users can disable the show selected report setting.

To disable show selected report

1. Click Options at the top of the Report List.

   The Options menu appears.

2. Clear the check beside Show Selected Report.

   Any report selected from the list is not displayed until Show Selected Report is re-enabled.

More information:

How to Create a Report (see page 261)
Edit a Report (see page 272)

Example: Run PCI Reports

The PCI Security Standards Council is an open global forum responsible for the development of the PCI Data Security Standard (PCI DSS) that includes requirements for security management, policies and procedures. Organizations that store, process or transmit cardholder data must comply with PCI DSS version 1.2, which details twelve requirements.

CA Enterprise Log Manager provides out-of-the-box PCI reports that you can view as soon as your system begins to collect and process event logs.

The examples in this section help you become familiar with the PCI reports and how to schedule and distribute them. The examples include references to the number associated with the PCI DDS Requirement that the report addresses.

More information:

View the List of Reports with the PCI Tag (see page 226)
Search for Reports for a Specific PCI DDS Control (see page 227)

View the List of Reports with the PCI Tag

You can begin your assessment of how to use CA Enterprise Log Manager reports to demonstrate PCI compliance by viewing the list of predefined reports that are tagged with the PCI tag.

To become familiar with reports with the PCI tag

1. Click the Queries and Reports tab and the Reports subtab.

   The Report Tag Filter and Report List appear.

2. Enter PCI in the Search field for the tag.

   The PCI tag appears.

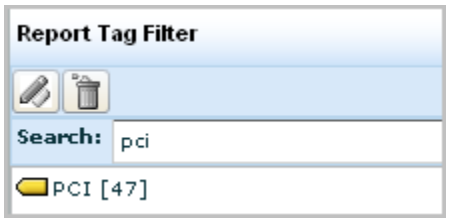

3. Review report list associated with the PCI tag.

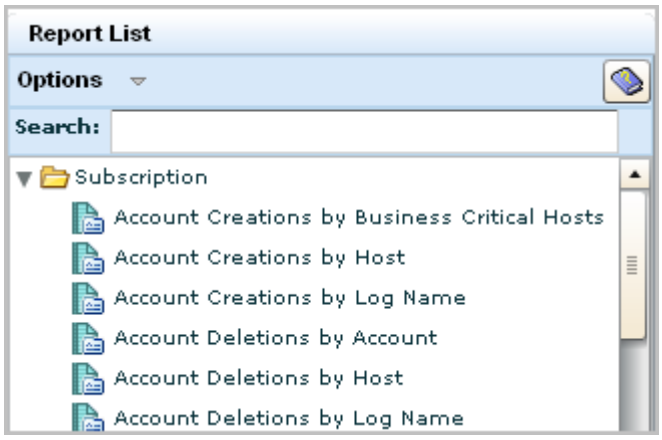

Search for Reports for a Specific PCI DDS Control

You can search for predefined reports using keywords relevant to specific PCI DDS controls. The following procedure covers a few examples.

Note: The referenced numbers are the number associated with the PCI DDS Requirement that the report addresses.

To display the list of reports relevant to specific PCI DDS controls

1. Click the Queries and Reports tab and the Reports subtab.
2. To locate the report that address changes to the firewall configuration (1.1.1), enter Firewall as the Search criteria.

   A list of reports similar to the following appears. Notice the one titled Firewall Configuration Changes.

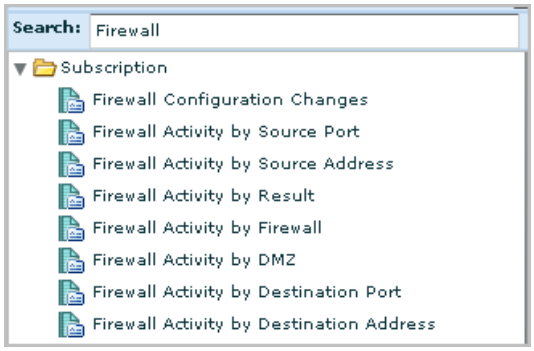

3. To locate the report that addresses changes to router configurations after you have verified synchronization (1.3.6), enter Router as the search criteria.

4. To locate reports that address password management (8.5), one of the strong access control measures, enter password as the Search criteria.

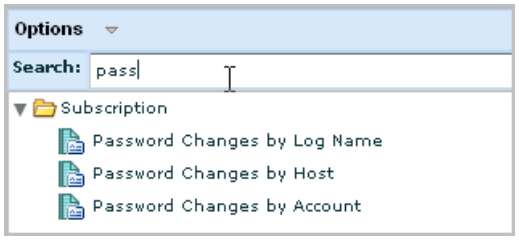

5. To locate reports that address additions, modifications, and deletions to user accounts (12.5.4), one of the measures for maintaining an information security policy, enter account as the search criteria.

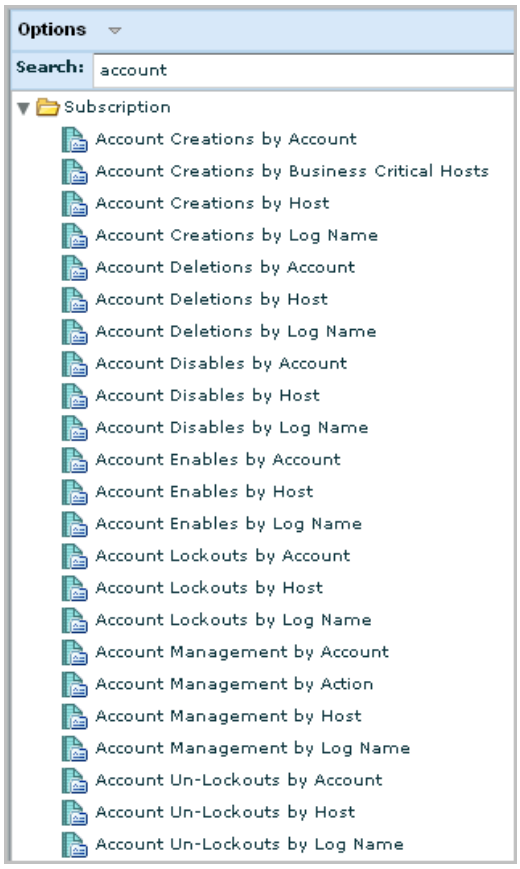

Work with a Single PCI Report

You can work with any report, including PCI reports, in the following ways:

- View the report by selecting the report name from the Report List.
- Print the report
- Schedule the report, with the option to email it to selected recipients.
- View the scheduled report job.
- View the generated report.

To view or act on a selected report

1. Click the Queries and Reports tab and the Reports subtab.
2. Select Show Selected Report in the Options drop-down list under Report List, if not already selected.

3. Select a report name from the report list.

   The resulting report displays the results of the underlying queries, which typically include a summary, the trend, and details, as well as report-specific queries.

4. To disable the loading of particular queries, select Cancel.

5. To print the displayed report, click Print Report in the right pane.

   When the Print dialog displays, select a printer and click Print.

6. To schedule the report to be generated for later viewing, click Schedule Report.

   The Schedule Report wizard appears with the displayed report in the Selected Reports area.

7. Enter a job name, for example, Resource Access by Host Report job.

   If you accept all defaults, the job is scheduled to run now with no recurrence, where the report is generated in PDF format with no email notification. The data is drawn from the current server, its federated peers and its federated descendants.

8. Click Save and Close.

9. View the scheduled job. Select the Scheduled Reports tab and then the Report Scheduling subtab.

   The job you just scheduled is shown.

Scheduled Jobs

	Job Name	Enabled	Server	Status	Recurrence	Scheduled Time ▼	Time Zone	Creator
☐	Resource Access by Host Report Job	true	calem_server	Schedule Expired	Now	Fri Aug 27 2010 3:30:38 P	America/New_York	admin

10. View the generated report.

    a. Select the Scheduled Reports tab and then the Generated Reports subtab.

    b. (Optional) Limit the displayed rows by selecting a recurrence other than All, a format other than All, or a Time span of the last hour.

    c. (Optional) Click Refresh.

11. After reviewing the generated report, you can modify the report job if you want to generate it on a recurring basis. Do the following:

    a. From the Reports Scheduling subtab, select the generated report, and click Edit.

    b. Select the Schedule Jobs step and select the option for the frequency of occurrence.

    c. Click Save and Close.

Prompts

A prompt is a special type of query that displays results based on the value you enter and the CEG fields you select. Rows are returned only for events where the value you enter appears in one or more of the selected CEG fields.

You can take any of the following actions on prompt query results:

- Select Show raw events to replace the view of refined events with the corresponding raw event and the time it occurred.
- Enter a string in the Match field and select Go to filter the display to rows containing data that matches your entry.
- Select an Export Query Data option to export the query results to a PDF document, an Excel spreadsheet or an XML file.
- Select Result Conditions to filter the display by a specified date range, set the limit for returned rows, or change the granularity of the displayed time. Or, reset result conditions to the default.
- Select Show/Edit Local Filter to specify quick filters or advanced filters.
- Print the query to a selected local printer.
- Refresh the query data manually or by selecting Auto Refresh.

Use the Connector Prompt

Each connector that is configured on an agent collects raw events from a specific event source and sends the events to the event log store on a CA Enterprise Log Manager collection server. The event refinement process converts raw events to refined events and archives them to the reporting CA Enterprise Log Manager server. The connector prompt queries for events on the reporting server that were collected as raw events by connectors with the name you specify. Connectors can have a default name or a user-defined name. You copy the name of the connector to use and paste it in the field of the connector prompt and click Go to display the prompt query results.

Use the connector prompt to:

- View events from all connectors based on the same integration. This is possible if you accept the default connector name when deploying connectors.
- Verify a new connector is retrieving events. If multiple agents have connectors with the name you specify, enter the agent name in the Match field to limit the query results to events retrieved by the new connector.

To copy the name of an active connector

1. Click the Administration tab.

   The Log Collection Explorer is displayed.

2. Click Agent Explorer.

   The Agent Status Monitor appears, where one column lists connector names.

3. Right-click the connector you want to use in the prompt query and select Copy Connector Name.

To use the Connector prompt

1. Select Queries and Reports.

   The Query List displays the Prompts folder, the Subscription folder, and possibly a Users folder.

2. Expand Prompts and select Connector.

   The Connector prompt displays the Connector field and the following CEG field, which must remain selected for the prompt to function:

   agent_connector_name

   Is the name of a connector.

3. Right-click in the Connector field and select Paste.

   The connector name you copied from the Agent Status Monitor appears in the Connector field.

4. Click Go.

   Results of the connector prompt query appear.

5. Use the following descriptions to interpret the query results:

   CA Severity

   Indicates the severity of the event, where the values in increasing order of severity include: Information, Warning, Minor Impact, Major Impact, Critical, and Fatal.

   Date

   Indicates when the event occurred.

   Category

   Identifies the high-level category of the corresponding event action. For example, System Access is the category for the Authentication action.

   Action

   Identifies the action, where possible actions are determined by the class of the event.

Agent Name

Identifies the agent on which the connector is running.

Host

Identifies the event source host from which the connector is collecting events.

Performer

Identifies the source actor of the event, that is, the identity that initiated the action. The performer can be expressed as the source username or source process name.

Account

Identifies the username of the account used for authentication when the connector attempts to connect to the host with the event source from which raw events are collected. This is typically a low-privileged account. The credentials for this account are configured on the event source and also on the log sensor of the connector.

Result

Specifies a code for the event result of the corresponding action, where S means Success, F means Failure, A means Accepted, D means Dropped, R means Rejected, and U means Unknown.

Connector Name

The name of the connector entered in the prompt filter field.

6. (Optional) Select Show raw events.

The first event collected by a new connector is for the action System Startup and ends with: result_string=<connector name> Connector Started Successfully

Use the Host Prompt

The host prompt queries for events where the hostname you specify appears in the selected CEG fields of the refined event. When raw event data is refined, event details can include several different CEG host names. Consider this scenario:

1. The event initiator on source_hostname attempts an act, event_action, on a target residing on dest_hostname.

   Note: Source_hostname and dest_hostname can be different hosts or the same host.

2. This event is recorded in a repository on event_source_hostname.

   Note: Event_source_name can be a different host than either source_hostname or dest_hostname or can be colocated.

3. A CA Enterprise Log Manager agent installed on agent_hostname makes a copy of the event recorded on event_source_hostname.

   Note: Agent_hostname is the same as event_source_name in agent-based log collection but is different in agentless and direct log collection.

4. The CA Enterprise Log Manager agent on agent_hostname transmits the copy of the event in event_logname to a CA Enterprise Log Manager collection server.

To use the Host prompt

1. Select Queries and Reports.

   The Query List displays the Prompts folder and one or more folders for other queries.

2. Expand Prompts and select Host.

   The Host prompt appears.

3. Enter the name of the host on which to base this query.

4. Select the fields on which to query for data matching your host name entry.

   source_hostname

   Is the name of the host where the event action was initiated.

   dest_hostname

   Is the name of a host that is the destination or target of the action.

   event_source_hostname

   Is the name of a host that records the event when the event occurs.

   For example, you can deploy a connector based on WinRM to collect events from the Event Viewer on a Windows Server 2008 host. To select events retrieved from a given Windows Server 2008 host, enter the hostname of that server and select this field.

receiver_hostname

Is the same as agent_hostname.

agent_hostname

Is the name of the host where a CA Enterprise Log Manager agent is deployed.

5. Click Go.

   Results of the host prompt query appear.

6. Use the following descriptions to interpret the query results:

   CA Severity

   Indicates the severity of the event, where the values in increasing order of severity include: Information, Warning, Minor Impact, Major Impact, Critical, and Fatal.

   Date

   Indicates when the event occurred.

   Source User

   Identifies the name of the user on source_hostname who initiated the event action.

   Result

   Specifies a code for the event result of the corresponding action, where S means Success, F means Failure, A means Accepted, D means Dropped, R means Rejected, and U means Unknown.

   Agent Host

   Identifies the name of the host where the CA Enterprise Log Manager agent who collected the event is installed.

   Receiver Host

   The same as agent host.

   Category

   Identifies the high-level category of the corresponding event action. For example, System Access is the category for the Authentication action.

   Action

   Identifies the event action performed by the source user.

   Log Name

   Identifies the log name used by the connector that collected the event. All connectors based on the same integration transmit events in a log file with the same log name.

Use the IP Prompt

The IP prompt queries for events where the IP address you specify appears in the selected CEG fields of the refined event. When raw event data is refined, event details can include several different CEG IP addresses. Consider this scenario:

1. The event initiator on source_address attempts an act, event_action, on a target residing on dest_address.

   Note: Source_address and dest_address can be different or the same.

2. This event is recorded in a repository on event_source_address.

   Note: Event_source_address can be different from either source_address or dest_address or can be the same as one or both.

3. A CA Enterprise Log Manager agent installed on agent_address makes a copy of the event recorded on event_source_address

   Note: Agent_address is the same as event_source_address in agent-based log collection but is different in agentless and direct log collection.

4. The agent on agent_address transmits the copy of the event in event_logname to a CA Enterprise Log Manager collection server.

To use the IP prompt

1. Select Queries and Reports.

   The Query List displays the Prompts folder and one or more folders for other queries.

2. Expand Prompts and select Host.

   The IP prompt appears.

3. Enter the IP address on which to base this query.

4. Select one or more of the following fields to query for data matching your IP address entry.

   source_address

   Is the IP address of the host where the action was initiated.

   dest_address

   Is the IP address of a host that is the destination or target of the action.

   event_source_address

   Is the IP address of a host that records the raw event when the event occurs.

   For example, you can deploy a connector based on WinRM to collect events from the Event Viewer on a Windows Server 2008 host. To select events retrieved from a given Windows Server 2008 host, enter the IP address of that server and select this field.

receiver_hostaddress

Is the same as agent_address.

agent_address

Is the IP address of a host where a CA Enterprise Log Manager agent is deployed.

5. Click Go.

   Results of the IP prompt query appear.

6. Use the following descriptions to interpret the query results:

   CA Severity

   Indicates the severity of the event, where the values in increasing order of severity include: Information, Warning, Minor Impact, Major Impact, Critical, and Fatal.

   Date

   Indicates when the event occurred.

   Result

   Provides a code for the result of the corresponding action, where the displayed letter has the following meaning: S for success, F for failure, A for Accepted, D for Dropped, R for Rejected, and U for Unknown.

   Destination Port

   Identifies the communication port on the destination host, the target of the event action.

   Source IP

   Identifies the IP address from which the event action was initiated.

   Destination IP

   Identifies the IP address of the host that was the target of the event action.

   Event Source IP

   Identifies the IP address of the host with the repository where the event was originally recorded.

   Agent IP

   Identifies the name of the host with the CA Enterprise Log Manager agent responsible for the collection of events from the event source.

Receiver IP

The same as Agent IP.

Category

Identifies the high-level category of the corresponding event action. For example, System Access is the category for the Authentication action.

Action

Identifies the event action.

Log Name

Identifies the log name used by the connector that collected the event

Use the Log Name Prompt

Each connector that is based on the same integration returns event logs collected from the event source to the CA Enterprise Log Manager collection server in a log file with a predefined name.The log name prompt queries for events involving the log name you specify.

Use the log name prompt to query for events transferred in a log file with the specified name. Each connector is based on an integration. Each integration uses a predefined log name. A query for a given log name returns results of events collected by different agents that use connectors based on the same integration or similar integrations.

A variety of conventions are used for naming logs:

- Integration name. CA Federation is the log name for the CA_Federation_Manager integration.
- Product name. McAfee Vulnerability Manager is the log name for both McAfee_VM and McAfee_VM_CM. MS AD Rights Management Services is the log name for both Microsoft_Active_Directory_RMS and Microsoft_Active_Directory_RMS_ODBC.
- Vendor name: Oracle is the log name for Oracl10g, Oracle9i, Oracle_AppLog, and Oracle_Syslog.
- Log type: Unix is the log name for the following integrations:AIX_Syslog, HPUX_Syslog, Linux_Syslog, SLES_Syslog, and Solaris_Syslog.

Some log names are reused as new releases or platforms are added. For example, NT-Security is the log name for security logs for the following integrations: NTEventLog, Windows2k8, and WinRM.

To use the Log Name prompt

1. Select Queries and Reports.

   The Query List displays the Prompts folder and one or more folders for other queries.

2. Expand Prompts and select Log name.

   The Log name prompt filter appears with the following field:

   event_logname

   Is the name of a log file associated with a specific integration.

3. Select the log name used to transmit events you want to view and click Go.

   Results of the log name prompt query appear.

4. Use the following descriptions to interpret the query results:

   CA Severity

   Indicates the severity of the event, where the values in increasing order of severity include: Information, Warning, Minor Impact, Major Impact, Critical, and Fatal.

   Date

   Indicates when the event occurred.

   Category

   Identifies the high-level category of the corresponding event action. For example, System Access is the category for the Authentication action.

   Action

   Identifies the event action performed by the corresponding performer.

   Host

   Identifies the event source host from which the connector is collecting events.

   Performer

   Identifies the source actor of the event, that is, the identity that initiated the action. The performer can be expressed as the source username or source process name.

Account

Identifies the username of the account used for authentication. When the connector attempts a connection to the event source, authentication occurs. Authentication typically uses a low-privileged account. During connector deployment, the administrator configures credentials for this account on the event source and then identifies this account on the log sensor.

Result

Specifies a code for the event result of the corresponding action, where S means Success, F means Failure, A means Accepted, D means Dropped, R means Rejected, and U means Unknown.

Log Name

The log name entered in the prompt filter field.

Use the Port Prompt

The port prompt queries for events where the port you specify appears in the selected CEG fields of the refined event. When raw event data is refined, event details can include several different CEG port numbers. Consider this scenario:

1. The event initiator on the source host uses the outbound source_port communication port for initiating the event action on a target residing on the destination host through the inbound dest_port communications port.

   Note: Source_port and dest_port are the same for local events. Otherwise, they are host-specific.

2. This event is recorded in a repository on the event source.
3. A CA Enterprise Log Manager agent makes a copy of the event recorded on the event source.
4. The agent transmits the copy of the event through the outbound port, receiver_port, to a CA Enterprise Log Manager collection server.

   Note: The agent uses port 17001, by default, to secure communications to the CA Enterprise Log Manager collection server.

To use the Port prompt

1. Select Queries and Reports.

   The Query List displays the Prompts folder and one or more folders for other queries.

2. Expand Prompts and select Port.

   The Port prompt appears.

3. Enter the port number on which to base this query.

4. Select the fields on which to query for data matching your port number entry:

   source_port

   Is the communications port used for initiating the action.

   dest_port

   Is the communication port on the destination host that is the target of the action.

   receiver_port

   Is the port that the agent uses to communicate with the CA Enterprise Log Manager collection server.

5. Click Go.

   Results of the port prompt query appear.

6. Use the following descriptions to interpret the query results:

   CA Severity

   Indicates the severity of the event, where the values in increasing order of severity include: Information, Warning, Minor Impact, Major Impact, Critical, and Fatal.

   Date

   Indicates when the event occurred.

   Source IP

   Identifies the IP address of the host from which the event action was initiated.

   Result

   Specifies a code for the event result of the corresponding action, where S means Success, F means Failure, A means Accepted, D means Dropped, R means Rejected, and U means Unknown.

   Source Port

   Identifies the outbound port used for initiating the action.

   Destination Port

   Identifies the inbound port on the destination host.

Receiver Host

Identifies the outbound port on the agent used to send event logs to the CA Enterprise Log Manager server.

Category

Identifies the high-level category of the corresponding event action. For example, System Access is the category for the Authentication action.

Action

Identifies the event action.

Log Name

Identifies the log name used by the connector that collected the event.

Use the User Prompt

Each event expresses information about two actors: the Source and the Destination.

- The Source actor initiates the action that causes the event.

  The source actor can be a user, source_username, or a process, source_processname.

- The Destination or "dest" actor is the target of the action.

  The destination actor can be a user, dest_username, or an object, dest_objectname.

The User prompt queries for events where the actor you specify appears in the selected CEG fields of the refined event. Consider this scenario:

1. The source actor, source_username or source_processname attempts an action on the target actor, destination_username or a destination_objectname.
2. This event is recorded in a repository on the event source.
3. A CA Enterprise Log Manager agent makes a copy of the event recorded on the event source and transmits it to a CA Enterprise Log Manager server.

To use the User prompt

1. Select Queries and Reports.

   The Query List displays the Prompts folder and one or more folders for other queries.

2. Expand Prompts and select User.

   The User prompt appears.

3. Enter the name of the user on which to base this query.

4. Select the fields on which to query for data matching your user name entry.

   source_username

   Is the name of the user that initiated the event action.

   dest_username

   Is the name of user that is the target of the action.

   source_objectname

   Is the name of the object involved in the action referenced in event information.

   dest_objectname

   Is the name of the object that is the target of the action.

5. Click Go.

   Results of the User prompt query appear.

6. Use the following descriptions to interpret the query results:

   CA Severity

   Indicates the severity of the event, where the values in increasing order of severity include: Information, Warning, Minor Impact, Major Impact, Critical, and Fatal.

   Date

   Indicates when the event occurred.

   Destination Host

   Identifies the name of the host with the user who was the target of the event action.

   Result

   Specifies a code for the event result of the corresponding action, where S means Success, F means Failure, A means Accepted, D means Dropped, R means Rejected, and U means Unknown.

Source User

Identifies the user who initiated the event action.

Source Object

Identifies the object on the source host that was involved in the event action.

Destination User

Identifies the user who was the target of the event action.

Destination Object

Identifies the object on the destination host that was involved in the event action.

Category

Identifies the high-level category of the corresponding event action. For example, System Access is the category for the Authentication action.

Action

Identifies the event action.

Log Name

Identifies the log name used by the connector that collected the event.

How to Create a Query

You can create custom queries using the Query Design wizard. When you create a query you must choose whether it applies to the event database or the incident database. A server's event database stores information for all events received by that server. A server's incident database stores information on incidents and elements of their component events as specified by correlation rules.

You can also delete custom queries and export query information, or copy a subscription query to create a custom query and then edit that query using the query design wizard. Only users with the Administrator or Analyst roles can create, delete, or edit queries.

Creating a query using the query design wizard involves the following steps:

1. Opening the query design wizard.
2. Adding identity and tag details.
3. Selecting query columns.
4. (Optional) Setting query conditions and filters.
5. Setting date range and result conditions.
6. (Optional) Choosing visualization options for the query display.
7. (Optional) Adding drill-down values for the query.

More information:

Open Query Design Wizard (see page 245)
Using Advanced Filters (see page 408)
Create a Query Display Visualization (see page 257)
Add a Drill Down Report (see page 258)

Open Query Design Wizard

To create a new custom query, create a copy of a query, or edit an existing query, you must open the query design wizard.

To open the query design wizard

1. Click the Queries and Reports tab, then the Queries subtab.

   The Query List appears.

2. Click Options and select New.

   The query design wizard appears.

   When using the wizard:

   - Click Save to save the query without closing the wizard.
   - Click Save and Close to save the query and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

More information:

How to Set Result Conditions (see page 410)
Create a Query Display Visualization (see page 257)
Add a Drill Down Report (see page 258)

Add Query Details

The first step in creating a query is entering identifying information and setting any tags you want to include.

To add a new query

1. Open the query design wizard.
2. Type a required query name, and optional short name for use in reports. The short name appears in the report's individual query pane when the query is included in a report.
3. Select the database you want your query to apply to :

   Event

   Applies the query to the event database, which stores all raw and refined event information received by the current server, or available through federation.

   Incident

   Applies the query to the incident database, which stores incidents created by the event correlation system, and event information used to create those incidents. The specific components of an event that are used to create an incident, and thus stored in the incident database, are set by the correlation rule.
4. Type any design notes you want in the Description entry field.

   Note: We recommend using this field for information about the query structure. For example, it could contain an explanation of why the query contains certain fields and function.
5. Select one or more tags that you want your query to be associated with using the Tags shuttle control.
6. (Optional) To add a custom category tag, enter a tag name in the Add Custom Tag entry field, and click the Add Tag button.

   The custom Tag appears, already selected, in the Tags shuttle control.
7. (Optional) To add one or more nested custom tags, select a tag, or type the name of the parent category tag, followed by a backslash, followed by the name of the child tag, then click Add Tag. For example, you could type: "Regulations\Industry Standards". You can add additional tags, maintaining the format: a\b\c and so on.

   Note: If you delete one of the custom nested tags, all the custom tags in which it is nested are also deleted, including the parent tag. If you nest a custom tag inside a subscription tag, and then delete it, only the custom tags are deleted.

   When you complete the process, the new tags appear in the list, with the nested custom tags visible when you expand the parent tag.
8. Click the appropriate arrow to advance to the query design step you want to complete next, or click Save and Close.

If you click Save and Close, the new query appears in the Query List, otherwise the query design step you choose appears.

Add Query Columns

To create a query, write a SQL statement that retrieves the event information you want from the event log store. The query design wizard helps automate this process.

To create a query SQL statement

1. Open the query design wizard.
2. Enter the name and tag, if not already specified, then advance to the Query Columns step.
3. (Optional) Select the Unique events only check box.
4. Set the CEG columns you want to query by dragging them from the list of Available Columns on the left into the Column field of the Selected Columns pane. They appear in the query display in the order in which they are entered.
5. (Optional) Select the settings you want for each column, including:

   Display Name

   Lets you enter a different name for the column, when it is displayed in Table or Event Viewer format. If you enter no Display Name, the native field name is used as the column name, "event_count" for example.

   Function

   Lets you apply one of the following SQL functions to the column values:

   - COUNT - returns the total number of events.
   - AVG - returns the average of the event_count values. This function is only available for event_count fields.
   - SUM - returns the sum of the event_count values. This function is only available for event_count fields.
   - TRIM - Removes any spaces in the queried text string.
   - TOLOWER - Converts the queried text string to lowercase.
   - TOUPPER - Converts the queried text string to uppercase.
   - MIN - returns the lowest event value.
   - MAX - returns the highest event value.
   - UNIQUECOUNT - returns the number of unique events.

Group Order

Sets the query display to show the selected columns grouped by the selected attribute. For example, you can set the query to group events by source name. You can control the order in which it is applied to various columns. If the first column values are identical, the second are applied. For example, you can group multiple events from the same source by username.

Sort Order

Controls the order in which the selected value is sorted. You can control the order in which it is applied to various columns. If the first column values are identical, the second are applied.

Descending

Sets the column values to display in descending order (highest to lowest value) rather than the default ascending order.

Not Null

Controls whether the row is displayed in a table or Event Viewer if it contains no value. Selecting the Not Null check box removes the row from the query result if it contains no displayable value.

Visible

Controls whether the column is visible in a table or Event Viewer format. You can use this setting to make the column data available in the details view without showing it in the display itself.

Note: If you select a Function except TRIM, TOLOWER, TOUPPER or a Group Order setting for a column, you must select the same setting for other columns too. Otherwise, CA Enterprise Log Manager displays error messages.

6. (Optional) Use the up and down arrows at the top of the Selected Columns pane to change the column order as needed.
7. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

   If you click Save and Close, the new query appears in the Query List, otherwise the Query Design step you choose appears.

Set Query Filters

You can filter the information returned by your query using simple or advanced filters. Simple filters let you create single-term filter statements easily and quickly. Advanced filters let you build more complex SQL language statements, including nested statements.

To set query filters

1. Open the query design wizard.
2. Enter the name and tag, if not already specified, then advance to the Query Filters step.

   The query filters dialog appears, displaying the Simple Filters Tab.
3. Create any simple filters you want, to search for stated CEG field values.
4. (Optional) Click the Advanced Filters tab.
5. (Optional) Create any advanced filters you want.
6. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

   If you click Save and Close, the new query appears in the Query List, otherwise the Query Design step you select appears.

More information:

Create a Simple Event Filter (see page 479)
Create an Advanced Event Filter (see page 482)
Using Advanced Filters (see page 408)

Create a Simple Event Filter

You can create simple filters to set search parameters for common values. For example, you could set the Ideal Model field to "Content Management" to identify all events with that value in the Ideal Model CEG field. Simple filters are used by many features, including event and incident queries, suppression and summarization rules, and event forwarding rules.

To create a simple filter

1. Select the check box for any simple filter field or fields you want to define, and select a value from the drop-down list, or enter the value you want in the text entry field.
2. Click Save when you have added all the simple filters you want.

More information

Using Advanced Filters (see page 408)
Create an Advanced Event Filter (see page 482)

Using Advanced Filters

You can use SQL-based advanced filters to qualify any function that queries the event or incident databases, including narrowing queries, or adding additional qualifications to simple filters. The Advanced Filters interface helps you create the appropriate filter syntax by providing a form for entering logic columns, operators and values according to your filtering requirements.

Note: This section contains a brief overview of the SQL terms used in advanced filters. To use advanced filters to their full potential you need a thorough understanding of SQL and the Common Event Grammar.

The following SQL terms join multiple filter statements:

And

Displays the event information if *all* the joined terms are true.

Or

Displays the event information if *any* of the joined terms are true.

Having

Refines the terms of the main SQL statement by adding a qualifying statement. For example, you could set an advanced filter for events from specified hosts, and add a "having" statement to return only events of a specified severity level from those hosts.

The following SQL operators are used by advanced filters to create the basic conditions:

Relational Operators

Include the event information if the column bears the appropriate relation to the value you enter. The following relational operators are available:

- Equal to
- Not Equal to
- Less than
- Greater than
- Less than or equal to
- Greater than or equal to

For example, using *Greater than* would include the event information from your chosen column if its value is greater than the value you set.

Like

Includes the event information if the column contains a pattern you enter, using % to set the pattern you want. For example, L% would return any values beginning with L, %L% would return any values with L included as neither first nor last letter.

Not like

Includes the event information if the column does not contain the pattern you specify.

In set

Includes the event information if the column contains one or more of the values in the quote-delineated set you enter. Multiple values in the set must be comma-separated.

Not in set

Includes the event information if the column does not contain one or more of the values in the quote-delineated set you enter. Multiple values in the set must be comma-separated.

Matches

Includes any event information that matches one or more of the characters that you enter, allowing you to search for key words.

Keyed

Includes any event information that is set as a key value during Report Server configuration. You can use key values to set business relevance or other organizational groups.

Not Keyed

Includes any event information that is not set as a key value during Report Server configuration. You can use key values to set business relevance or other organizational groups.

Create an Advanced Event Filter

Advanced filters are used by many features, including query creation, report scheduling,alert jobs, and local and global filters.

To create an advanced filter

1. If you are creating a scheduled report job or action alert job, click the Events or Incidents tab to set the appropriate filter type. Since a report or alert job may contain both event and incident queries, you can set the filter types separately.
2. Click New Event Filter.

   The first row of the event filter table becomes active, and its Logic and Operator columns are populated with the default values "And" and "Equal to" respectively.
3. Click the Logic cell and change the logic value as needed.
4. Click the Column cell, and select the event information column you want from the drop-down menu.
5. Click the Operator cell, and select the operator you want from the drop-down menu.
6. Click the Value cell, and enter the value you want.
7. (Optional) Click the open and closed parentheses cells and enter the number of parentheses you need.
8. (Optional) Repeat steps 1 through 6 as needed to add additional filter statements.
9. Click Save when you have entered all the filter statements you want.

More information:

Create a Simple Event Filter (see page 479)
Using Advanced Filters (see page 408)

How to Set Result Conditions

You can set a date range and other result conditions for the query, including row limits and base display time period. Result conditions can be altered at any time up to the query's run time, making them a useful way to modify queries without altering the base query or its filters.

If you are creating a scheduled report job or action alert job, you can set result conditions for both event and incident queries that make up the job as needed.

You can set the following types of result conditions:

- Date range conditions governing the query's search period
- Display conditions, such as maximum rows
- Grouped Event conditions, such as the most recent grouped events after a given date, or grouped events containing a set number of events.

  Note: If you do not group at least one column when creating a query, users will not be able to edit result conditions from the query display.

Set a Time or Date Range

You can set a time or date range condition for your query. This improves the efficiency of your query by narrowing the portion of the event log store it must search.

You can use a predefined time range, or create a custom time range. For a custom time range to work properly you must set both a beginning and end time. If you only set a single time parameter, it is expressed as a "Where" clause in the query SQL.

To set result conditions

1. Open the result conditions dialog.
2. If you are creating a scheduled report job or action alert job, click the Events or Incidents tab to set the appropriate filter type. Since a report or alert job may contain both event and incident queries, you can set the filter types separately.
3. Select a predefined time range from the drop-down list. For example, if you want to view events received in the last day, select "previous day".

   Note: If you are creating an action alert or scheduled report, the interface displays the following default time ranges:

   - Action Alert: the previous 5 minutes
   - Scheduled Report: the previous 6 hours
4. (Optional) Create a custom time range using the following substeps:

   a. Click Edit beside the 'Dynamic End Time' entry field in the Date Range Selections area. This lets you set the end of the time period you want the query to search.

      The Dynamic Time Specification dialog appears.

   b. Select the reference time you want to base the parameter on, and click Add.

   c. Select the time parameter you want, and click Add. You can add multiple time parameters.

   d. When you are finished adding parameters, click OK.

      The Dynamic Time Specification dialog closes, and the values you choose appear in the 'Dynamic End Time' area. If you use multiple parameters, they form a complete time statement, with each parameter referring to the first. For example, adding the 'Start of the Month,' and 'Day of the Week - Tuesday' values in the 'Dynamic End Time' area will end your query on the first Tuesday of the month.

      Note: When using the 'Number of' values, such as 'Number of days' or Number of hours' you must enter a *negative* number to set a time in the past. Using a positive number will set a future end time, and cause the query to continue sending results as long as at least one qualified event is in the log store.

For example, adding the 'now,' and 'number of minutes -10' values to the 'Dynamic Start Time' area starts your query 10 minutes before the selected end time.

e. Repeat step 2 in the 'Dynamic Start Time' area to set the beginning of the time period you want the query to search.

If you do not enter a date range, the query is applied all events in the log store. If you enter an invalid date range, your query might not return any results.

5. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

   If you click Save and Close the new query appears in the Query List, otherwise the Query Design step you choose appears.

More information:

How to Set Result Conditions (see page 410)
Set Display and Group Conditions (see page 413)

Set Display and Group Conditions

You can set conditions that allow you to control the query display and conditions that search for events based on how they are grouped.

To set display and group conditions

1. Open the result conditions dialog.
2. If you are creating a scheduled report job or action alert job, click the Events or Incidents tab to set the appropriate filter type. Since a report or alert job can contain both event and incident queries, you can set the filter types separately.
3. Use the Results check boxes to enable any of the following display qualifications you want:

   Default Query Limit

   This value is only available for Action Alerts and Report Scheduling. It sets the alert or report job to use the row limit of the individual queries in the job. If you select any other results value when creating a job, CA Enterprise Log Manager overrides the row limits in the component queries.

   Row Limit

   Sets the maximum number of event rows that the query displays, starting with the most recent events.

   No Limit

   Sets the query to retrieve all events that match its filter. This can include a large number of events so you should plan the query accordingly.

   Show Other

   Indicates the presence of other results that are not displayed due to the row limit, allowing you to compare the selected events in the context of all events of that type. For example, if you choose a row limit of 10 for your event viewer display and select show other, events beyond 10 are displayed as a single entry titled Other, showing all remaining events. This setting is only effective when row limit is selected.

   Time Granularity

   Sets the detail level of the time period field used in the query display.
4. Use the Result Conditions to query for various types of grouped event conditions. For example you could set your query to search for the latest grouped event after a selected date, or a certain number of grouped events. A grouped event is a refined event for which you have set a Function and Group Order in the Query Creation step.

   The group conditions use the same time statement system as the time range fields.
5. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

If you click Save and Close the new query appears in the Query List, otherwise the Query Design step you choose appears.

More information:

How to Set Result Conditions (see page 410)

Create a Query Display Visualization

To create a new query display you must set the Visualization details that control how the event information appears.

To create a query display visualization

1. Open the query design wizard.
2. Enter the name and tag, if not already specified, then advance to the Visualization step.
3. Choose whether you want your query display to use an Event Viewer or Chart.

   If you choose Event Viewer, the visualization step is complete. The event columns appear in the Event Viewer display in the order in which you placed them during the Query Columns construction step.
4. If you choose a Chart, you can select one or more chart types. Selecting multiple chart types allows users to toggle back and forth between them in the report display. The up and down arrows that appear beside each type control the order in which they appear in the Change Visualization menu.

   Note: Table is always available as a visualization even if you do not add it during this step.
5. Select the event you want to appear as the X (horizontal) Axis from the column drop-down, enter label text if you want any to appear, and select one of the following options from the display type menu:

   - Category - Use this option for string or text value columns, such as source_username.
   - Linear - Use this option for numeric values, such as event_count. When the values are widespread, you can use the Log Axis check box to allow the axis to be logarithmic
   - Datetime - Use this option to display time values in the local date/time.
6. Repeat Step 4 using the Y-Axis Settings menus to set the Y (vertical) Axis column, label, and type options.
7. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

   If you click Save and Close the new query appears in the Query List, otherwise the Query Design step you choose appears.

Add a Drill Down Report

You can add one or more drill down reports to your query. Drill down reports let users click a query display element and display another related report.

To add a drill down report

1. Open the query design wizard.
2. Enter the name and tag, if not already specified, and then advance to the Drill Down step.
3. Click Add Drilldown.
4. Enter the name or browse for the report you want to make available as a drill down item.
5. Select one or more available parameters on which to focus the drill down report and move them to the Selected Parameters list. The drill down reports use the selected parameters to preserve your query focus.
6. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

   If you click Save and Close the new query appears in the Query List, otherwise the Query Design step you choose appears.

Edit a Query

You can edit existing custom queries. You cannot edit a subscription query; however, you can copy a subscription query and edit your copy. If you edit a query, the changes you make affect any reports using that query.

To edit a query

1. Click the Queries and Reports tab and the Queries subtab.

   The Query Tag Filter list and the Query List appear.
2. Expand the User folder in the Query List and select the query you want to edit.
3. Click Options at the top of the list, and select Edit.

   The query design wizard appears, populated with the specifications of the query you selected.
4. Make the changes you want, and click Save.

Delete a Custom Query

You can delete a custom query. You cannot delete a subscription query.

To delete a query

1. Select the query you want to delete.
2. Click Options at the top of the list, and select Delete.

   A confirmation dialog appears.
3. Click Yes.

   The deleted query is removed from the Query List.

More information:

Disable Show Selected Query (see page 259)
Edit a Query (see page 258)
Exporting and Importing Query Definitions (see page 260)

Disable Show Selected Query

You can set your query list so that you can make changes without loading queries. Normally, selecting a query from the list displays it in the details window.

Disabling this default mode saves time by letting you select a query from the list and edit it immediately, without waiting for it to display. This is especially useful if you have multiple queries to edit and already know what changes you plan to make.

Note: Because only users with the Administrator or Analyst mode can create or edit queries, only these users can disable the show selected query setting.

To disable show selected query

1. Click Options at the top of the Query List.

   The Options menu appears.
2. Clear the check beside Show Selected Query.

   Any query selected from the list is not displayed until Show Selected Query is re-enabled.

Exporting and Importing Query Definitions

You can export and import details of custom queries for use in other management servers. This lets you transfer successful custom queries between CA Enterprise Log Manager environments, or from a test to a live environment.

More information:

Import Query Definitions (see page 261)
Export Query Definitions (see page 260)

Export Query Definitions

You can export the details of user-created queries for use in other management servers. The export is saved as an XML file.

To export query details

1. Click the Queries and Reports tab, and then click the Queries subtab.

   The Query List appears.

2. Click Options at the top of the list, and select Export.

   The Export User Query Definitions dialog appears, displaying available user-created reports.

3. Select the query or queries you want to export using the shuttle control, and click Export.

   An export dialog appears.

4. Enter or browse for the location you want to save the XML export files, and click Save.

   The query files are saved to your chosen location and a confirmation dialog appears.

5. Click OK, and then Close.

   The Export User Query Definitions dialog closes.

More information:

Exporting and Importing Query Definitions (see page 260)
Import Query Definitions (see page 261)

Import Query Definitions

You can import query definition XML files for use in the local management server.

To import report details

1. Click the Queries and Reports tab, and then the Queries subtab.

   The Query List appears.

2. Click Options at the top of the list, and select Import.

   The Import File dialog opens

3. Enter or browse for the location of the file you want to import, and click OK.

   The Import Results window appears.

4. Click Import Another File to repeat step 3, or Close.

   The Import Results window closes.

More information:

Exporting and Importing Query Definitions (see page 260)
Export Query Definitions (see page 260)

How to Create a Report

You can create custom reports for your environment, either by using the process outlined in this section to create an entirely new report, or by using a pre-defined report as a template. You can view custom reports or set them as scheduled report templates.

You can also edit or delete custom reports, and export report information. You can perform these customization tasks only if you are logged in as a user with the Administrator or Analyst roles.

The process of creating a new report using the report design wizard has the following steps:

1. Opening the report design wizard.
2. Adding report details - Naming the new report and assigning category tags.
3. Designing a report layout - choosing which queries are included in the report and how they will be displayed.

Open Report Design Wizard

To create a new custom report from scratch or based on an existing report, you must open the report design wizard.

To open the report design wizard

1. Click the Queries and Reports tab, then the Reports subtab.

   The Reports List appears.

2. Click Options, and then select either New or Copy.

   The Report Design wizard appears.

   When using the wizard:

   - Click Save to save without closing the wizard.
   - Click Save and Close to save the report and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Add Report Details

You can create a new report from scratch or from a copy of an existing report. When you create a report, you name it and add any subscription or custom tags you want to associate with it.

To add report details

1. Open the report design wizard.
2. Enter a report name. You may also enter optional description information for reference.
3. Select one or more tags that you want the report to be associated with using the Tags shuttle control.
4. (Optional) To add a custom category tag, enter a tag name in the Add Custom Tag entry field, and click Add Tag.

   The custom Tag appears as in the Selected Tags list.
5. (Optional) To add one or more nested custom tags, select a tag, or type the name of the parent category tag, followed by a backslash, followed by the name of the child tag, then click Add Tag. For example, you could type: "Regulations\Industry Standards". You can add additional tags, maintaining the format: a\b\c and so on.

   Note: If you delete one of the custom nested tags, all the tags in which it is nested are also deleted, including the parent tag. If you nest a custom tag inside a subscription tag, and then delete it, only the custom tags are deleted.

   When you complete the process, the new tags appear in the list, with the nested custom tags visible when you expand the parent tag.
6. Advance to the Layout step or click Save and Close if at least one query has already been selected.

Design Report Layout

You can design your report structure by specifying the grid size and dimensions and then selecting the queries to display in each section of the grid.

To design a report layout

1. Open the report design wizard. If this is a new report, enter a name, select a tag, and advance to the Layout step.
2. Select or enter the number of rows and columns you want to appear in your report, using the Grid Rows and Columns areas in the Report Layout pane. These settings control the number of query display areas the report contains. You may include up to 10 rows and/or columns.

   The appropriate number of rows, columns, and corresponding query displays appears in the report layout pane.

   Note: You can use the arrows at the right side and bottom of individual query display areas to expand or shrink them horizontally or vertically as needed.
3. (Optional) Enter or select a minimum pixel size for the query display areas in the Min. Width and Min. Height areas.
4. Drag the query you want to display in each display area from the Query List to the appropriate area in the report layout.
5. (Optional) Click Edit at the top of each query display area to edit the query you have placed there or create a new custom query.
6. Click Save and Close.

   The Report Design wizard closes. The new report appears in the Report List under the User folder.

Example: Create a Report from Existing Queries

You can create custom reports composed of predefined queries and tailor it to your specifications.

To create a report from existing queries

1. Identify the queries to include in the custom report.
   a. Click the Queries and Reports tab and the Queries subtab, if not displayed.
   b. Enter a key word or key phrase in the Search field to display the queries with the content from which you want to make a selection. For example, enter critical hosts trend.
   c. Note the names of the queries you want to include in the custom report. For example, you can define a report of the trends associated with business critical hosts from those listed in the following illustration, for example, the ones for System Access, Resource Access and Account Creations.

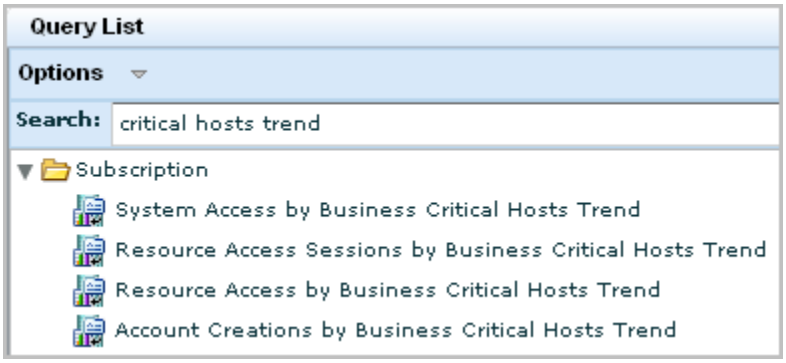

2. For the first query to include in the report, create a copy and add a custom tag.
   a. Select a query and select Copy from the Options drop-down list.
   b. Rename the query and enter a custom tag to add. For example, rename Copy of System Access by Business Critical Hosts Trend to Custom System Access by Business Critical Hosts Trend.

c. Add a custom tag. For example, enter Critical_Assets_Trend and click Add Tag.

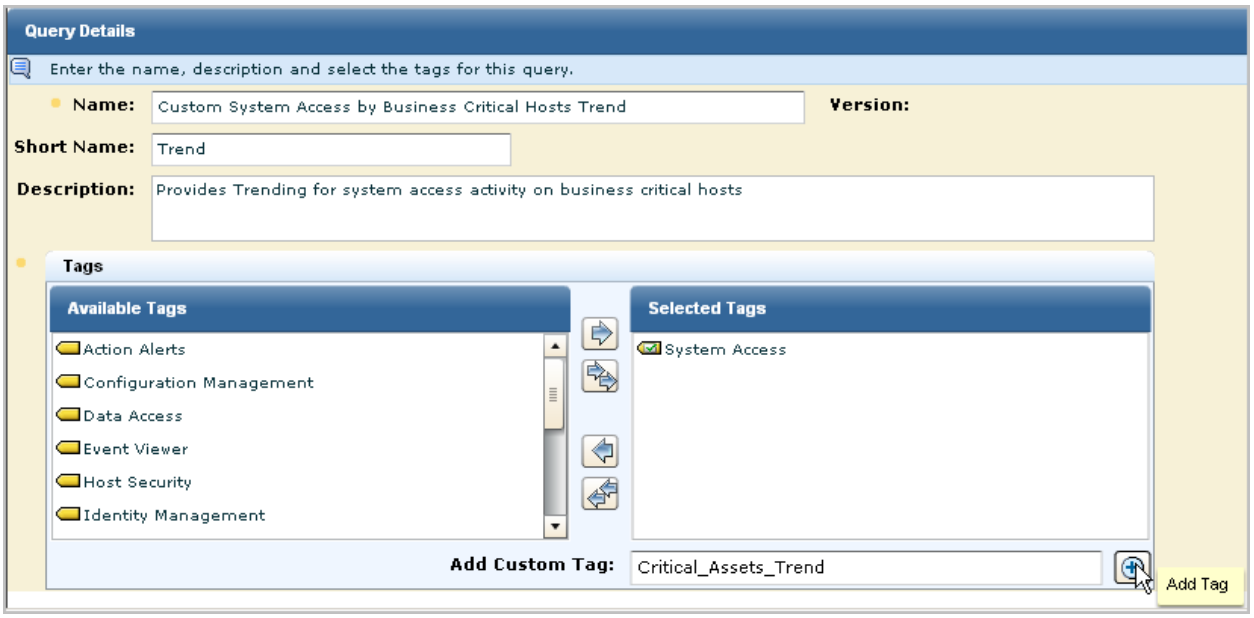

d. Click the move button to move the preselected tag to the Available Tags area. For example, move System Access. The only tag selected is the one you added.

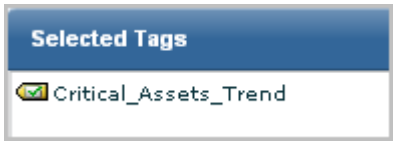

e. Click Save and Close.

3. For the other queries to include in the report, create a copy and select the custom tag you created.

   a. Select a query and select Copy from the Options drop-down list.

   b. Rename the query and select the new custom tag. For example, rename Copy of Resource Access by Business Critical Hosts Trend to Custom Resource Access by Business Critical Hosts Trend, move Critical_Assets_Trend to the Selected Tags list and remove the preselected tag.

   c. Click Save and Close.

   The copied queries display under User:

4. If the queries are associated with keyed list, define the values for that keyed list.
5. Initiate the report creation process as follows:
   a. Click the Queries and Reports tab and then the Reports subtab.
   b. Select New from the Options drop-down list under the Report List.

      The Report Design wizard appears.

      Add the custom tag, Critical_Assets_Trend.
6. Design the report layout.

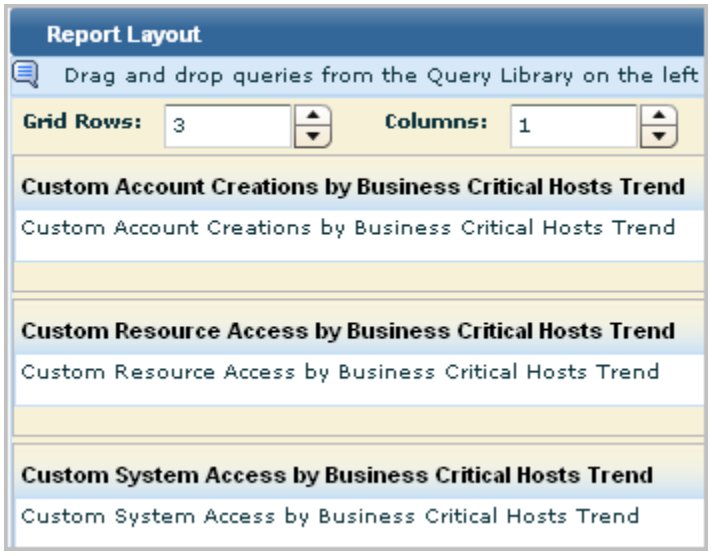

7. Click Save and Close.
8. Schedule the report based on the custom tag you created.
9. View the report.

Note: It is good practice to examine any new report to verify that it provides the desired information in the best possible way.

Example: Set Up Federation and Federated Reports

You can collect logs from high-volume, geographically separate data centers and set up reporting so that distributed data is queried from just one of the data centers.

Consider an example scenario where the two high-volume data centers are located in New York and Virginia, where New York is the corporate headquarters. Each data center has a collection server that collects and processes incoming event logs and sends them to its reporting server. The reporting server handles queries, alerts, and reports. Most queries, alerts, and reports target event data collected through agents; consolidating data from these event sources requires federation among reporting servers and collection servers.

Some queries, alerts, and reports target self-monitoring events generated by CA Enterprise Log Manager servers; consolidating this type of data requires inclusion of the management server in the federation. If consolidating self-monitoring event data is not desired, the management server can be excluded from the federation. Self-monitoring events from this server can be monitored with non-federated local reports. For simplicity, the management server is excluded in this federation; inclusion could be achieved by creating a meshed federation between the NY-Reporting-ELM and Management-ELM.

The server names are as follows:

- Management-ELM
- NY-Collection-ELM
- NY-Reporting-ELM
- VA-Collection-ELM
- VA-Reporting-ELM

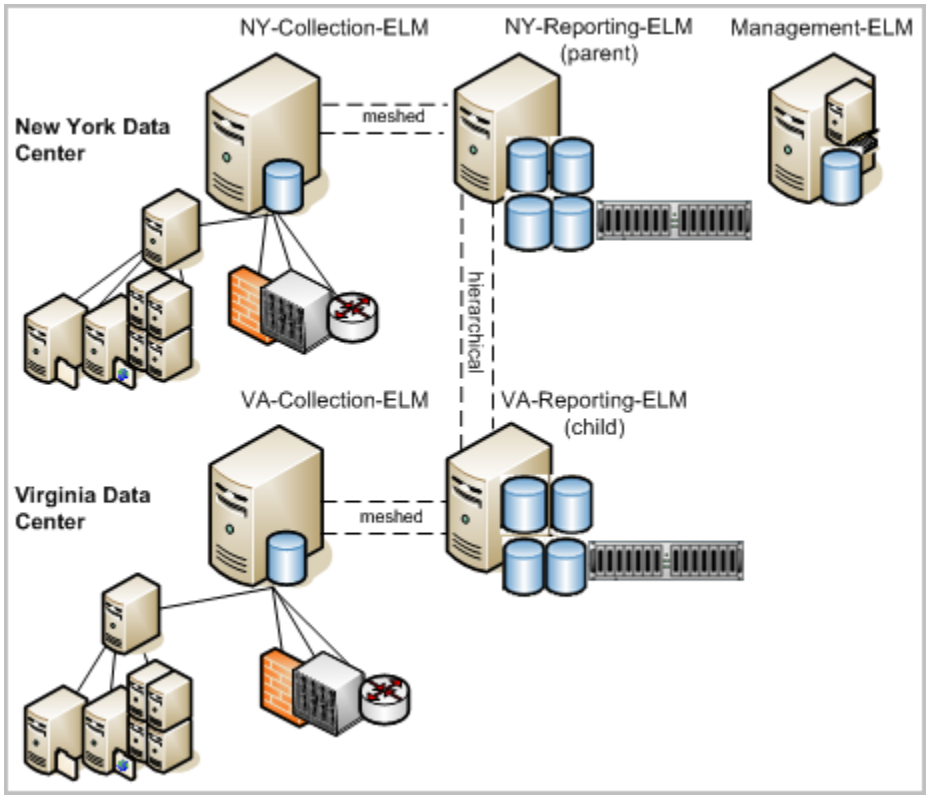

Assume the Administrator in New York wants all reports and alerts that are run from the New York site to include data from the Virginia site, but wants all reports and alerts run from the Virginia site to include only locally collected data.

The following example shows how to federate the servers and configure reporting to meet the criteria for this scenario. Procedures for configuring auto-archiving are not included in this example, but auto-archiving should be configured for any high-volume architecture.

1. Log into a CA Enterprise Log Manager with Administrator credentials.
2. Click the Administration tab and select the Services subtab.
3. Create a hierarchical federation, where NY-Reporting-ELM is the parent and VA-Reporting-ELM is the child as follows:
   a. Expand the Event Log Store service, and then select the server name that is to be the parent in the hierarchical federation, in this case, NY-Reporting-ELM.

   b. Select VA-Reporting-ELM from the available federation children list and move it to the selected list.

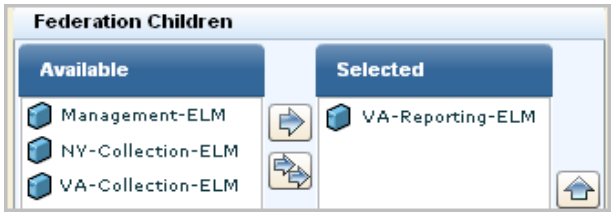

4. Create a meshed federation between the NY-Reporting-ELM and the NY-Collection-ELM as follows, where each is a child of the other:
   a. Select the NY-Reporting-ELM from the Event Log Store list.
   b. Select NY-Collection-ELM from the available federation children and move it to the selected list.
   c. Select the NY-Collection-ELM from the Event Log Store list.
   d. Select NY-Reporting-ELM from the available federation children and move it to the selected list.

5. Create a meshed federation between the VA-Reporting-ELM and the VA-Collection-ELM as follows, where each is a child of the other:
   a. Select the VA-Reporting-ELM from the Event Log Store list.
   b. Select VA-Collection-ELM from the available federation children and move it to the selected list.
   c. Select the VA-Collection-ELM from the Event Log Store list.
   d. Select VA-Reporting-ELM from the available federation children and move it to the selected list.
6. Configure global report server settings and local overrides for VA-Reporting-ELM as follows. Geographically distant servers often use different mail servers.
   a. Select Alerting Service on the Service List
   b. Configure global or local settings as needed for mail server options from the NY-Reporting-ELM node.
   c. If you plan to email reports, select Report Server and then the NY-Reporting-ELM node.
   d. Set global or local PDF format options, or report options related to report and alert retention.
7. For each report scheduled to run from NY-Reporting-ELM, do the following:
   a. Select the Scheduled Reports tab and the Report Scheduling tab.
   b. Click Schedule a Report.
   c. Select the report to schedule and complete steps 2, 3, 4, and 5 as needed.
   d. Click the Server Selection step, select NY-Reporting-ELM from the available servers list and move it to the selected servers list and then accept the default, Yes, for federated query.
   e. Click Save and Close.

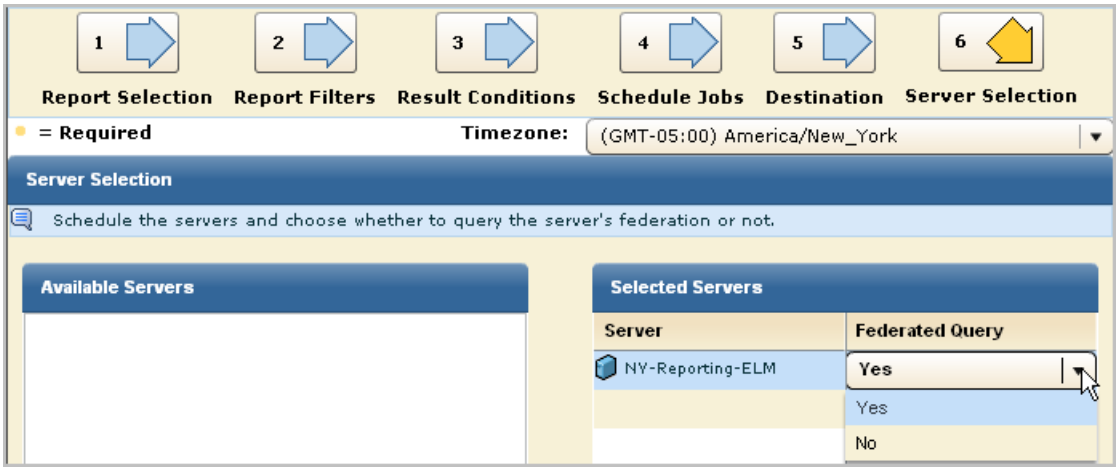

The resulting reports include data from NY-Reporting-ELM, its peer, NY-Collection-ELM, its child, VA-Reporting-ELM, and its child's peer, VA-Collection-ELM.

Note: A federated query run from VA-Reporting-ELM includes data from VA-Reporting-ELM and its peer VA-Collection-ELM. It does not include data from NY-Reporting-ELM, since this server is its parent in the hierarchical federation.

Edit a Report

You can edit a custom report.

Note: You can disable the Show Selected Report option when editing multiple reports. This lets you select and edit reports without waiting for them to display in the details pane.

To edit a report

1. Select the report you want to edit from the Report List.
2. Click Options at the top of the list, and select Edit.

   The Report Design wizard appears, populated with the specifications of the report you selected.
3. Make the changes you want and then click Save and Close.

   The edited report appears in the Report List under User folder.

Delete a Custom Report

You can delete a custom report. You cannot delete a subscription report.

To delete a custom report

1. Select the custom report you want to delete from the Report List.
2. Click Options at the top of the list, and select Delete.

   A confirmation dialog appears.
3. Click Yes.

   The deleted report is removed from the Report List.

More information:

How to Create a Report (see page 261)
Edit a Report (see page 272)

Example: Delete Daily Reports More Than 30 Days Old

You can implement policies on report retention through the global configuration of report servers. You can set a different retention policy for each schedule report recurrence, that is,

- One-time Reports Retention
- Daily Reports Retention
- Weekly Reports Retention
- Monthly Reports Retention
- Yearly Reports Retention

You must change the default of Never Runs for the Reports Retention utility to a frequency. Be sure the frequency with which you set the utility to run is often enough to do the deletions at the frequency you configure. For example, if you want to delete your daily reports 1 day after they run, and you schedule daily reports to be run at 6 a.m. and 6 p.m., you would set the reports retention utility to run every 12 hours at the minimum.

Example: Delete all Daily Reports Older Than 30 Days

1. Click the Administration tab and the Services subtab.

   The Service List shows services by service.

2. Click Report Server

   The Global Service Configuration: Report Server appears.

3. Use the following guidelines to complete this configuration:

   - To automate the deletion of all daily reports 30 days after they are generated, set the daily report retention to delete after 30 days.
   - Be sure to set the Reports Retention utility to run every specified number of hours, days, or weeks.

4. Click Save.

Export Report Definitions

You can export the details of user-created files for use in other management servers. The export is saved as an XML file. An exported report definition includes the definitions for all the queries in that report.

To export report details

1. Click the Queries and Reports tab, and then the Reports subtab.

   The Report List appears.

2. Click Options at the top of the list, and select Export.

   The Export User Definitions dialog appears, displaying available user-created reports.

3. Select the report or reports you want to export using the shuttle control, and click Export.

   An export dialog appears.

4. Enter or browse for the location you want to save the XML export files, and click Save.

   The Report files are saved to your chosen location and a confirmation dialog appears.

5. Click OK, and then Close.

   The Export User Report Definitions dialog closes.

More information:

Import Report Definitions (see page 275)

Import Report Definitions

You can import report definition XML files for use in the local management server.

To import report details

1. Click the Queries and Reports tab, and then the Reports subtab.

   The Report List appears.

2. Click Options at the top of the list, and select Import.

   An Import File dialog opens

3. Enter or browse for the location of the file you want to import, and click OK.

   The Import Results window appears.

4. Click Import Another File to repeat step 3, or Close.

   The Import User Report and Query and Query Definitions window closes.

More information:

Export Report Definitions (see page 274)

Preparing to Use Reports with Keyed Lists

All reports are built from one or more queries. Some queries used in predefined reports are designed to select all values from a given table where a certain attribute field contains a value used as criteria for compiling the list of key values. For example, an assets table has an IsCritical field. A query that selects all asset names from the asset table where IsCritical equals Yes, would select only the names of critical assets. These names could be returned to CA Enterprise Log Manager to refresh the values for the Critical_Assets key.

Preparing to use predefined reports with keyed lists involves the following tasks:

- (Optional) Enabling dynamic values import, if you use CA IT PAM.
- Creating keyed lists for predefined keys that have no predefined values.
- Customizing keyed lists for predefined keys that have predefined values.
- Maintaining keyed lists used in the predefined reports you want to use. Update each keyed list with current values.

In addition, you can add new keys for custom reports that use keyed lists and then add values for each new key. You can also add values to the Business_Critical_Sources key and the ELM_System_Lognames key for on demand queries of your own design.

Enabling Dynamic Values Import

The procedures required to enable dynamic values import apply only to CA IT PAM users.

If you use CA IT PAM and have existing tables or spreadsheets where you keep lists of files, databases, hosts, and users, for example, you can leverage this data. You can create a process that reads the table or file, selects the values pertinent to the key, and returns those values to the CA Enterprise Log Manager values list for that key.

To import dynamic values

1. Create a process in CA IT PAM for each key values list that you want to generate on demand.

   Note: If any process reads a database table, install a CA IT PAM agent on the server with the SQL Server 2005 database.

2. Configure CA IT PAM integration for dynamic values in CA Enterprise Log Manager.

More information:

Create a CA IT PAM Process to Generate a Values List (see page 277)
Configure CA IT PAM Integration for Dynamic Values (see page 277)

About Dynamic Values Processes

A dynamic values process is a CA IT PAM process that you can invoke to populate or update the values list for a selected key that is used in reports or alerts. The assumption is that you are already storing master lists of the files, databases, hosts, users, and so forth that make up your work environment and that this master list is designed with attributes that allow you to query for sets of values of interest. If you use CA IT PAM, you can create processes that can be invoked to run the queries that return the data to CA Enterprise Log Manager for use as key values in the reports and alerts based on keys. Being able to dynamically create a values list is a useful way to keep volatile key list updated with current values.

Create a CA IT PAM Process to Generate a Values List

You can create a process in CA IT PAM for each key values list that you want to be able to generate on demand. Use your CA IT PAM documentation for details on process creation. Each process must meet CA Enterprise Log Manager requirements regarding InputKey, the ValueList and FaultString local process parameters, and the Success and Failure calculation operators.

Use the following guidelines:

- The process must accept the selected key as InputKey.
- The process must define the following two local process parameters:
  - ValueList retrieves the value list
  - FaultString retrieves the error string

  Note: CA Enterprise Log Manager requires these exact parameter names be used as output interface parameters.
- The process must contain the following two calculation operators:
  - Success calculation operator: Process.ValueList =<variable containing comma-separated list of values>
  - Failure calculation operator: Process.FaultString =<variable containing error message>

If you create a script, consider these additional guidelines:

- If your script selects columns from a database table, an IT PAM agent must exist on the server where the SQL Server 2005 is installed. SQL Servers must be listed in your domain under All Touchpoints. The SQL Server containing key data must display the agent name for the SQL Server.
- The inline script must run the sqlcmd utility to retrieve the desired list.

More information:

Configure CA IT PAM Integration for Dynamic Values (see page 277)

Configure CA IT PAM Integration for Dynamic Values

You can configure CA IT PAM integration to leverage either or both of the following types of CA IT PAM processes:

- Event/alert output process--a process that invokes processing on a third-party system such as a help desk product
- Dynamic values process--a process that accepts an input key and returns current values for that key as a comma separated values (*.csv) file

Configuration for either purpose requires the ability to launch and log in to CA IT PAM. Gather the following values:

- Fully qualified host name, or IP address, of the CA IT PAM server
- Port (default is 8080)
- User name and password that CA Enterprise Log Manager is to use to log in to CA IT PAM

Configuration of CA IT PAM for dynamic values lets you import the list of values that are dynamically generated by the configured dynamic values process. The import is done when setting up or refreshing keyed values used in certain reports and alerts.

The following procedure addresses both the common settings and the one specific to dynamic values.

To configure CA IT PAM integration for the dynamic values process

1. Click the Administration tab and the Services subtab.
2. Click Report Server

   The Global Service Configuration: Report Server appears.
3. Scroll to the IT PAM area.
4. Make the following entries to enable CA IT PAM access:
   a. Enter the fully qualified host name of the server on which CA IT PAM is installed
   b. Accept the default port number, 8080
   c. Enter valid login credentials for CA IT PAM
5. Enter a process path in the Dynamic Values Process field.

   This process path becomes the default when importing dynamic values.
6. Click Save.

   The following message appears: Confirmation: Configuration changes saved successfully.

Approaches to Maintaining Keyed Lists

Keyed lists are used in some predefined reports and in some predefined queries tagged as appropriate for action alerts. If you plan to use these reports or create alerts that use these queries, you can use any combination of the following approaches to maintaining your keyed lists.

- You can add key values directly for any selected key. You can also select any key value and edit or delete it.
- You can import key values stored in a CSV list. Or, you can export the current list of values as a CSV file, update that file, and then import the updated file to populate the values list.
- You can run a CA IT PAM process that dynamically generates a current list and returns the values as a CSV file that populates the values list.

If you plan to create custom reports that use a keyed list, you can add a custom key and then add or import its values.

You can identify the keyed list or lists used in a query and then update that list before scheduling a report or alert that includes that query.

More information:

Create a Keyed List (see page 280)
Update a Keyed List Manually (see page 281)
Update a Keyed List with Export/Import (see page 282)
Example: Update a Keyed List with a CSV File (see page 283)
Using Queries Tagged as Action Alert (see page 292)

Create a Keyed List

Keyed lists let you create a group and assign values to it. You can then query the group name and any of the values in the group will return a positive result. You can assign values individually or import them from a .csv file. You can create custom keyed lists or add values to predefined lists.

For example, some queries for Privilege Grant reports search for a key value named "Privileged_Groups". When a query includes this value, it returns all rows where that field contains any of the values specified in the group.

To create a Keyed List

1. Click the Administration tab, the Library subtab, and the Keyed List folder.
2. Click New for a new list, or select the list to which you want to add values.

   Keyed List Details appears in the right pane.
3. Type a name and description for the keyed list, and select a type.
4. If you want to import values, go to Step 8. If you want to add values individually, go to Step 5.
5. Click Add at the top of the Keyed List Values table.

   A highlighted row appears in the User column.
6. Click the row, and type a value.
7. Repeat step 4-5 to add additional values.
8. Click Import at the top of the List Details area, browse for the file containing the values you want to import, and click OK.

   The values appear in the Values area.

   Note: You can only import .csv files that do not contain special characters.
9. When you have added all the values you want, click Save.

   The new list appears in the User folder of the Keyed Lists folder.

Update a Keyed List Manually

You can update the values in a keyed list in several ways. One way is to add, edit, and delete values manually.

To update a keyed list manually

1. Click the Administration tab, the Library subtab, and the Keyed List folder.
2. Expand the Keyed List folder, and select the keyed list you want to update.
3. To add a value to the keyed list:
   a. Select the key to which you want to add a value.
   b. Click Add Value.
   c. Enter the name of the value in the Name field and click OK.

      The added value appears in the Values list for the selected key.
   d. Repeat these steps for each value to add.
4. To delete a value in a keyed list:
   a. Select the key with an unneeded value.
   b. Select the value to deleted and click Remove Value

      A confirmation message appears.
   c. Click OK.

      The value is deleted from the Values list of the selected key.
   d. Repeat these steps for each value to delete.
5. To edit a value in the keyed list:
   a. Select the key with the value to modify.
   b. Select the value to modify and click Edit Value.
   c. Edit the entry in the Name field and click OK.

      The value is displayed with the modified name in the Values list of the selected key.
   d. Repeat these steps for each value to edit.
6. Click Save.

   The values for the selected keys are updated.

Update a Keyed List with Export/Import

If you store values that correspond to a key in an Excel spreadsheet, you can save that spreadsheet as a comma-separated values list (*.csv) and populate the Keyed List for the selected with an import.

You can update keyed list values you store in a CSV file in the following ways:

- If the CSV file contains current values for a given key and the displayed Values list is not current, you can import the values directly from the CSV file.
- If you want to create a CSV file or update one with obsolete values, use an export, edit, import sequence.

To update a keyed list with export or import

1. Click the Administration tab, the Library subtab, and the Keyed List folder.
2. Expand the Keyed List folder, and select the keyed list you want to update.
3. To update values for a selected key from a CSV file that contains current values:
   a. Select the key in the Key Values list that you want to update.
   b. Click Import Values on the Values list toolbar.

      The Import file dialog appears.
   c. Click Browse, and navigate to the location where the CSV file containing the values for the selected key is saved.
   d. Select the file to import and click Open, and then Click OK.

   The Values list is updated with the values from the CSV file.
4. To update the values for a selected key where the CSV file either does not exist or is not current:
   a. Select the key in the Key Values list that you want to update.
   b. In the Values toolbar, click Export Values, navigate to the location where you want to save the CSV file, and click Save.

      A success confirmation appears.
   c. Click OK.
   d. Navigate to the exported file, open the spreadsheet, and modify or delete existing columns as required. Scroll to display the last column, and add new entries. Then, save the file as a CSV file.
   e. Select the same key and click Import Values.
   f. Click Browse, select the file you saved, and click Open.
   g. Click OK.

   The file is uploaded. You can scroll to the bottom of the Values list to confirm your new entry is present.

Example: Update a Keyed List with a CSV File

You can supply values for keyed lists in the following three ways:

- Enter the key values manually
- Import the key values from a CSV file
- Import the key values from a specified CA IT PAM process

Use the following example as a guide to updating the values in any user-defined keyed list where the values are stored in an Excel spreadsheet saved as a comma-separated values list (*.csv).

To update a keyed list with a CSV file

1. Click the Administration tab, the Library subtab, and the Keyed List folder.
2. Expand the Keyed List folder, and select the keyed list you want to update, such as Default_Accounts, and click Export Values.

   An Export dialog appears with file.csv as the default filename.
3. Select the directory where you want to save the exported file. Change the file name, for example, Default_Accounts.csv and click Save.

   A confirmation message appears.
4. Click OK.
5. Browse to the exported .csv file, open it and scroll to display the last column, and add the entry you want to include. Optionally, delete the column for any default entry you want to remove from the keyed list for Default_Accounts.
6. Save and close the .csv file and return to the CA Enterprise Log Manager interface.
7. Click Import Values for the list you want to update here, the Default_Accounts keyed list.
8. Click Browse, select the file you saved, and click Open.
9. Click OK.

   The file is uploaded. Scroll to the bottom of the Values list to confirm that your new entry is present.

Update a Keyed List with a Dynamic Values Process

If you use CA IT PAM processes to generate a list of values associated with a key used in CA Enterprise Log Manager queries, run the IT PAM dynamic values process from CA Enterprise Log Manager and update the values for a given key. Importing saves you the time of manually entering all the values for a given key. When values for one of your keys change, you can refresh them in CA Enterprise Log Manager by selecting the key and repeating the import of dynamic values.

Configure CA IT PAM integration for dynamic values before attempting to import keyed list values from CA IT PAM.

To import values for a keyed list from CA IT PAM

1. Click the Administration tab, the Library subtab, and the Keyed List folder.
2. Expand the Keyed List folder, and select the keyed list you want to update.
3. To create a key for the values to import:
   a. Click Add at the top of the keyed values table.

      The first available row in the User column is selected
   b. Click the row, and type the name of the new key.
   c. Click Save.
4. To refresh dynamic values for an existing key:
   a. Select the key.
   b. Click Import Dynamic Values list at the top of the details pane.

      The Import Dynamic Values dialog appears.
   c. Enter the name of the IT PAM Dynamic Values Process that generates the values for the selected key, and then click OK.

      The associated CA IT PAM process is run, a file with the results is returned, and values for the selected key are refreshed.
   d. Click Save.

More information:

Enabling Dynamic Values Import (see page 276)
About Dynamic Values Processes (see page 276)
Create a CA IT PAM Process to Generate a Values List (see page 277)
Configure CA IT PAM Integration for Dynamic Values (see page 277)

Determine Keyed List Usage for a Query

It is good practice to keep keyed lists updated with current values. To update a keyed list used in a particular report or alert, first identify the queries used in the report or alert. Then, determine the keyed list used in the source query or query. Queries that use a keyed lists often reference the keyed list name in the query name. For example, there are queries with "Default Accounts" or "Privileged Group" in the query name.

To determine keyed list usage for a query

1. Open a copy of the query you want to check for keyed list usage in the query design wizard.
2. Click the Query Filters step and then click the Advanced Filters tab.
3. A query using a keyed list has a filter with the operator Keyed. The value is the name of the keyed list Default_Accounts for example.
4. Click Cancel. The query copy is not saved.

Creating Keyed Values for Predefined Reports

Some predefined keys that are used in predefined reports have no predefined values. To use these reports effectively, you must supply values for the respective keyed lists. You can also add custom values to keyed lists *with* predefined values.

Examples of Keyed Lists that have no predefined values include:

- Critical_Assets
- DMZ_Hosts
- EPHI_Database
- Business_Critical_Sources

You can add values to any keyed lists manually or by import.

More information:

Update a Keyed List Manually (see page 281)
Update a Keyed List with Export/Import (see page 282)
Example: Update a Keyed List with a CSV File (see page 283)

Create Keyed Values for Critical_Assets

This topic provides an example of adding custom values to a keyed list that has none provided by default. You can follow this example to add values to other existing keyed lists.

You can use certain reports and queries to monitor activities by your business critical hosts. To do this, you must first identify these hosts as values in the key-value list for Critical_Assets.

Reports that use the Critical_Assets list include the following:

- Account Creations by Business Critical Hosts
- Failed Login on Business Critical Hosts
- Resource Access Sessions by Business Critical Hosts
- Resource Access by Business Critical Hosts
- System Access by Business Critical Hosts

Similar reports for CA Access Control, CA Identity Manager, and CA SiteMinder use the Critical_Assets keyed list, for example: CA Access Control - Account Creations by Business Critical Hosts.

Queries that use the Critical_Assets list include the following:

- (>5) Logins by Admin Accounts on Critical Systems during Night for Last 1 Day
- (>5) Logins by Admin Accounts on Critical Systems during Weekends for Last 1 Week
- System Exception Activity...

If you create a custom query on critical assets, define the filter as follows:

Column	Operator	Value
dest_hostname	Keyed	Critical_Assets

To define a filter for other keyed lists, replace the value with the list value you want. For example, you could set the filter value to EPHI_Database to filter for hostnames belonging to that keyed list.

To create keyed values for Critical_Assets

1. Click the Administration tab, the Library subtab, and the Keyed List folder.
2. Expand the Keyed List folder, and select Critical_Assets.
3. Take one of the following actions to create this list:
   - Click Add Value and enter each new value to include in the keyed list.
   - Create an Excel spreadsheet with one row, where each column is a single value. Save it as a csv file. Click Import Values to import your edited list.
   - If the values for this key are dynamically generated by the CA IT PAM dynamic values process, click Import Dynamic Values List.
4. Click Save.

   Reports using this keyed list that are generated by scheduled jobs begin reflecting data for the updated values.

Customize Keyed Values for Administrators

This topic provides an example of adding custom values to a predefined keyed list that has some values already set. You can follow this example to add values to other existing keyed lists.

You can use predefined reports and their associated queries to monitor activities by your administrators. Predefined values include Administrator, root, sa, and admin. To customize the list, identify other accounts in your environment that have admin privileges as values in the key-value list for Administrators.

If you create a custom query that uses this key, define the filter as follows:

Column	Operator	Value
dest_username	Keyed	Administrators

To define a filter for other keyed lists, replace the value with the list value you want. For example, you could set the filter value to EPHI_Database to filter for hostnames belonging to that keyed list.

To customize keyed values for Administrators

1. Click the Administration tab, the Library subtab, and the Keyed List folder.

   A list of keys to which you add user-defined values is displayed at the bottom of the main pane.

2. Select the key, Administrators.

   The predefined values appear.

3. Take one or more of the following actions to update this list:

   - Update the list manually:
     - Click Add Value and enter a new value to include in the keyed list.
     - Select a value and click Remove Value to delete the value from the list.
     - Select a value, click Edit Value, modify the value and click OK.
   - Update the list with export/import:
     a. Click Export Values to export the current list.
     b. Open the exported list, edit the list to change its values, and save the file.
     c. Click Import Values to import your edited list.
   - Click Import Values to import the values in an updated csv file.
   - If the values for this key are dynamically generated by the configured CA IT PAM dynamic values process, click Import Dynamic Values List.

4. Click Save.

   Reports using this keyed list that are generated by scheduled jobs begin reflecting data for the updated values.

View a Report Using a Keyed List

You can view the results of a report before scheduling it to be generated. Certain predefined reports use keyed lists, where the key is predefined but the values are user-defined. Once you add or import values for a key, it is a good practice to view the report using the keyed list.

To view a report using a keyed list

1. Click the Queries and Reports tab and the Reports subtab.
2. Select a report that uses a keyed list.
3. View the results.

Chapter 10: Action Alerts

This section contains the following topics:

About Action Alerts

Action Alerts are specialized reports that generate an event when their query conditions are fulfilled. They can help you monitor your environment - allowing automatic notifications for a wide variety of situations and occurrences. For example, you can set action alerts to deliver event trend information, track disk space usage, or deliver notifications when failed access thresholds are exceeded.

Action alerts are a good way to sift through mountains of collected data for those few events on which you need to act right now. You can use action alerts to notify you about almost anything that happens in your log collection network. You can create alerts to let you know about spikes in inbound or outbound traffic, traffic on specific ports, access of certain privileged resources, configuration changes to various network entities like firewalls, databases, or key servers, and so forth.

You can create action alerts in the following ways:

- Using the action alert wizard
- From a query display
- Using a custom-created query

Scheduling options are a significant part of creating an alert, so you have control over how long and how often your alert job runs.

Using Queries Tagged as Action Alert

CA Enterprise Log Manager provides a number of queries with the tag, Action Alerts. To view the list of queries tagged Action Alerts, click the Queries and Reports tab, Queries subtab, and select the Action Alerts tag. The queries with this tag appear in the Query List. When you move your cursor over a query name, its tag or tags display.

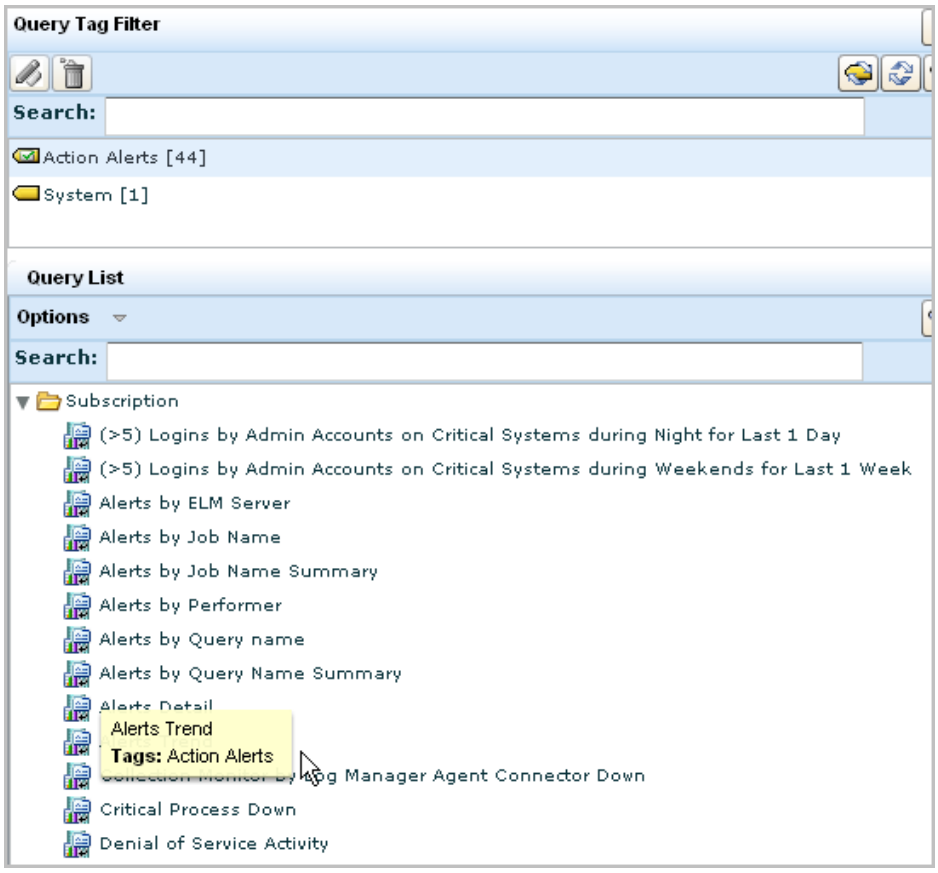

Before you schedule action alerts from these queries, you can get more information about what each of these queries do. To view a description and details on a query such as Low Available Disk Space, select that query form the query list, then move your cursor over the query name.

A summary of the query appears, including a description, its filters, and the query conditions.

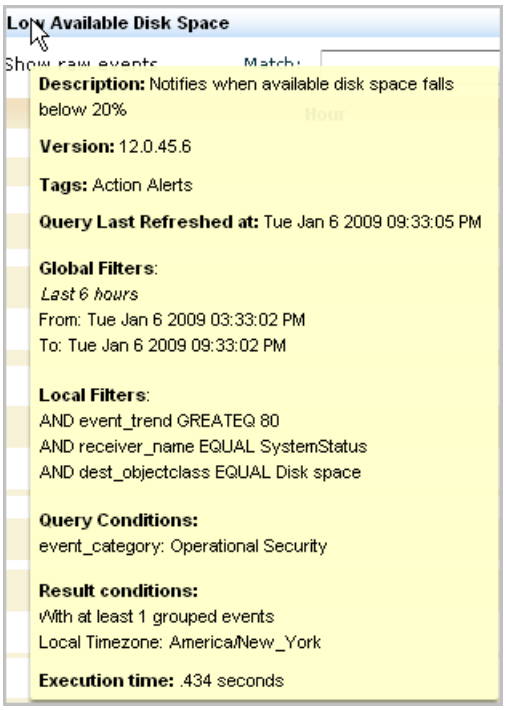

.

You can either schedule the query as is or you can copy the query to a new name and customize it to your requirements. For example, you can generate an alert when available disk space falls below 25 percent instead of 20 percent. You can create a user-defined query based on the predefined query and then select it for your action alert.

Note: Before using the queries with Privileged Group or Default Account in the title, consider adding your own keyed values for the corresponding keyed lists.

Identifying Other Queries to Use for Alerts

There are queries that are not tagged as Action Alerts that are good candidates for including in a scheduled action alert because they retrieve only events evaluated as severe.

For example, Security Log Cleared by Host Detail retrieves all events where the event action is Security Log Clear. The only tag for this query is Operational Security.

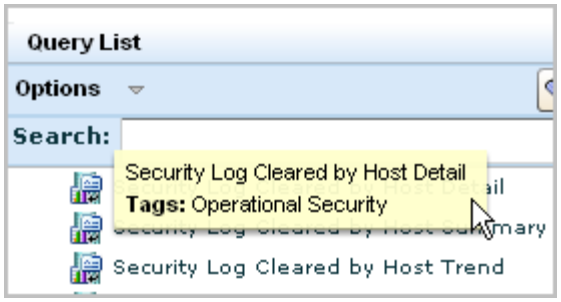

The action, Security Log Clear, is listed in the CEG. The CEG defines the following two event types with a security level mapped to 6, which is severe.

Category	Class	Action	Result	Security Level
Operational Security	Security Log Activity	Security Log Clear	Success	6
Operational Security	Security Log Activity	Security Log Clear	Failure	6

It is a good practice to schedule an alert with this query.

More information:

Identify the Simple Filter for Severe Events (see page 296)

Customizing Queries for Action Alerts

Alerts are designed to notify the appropriate person, process, or product when a severe event occurs. When attempting to identify queries on which to base alerts, consider queries designed to retrieve events with a high security level.

After you identify the definitions for severe events, you can identify the queries that retrieve severe events. If queries do not exist, you can create them.

Consider the following process:

1. Identify the event types that CA considers very severe, where event types are defined by category, class, action, and result.
2. Identify predefined queries that are designed to retrieve only such events.
3. Identify predefined queries that are designed to retrieve events that would include severe events, but could be customized to include only severe events.
4. Create custom queries where predefined queries do not exist.
5. Schedule alerts to run these queries frequently.

More information:

Identify the Simple Filter for Severe Events (see page 296)
Customize Queries to Retrieve Only Severe Events (see page 299)
Create a Query to Retrieve Only Severe Events (see page 297)

Identify the Simple Filter for Severe Events

Events vary in severity from informational to fatal. CA assigns a value between 2 and 7 to indicate the severity of events based on the CEG model of Category, Class, Action and Result. Severity 7 is assigned to system shutdown events. Severity 6 is assigned to events with high security implications or that need immediate attention.

If you plan to create custom queries or to customize predefined queries for use in alerts, it is a good idea to examine the CEG model definitions of severe event types. The model definition is the basis for simple filters. That is, you can create queries that retrieve events based on your specification of their event category, event class, event action, and event result.

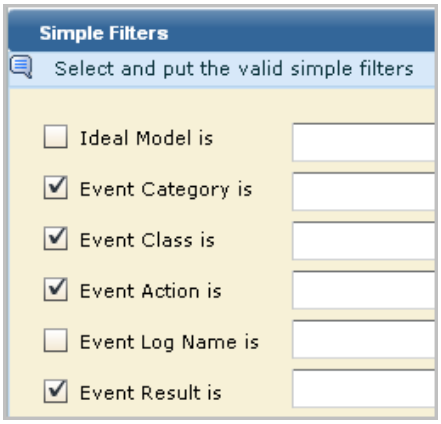

To identify the simple filter for severe events

1. Click the Help link.
2. Expand Common Event Grammar, and select Security Level Assignment.
3. Copy the table to a spreadsheet and sort by Security Level from highest to lowest.

   The resulting table lists event types beginning with the most severe based on CA Security Level assignment.

   An example follows. Your results will reflect the current CEG definitions.

Category	Class	Action	Result	Security Level
Operational Security	System Activity	System Shutdown	Success	7
Operational Security	System Activity	System Shutdown	Failure	7
Configuration Management	Configuration Management	Configuration Error	Success	6
Data Access	Object Management	Control File Creation	Success	6
Host Security	Antivirus Activity	Scan Error	Success	6

Category	Class	Action	Result	Security Level
Host Security	Antivirus Activity	Virus Clean	Failure	6
Host Security	Antivirus Activity	Virus Detected	Success	6
Host Security	Antivirus Activity	Virus Quarantine	Failure	6
Host Security	IDS/IPS Activity	Signature Violation	Success	6
Network Security	Signature Violation Activity	Signature Violation	Success	6
Operational Security	System Activity	System Startup	Failure	6
Operational Security	Security Log Activity	Security Log Clear	Success	6
Operational Security	Security Log Activity	Security Log Clear	Failure	6
System Access	Authentication Activity	Authentication Fallback	Failure	6
System Access	Authentication Activity	Authentication Start	Failure	6

Create a Query to Retrieve Only Severe Events

You can create a query from scratch if you do not find a predefined query that retrieves the types of events you want to be notified about. Consider the following types of severe even types:

Category	Class	Action	Result	Security Level
Host Security	Antivirus Activity	Virus Quarantine	Failure	6
Host Security	IDS/IPS Activity	Signature Violation	Success	6
Network Security	Signature Violation Activity	Signature Violation	Success	6

Example: Create a query to retrieve only virus quarantine failures

Assume, for example, that you want to be notified of any virus quarantine failure. Perhaps the keyword quarantine does not appear in the query list. If such were the case, you can create the query you need and then schedule an alert that runs the query.

To create a query to retrieve virus quarantine failures

1. Click Queries and Reports.
2. Under Query List Options, select New.

   Query Design wizard appears with the Details step displayed.
3. Enter a name.

   For example, enter Alert: Virus Quarantine Failure
4. Enter a custom tag.

   For example, enter Virus Quarantine
5. Click the Query Columns step and add the desired columns.
6. Click the Query Filters step.
7. Enter a simple filter based on the CEG entry for the event.

   For example, select Host Security for category, Antivirus Activity for Class, Virus Quarantine for action, and F for result.

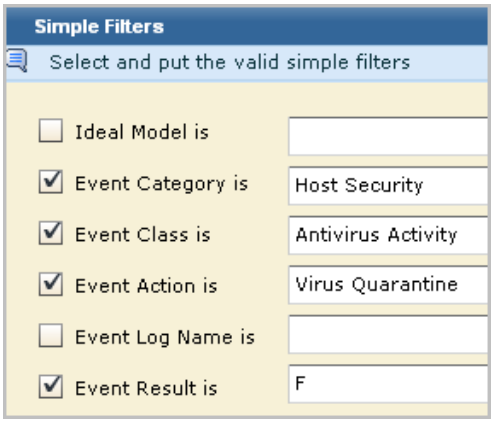

8. Select the Result Conditions step and select Last 5 minutes from the Predefined Ranges drop-down, to ensure timely alerting.
9. Click Save and Close.

Customize Queries to Retrieve Only Severe Events

Predefined queries that are not tagged as action alerts are designed for reports. It is appropriate for reports to contain data reflecting events of all levels of severity. You can customize selected queries to retrieve only severe events. To do this, you identify a query that retrieves severe events along with less severe events, copy it, enter filters that ensure retrieval of only the severe event, and save it for selection in an alert.

Before you begin, have at hand your spreadsheet that lists the definitions of severe events. This example is based on the following CEG information:

Category	Class	Action	Result	Security Level
Operational Security	System Activity	System Shutdown	Success	7
Operational Security	System Activity	System Shutdown	Failure	7

The query to customize retrieves events for both system shutdown and system startup.

To customize a query to retrieve only severe events

1. Click the Queries and Reports tab.
2. Select a query tag filter that matches the Category of a severe event.

   For example, select Operational Security.

3. Review the query list for queries with names containing keywords found in the Class or Action for the identified event type.

   For example, the keywords System Shutdown appear in queries beginning with the phrase System Startup or Shutdown by Host.

4. Copy the query System Startup or Shutdown by Host Detail. Highlight the query and select Copy from the Options drop-down list.

5. Click Query Filters and compare the default with the table entries for the severe event type.

   For this query, only Operational Security is selected.

6. Refer to the table for values to enter for Class and Action.

   For example, select System Activity for the Class and System Shutdown for the action.

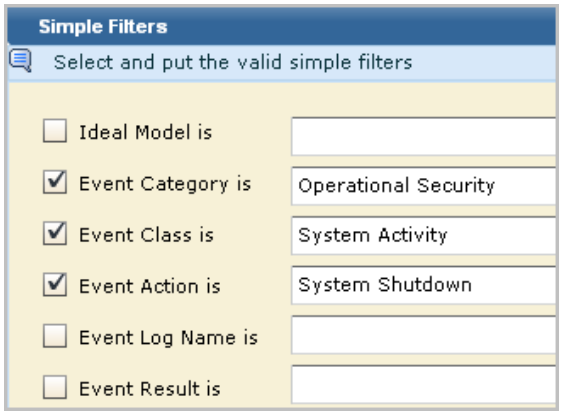

7. Select the Advanced Filters tab to determine whether modification is needed.

   Click delete for each line since the filter event_action is equal to system startup or shutdown is not pertinent to this custom query.

8. Replace that with a filter for the result.

   For example, create a filter where event_result is equal to either success or failure.

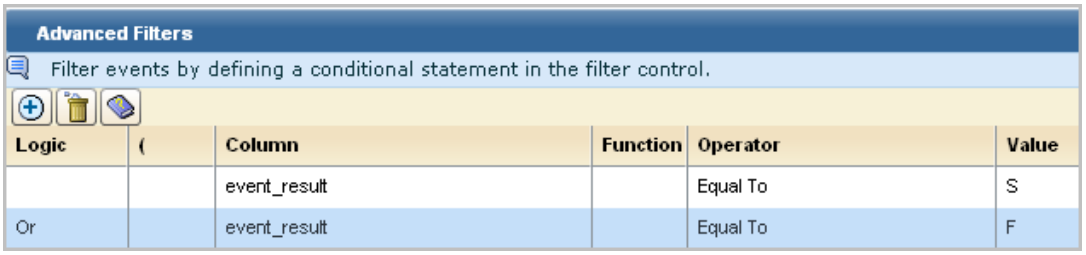

9. Click Details and name the query in a way that indicates you want to use it for an alert.

   For example, enter Alert: System Shutdown by Host Detail as the name. Change the description accordingly.

10. Click Result Conditions. For severe conditions, consider querying frequently.

    For example, select the predefined range for the last 5 minutes to run the query every 5 minutes for the occurrence of this severe event.

11. Click Save.

    You can create an alert with this query to notify a person, product, or process of a system shutdown success or failed attempt. (Product notification is done through SNMP traps; process notification is done through IT PAM event/alert output.)

Candidate Queries for Modification

Consider modifying selected predefined queries for use with alerts. To customize the query, add the simple filter based on the CEG analysis. Set the Date Range Selection with the Predefined Range, Last 5 minutes to ensure immediate notification. A few examples follow:

Query for Successful Configuration Error

1. Copy Configuration Error Activity Detail.

   This query returns successes as well as failures. Only successes are needed.

2. Set the simple filter as follows:

Category	Class	Action	Result	Security Level
Configuration Management	Configuration Management	Configuration Error	Success	6

3. Save as Alert: Successful Configuration Error

Query for Successful Control File Creation

1. Copy Data Manipulation Activity Detail

   This query retrieves all data access actions.

2. Set the simple filter as follows:

Category	Class	Action	Result	Security Level
Data Access	Object Management	Control File Creation	Success	6

3. Save as Alert: Successful Control File Creation

Query for Antivirus Scan Failure

1. Copy Virus Activity by Action

   This query filters for all Antivirus host security actions.

2. Use the following definition as a guide:

Category	Class	Action	Result	Security Level
Host Security	Antivirus Activity	Scan Error	Success	6

3. Define the simple filter as follows:

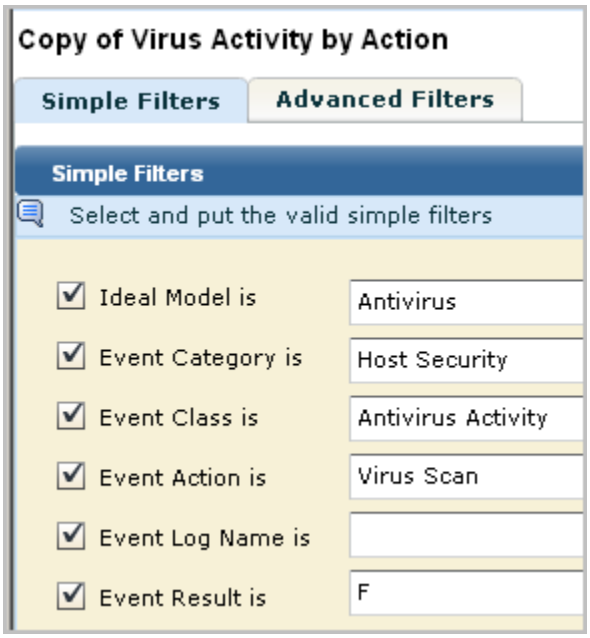

4. Save as Alert: Virus Scan Failed

Query for Virus Cleaning Failure

You can use the predefined query Virus Detection or Cleaning Activity Detail to retrieve both actions with either success or failure results. This may be sufficient for your needs. Optionally, you can create two separate queries based on this query where you specify the result as indicated on the CEG table for severe events.

1. Copy Virus Detection or Cleaning Activity Detail.
2. Create a simple filter to specify result of failure.

Category	Class	Action	Result	Security Level
Host Security	Antivirus Activity	Virus Clean	Failure	6

3. Remove the Advanced Filter.
4. Save as Alert: Virus Cleaning Failure

Query for Successful Detection of a Virus

You can use the predefined query Virus Detection or Cleaning Activity Detail to retrieve both actions with either success or failure results. This may be sufficient for your needs. Optionally, you can create two separate queries based on this query where you specify the result as indicated on the CEG table for severe events.

1. Copy Virus Detection or Cleaning Activity Detail.
2. Create a simple filter to specify result of success with just the detection activity.

Category	Class	Action	Result	Security Level
Host Security	Antivirus Activity	Virus Detected	Success	6

3. Remove the Advanced Filter.
4. Save as Alert: Virus Detected

Action Alert Considerations

You can view the results of any action alert from CA Enterprise Log Manager without any special configuration. Additionally, an action alert can be sent to the following destinations:

- RSS feed
- Email recipients
- SNMP trap destinations, such as CA Spectrum or CA NSM
- A CA IT PAM event/alert output process

Administrators configure these destinations from the Administration tab, Services subtab, under either Global Configuration or Global Service Configuration: Report Server.

Ensure these destinations are configured as follows before attempting to schedule an alert.

- To use the Feed Reader, ensure the checkbox is cleared for Viewing Action Alerts Require Authentication in the Global Configuration.

  The RSS Feed URL follows, where *elmhostname* is the host name of the CA Enterprise Log Manager server:

  https://{*elmhostname*}:5250/spin/calm/getActionQueryRssFeeds.csp

- To send alerts to email recipients, ensure the Email Settings section is configured in the Global Service Configuration: Report Server.
- (Optional) To send alerts to SNMP destinations, ensure SNMP Configuration section is configured in Global Service Configuration: Report Server.
- To send alerts to the CA IT PAM event/alert output process, ensure the IT PAM section is configured in Global Service Configuration: Report Server. (The only value not required for alerts is that for the Dynamic Values process.)

When you specify result conditions for an action alert, consider the following:

- Use the preset dynamic start time and dynamic end time for the predefined ranges.
  - The predefined range, last 5 minutes, is set with the dynamic end time to 'now', '-2 minutes' and the dynamic start time to 'now', '-7 minutes.' This default range and the other predefined time ranges allow adequate time for the events to be saved to the database.

    Note: Do not change the dynamic end time to *'now'* or to *'now', '-1 minutes'*. This change from the predefined value can cause incomplete data to be displayed when the URL is launched from the destination. For example, if event count is one of the values, the displayed count when viewed from the URL may be less than the displayed count when viewed from CA Enterprise Log Manager.
- Extend the dynamic end time if incomplete data is displayed with the default setting. For example, set it to 'now', '-10 minutes'

When you create an action alert schedule, consider the following:

- The recurrence interval is the frequency with which the query is run. Therefore, a recurrence interval of 5 minutes means the query is run every five minutes, or 12 times per hour. An action alert is generated only if the query returns results when it is run.
- Set the recurrence interval based on how time critical it is for you to respond when the tested condition occurs.
  - If you need to take immediate action to remedy the condition, set the recurrence interval to a high frequency so you can be notified as soon as possible.
  - If the condition is one that you want to track, but not one that requires intervention, set the recurrence interval to a low frequency.
- Avoid setting the recurrence interval to a high frequency, such as every five minutes, if your CA Enterprise Log Manager server time is not synchronized with your NTP server.

  Important! Your CA Enterprise Log Manager server time must be synchronized with your NTP server to ensure complete results are returned when the query is set to run at a high frequency.

Consider the following filtering options:

- To use the filters that are defined with the included queries, no action is needed.
- To apply additional filters to the queries included in an alert, define them in the Alert Filters step.
- To apply the same set of filters to multiple alert jobs, use a profile.

Before you configure thresholds for action alerts on a CA Enterprise Log Manager report server, consider the following:

- To keep the RSS feed to a reasonable size, set the maximum number of alerts to allow. The more frequent the recurrence intervals on enabled alerts, the faster the feed will fill up if the query or queries return results.
- To ensure that the RSS feed does not retain alerts longer than the data is of interest, specify the action alert retention to the maximum age in days of the oldest record to retain.
- Consider how often you want to check the RSS feed for alerts. That frequency can help you plan how long to keep records.
- If you want the RSS feed to display the most recent result of all jobs at all times, configure the retention values such that infrequently run alerts don't get deleted because they are older than frequently run alerts that fill the queue to capacity.

More information:

Configure Action Alert Retention (see page 397)
Example: Create an Action Alert for Low Disk Space (see page 388)

Working with CA IT PAM Event/Alert Output Processes

Working with CA IT PAM event/alert output processes that are integrated with CA Enterprise Log Manager involves some combination of the following tasks:

- Import the sample event/alert output process
- Create event/alert output processes in CA IT PAM that meet integration requirements
- Configure CA IT PAM integration and specify the default event/alert output process
- Run event/alert output processes from selected query results
- Schedule alerts that run a CA IT PAM process per row
- Schedule alerts that run a CA IT PAM process per query

More information:

Import the Sample Event/Alert Output Process (see page 315)
Guidelines for Creating an Event/Alert Output Process (see page 322)
Example: Run an Event/Alert Output Process with Selected Query Results (see page 328)
Example: Send an Alert that Runs an IT PAM Process Per Row (see page 334)
Example: Send an Alert that Runs an IT PAM Process Per Query (see page 338)

About CA IT PAM Event/Alert Output Processes

CA Enterprise Log Manager detects events that require intervention. You can generate alerts as soon as unwanted events occur. Integration with CA IT PAM makes it possible for an alert to run an event/alert output process. Event/alert output processes are designed to invoke appropriate remedial actions by other products. That is, event/alert output processes are CA IT PAM processes that command other products to take specified actions on specified objects.

CA Enterprise Log Manager, CA IT PAM, and third-party products work together to protect your environment. CA Enterprise Log Manager automates the detection of unwanted events and the IT PAM event/alert output process invokes other products to take the appropriate series of responses.

Integration involves configuring the connection to the CA IT PAM server, specifying the process to run, and specifying the process parameters with default values.

Running the CA IT PAM process can be done on demand from a displayed query result (row) or through scheduled alerts. In both cases, parameter values such as summary and description can be tailored to provide supporting details to the destination product of the CA IT PAM process.

More information:

Architecture Supporting CA IT PAM Integration (see page 308)
Process of Working with Event/Alert Output Processes (see page 309)
How CA IT PAM Integration Works (see page 311)
Example: Data Flow for Event/Alert Output Processing (see page 313)

Architecture Supporting CA IT PAM Integration

You need the following network components to run a CA IT PAM event/alert output process from an alert:

- A working CA Enterprise Log Manager environment, for example:
  - Agents with connectors that capture raw events from event sources
  - CA Enterprise Log Manager collection servers that refine the raw events and send them to reporting servers
  - CA Enterprise Log Manager reporting servers that process scheduled alerts and on demand queries
- A CA IT Process Automation Manager r2.1 (CA IT PAM) server configured with processes that invoke another product to perform a routine remediation action
- A server with a product used by the CA IT PAM process, for example, a server with a help desk product

Process of Working with Event/Alert Output Processes

An overview of the work flow for leveraging a CA IT PAM event/alert output process follows:

1. Determine whether to set up CA IT PAM integration with or without the sample process. The advantage of using the sample process is that it lets you see results right away. You can defer updating your own process until you become familiar with integration results. Using the sample process requires CA Service Desk.
2. Do one or both of the following:
   - Import the sample process and specify the CA ServiceDesk connection parameters
   - Create event/alert output processes that meet requirements of CA Enterprise Log Manager integration
3. Gather details for CA IT PAM integration from the sample process or the process you created.
4. Configure CA IT PAM integration for event/alert output.
5. Ensure that users who monitor event/alert output process results at the third-party product have user accounts in CA Enterprise Log Manager and know the credentials with which to log in. You can assign the role of Auditor to such accounts.

   Note: When users log in, all they can do is view the page with the associated query results.
6. Prepare to automate the running of an event/alert output process:
   a. Identify the query or queries that return data on which the third-party product can take action according to the configured CA IT PAM process.
   b. If the query uses a keyed list, ensure the keyed list is populated with the values you need.
   c. Run the event/alert output process on the query results, and verify that the process runs successfully.

7. Schedule an action alert using the documented procedure and the following guidelines.
   a. On the Alert Selection step:
      - Type a job name.
      - Verify selection type is Queries.
      - Select the query or queries you identified during planning.
   b. On the Destination step, select the IT PAM Process tab and specify event/alert output details as follows:
      - Select the queries on which to base the alert.
      - Specify whether to run the process once per query that returns results or once per returned row.
      - Specify IT PAM process parameter values. You can include field values and text for the Summary and Description parameter values only if running the process per row.
   c. Specify details for the remaining steps as with any action alert you schedule, then save and close the wizard.
8. Monitor the results:
   a. Verify the Action Alert Jobs list includes this job.
   b. Monitor self-monitoring events, Event Notification action, to verify that the result of running the IT PAM process was successful.
   c. (Optional) Log on to the third-party product that responded to the event/alert output information from CA Enterprise Log Manager that was passed to it by the IT PAM process.

More information:

Import the Sample Event/Alert Output Process (see page 315)
Guidelines for Creating an Event/Alert Output Process (see page 322)
Example: Run an Event/Alert Output Process with Selected Query Results (see page 328)
Design Queries for Events to Send to the Event/Alert Output Process (see page 333)
Set Notification Destinations (see page 384)
Example: Send an Alert that Runs an IT PAM Process Per Row (see page 334)

How CA IT PAM Integration Works

Assume the following setup has occurred:

- You have configured CA IT PAM on the Report Server configuration page and specified the event/alert output process to run.
- You have scheduled an alert with CA IT PAM as a destination and specified to run the process once per row. For parameters that allow entry of summary and description statements, you entered statements that included CEG fields.
- You have scheduled another alert with CA IT PAM as a destination and specified to run the process once per query. For parameters that allow entry of summary and description statements, you entered literal text.

The end-to-end process involves actions by multiple sources:

- The generation of raw events by event sources
- The collection, and refinement of events by CA Enterprise Log Manager
- The generation of alerts when refined events meet query criteria by CA Enterprise Log Manager
- The sending of event and alert output by CA Enterprise Log Manager to CA IT PAM
- The running of the configured event/alert output process by CA IT PAM on a third-party system
- One of the following:
  - An evaluation of data by a user of the third party system who determines the correct action and takes it.
  - The automated response by that third-party system to the occurrence of the events.

A summary of the processing follows:

1. Event sources generate raw events.
2. Agents collect some of these raw events based on their connectors and transfer the raw events to a collection server.
3. The collection server normalizes and classifies the raw events and transfers the refined events to a reporting server.

   For example, when a configuration change is made on any system, a log is created and classified as a configuration change. The event captures the time of the change, the host where the change was made, the user who performed the change, and the result of the change attempt.
4. The reporting server runs the queries selected for each scheduled alert.

5. When refined events meet the query criteria, the reporting server generates an alert and transfers the following information to CA IT PAM:
   - Alert details
     - Displayed process parameters and their values
     - CEG fields sent for undisplayed process parameters
   - Event details
     - For per row, event details are conveyed by the entries in the fields available for summary and description statements, where users describe the event with the CEG field variables composing the query selected for the alert.
     - For per query, event details are conveyed with a URL to a CA Enterprise Log Manager page that displays event details at the row level.
6. If the send is successful, CA IT PAM continues processing as defined in the configured event/alert output process.
7. If the third party product is CA Service Desk and the process is the sample event/alert output process, the following occurs:
   - A help desk ticket is opened and assigned a number. Fields on the ticket are populated with the parameter values from the alert definition. If a URL is received, it is displayed with the summary statement.
   - CA Service Desk returns the ticket number to CA IT PAM
8. CA IT PAM passes the ticket number back to CA Enterprise Log Manager
9. CA Enterprise Log Manager displays the ticket number as a self-monitoring event.

Example: Data Flow for Event/Alert Output Processing

The arrows on the following diagram illustrate the data flow:

- From collection servers to reporting servers
- From reporting servers to CA IT PAM
- From CA IT PAM to the product to which the CA IT PAM process sends the CA Enterprise Log Manager output, for example, CA Service Desk.

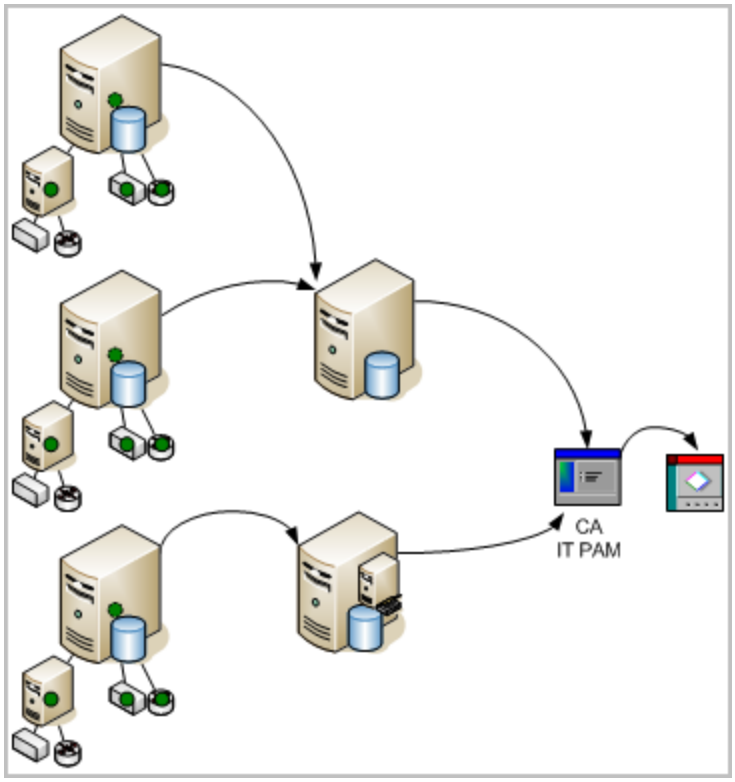

When CA Enterprise Log Manager receives notice that the send was successful, it polls CA IT PAM for the status of the process that was run. As soon as CA IT PAM sends the status update, CA Enterprise Log Manager creates a self-monitoring event with the result. The processing sequence follows:

1. CA IT PAM notifies CA Enterprise Log Manager whether the process that was run succeeded or failed.
2. CA Enterprise Log Manager generates a notification creation self-monitoring event with the received result.

Consider the example where the CA IT PAM process creates a help desk ticket with the process parameter values and the event data retrieved by the query. The arrows on the following diagram illustrate the following data flow:

- From the help desk product to CA IT PAM
- From CA IT PAM to the source CA Enterprise Log Manager reporting servers.

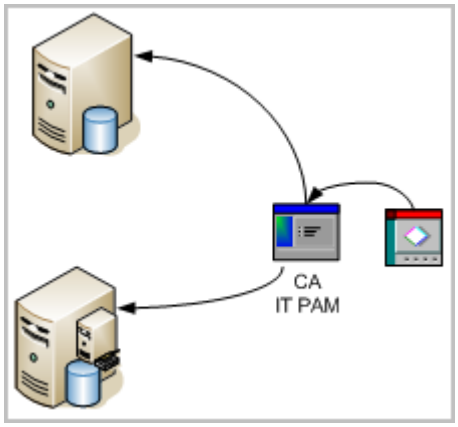

Import the Sample Event/Alert Output Process

To let you test CA IT PAM integration right away and practice the configuration procedure with known values, CA provides a sample process for this purpose. It is on the DVD with the application. Use of this sample IT PAM process assumes you are using CA Service Desk as your help desk application.

You can then configure CA IT PAM in CA Enterprise Log Manager and test running this sample CA IT PAM process with query results you select. After you become familiar with how CA Enterprise Log Manager operates with CA IT PAM, you can ensure compliance of your own process and substitute those values in the CA IT PAM configuration for your production integration.

To import a sample process and test IT PAM integration

1. Launch CA IT PAM and log on.
2. Launch the ITPAM Client.

3. Import the sample IT PAM process, EventAlertOutput.xml, provided on the application DVD under CA/ITPAM. This sample has all the required values defined.
   a. Select File, Open Library Browser.
   b. Click Folders in the left pane, and at the root folder, click Import.

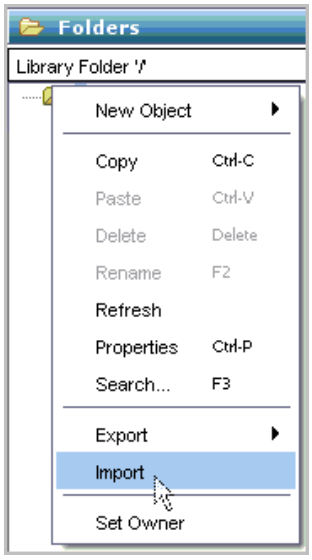

c. Select the sample IT PAM process, EventAlertOutput.xml, from the extracted iso image and click Open.

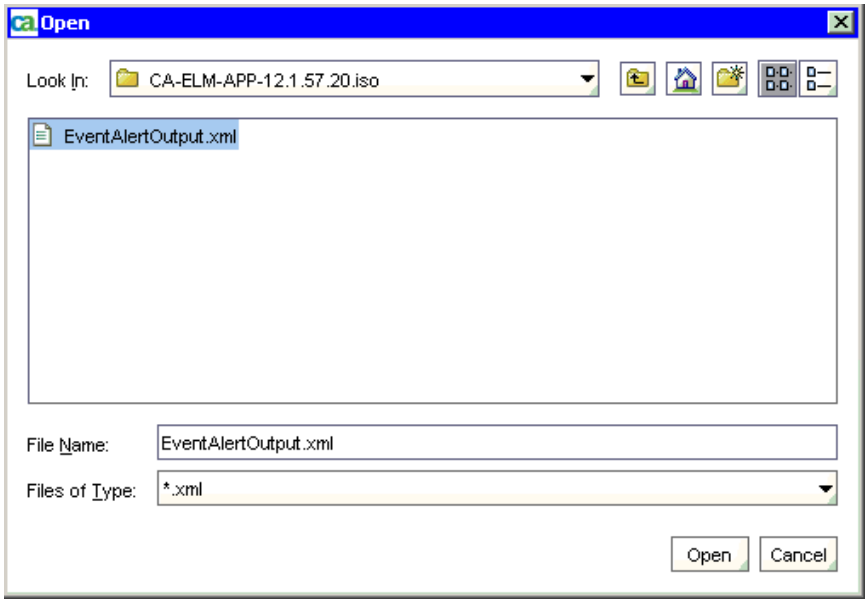

d. Select both options on the Import Object dialog and click OK.

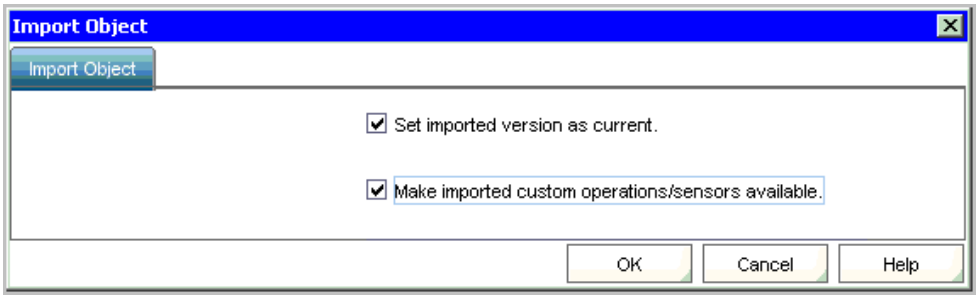

The resulting display shows the exact name and path. For example, the name is EventAlertOutput and the path is /CA_ELM/.

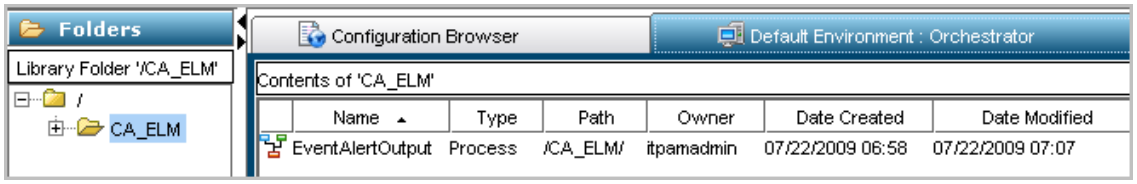

4. Specify the Service Desk connection parameters.
   a. Click the ServiceDesk Connect Parameters tab for Request_Create to view the ServiceDesk Connect Parameters.
   b. Use the following syntax for specifying the Service Desk URL:

      "http://<server name>:8080/axis/services/USD_R11_WebService"

   c. Enter valid login credentials to the Service Desk for Service Desk User ID and Password.
5. (Optional) Test the imported process to ensure that it works as a standalone process.
6. Close the ITPAM Client, then click Sign Out to exit CA IT PAM.

View the Sample Event/Alert Output Process

If you import the sample event/alert output process, you can examine its design in CA IT PAM. Use the following guidelines to become familiar with CA Enterprise Log Manager requirements in the context of the sample process. During this walk-through, you will see where to define web service connect parameters and how the calculation operators are defined. In addition, you will notice product-specific requirements. For example, configuring CA Service Desk as the third party product requires use of the Request_Create operator from the CA Service Desk Module and a precalculation operator that maintains values for severity and priority.

To become familiar with the sample event/alert out process

1. Display the model of your target process.
   a. Launch CA IT PAM and log in.
   b. Click ITPAM Client.
   c. From the File menu, select Open Library Browser.
   d. From the Folders tab, select the library folder containing the model for your target process.

      The name of your process and path appear in the main pane.

   e. Double-click the row containing your process name and path.

A model similar to the following appears. This example model contains minimal requirements for CA Enterprise Log Manager.

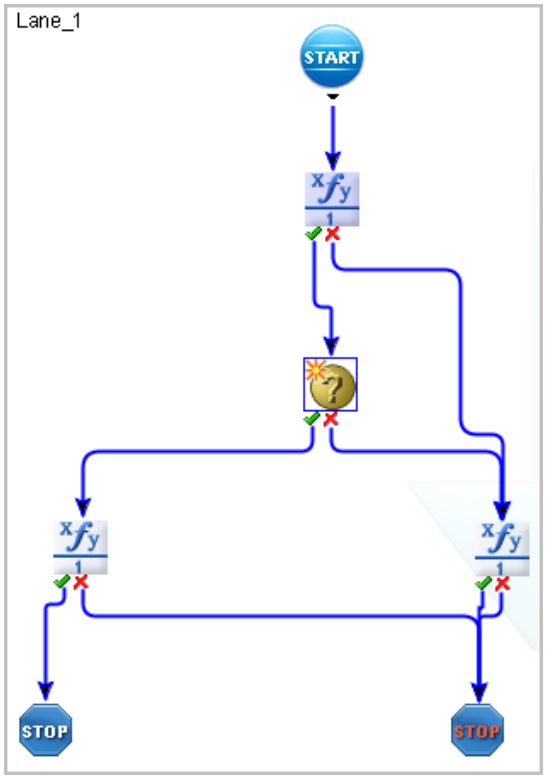

2. Notice how the ServiceDesk Basic Parameters meet CA Enterprise Log Manager requirements.

   a. Double click the Request_Create_1 icon.

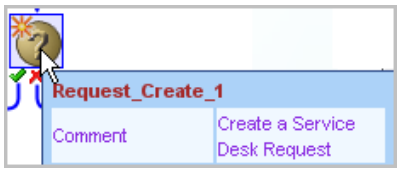

   The Request_Create operator passes the data returned by the action alert query to your target product (application). A similar operator is required for any process that is to be run from CA Enterprise Log Manager.

b. Under ServiceDesk Basic Parameters, notice that local process parameters are specified with the following syntax:

BasicParameter = Process.LocalParameter

Note: Local process parameters are the Event/Alert Output Process Parameters you add to CA Enterprise Log Manager when you configure CA IT PAM.

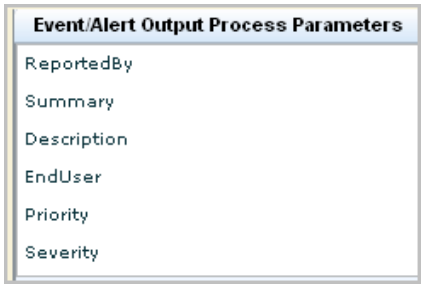

c. Since the target application is the CA Service Desk product, the following local process parameters are defined as described on the following table:

ServiceDesk Basic Parameter	Local Parameter	Service Desk Field	Notes
Request Creator ID	Process.ReportedBy	Assignee,Reported By	A valid "Contact" in CA Service Desk
Summary	Process.Summary	Summary	(Leave blank)
Description	Process.Description	Description	(Leave blank)
Customer ID	Process.EndUser	Affected End User	A valid "Contact" in CA Service Desk
Priority	Process.Priority	Priority	1-5
Severity	Process.Severity	Severity	1-5

The following example shows valid local parameters for ServiceDesk Basic Parameters. The entries are case-sensitive. That is, Process.ReportedBy must be entered exactly as shown with a capital "R" and a capital "B" for example.

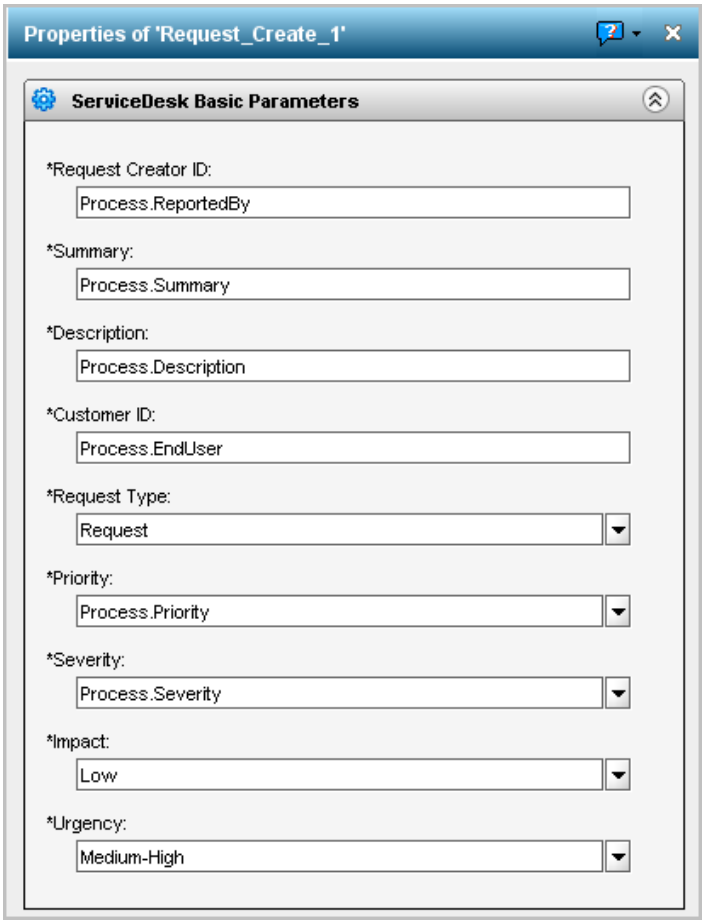

3. Click the ServiceDesk Connect Parameters tab for Request_Create to view the ServiceDesk Connect Parameters.
   - Service Desk URL:"http://<server name>:8080/axis/services/USD_R11_WebService"
   - Service Desk User ID:"<SD user>"
   - Password:"<SD password>"

4. Notice that for CA Service Desk, an adjustment is needed to ensure that the values of severity and priority that are entered in CA Enterprise Log Manager are correctly interpreted by CA Service Desk.

   a. A pre-calculation operator appears after Start and before the Create_Process operator. In the following example, it is named Fix_Sev_Pri.

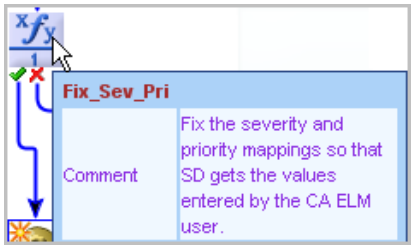

   b. Under Properties, Calculate, the following mappings are defined:

```
if (Process.Priority == 1) Process.Priority = "pri:504";
else if (Process.Priority == 2) Process.Priority = "pri:503";
else if (Process.Priority == 3) Process.Priority = "pri:502";
else if (Process.Priority == 4) Process.Priority = "pri:501";
else if (Process.Priority == 5) Process.Priority = "pri:500";

if (Process.Severity == 1) Process.Severity = "sev:800";
else if (Process.Severity == 2) Process.Severity = "sev:801";
else if (Process.Severity == 3) Process.Severity = "sev:802";
else if (Process.Severity == 4) Process.Severity = "sev:803";
else if (Process.Severity == 5) Process.Severity = "sev:804";
```

5. Notice that the following return value, or output interface, parameters are formatted as required by CA Enterprise Log Manager:

   - ResultString
   - FaultString

6. View the calculation operator for request creation success. This format must be used in any event/alert output process to be run from CA Enterprise Log Manager.

   a. Click the icon for the calculation operator for request creation success.

   b. Select the Calculate tab and click ... in the source code field.

   c. Notice how the success calculation operator is defined in the source code:

```
Process.ResultString = "Request " + Request_Create_1.newRequestNumber + " created in CA Service Desk.";
```

7. View the calculation operator for failure. This format is required for any event/alert output process to be run from CA Enterprise Log Manager.
   a. Click the icon for the calculation operator for failure.
   b. Select the Calculate tab and click ... in the source code field.
   c. Notice how the failure calculation operator is defined in the source code, where the Process.FaultString maps to the appropriate SOAP variable:

      ```
      Process.FaultString = Request_Create_1.SoapErrorResponse;
      ```

Guidelines for Creating an Event/Alert Output Process

Certain guidelines must be satisfied for a CA IT PAM process to run from CA Enterprise Log Manager. Before you attempt to run a CA IT PAM process from CA Enterprise Log Manager, verify that the process includes the following:

- Web Service Connect Parameters.
- A success calculation operator that maps Process:ResultString to a statement with literals and variables that expresses the response from the third-party product.
- A failure calculation operator that maps Process:FaultString to the appropriate SOAP response variable.

If your target IT PAM process is for a third-party help desk product, verify that the process also includes the following:

- The product-specific operator.

  For example, a process that targets the BMC Remedy Module would be defined with the Create_Help_Desk_Case operator.
- Product- specific parameters that are mapped to local process parameters: ReportedBy, Summary, Description, EndUser, Priority, and Severity.

  For example, a process that targets the BMC Remedy module would map local parameters to the HelpDesk Create Case Parameters.

Typically, a CA IT PAM process includes only the default process parameters, each of which is mapped to a field in the third-party product. Optionally, you can add CEG fields as process parameters for a given process. The following example shows the following CEG fields in the dataset:

- event_severity
- event_count
- event_datetime

Each basic parameters is mapped to a Service Desk field. For example, the ReportedBy process parameter is mapped to the CA Service Desk field named Assignee. When CEG fields are added as process parameters, they can be referred to as values in a basic parameter. For example, the value for the CEG field event_datetime can be defined to appear in the Description field in CA Service Desk by default. This is achieved by adding the Process.event_datetime in the Description field of the Service Desk Basic Parameters.

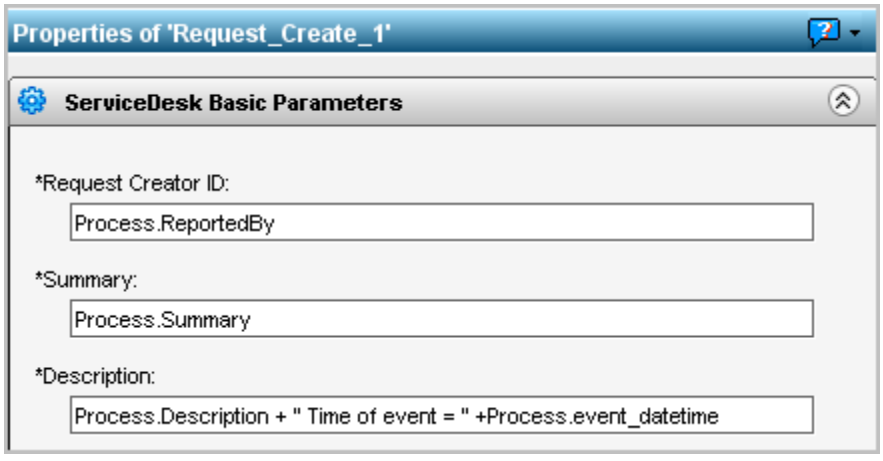

When you create an alert that runs this process, examine the CEG fields listed under Send field values as parameters. If any listed parameter is a CEG field that you defined as a process parameter, select that field. Consider the following examples:

- All three CEG fields defined in the dataset are displayed for the query System Event Count by Event Action. Therefore, you would select all three to send as parameters to CA IT PAM.

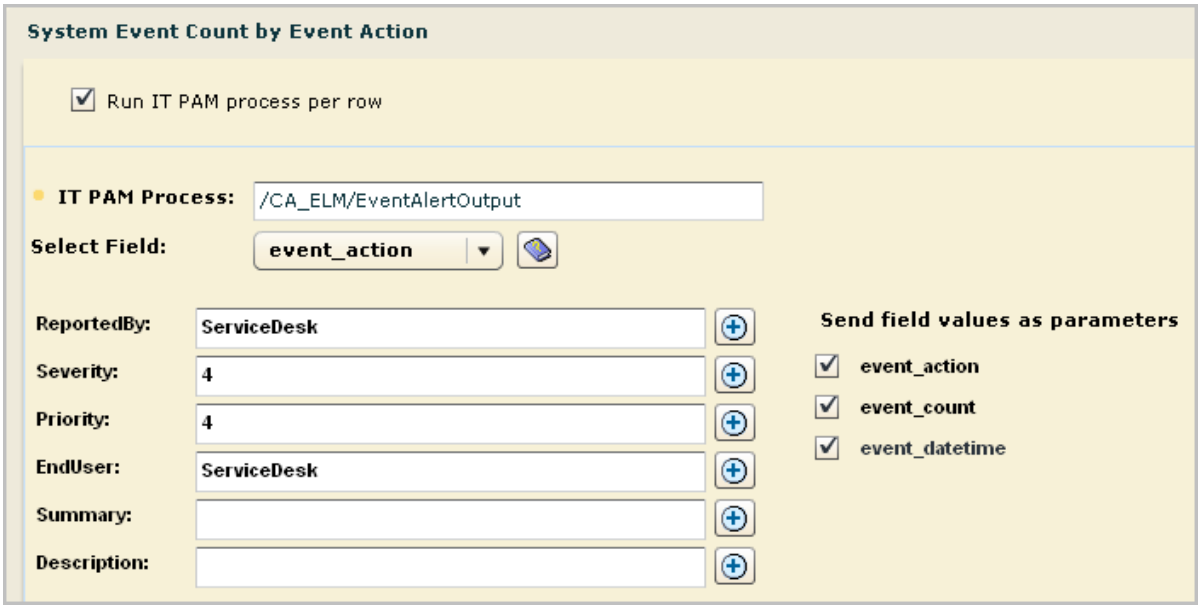

- Two of the three CEG fields defined in the dataset are displayed for the query >5 Logins by Admin Accounts. You would select those two to send as parameters to CA IT PAM.

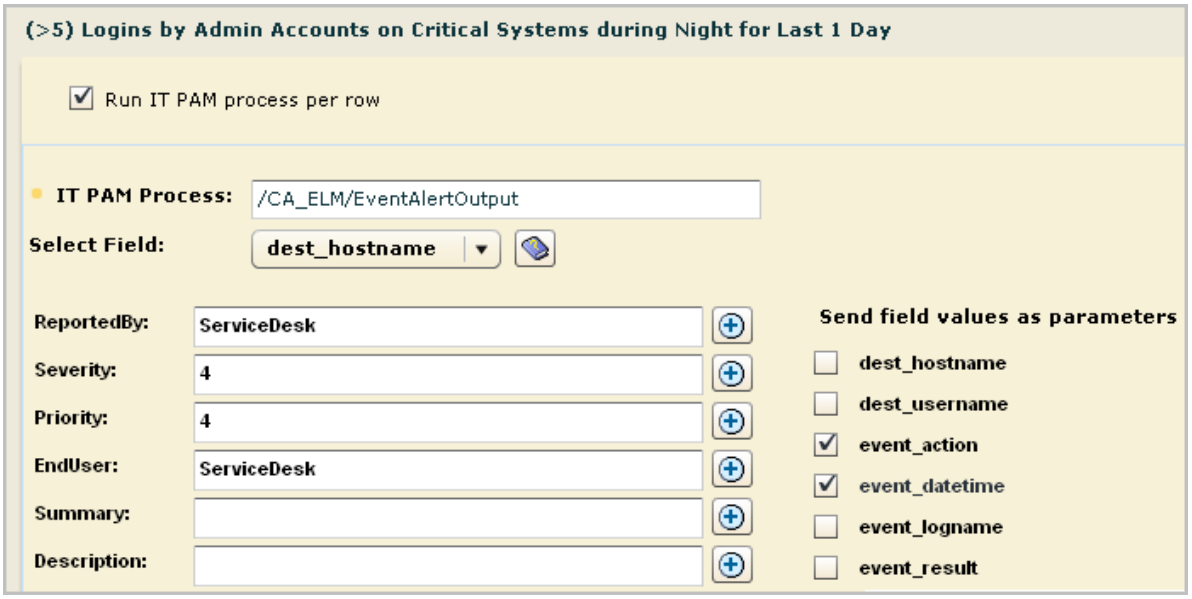

More information:

View the Sample Event/Alert Output Process (see page 317)

Gather Details for CA IT PAM Integration

Most of the details required for CA IT PAM integration are part of the CA IT PAM product and process configurations. You can launch CA IT PAM and search for the details as you need them for configuration or you can gather the details first, record them, and then quickly configure CA IT PAM by entering the values you recorded.

You can reference either the sample processes you imported or your own processes that you have modified to meet CA Enterprise Log Manager requirements.

To gather details for CA IT PAM integration

1. Log on to your local CA IT PAM server and verify it is CA IT Process Automation Manager 2.1.
2. Click the ITPAM Client link.
3. Gather details for the first four fields of the IT PAM configuration.
   a. Click Configuration Browser
   b. Click the Properties tab.
   c. Record the Server Name value as your value for IT PAM Server.
   d. Accept port 8080 as the IT PAM port.
   e. Obtain login credentials for CA Enterprise Log Manager from the CA IT PAM administrator and record them for Username and Password.

IT PAM Configuration Field	Description	Your Value
IT PAM Server	The fully qualified host name of the server where CA IT PAM is installed. This value appears in the Server Name field on the Properties tab of the Configuration Browser	
IT PAM Port	Port 8080 is the default This value appears in the Domain URL on the Properties tab of the Configuration Browser.	8080

IT PAM Configuration Field	Description	Your Value
Username	The user ID that CA Enterprise Log Manager is to use to log into IT PAM and run a process. Obtain from your CA IT PAM administrator Example: itpamadmin	
Password	The password associated with the Username. Obtain from your CA IT PAM administrator.	

4. Record the process path and names of the processes you plan to run from CA Enterprise Log Manager.
   a. From the File menu of the ITPAM client, select Open Library Browser
   b. In the Folders tab, select the library folder containing the event/alert output process.
   c. Record the path and name of the process for Event/Alert Output Process.
   d. If different, select the library folder containing the process that returns current values for a specified key.
   e. Record the path and name for Dynamic Values Process.

IT PAM Process-Specific Field	Description and Example	Your Value
Event/Alert Output Process	Path and process name. Identifies the process designed to pass details configured with the alert or a URL to an external product such as CA Service Desk. Example: /CA_ELM/EventAlertOutput	
Dynamic Values Process	Path and process name. Identifies the process designed to collect values for the input key and return them for parsing into a csv file. Example: /CA_ELM/ValuesList	

5. Collect event/alert output process parameters:
   a. Double-click the Event Alert Output process you referenced to open the process.
   b. On the Main Editor tab, click the Request_Create icon to display properties.
   c. Display the ServiceDesk Basic Parameters.
   d. Record those parameters prefixed by Process: in the first column below if they do not exactly match what is shown
   e. Click the Dataset tab.
   f. Click each parameter for the Local_Dataset and record its default value if any.

Event/Alert Output Process Parameters	Description and Example	Your Value
ReportedBy	A valid ServiceDesk user name.	
Summary	This text appears in the Service Desk request Summary field. For example “Request created from CA ELM”	---
Description	This text appears in the Service Desk request Description field.	---
EndUser	A valid ServiceDesk user name.	
Priority	Sets the default priority. If no default is configured, record a value between 1 and 5. Example: 3	
Severity	Sets the default severity. If no default is configured, record a value between 1 and 5. Example: 4	

Example: Run an Event/Alert Output Process with Selected Query Results

All users are authorized to run a CA IT PAM process on demand. You can run the configured CA IT PAM event/alert output process with selected query results for any of the following purposes:

- To perform an on demand event/alert output process based on current needs.
- To test the processing results before creating a scheduled alert for this query with the CA IT PAM process as a destination.

You can run a CA IT PAM process from a displayed query result row. This assumes the results are displayed as a table rather than a chart. You can display query result rows in any of the following ways:

- Select a query from the query list that returns results.
- Select a report from the report list, select a query that returns results.
- Enter a prompt that returns results.

Note: The following topic assumes that a query result row displays when you select the query from the query list.

To become familiar with what data is returned for the CEG fields, see the *Common Event Grammar (CEG) Reference* guide in online help.

To run the configured CA IT PAM process manually based on a displayed query result row

1. Click the Queries and Reports tab and the Queries subtab.

   The query tag filter and the query list appear.

2. (Optional) Enter search criteria, such as default accounts, on the query list.

   Events that reflect logins by default accounts are good candidates for forwarding to your CA IT PAM event/alert output process.

3. Select the query from the query list for which you want to view results.

   As an alternative, you can display the Reports subtab, select an option from the Report List, switch to individual query view, and select the query from this view.

4. If the results display in a chart, select Change Visualization from the query name drop-down list and select Table.

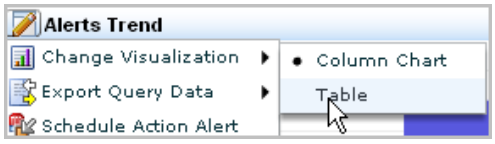

5. Select the query result row for which you want to run the CA IT PAM process.
6. Right-click this query result row and select Run IT PAM process from the drop-down list.

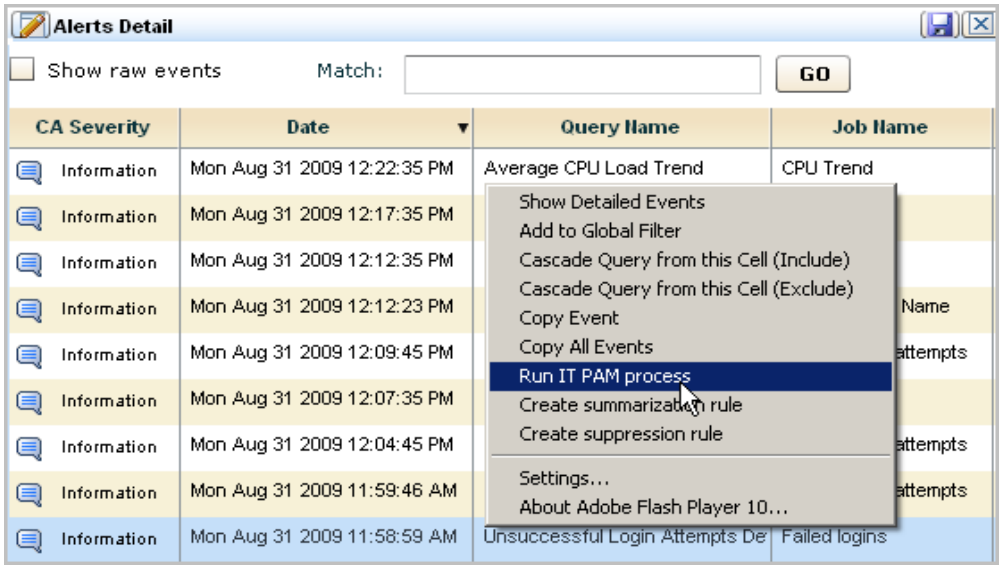

The Run IT PAM process dialog appears. It contains the process name and process parameters defined in the IT PAM configuration of the Report Server service. Additionally, it contains a Select Field drop-down list that allows you to enter variable data returned to the selected CEG field.

7. Complete the fields as follows:

   a. Review the default values shown for the displayed process parameters and identify any values that need to be changed.

      These parameters and their values are derived from the CA IT PAM integration configuration.

   b. To change the displayed default value, type the new value.

   c. To specify a variable value, select that CEG field from the Select Field drop-down list at the top of the dialog, then click Add Field next to the text box to which it applies.

   d. For any field that is blank, type a value, select a variable and add it, or type a sentence that includes selected variables.

      Example Summary: On (event_datetime), the (dest_username) account performed a (event_action) action on the (dest_hostname) host.

      Example Description: The action result (event_result), is logged in the (event_logname) log. The CA Severity is (event_severity).

   e. If the CA IT PAM process specifies parameters that refer to additional CEG fields, select these fields from the displayed list to send as parameters.

   An example follows. Your display may include other fields defined in the custom IT PAM event/alert output process.

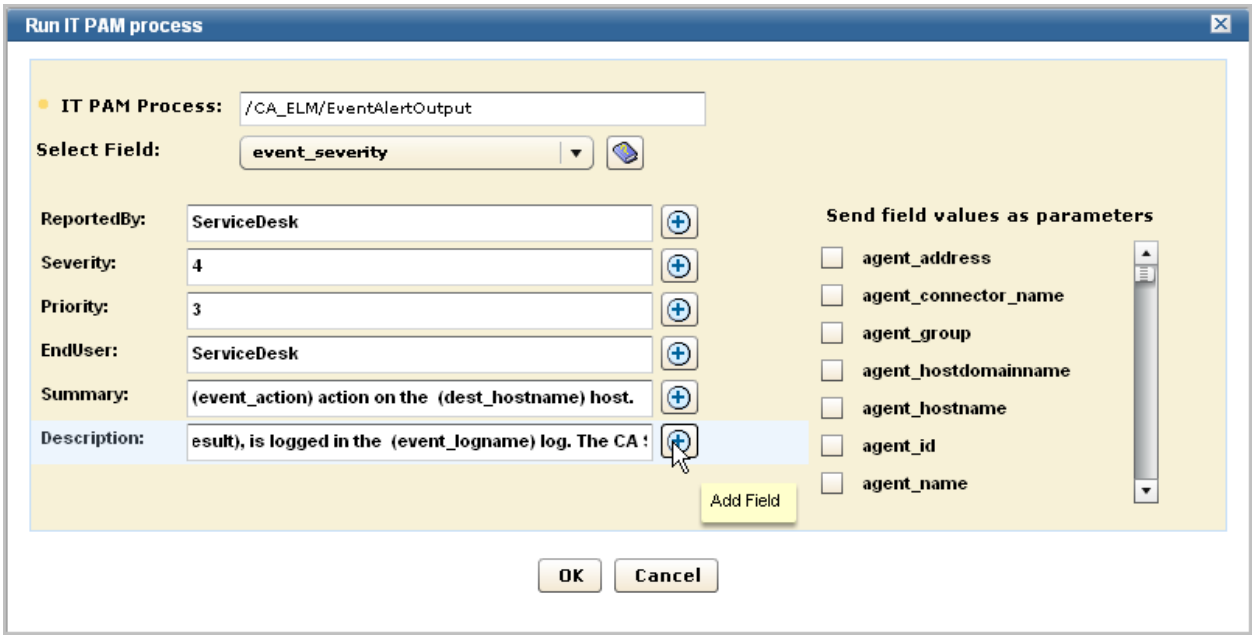

8. Click OK.

   The progress dialog appears, followed by a message indicating whether the CA IT PAM process ran successfully, and if so, the results of running the process.

   An example follows, where the result is Request 4590 created in Service Desk.

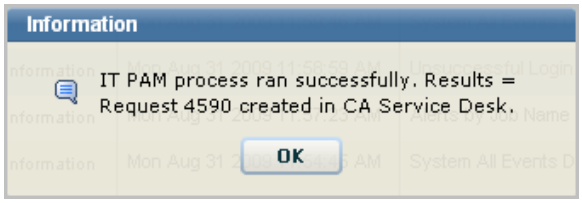

9. Click OK.

10. To see the results in CA Service Desk, log on and search for "Request" with the number in the message.

    For example, select Request and enter 4590.

11. Service Desk results similar to the following appear.

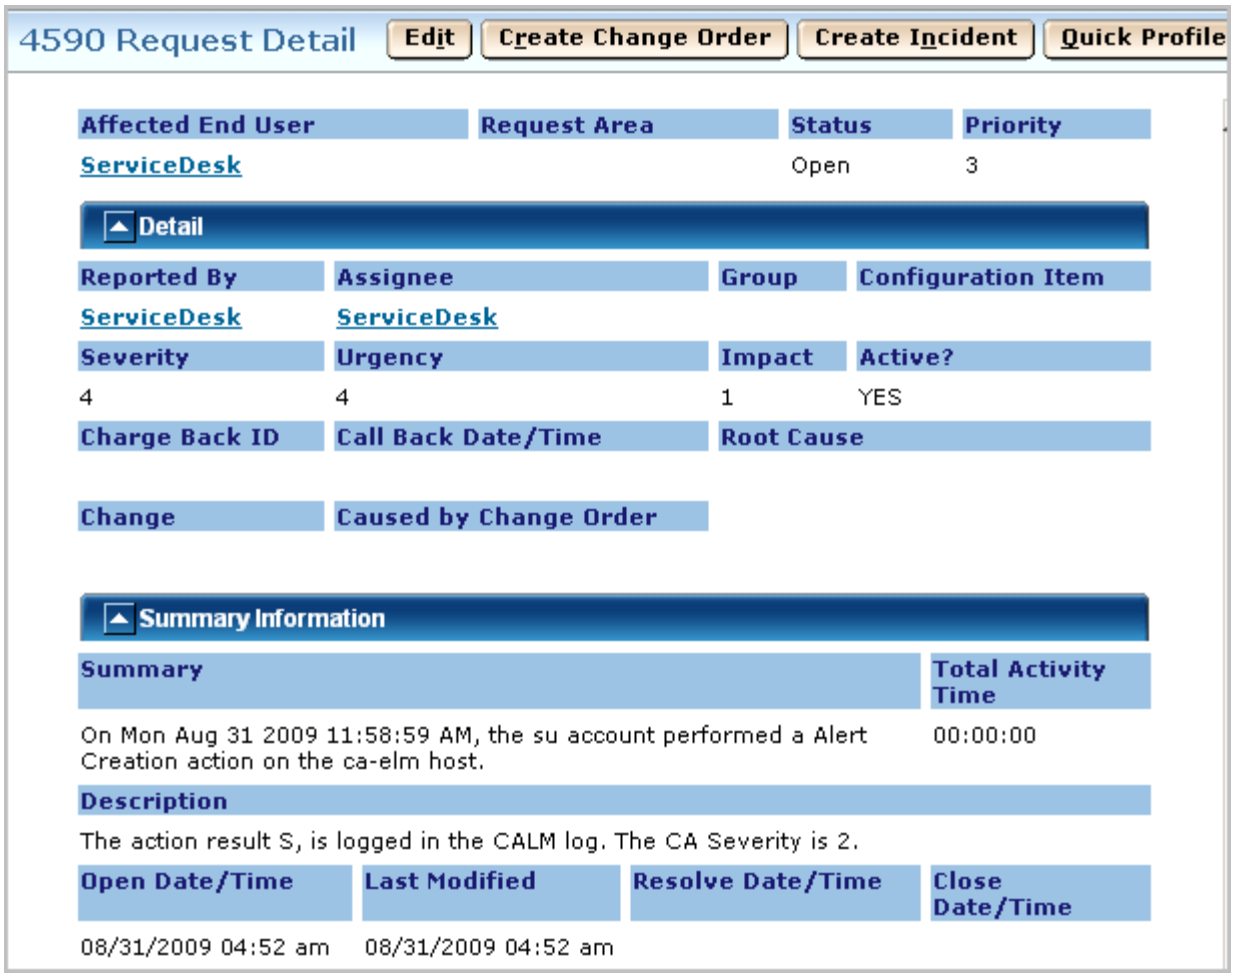

12. Compare the planned summary and description data determined in Step 7 with the summary and description data displayed under Summary Information. It includes the CA Severity data.

Design Queries for Events to Send to the Event/Alert Output Process

After you set up CA IT PAM integration, you can take the first step toward scheduling alerts that generate event/alert output--that of compiling a list of queries on which the alerts are to be based. These are typically queries for events that suggest a policy violation. You can take a combination of several approaches:

- Analyze currently scheduled alerts to identify any that should run the event/alert output process. For example, if the event/alert output process notifies a help desk application, you identify alerts that should open a help desk ticket.
- Analyze your policies to identify those where a violation could be traced back to a logged event, then create a query for such an event.
- Examine the results of other predefined queries to identify data that a third-party product, such as a help desk product, could use to take remedial action.
- If your CA IT PAM event/alert output process creates tickets in a third-party help desk product, review typical types of help desk tickets for causes that could be captured as event logs.

To identify or design queries on which to base alerts that run the CA IT PAM event/alert output process

1. For each event type requiring a help desk ticket, identify, modify, or create one or more queries that capture data for such an event.
   - Identify each predefined query that collects events on such conditions.
   - If a predefined query requires customization, copy the query and then tailor the copy to your needs.
   - If no predefined query exists to collect a particular type of event that requires help desk notification, create the query or queries you need.
2. For any query that is to search for an IT event where one of its fields can have any of several known values, use a predefined keyed list, customize a keyed list, or create a new keyed list. If the values for such a key exist in a csv file, import it. For a list generated by an IT PAM process, configure that process as the Dynamic Values process, create the key and then import the values from CA IT PAM.
3. Determine whether to run the CA IT PAM event/alert output process per query that returns results or per result row.

4. Test the query.
   a. Create the condition that produces the event you want to capture.
   b. Run the query or set of queries manually
   c. Evaluate whether the query results are sufficient for the help desk personnel to complete the needed follow-up.
   d. If not, modify the query or set of queries to provide the required information and retest.

This preparation ensures that when you schedule an alert that runs each such query or set of queries, the resulting event/alert output will contain the data required for resolution.

More information:

Customizing Queries for Action Alerts (see page 295)

Example: Send an Alert that Runs an IT PAM Process Per Row

You can send an alert that runs the CA IT PAM event/alert output process per row or per query. This example illustrates the procedure of running the process per row. It includes an example of what can be viewed for this type of alert by personnel working with both CA IT PAM and the third-party product to which CA IT PAM sends the details.

Prior to creating an alert to run an IT PAM process for a given query, it is a good practice to identify the CEG columns that return data. These columns are the ones to select when creating a summary and description statement for the alert.

Note: Copy the query and click the Query Columns step. For fields designed to be visible, notice the column name corresponding to the display name. For example, the CEG field used to populate the Account column is dest_username.

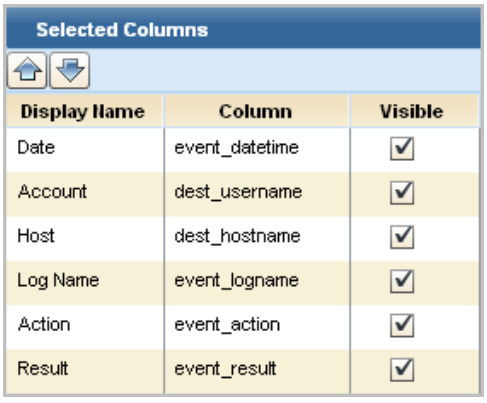

To create an alert when a default account member logs in successfully

1. Click the Alert Management tab and then click the Alert Scheduling subtab.
2. Click Schedule an Action Alert.

   The Schedule Action Alerts wizard appears.
3. Complete the Alert Selection step as follows:
   a. Enter the job name, for example, Default Account Logins.
   b. Click the Action Alerts tag.
   c. Select the Successful Login by Default Account in last 24 hours query and move it to the Selected Queries list.

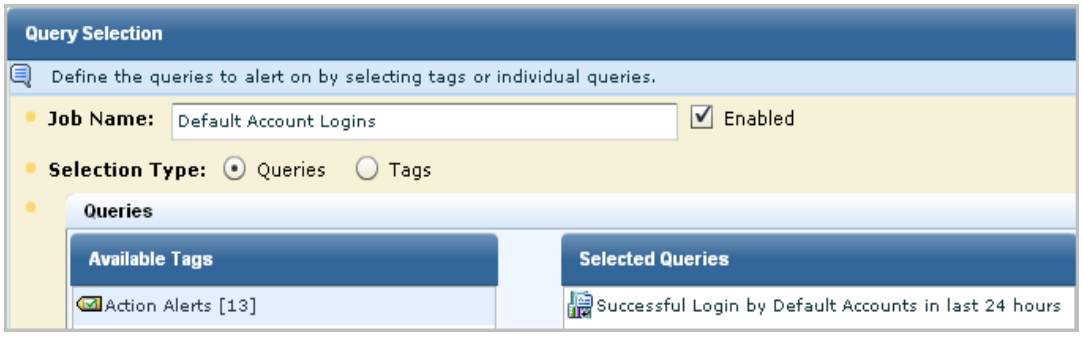

4. Select a date range for running the query and the maximum number of rows to display.
   a. Click Result Conditions.
   b. Select a date range such as 'now' and 'now' '-1 hours'
   c. Select result display parameters such as row limit of 10 and time granularity as event_datetime.
   d. Skip grouped events.
5. Define the schedule.
6. Define the alert data to pass to the IT PAM process along with the event data retrieved by the query.
   a. Click the Destination step.
   b. Select the IT PAM Process tab.
   c. Select Successful Login by Default Account in the last 24 hours.
   d. Select Run IT PAM process per row.
   e. If the configured IT PAM Process is not the one you want to run, change the path for IT PAM Process. The IT PAM process must contain the full path beginning with a forward slash (/).

f. (Optional) Create a summary statement with literal text and variables. Here, the variables are derived from CEG fields when the collected data for a row is refined. Following is an example summary statement using variables.

The (dest_username) account performed the (event_action) action on (dest_hostname)

The first statement is created as follows:

- Type the word, "The"
- Select dest_username from the Select Field drop-down list, then click + next to the Summary field.
- Type the phrase "account performed the"
- Select event_action from the Select Field drop-down list, then click + next to the Summary field.
- Type the phrase "action on"
- Select dest_hostname from the Select Field drop-down list, then click + next to the Summary field.

g. (Optional) Create a description with literal text and text derived from CEG fields. Select the desired field from the Select Field drop-down list and click +. For example:

The (event_logname) log shows the result of (event_result) on (event_datetime)

The(event_result) of the (event_action) is logged in the (event_logname) log.

The (event_logname) log shows the (event_action) action had a result of (event_result).

h. For Send field values as parameters, select each CEG field that the specified IT PAM process uses as a process parameter.

Note: Since the selected process does not use any CEG field names as parameters, no fields are checked in this example. To determine if a custom process uses such parameters, view the Dataset tab in the CA IT PAM event/alert output process.

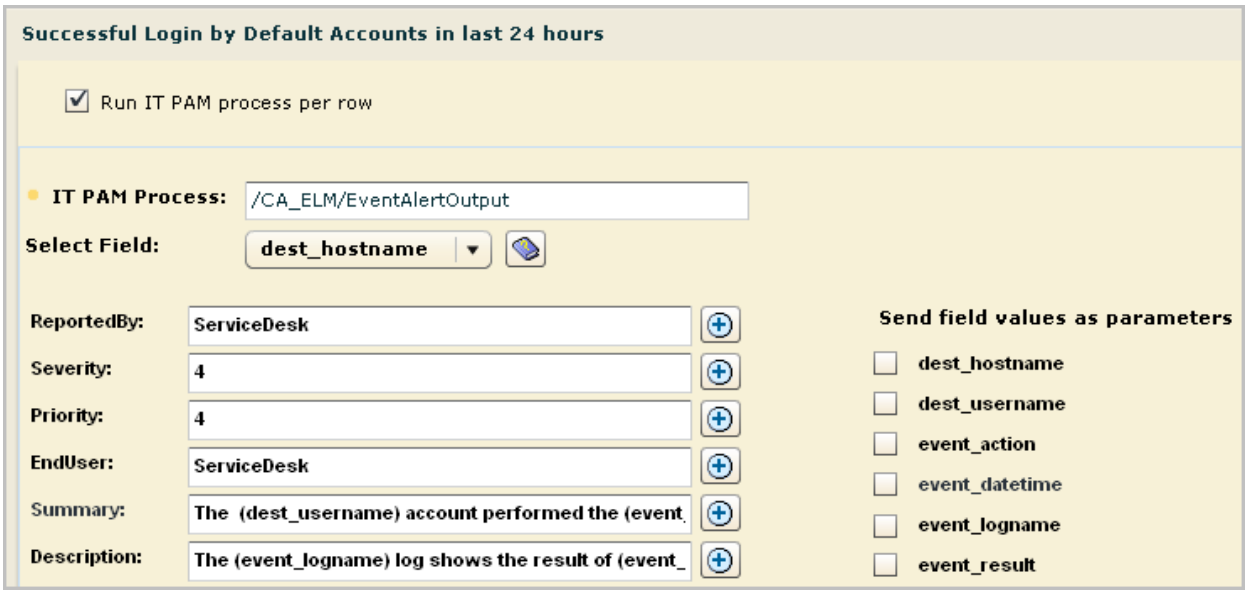

7. Select a Server.
8. Click Save and Close.
   The job appears on the Action Alert Jobs list.

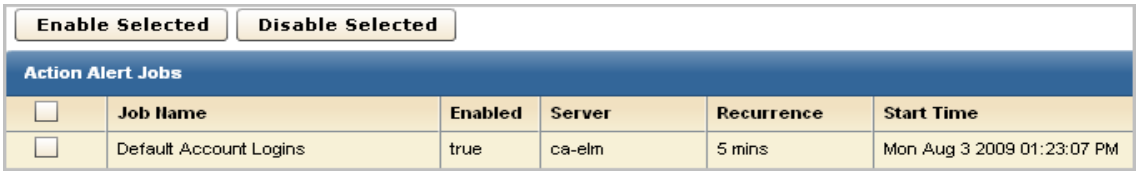

9. Click Alert Management, Self-Monitoring Events to view results. A partial view the information rows follows:

Action	Result	Result Description
Resource Modify	S	Update RSSFeed Alert Name [login attempts] on reportServer [ca-elm] recurrence [5] recurrenceType [Minutes] was Successful.
Alert Creation	S	Alert job [Default Account Logins] created successfully.
Alert Job Setup	S	Schedule Action Query Alert Name [Default Account Logins] on reportServer [ca-elm] was Successful.

10. Click the Alert Management tab, Action Alerts subtab. Select the alert you scheduled to view query results.

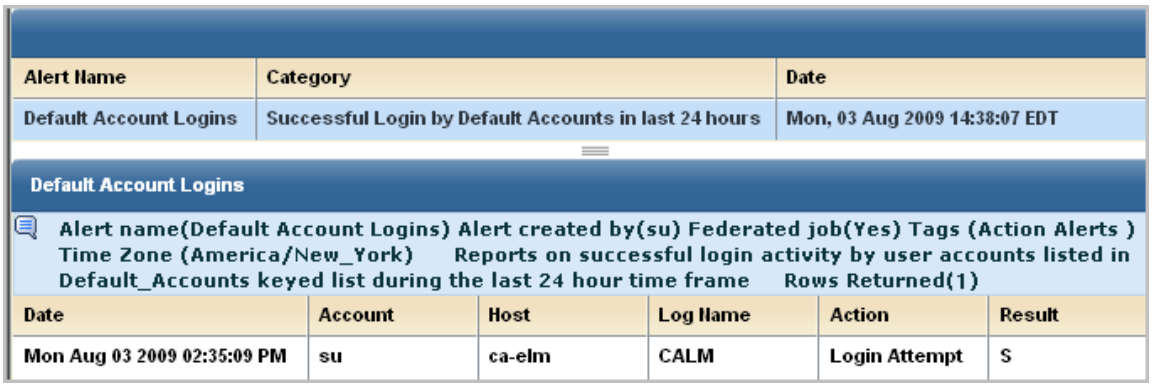

Alert Name	Category	Date
Default Account Logins	Successful Login by Default Accounts in last 24 hours	Mon, 03 Aug 2009 14:38:07 EDT

Default Account Logins

Alert name(Default Account Logins) Alert created by(su) Federated job(Yes) Tags (Action Alerts) Time Zone (America/New_York) Reports on successful login activity by user accounts listed in Default_Accounts keyed list during the last 24 hour time frame Rows Returned(1)

Date	Account	Host	Log Name	Action	Result
Mon Aug 03 2009 02:35:09 PM	su	ca-elm	CALM	Login Attempt	S

11. Check the self-monitoring event tab for results returned from CA IT PAM.

    A partial example of a success message follows, where this message appears in the self monitoring events for the Report Server. Notice the ticket number following Results =.

Action	Result	Result Description
Notification Creation	S	IT PAM process ran successfully. Results = [Request 631 created in CA Service Desk.]

12. (Optional) Review the results on CA Service Desk as follows:
    a. Log on to CA Service Desk.
    b. Select Request and enter the issue number.
    c. Click the request number link to review the issue detail and summary information.

More information:

Guidelines for Creating an Event/Alert Output Process (see page 322)

Example: Send an Alert that Runs an IT PAM Process Per Query

You can send an alert that runs the CA IT PAM event/alert output process per row or per query. This example illustrates the procedure of running the process per query. It includes an example of what can be viewed for this type of alert by personnel working with the third-party product to which CA IT PAM sent the details.

To send an alert that runs the CA IT PAM event/alert output process per query

1. Click the Alert Management tab and then click the Alert Scheduling subtab.
2. Click Schedule an Action Alert.

   The Schedule Action Alerts wizard appears.

3. Complete the Alert Selection step as follows:
   a. Enter the job name.
   b. Select a query.
4. (Optional) Select a date range for running the query and the maximum number of rows to display.
   a. Click Result Conditions.
   b. Select a date range such as 'now' and 'now' '-1 hours'
   c. Select result display parameters.
5. Define the schedule.
6. Define the alert data to pass to the IT PAM process along with the event data retrieved by the query.
   a. Click the Destination step.
   b. Select the IT PAM Process tab.
   c. Select the query to send

   d. If you want results reported by query, leave the Run IT PAM process per row blank.
   e. Optionally, type literal text in the Summary and Description fields.

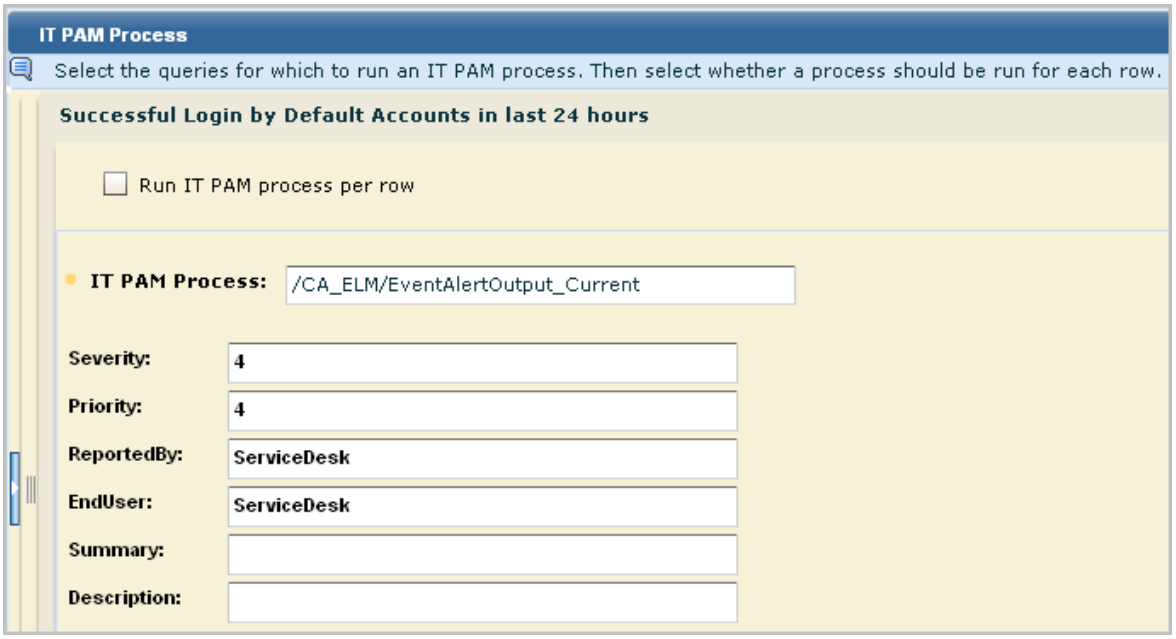

7. Select a Server.
8. Click Save and Close.
   The job appears on the Action Alert Jobs list.
9. Click the Alert Management tab, Action Alerts subtab. Select the alert you scheduled to view query results.
10. Check the self-monitoring event tab for the action, Notification Creation, with results returned from CA IT PAM. A success message includes the Request number created in the third-party application, if it is a help desk product.

Action	Result	Result Description
Notification Creation	S	IT PAM process ran successfully. Results = Request 2936 created in CA Service Desk.

11. (Optional) To see what the help desk personnel sees, review the results on CA Service Desk as follows:
    a. Log on to CA Service Desk.
    b. Select Request and enter the number displayed in the result description for Notification Creation. Click Go.

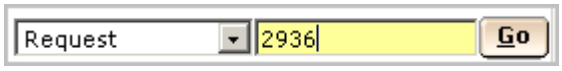

    c. Copy the URL displayed in the Summary Information section and paste it into your browser.

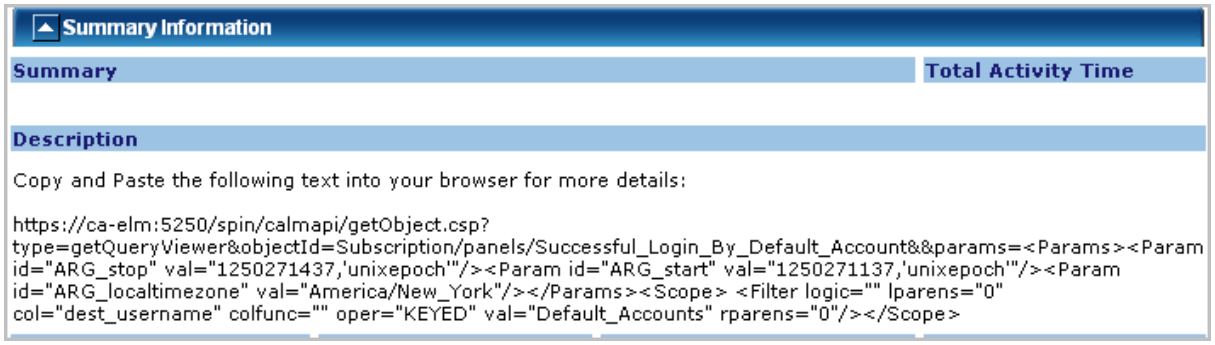

The CA Enterprise Log Manager logon dialog appears.

d. Log into CA Enterprise Log Manager. You can use an account with a low-privilege role such as Auditor.

The event data returned by the query is presented in the format of the default view of the query, that is, table or chart.

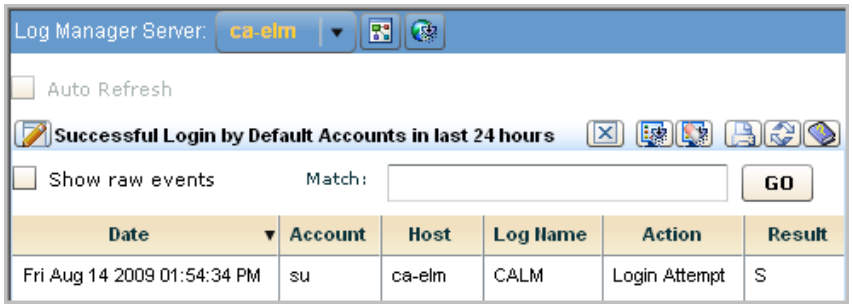

If the display is in table format, you can view raw event data.

More information:

Set Notification Destinations (see page 384)

Working with SNMP Traps

Fault management systems and network operations centers (NOCs) typically receive SNMP traps. You can send alerts to such systems as SNMP v2 traps or SNMP v3 traps, depending on the destination product.

The only required tasks for working with SNMP traps follow:

- Create a custom MIB for each action alert destined for CA NSM.
- Prepare the destination products to receive SNMP traps from CA Enterprise Log Manager.
- Schedule alerts with one or more SNMP trap destinations.

Configuring a default SNMP trap destination is optional.

More information:

Preparing CA Spectrum to Receive SNMP Traps from Alerts (see page 361)
Preparing CA NSM to Receive SNMP Traps from Alerts (see page 370)
Configure Integration with an SNMP Trap Destination (see page 128)
Example: Alerting CA Spectrum of Configuration Changes (see page 365)
Example: Alerting CA NSM of Configuration Changes (see page 373)

About SNMP Traps

SNMP is the acronym for Simple Network Management Protocol, an open standard for sending alert messages to a specified destination. There are three versions of SNMP: SNMPv1, SNMPv2, and SNMPv3. CA Enterprise Log Manager can use either SNMPv2 or SNMPv3 to alert one or more third-party management systems when an event that generates an alert occurs.

In CA Enterprise Log Manager, an alert is generated when a scheduled query returns results from the event log databases of recently refined events. A scheduled query can be configured with SNMP trap as a destination. Trap receivers, the destination management systems, can process traps at the rate of approximately 200 traps per second. Trap receivers typically listen on UDP port 162, the well-known port for snmptrap.

CA Enterprise Log Manager gives you the flexibility to create your own custom alerts to send as SNMP traps. For example, you can define alerts that send notification that a critical event has occurred. You can also define alerts for events such as configuration changes. You decide which alerts to send as SNMP traps.

Example Simple Filters for Alerts to Send as Traps

Events that negatively impact operations, such as shutdown of services, errors on devices, and deletion of resource, are of interest to Network Operations Centers (NOCs). You can generate action alerts when such events occur and route them to your NOC. You can create custom alerts for this purpose using Simple Filters in a custom query. Consider the following simple filter examples.

- Device error

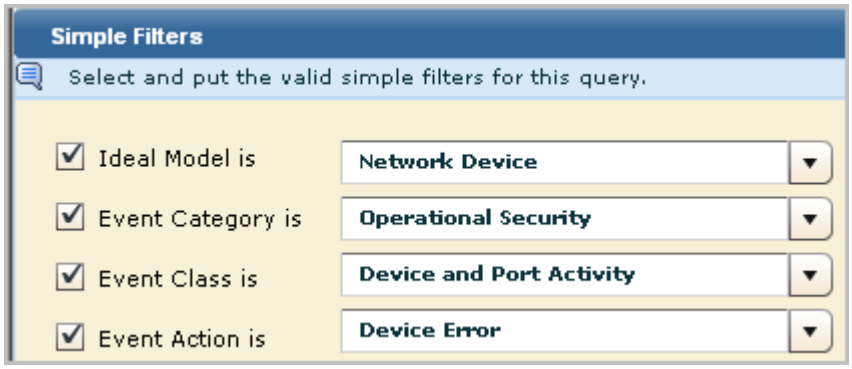

- Resource deletion

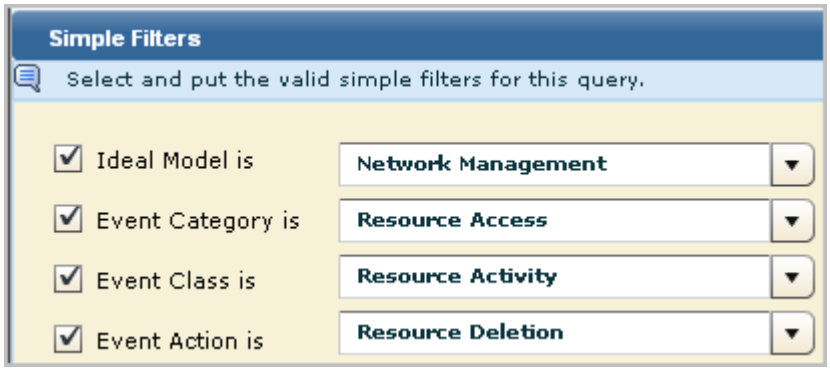

- System shutdown

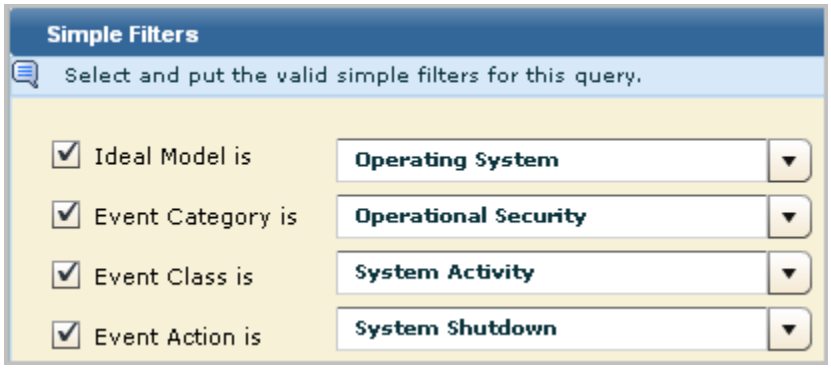

About MIB Files

SNMP traps are defined in either standard Management Information Base (MIB) files or enterprise-specific MIBs.

Each private enterprise on the MIB tree has a unique number that is preceded by the numbers of its parent nodes. IANA assigned CA, Inc. the private enterprise number 791. All data sent in SNMP traps by any CA application is associated with object IDs that begin with 1.3.6.1.4.1.791. The CA Enterprise Log Manager application that belongs to CA has 9845 as its identifier. All SNMP trap data sent by CA Enterprise Log Manager action alerts is associated with object IDs (OIDs) beginning with 1.3.6.1.4.1.791.9845.

CA Enterprise Log Manager provides one MIB file. The name of this MIB is CA-ELM.MIB. This MIB defines all the fields that can be sent by action alerts with one trap. That trap includes all CEG fields available in CA Enterprise Log Manager.

When an action alert is sent to an SNMP trap destination, the data that is sent includes a URL. The individual monitoring incoming traps can browse the URL sent by the action alert. Browsing the URL launches a CA Enterprise Log Manager page that displays query results in an easy to read format. This functionality makes the use of MIBs to interpret data sent as SNMP traps unnecessary.

The CA-ELM MIB Tree

You can view the structure of the CA-ELM.MIB file in the MIB tree form. CEG fields are defined under elmAlertVariables with unique SNMP object identifiers. For example, result_severity has an OID of 1.3.6.1.4.1.791.9845.2.88.

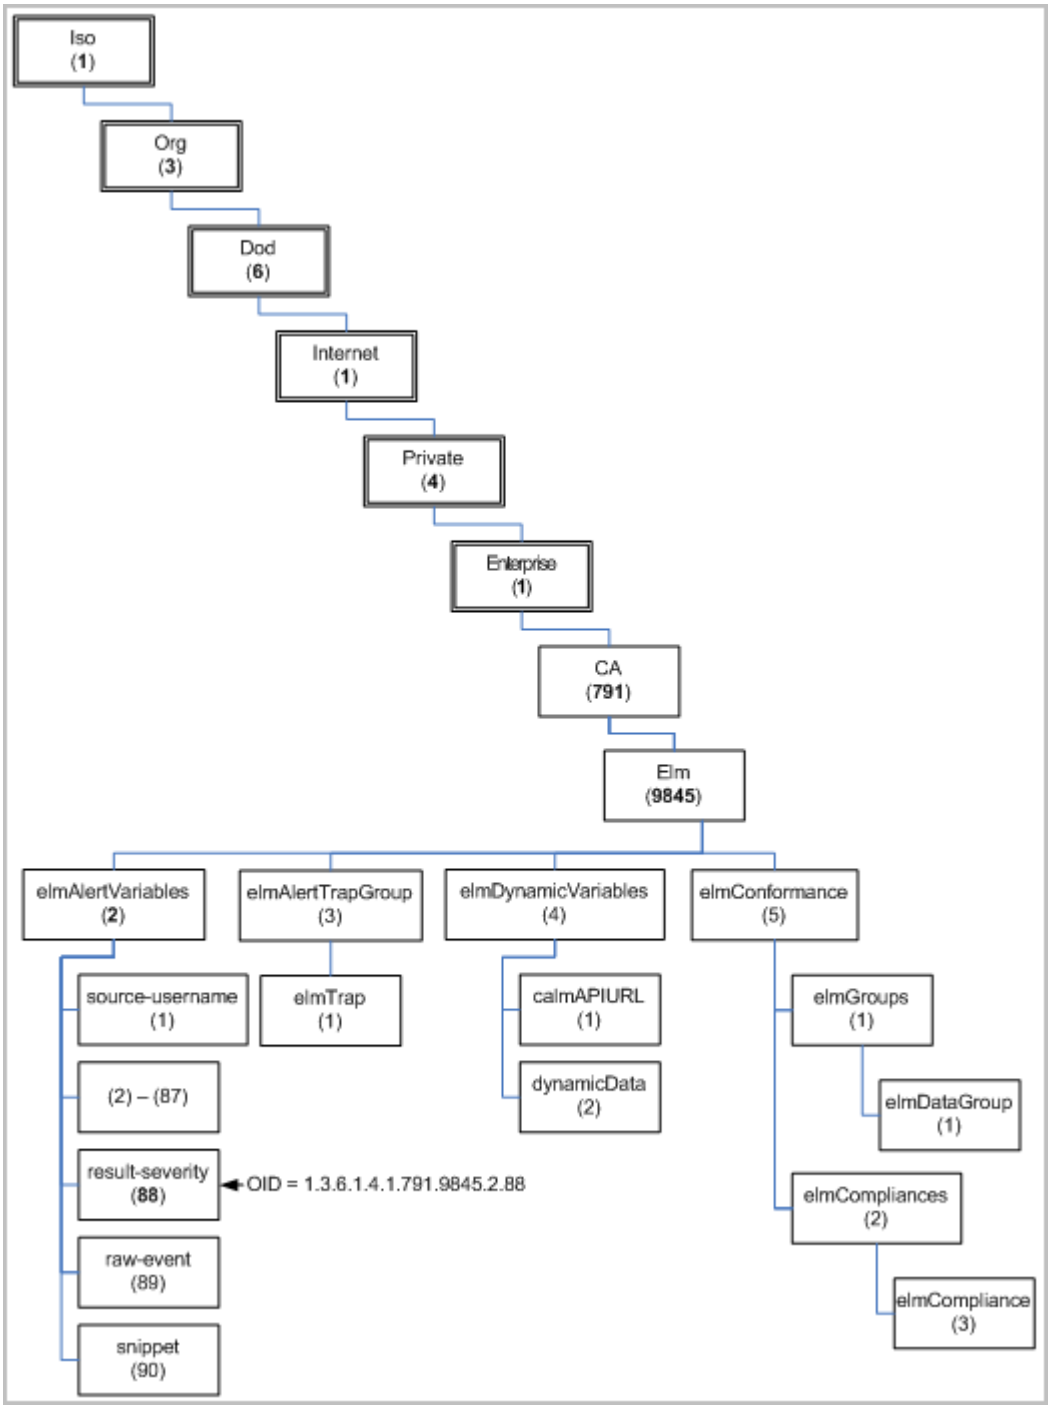

The CA-ELM.MIB File

The CA Enterprise Log Manager MIB file, CA-ELM.MIB, is on the installation DVD. The CA Enterprise Log Manager MIB is generated from the CEG source document, which contains the OIDs for each CEG field (elmAlertVariables).

The CA-ELM.MIB file begins with imports as follows:

```
CAELM-MIB DEFINITIONS ::= BEGIN

 IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE, Integer32, NOTIFICATION-TYPE
        FROM SNMPv2-SMI
    MODULE-COMPLIANCE, OBJECT-GROUP,NOTIFICATION-GROUP
        FROM SNMPv2-CONF                                    DisplayString
        FROM SNMPv2-TC;
```

The following representation is designed to show the structure of the CA Enterprise Log Manager MIB tree, where the top-level nodes include iso(1) org(3) dod(6) internet(1) private(4) enterprises(1). The actual CA-ELM.MIB is not formatted like this representation.

```
ca OBJECT IDENTIFIER::= { enterprises 791 }
        elm MODULE-IDENTITY...::= { ca 9845 }
                elmAlertVariables     ::= { elm 2 }
                        source-username    ::= { elmAlertVariables 1 }
                        source-domainname  ::= { elmAlertVariables 2 }
                        source-groupname   ::= { elmAlertVariables 3 }
            ...
                        result-severity    ::= { elmAlertVariables 88 }
                        raw-event          ::= { elmAlertVariables 89 }
                        snippet            ::= { elmAlertVariables 90 }
                elmAlertTrapGroup     ::= { elm 3 }
                        elmTrap            ::= { elmAlertTrapGroup 1 }
                elmDynamicVariables   ::= { elm 4 }
                        calmAPIURL         ::= { elmDynamicVariables 1 }
                        dynamicData        ::= { elmDynamicVariables 2 }
                elmConformance        ::= { elm 5 }
                        elmGroups          ::= { elmConformance 1 }
                                elmDataGroup   ::= { elmGroups 1 }
                        elmCompliances     ::= { elmConformance 2 }
                                elmCompliance  ::= { elmCompliances 3 }
```

The CA-ELM.MIB file defines one trap. That trap is defined as follows:

```
elmTrap NOTIFICATION-TYPE
   OBJECTS {
source-username,source-domainname,source-groupname,source-uid,source-gid,source-hostname,source-hostd
omainname,source-address,source-mac-address,source-port,source-processname,source-objectname,source-o
bjectattr,source-objectid,source-objectclass,source-objectvalue,dest-username,dest-domainname,dest-groupnam
e,dest-uid,dest-gid,dest-hostname,dest-hostdomainname,dest-address,dest-mac-address,dest-port,dest-objectna
me,dest-objectattr,dest-objectid,dest-objectclass,dest-objectvalue,agent-name,agent-address,agent-hostname,ag
ent-hostdomainname,agent-version,agent-id,agent-connector-name,agent-group,event-source-hostname,event-s
ource-hostdomainname,event-source-address,event-source-processname,receiver-name,receiver-hostname,rec
eiver-hostaddress,receiver-hostdomainname,receiver-port,receiver-time-gmt,receiver-timezone,receiver-version,e
vent-protocol,event-logname,event-euuid,event-count,event-summarized,event-duration,event-time-year,event-ti
me-month,event-time-monthday,event-time-weekday,event-time-hour,event-time-minute,event-time-gmt,event-da
tetime,event-year-datetime,event-month-datetime,event-day-datetime,event-hour-datetime,event-quarterhour-dat
etime,event-minute-datetime,event-timezone,event-sequence,event-trend,event-action,event-id,event-category,ev
ent-class,ideal-model,event-severity,event-result,result-string,result-signature,result-code,result-version,result-priori
ty,result-scope,result-severity,raw-event,snippet }
   STATUS  current
   DESCRIPTION
        "The ELM SNMP Trap."
   ::= { elmAlertTrapGroup 1 }
```

The elmAlertTrapGroup is 1.3.6.1.4.1.791.9845.3 and the elmTrap is defined by the next node. The default elmTrap ID is 1.3.6.1.4.1.791.9845.3.1. User-defined Custom Trap IDs have the range 1.3.6.1.4.1.791.9845.3.2 to 1.3.6.1.4.1.791.9845.3.999.

Important! The best practice for sending traps to CA Spectrum is to use the default elmTrap ID. The best practice for sending traps to CA NSM is to specify a Custom Trap ID that references an elmTrap ID in a custom MIB.

More information:

Object ID (OID) to CEG Field Mapping (see page 346)
Custom MIBs (see page 350)

Object ID (OID) to CEG Field Mapping

The following table shows the CEG field corresponding to each Object ID (OID) under elmAlertVariables in the MIB tree. This branch of the tree will grow as new fields are added to the CEG. Be sure to check for updates to the MIB and be sure the latest version is available to your SNMP trap destination products.

Object ID (OID)	CEG Field
1.3.6.1.4.1.791.9845.2.1	source-username
1.3.6.1.4.1.791.9845.2.2	source-domainname

Object ID (OID)	CEG Field
1.3.6.1.4.1.791.9845.2.3	source-groupname
1.3.6.1.4.1.791.9845.2.4	source-uid
1.3.6.1.4.1.791.9845.2.5	source-gid
1.3.6.1.4.1.791.9845.2.6	source-hostname
1.3.6.1.4.1.791.9845.2.7	source-hostdomainname
1.3.6.1.4.1.791.9845.2.8	source-address
1.3.6.1.4.1.791.9845.2.9	source-mac-address
1.3.6.1.4.1.791.9845.2.10	source-port
1.3.6.1.4.1.791.9845.2.11	source-processname
1.3.6.1.4.1.791.9845.2.12	source-objectname
1.3.6.1.4.1.791.9845.2.13	source-objectattr
1.3.6.1.4.1.791.9845.2.14	source-objectid
1.3.6.1.4.1.791.9845.2.15	source-objectclass
1.3.6.1.4.1.791.9845.2.16	source-objectvalue
1.3.6.1.4.1.791.9845.2.17	dest-username
1.3.6.1.4.1.791.9845.2.18	dest-domainname
1.3.6.1.4.1.791.9845.2.19	dest-groupname
1.3.6.1.4.1.791.9845.2.20	dest-uid
1.3.6.1.4.1.791.9845.2.21	dest-gid
1.3.6.1.4.1.791.9845.2.22	dest-hostname
1.3.6.1.4.1.791.9845.2.23	dest-hostdomainname
1.3.6.1.4.1.791.9845.2.24	dest-address
1.3.6.1.4.1.791.9845.2.25	dest-mac-address
1.3.6.1.4.1.791.9845.2.26	dest-port
1.3.6.1.4.1.791.9845.2.27	dest-objectname
1.3.6.1.4.1.791.9845.2.28	dest-objectattr
1.3.6.1.4.1.791.9845.2.29	dest-objectid
1.3.6.1.4.1.791.9845.2.30	dest-objectclass
1.3.6.1.4.1.791.9845.2.31	dest-objectvalue
1.3.6.1.4.1.791.9845.2.32	agent-name

Object ID (OID)	CEG Field
1.3.6.1.4.1.791.9845.2.33	agent-address
1.3.6.1.4.1.791.9845.2.34	agent-hostname
1.3.6.1.4.1.791.9845.2.35	agent-hostdomainname
1.3.6.1.4.1.791.9845.2.36	agent-version
1.3.6.1.4.1.791.9845.2.37	agent-id
1.3.6.1.4.1.791.9845.2.38	agent-connector-name
1.3.6.1.4.1.791.9845.2.39	agent-group
1.3.6.1.4.1.791.9845.2.40	event-source-hostname
1.3.6.1.4.1.791.9845.2.41	event-source-hostdomainname
1.3.6.1.4.1.791.9845.2.42	event-source-address
1.3.6.1.4.1.791.9845.2.43	event-source-processname
1.3.6.1.4.1.791.9845.2.44	receiver-name
1.3.6.1.4.1.791.9845.2.45	receiver-hostname
1.3.6.1.4.1.791.9845.2.46	receiver-hostaddress
1.3.6.1.4.1.791.9845.2.47	receiver-hostdomainname
1.3.6.1.4.1.791.9845.2.48	receiver-port
1.3.6.1.4.1.791.9845.2.49	receiver-time-gmt
1.3.6.1.4.1.791.9845.2.50	receiver-timezone
1.3.6.1.4.1.791.9845.2.51	receiver-version
1.3.6.1.4.1.791.9845.2.52	event-protocol
1.3.6.1.4.1.791.9845.2.53	event-logname
1.3.6.1.4.1.791.9845.2.54	event-euuid
1.3.6.1.4.1.791.9845.2.55	event-count
1.3.6.1.4.1.791.9845.2.56	event-summarized
1.3.6.1.4.1.791.9845.2.57	event-duration
1.3.6.1.4.1.791.9845.2.58	event-time-year
1.3.6.1.4.1.791.9845.2.59	event-time-month
1.3.6.1.4.1.791.9845.2.60	event-time-monthday
1.3.6.1.4.1.791.9845.2.61	event-time-weekday
1.3.6.1.4.1.791.9845.2.62	event-time-hour

Object ID (OID)	CEG Field
1.3.6.1.4.1.791.9845.2.63	event-time-minute
1.3.6.1.4.1.791.9845.2.64	event-time-gmt
1.3.6.1.4.1.791.9845.2.65	event-datetime
1.3.6.1.4.1.791.9845.2.66	event-year-datetime
1.3.6.1.4.1.791.9845.2.67	event-month-datetime
1.3.6.1.4.1.791.9845.2.68	event-day-datetime
1.3.6.1.4.1.791.9845.2.69	event-hour-datetime
1.3.6.1.4.1.791.9845.2.70	event-quarterhour-datetime
1.3.6.1.4.1.791.9845.2.71	event-minute-datetime
1.3.6.1.4.1.791.9845.2.72	event-timezone
1.3.6.1.4.1.791.9845.2.73	event-sequence
1.3.6.1.4.1.791.9845.2.74	event-trend
1.3.6.1.4.1.791.9845.2.75	event-action
1.3.6.1.4.1.791.9845.2.76	event-id
1.3.6.1.4.1.791.9845.2.77	event-category
1.3.6.1.4.1.791.9845.2.78	event-class
1.3.6.1.4.1.791.9845.2.79	ideal-model
1.3.6.1.4.1.791.9845.2.80	event-severity
1.3.6.1.4.1.791.9845.2.81	event-result
1.3.6.1.4.1.791.9845.2.82	result-string
1.3.6.1.4.1.791.9845.2.83	result-signature
1.3.6.1.4.1.791.9845.2.84	result-code
1.3.6.1.4.1.791.9845.2.85	result-version
1.3.6.1.4.1.791.9845.2.86	result-priority
1.3.6.1.4.1.791.9845.2.87	result-scope
1.3.6.1.4.1.791.9845.2.88	result-severity
1.3.6.1.4.1.791.9845.2.89	raw-event

Custom MIBs

You can create custom MIB files from the provided boilerplate text by adding selected varbinds from the CA-ELM.MIB file content. A custom MIB file for a single alert contains a subset of the contents of the CA-ELM.MIB file. A custom MIB for an alert differs from CA-ELM.MIB in these ways:

- The custom MIB defines only the fields sent by that alert.
- The custom MIB defines a trap that lists these fields in the sequence in which they are sent.
- The custom MIB trap is defined with the OID, 1.3.6.1.4.1.791.9845.3.x, where x is a value between 1 and 999.

  Note: An alert that uses a custom MIB specifies this OID as the value for Custom Trap ID.
- A custom MIB includes the dynamicData varbind *only if* the query includes calculated fields.

  Calculations can be applied to any field. The event_count field is an example field to which calculations are commonly, but not always, applied. Event_count in the query, System Event Count by Event Action, is a calculated field; it is calculated with Sum. To determine whether a field is calculated, examine the query where the field is defined. An example of a definition of event_count where it is a calculated field follows:

  System_Event_Count_By_Event_Action.xml: <Column columnname="event_count" datatype="I" displayname="Count" **functionname="sum"** resultname="event_count" sortdesc="true" sortorder="1" visible="true"/>

Boilerplate Text for a Custom MIB

Boilerplate text for a custom MIB follows. If you start a custom MIB with this example, you can replace or insert custom data in locations indicated with the string ###. In sections where you modify data, you can, optionally, modify the description.

```
CAELM-MIB DEFINITIONS ::= BEGIN
 IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE, Integer32, NOTIFICATION-TYPE
        FROM SNMPv2-SMI
MODULE-COMPLIANCE, OBJECT-GROUP,NOTIFICATION-GROUP
        FROM SNMPv2-CONF                                    DisplayString
       FROM SNMPv2-TC;
```

```
elm MODULE-IDENTITY
   LAST-UPDATED "200907050600Z"
   ORGANIZATION "CA"
   CONTACT-INFO
      "100 Staples drive
      Framingham MA"
   DESCRIPTION
      "Contains objects describing data for ELM events"
   REVISION "200907050600Z"
   DESCRIPTION
      "Custom MIB <###>."
   ::= { ca 9845 }

ca OBJECT IDENTIFIER ::= {enterprises 791}
elmAlertTrapGroup OBJECT IDENTIFIER ::= { elm 3 }
elmAlertVariables OBJECT IDENTIFIER ::= { elm 2 }
elmDynamicVariables OBJECT IDENTIFIER ::= { elm 4 }
elmConformance OBJECT IDENTIFIER ::= { elm 5 }
elmGroups     OBJECT IDENTIFIER ::= { elmConformance 1 }
elmCompliances OBJECT IDENTIFIER ::= { elmConformance 2 }


<### Insert elmAlertVariable varbind for each query field ###>

<### Insert the following dynamicData varbind only if query includes calculated fields ###>
dynamicData OBJECT-TYPE
   SYNTAX  DisplayString
   MAX-ACCESS  read-only
   STATUS     current
   DESCRIPTION
" This field contains all the elm dynamic variables and data in name=value format."
   ::= { elmDynamicVariables 2 }

calmAPIURL OBJECT-TYPE
   SYNTAX  OCTET STRING
   MAX-ACCESS  read-only
   STATUS     current
   DESCRIPTION
"The OPEN API URL which points to the query result."
   ::= { elmDynamicVariables 1 }

elmTrap NOTIFICATION-TYPE
   OBJECTS { <### insert list of query fields with hyphens ###> }
   STATUS  current
   DESCRIPTION
         "The ELM SNMP Trap."
   ::= { elmAlertTrapGroup <### insert custom trap ID node number ###> }
```

```
elmCompliance MODULE-COMPLIANCE
   STATUS  current
   DESCRIPTION
         "The compliance information."
   MODULE  -- this module
      GROUP      elmDataGroup
      DESCRIPTION
         "This group is mandatory."
   ::= { elmCompliances 3 }
-- units of conformance

elmDataGroup    OBJECT-GROUP
   OBJECTS { <### insert list of query fields with hyphens ###> }
   STATUS  current
   DESCRIPTION
         "A collection of objects providing information specific to
         ELM data."
   ::= { elmGroups 1 }
END
```

Example: Create Custom MIB 33 for the Average CPU Load Trend Query

Create a custom MIB for each query sent to CA NSM as an SNMP trap. Each such query is associated with a custom trap ID. The custom MIB defines the fields selected to include in the trap in the order displayed in the action alert.

Consider the example where the query selected for the action alert is Average CPU Load Trend. The selected fields are event_datetime and event_trend.

The Custom Trap ID is 1.3.6.1.4.1.791.9845.3.33.

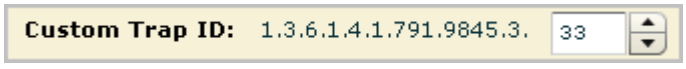

To create a custom MIB for the custom trap ID ending in 33

1. Open a copy of CA-ELM.MIB for the purpose of copying text to your custom MIB.
2. Open an editor, copy the boilerplate text for custom MIB, and save the file as a new name. For example, save it as Custom MIB n.mib, where n is 33, the final node of the Custom Trap ID specified for the query in the action alert.

3. (Optional) Under elm MODULE-IDENTITY, replace <###> with 33. For example:

   Custom MIB 33."

4. Replace the following boilerplate text with text from CA-ELM.MIB

   <### Insert elmAlertVariable varbind for each query field in trap sequence ###>

   Copy the elmAlertVariable varbinds for event_datetime and then for event_trend. These varbinds must appear in the MIB in the same sequence that they are sent in the SNMP trap. For example:

```
event-datetime OBJECT-TYPE
   SYNTAX  DisplayString
   MAX-ACCESS  read-only
   STATUS      current
   DESCRIPTION
"The calendar date and time expressed in the event information"
   ::= { elmAlertVariables 65 }

event-trend OBJECT-TYPE
   SYNTAX  Integer32
   MAX-ACCESS  read-only
   STATUS      current
   DESCRIPTION
"Trending data for this event."
   ::= { elmAlertVariables 74 }
```

5. Because neither of the fields in this query are calculated fields, delete the following boilerplate text:

```
<### Insert the following dynamicData varbind only if query includes calculated fields ###>
dynamicData OBJECT-TYPE
   SYNTAX  DisplayString
   MAX-ACCESS  read-only
   STATUS      current
   DESCRIPTION
" This field contains all the elm dynamic variables and data in name=value format."
   ::= { elmDynamicVariables 2 }
```

6. Replace the following boilerplate text under elmTrap:

   OBJECTS { <### insert list of query fields with hyphens ###> }

   with the list of selected query fields, as follows:

   OBJECTS { event-datetime,event-trend }

7. Replace the following boilerplate text under elmTrap:

   ::= { elmAlertTrapGroup <### insert custom trap ID node number ###> }

   with the following:

   ::= { elmAlertTrapGroup 33 }

8. Replace the following boilerplate text under elmDataGroup:

   OBJECTS { <### insert list of query fields with hyphens ###> }

   with the following:

   OBJECTS { event-datetime,event-trend }

9. Save the file.

Example: Custom MIB 33

The following example is a custom MIB developed for an action alert sent as an SNMP trap with the Custom Trap ID ending in 33. The custom trap ID was 1.3.6.1.4.1.791.9845.3.33. The selected query was Average CPU Load Trend and the fields selected to be sent in the SNMP trap are event_datetime, and event_trend.

```
CAELM-MIB DEFINITIONS ::= BEGIN
 IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE, Integer32, NOTIFICATION-TYPE
        FROM SNMPv2-SMI
MODULE-COMPLIANCE, OBJECT-GROUP,NOTIFICATION-GROUP
        FROM SNMPv2-CONF                                    DisplayString
       FROM SNMPv2-TC;

elm MODULE-IDENTITY
  LAST-UPDATED "200907050600Z"
  ORGANIZATION "CA"
  CONTACT-INFO
    "100 Staples drive
    Framingham MA"
  DESCRIPTION
    "Contains objects describing data for ELM events"
  REVISION "200907050600Z"
  DESCRIPTION
    "Custom MIB 33."
  ::= { ca 9845 }

ca OBJECT IDENTIFIER ::= { enterprises 791}
elmAlertTrapGroup OBJECT IDENTIFIER ::= { elm 3 }
elmAlertVariables OBJECT IDENTIFIER ::= { elm 2 }
elmDynamicVariables OBJECT IDENTIFIER ::= { elm 4 }
elmConformance OBJECT IDENTIFIER ::= { elm 5 }
elmGroups     OBJECT IDENTIFIER ::= { elmConformance 1 }
elmCompliances OBJECT IDENTIFIER ::= { elmConformance 2 }
```

```
event-datetime OBJECT-TYPE
   SYNTAX  DisplayString
   MAX-ACCESS  read-only
   STATUS      current
   DESCRIPTION
"The calendar date and time expressed in the event information"
   ::= { elmAlertVariables 65 }

event-trend OBJECT-TYPE
   SYNTAX  Integer32
   MAX-ACCESS  read-only
   STATUS      current
   DESCRIPTION
"Trending data for this event."
   ::= { elmAlertVariables 74 }

calmAPIURL OBJECT-TYPE
   SYNTAX  OCTET STRING
   MAX-ACCESS  read-only
   STATUS      current
   DESCRIPTION
"The OPEN API URL which points to the query result."
   ::= { elmDynamicVariables 1 }

elmTrap NOTIFICATION-TYPE
   OBJECTS { event-datetime,event-trend }
   STATUS  current
   DESCRIPTION
        "The ELM SNMP Trap."
   ::= { elmAlertTrapGroup 33 }

elmCompliance MODULE-COMPLIANCE
   STATUS  current
   DESCRIPTION
        "The compliance information."
   MODULE  -- this module
      GROUP       elmDataGroup
      DESCRIPTION
         "This group is mandatory."
   ::= { elmCompliances 3 }
-- units of conformance
```

```
elmDataGroup    OBJECT-GROUP
   OBJECTS { event-datetime,event-trend }
   STATUS  current
   DESCRIPTION
        "A collection of objects providing information specific to
        ELM data."
   ::= { elmGroups 1 }
END
```

Example: MIB Tree for Custom MIB 33

The MIB tree for a custom MIB differs from the MIB tree for CA-ELM.MIB in the following ways:

- The objects under elmAlertVariables are limited to the fields in the query. Consider the example where the selected fields are:
  - event_datetime
  - event_trend
- The elmTrap that contains only the query fields replaces elmTrap (1) in CA-ELM.MIB, which contains all CEG fields. Consider the example where the elmTrap is 33. This elmTrap identifier references the final node of the Custom Trap ID, which is 1.3.6.1.4.1.791.9845.3.33.

A depiction of Custom MIB 33 in MIB tree format follows, where the highlighted blocks indicate differences between this custom MIB and the CA-ELM.MIB. The custom MIB defines only two fields under elmAlertVariables. The custom elmTrap includes only the two query fields and has a unique number, 33. The elmDataGroup includes only the two query fields.

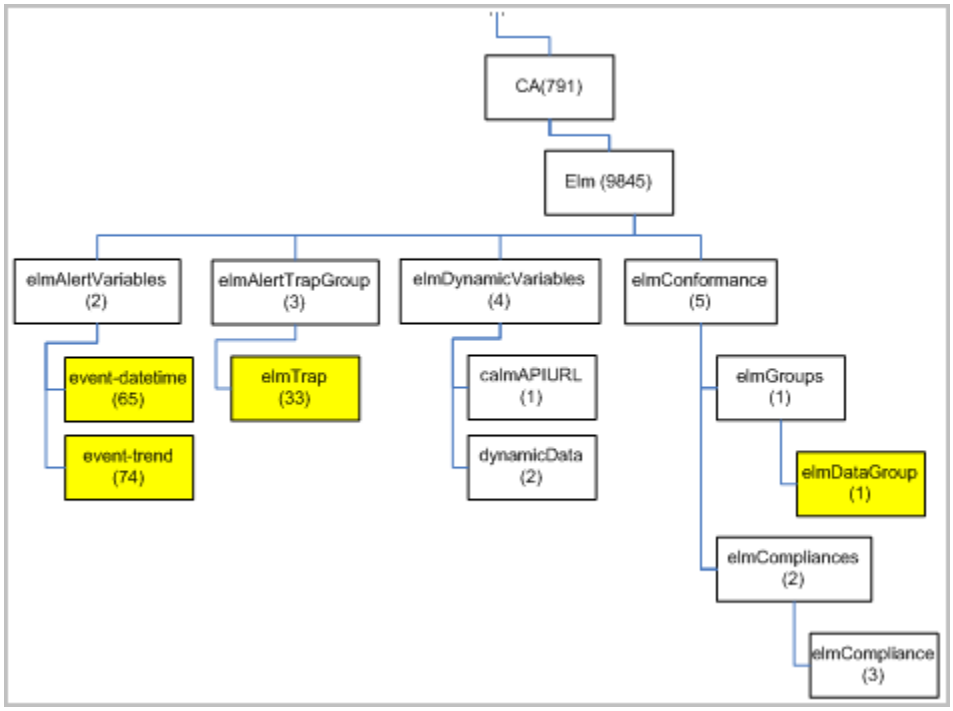

MIB Usage Considerations

For a system to understand an SNMP trap it receives from CA Enterprise Log Manager using MIBs, it must know what the composing OIDs define. Requirements follow:

- Import and compile the CA-ELM.MIB file; keep this file updated.
- (CA NSM only) Create, import, and compile one custom MIB for each scheduled alert sent as an SNMP trap.

  Note: The same custom MIB can be used for alerts based on queries that send the same fields in the same order as a trap.

  For example, the following queries all return values for the fields dest_hostname and event_count.

  - Five Failed System Access by System
  - Failed Restore Activity in Last Seven Days by Host Summary
  - Failed Backup Activity in Last Seven Days by Host Summary
  - Excessive (25) Configuration Failures in Last Hour
  - Excessive (25) Configuration Changes in Last Hour
  - Excessive (25) SU Activity by Host in Last Hour
  - Critical or Fatal Events on Critical Host Summary

  If you create separate alerts based on these queries, those alerts would specify the same Custom Trap ID and would be interpreted with the same custom MIB.

Individuals who monitor the SNMP traps received at the destination product can interpret traps sent from CA Enterprise Log Manager in two ways:

- Launch the SNMP trap results page from the URL sent in the trap.
- Use an application that references imported MIBs.

Process of Working with SNMP Traps

Using SNMP traps involves the following procedures:

1. Prepare CA Enterprise Log Manager to send SNMP traps.
   - Configure the default SNMP trap destination.
   - Identify the IP address and port of each additional SNMP trap destination that you can specify when sending alerts as SNMP traps.
   - Identify alerts where query results would be of interest to CA Spectrum, CA NSM, or another SNMP trap receiver.

2. Prepare the SNMP trap destination products to receive SNMP traps from CA Enterprise Log Manager
   - If CA Spectrum is to be a destination:
     - Create an Event Model according to Spectrum documentation. Without an event model, you cannot view trap results at the destination.
     - Configure CA Spectrum to receive SNMP v3 traps.
   - If CA NSM is to be a destination:
     - Install NSM r11.2 GA build on a Windows Server 2003 EE SP1 and apply the patch to update the aws_snmpex.dll file.
     - Configure CA NSM to receive SNMP traps, including SNMP v3 traps.
3. (Optional) Prepare the SNMP trap destination to interpret SNMP traps from CA Enterprise Log Manager with MIBs.
   - Download the CA Enterprise Log Manager MIB to a location accessible from your SNMP trap destination product.

     Note: The CA-ELM.MIB is delivered on the installation DVD. You can download the latest version of this MIB from the CA Enterprise Log Manager product page.
   - Import and compile the CA-ELM.MIB file.
   - (CA NSM only.) For each alert to be sent as an SNMP trap, create a custom MIB with a trap defined with 1.3.6.1.4.1.791.9845.3.x, where x is a whole number equal to or less than 999. Import and compile all custom MIBs.

   Important! This step is optional because traps received from CA Enterprise Log Manager can be interpreted by launching the trap results page from the URL sent in the trap.
4. Schedule alerts with SNMP trap destinations.
5. Verify that the alert was successfully sent as an SNMP trap.
6. (Optional) Monitor results of sent SNMP traps from the trap destination.
   - View the SNMP trap results at the trap destination.
   - Launch the URL to view the data sent by the alert in chart or graph format.

More information:

Configure Integration with an SNMP Trap Destination (see page 128)
Preparing CA Spectrum to Receive SNMP Traps from Alerts (see page 361)
Preparing CA NSM to Receive SNMP Traps from Alerts (see page 370)
Send SNMPv2 Traps to CA Spectrum (see page 366)

Configure Integration with an SNMP Trap Destination

Configure SNMP integration as part of the Global Service Configuration for Report Server. The configuration is the IP address and port of one SNMP trap destination.

You can configure SNMP integration either before or after preparing the destination product to receive and interpret SNMP traps from CA Enterprise Log Manager.

When you create an alert destined for an SNMP trap recipient, you can specify one or more destinations. This configuration serves as the default. This default applies to all servers listed under Report Server.

To configure SNMP integration

1. Click the Administration tab and the Services subtab.
2. Click Alerting Service.
3. Enter the IP address or host name of the destination server for the SNMP traps.
4. Accept the default port, 162, or change it.
5. Click Save.

Preparing CA Spectrum to Receive SNMP Traps from Alerts

You can send alerts in the form of SNMP traps from CA Enterprise Log Manager to any destination in your network that receives and interprets traps. Each trap receiver product has its own requirements.

Prepare CA Spectrum to receive traps from CA Enterprise Log Manager action alerts by:

- Creating a CA Spectrum southbound gateway integration, an integration point that can support any incoming alert data stream format from a third-party system, including SNMP traps such as those generated by CA Enterprise Log Manager.
- Creating a model of the CA Enterprise Log Manager node to enable receipt of SNMP v3 traps.
- Downloading the CA Enterprise Log Manager MIB.
- Importing the CA Enterprise Log Manager MIB into CA Spectrum.

The process for creating a southbound gateway integration is fully documented in the Spectrum *Southbound Gateway Toolkit Guide*. Creating a southbound gateway integration includes mapping SNMP traps to CA Spectrum events in an AlertMap file and defining the required models. The southbound gateway integration point accepts alert data from third-party systems and displays it within OneClick.

After downloading the MIB file from the CA Enterprise Log Manager product page on Support Online or retrieving it from the installation DVD, you can import it into CA Spectrum. For more information about using the MIB Tool for importing in CA Spectrum OneClick, see the *CA Spectrum Device Management User Guide*.

More information:

Configure CA Spectrum to Accept SNMP v3 Traps (see page 362)
Download the CA Enterprise Log Manager MIB (see page 364)
Import the CAELM-MIB into CA Spectrum (see page 364)

Configure CA Spectrum to Accept SNMP v3 Traps

Before you can send SNMP V3 traps from CA Enterprise Log Manager to CA Spectrum, you must create a model of the CA Enterprise Log Manager appliance in CA Spectrum. SNMP v3 traps are then directed to the CA Enterprise Log Manager node that you modeled.

To create a model that enables Spectrum to receive SNMP v3 traps from action alerts

1. Log on to the Windows server where CA Spectrum is installed.
2. Access the Spectrum OneClick console:
   a. From the Start menu, click All Programs, CA, SPECTRUM Control Panel.

      The SPECTRUM Control Panel appears with a Status indicator at the bottom of the screen.

   b. If Status does not display RUNNING, click Start SpectroSERVER under Process Control.

   c. When Status displays RUNNING, click OneClick Administration.

      OneClick Administration - SPECTRUM Control Panel appears with Host as localhost and Port as 80.

   d. Click OK

      A login dialog appears.

   e. Provide your credentials.

      The SPECTRUM NFM OneClick page appears.

   f. Click Start Console.

      The Login - SPECTRUM OneClick login dialog appears to connect you to SPECTRUM OneClick on local host.

   g. Click OK

      The Console - SPECTRUM OneClick appears with a Navigation pane, a Contents pane, and a Component Detail pane.
3. On the Explorer tab in the Navigation pane, expand the top-level node and select Universe.

   The Contents and Component Detail pane titles display Universe of type Universe.
4. On the Contents pane, click the Topology tab.

   The second button on the tab lets you create a new model by type and add it to this view.
5. Click Create a new model.

   The Select Model Type - SPECTRUM OneClick dialog appears.

6. Click the All Model Types tab
7. Type a string in the Filter field. For example, type gn.

   Model types beginning with Gn appear in the list.
8. Select the desired model type and click OK. For example, select GnSNMPDev and click OK.

   The Create Model of Type <selected model type> opens.
9. Complete the Create Model of Type dialog as follows:
   a. Enter the host name of a CA Enterprise Log Manager server in the Name field.
   b. Enter the IP address of the same server in the Network Address field.
   c. Enter a port in the Agent Port field, if the default 161 is not what you want. For example, enter 162.
   d. Select SNMP v3 as the SNMP Communication option.
   e. Click Profiles.

   The Edit SNMP v3 Profiles window appears with a list of existing profiles, if any.
10. To add a profile, follow these steps:
    a. Type the profile name and type the User ID.
    b. Since this is for SNMP v3, select Authentication with Privacy as the authentication type.
    c. In the next four fields, type an 8-character authentication password twice and type an 8-character privacy password twice.
    d. Click Add to add the profile to the list.
    e. Click OK.

    The profile you added appears first in the V3 Profile drop-down list on the Create Model of Type dialog.
11. Select Discover Connections and click OK.

    The Creating Model progress indicator appears. When processing completes, the created model appears on the Topology tab as a graphic with the host name that you entered and the model type you selected.

Download the CA Enterprise Log Manager MIB

You can download the MIB file from the CA Enterprise Log Manager product page on Support Online or you can retrieve it from the installation DVD. After downloading the CA Enterprise Log Manager MIB, you can import/compile it into each product you configure as an SNMP trap destination.

To download the CA Enterprise Log Manager MIB

1. Log on to the server where you have installed CA Spectrum
2. Launch CA Support Online and log on.
3. Access the CA Enterprise Log Manager product page.
4. Download the CA Enterprise Log Manager MIB file to your network.
5. If you plan to send SNMP traps to CA Spectrum, import the CA Enterprise Log Manager MIB into CA Spectrum.
6. If you plan to send SNMP traps to CA NSM, import the CA Enterprise Log Manager MIB into CA NSM. Refer to the CA NSM documentation for the procedure.

Import the CAELM-MIB into CA Spectrum

Before you send SNMP traps from CA Enterprise Log Manager to CA Spectrum, you can import and compile the CA Enterprise Log Manager MIB using the CA Spectrum OneClick MIB Tools.

Note: The SNMPv2 MIBs referenced in the CA-ELM.MIB are preloaded in CA Spectrum.

To import the CA-ELM.MIB into CA Spectrum

1. Log on to CA Spectrum.
2. Launch the OneClick Console.
3. Click Tools, Utilities, MIB Tools.

   The MIB Tools: Add MIB dialog opens.
4. Click Browse, navigate to the location where you downloaded CA-ELM.MIB, and select this file.
5. Click Compile.

   A success message indicates that the CA Enterprise Log Manager MIB is successfully stored in the following directory on the OneClick web server:

   <*$SPECROOT*>/MibDatabase/userContrib

6. Close the MIB Tools: Add MIB dialog.

   CAELM-MIB is added to the Navigation bar under CA.

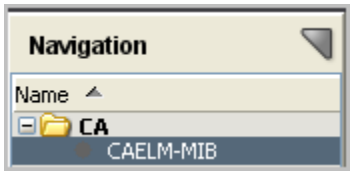

   In the hierarchy, cai expands to display elm with its subordinate tree objects and their associated OIDs.

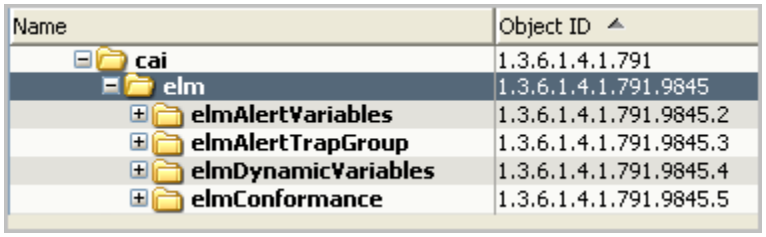

Example: Alerting CA Spectrum of Configuration Changes

Before you send SNMP traps to CA Spectrum for the first time, it is a good practice to identify the queries that return results pertinent to this destination. When you schedule your first alert with Spectrum as a destination, you may want to track the progress and compare the results displayed in CA Enterprise Log Manager with those that appear in the CA Spectrum interface. Once sending traps to CA Spectrum becomes routine, you may not ever take these preparation and follow-up steps again.

The following example is designed to walk you through the initial process, including:

- Preparing to send SNMP Traps to CA Spectrum
- Sending SNMP traps to CA Spectrum
- Verifying SNMP traps were sent successfully
- Viewing the SNMP traps received by CA Spectrum

More information:

Send SNMPv2 Traps to CA Spectrum (see page 366)
Track the Alert Job Progress (see page 368)
View SNMP Traps on CA Spectrum (see page 369)

Send SNMPv2 Traps to CA Spectrum

The following example shows how to create an alert that notifies CA Spectrum of configuration changes with SNMPv2 traps.

To send SNMPv2 traps to CA Spectrum

1. Open the Alert Scheduling wizard.
   a. Click the Alert Management tab and the Alert Scheduling subtab.
   b. Click the Schedule an Action Alert button.
2. Complete the Alert Selection step.
   a. Type a job name; this is required for any alert.
   b. Verify selection type is Queries.

      Selection of SNMP trap destinations is not allowed for alerts based on tags.
   c. If the queries you want to select are tagged Action Alerts, click the Action Alerts tag to filter the displayed list.
   d. Select the query or queries you identified.

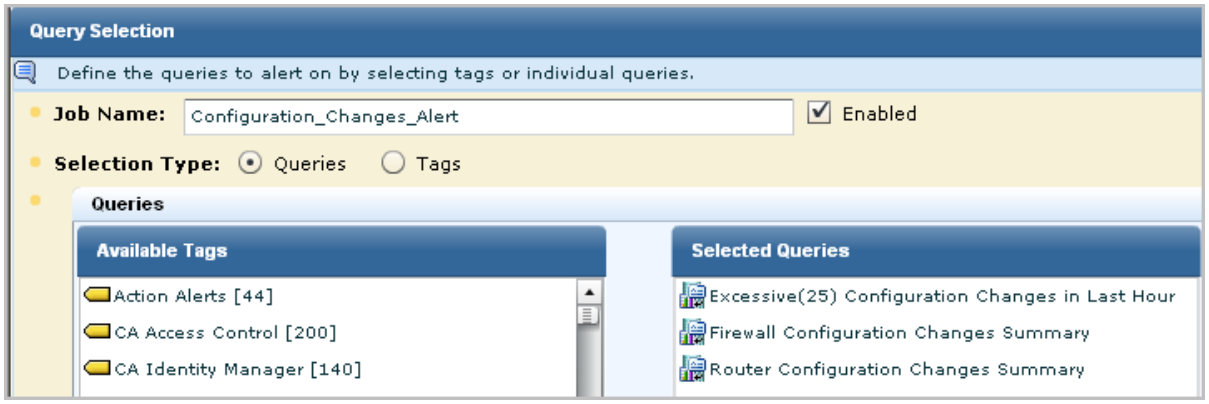

3. (Optional) Complete the Alert Filters, Result Conditions, and Schedule Jobs steps as documented in the online help for this wizard.

4. Set the SNMP trap details.

   a. Click the Destination step.

   b. Click the SNMP Trap tab.

      The configured SNMP Trap destination and the queries selected in step 1 of the wizard appear.

      Note: By default, the SpectroSERVER listens on the standard SNMP trap port 162. If changed, the port must match the snmp_trap_port parameter in the SPECTRUM .vnmrc file located in the SS directory.

   c. (Optional). To send the trap to up to nine servers in addition to the configured destination server, click the Add button and enter the IP address and port of the server.

   d. For a query where you want all fields included in the trap, just select the query.

      All fields of a selected query are selected by default. The name of the selected query appears above the field list.

   e. For a query where you want selected fields included in the trap, select the query and clear the fields that are not to be sent.

f. Select the SNMP Version supported by the selected trap destination for traps received from applications.

Note: Some trap destination accept Version 3 traps sent directly by devices, but only Version 2 traps from applications that collect events from devices. For this example, we accept Version 2.

5. Select the server and specify whether the query should return results from just selected server(s) or from this server and all of its child (if hierarchical) or peer (if meshed) federated servers.

6. Click Save and Close.

   The job appears on the Action Alert Jobs list. Unless you cleared the Enabled check box on the first step of the wizard, it is displayed as enabled (true in the Enabled column). An abbreviated example follows:

Action Alert Jobs

	Job Name	Enabled	Server	Recurrence	Start Time	End Time	Time Zone	Creator
☐	Configuration_Changes_Alert	true	ca-elm	5 mins	Thu Jun 25 2009 01:59:28 PM		America/New_York	su

Track the Alert Job Progress

You can view results returned by the queries selected for the alert you created. The results displayed for the example Configuration_Changes_Alert are displayed in CA Enterprise Log Manager under the headings Host and Count.

1. Select the Alert Management tab, and the Action Alert subtab.
2. Click the name of the alert that you scheduled.
3. View the results for that alert.

   Example results follow:

Alert Name	Category	Date
Configuration_Changes_Alert	Excessive(25) Configuration Changes in Last Hour	Thu, 25 Jun 2009 14:59:28 EDT

Configuration_Changes_Alert

Alert name(Configuration_Changes_Alert) Alert created by(su) Federated job(Yes) Tags (Action Alerts) Time Zone (America/New_York) Lists Excessive Configuration Changes (more than 25) in last hour. Rows Returned(1)

Host	Count
ca-elm	2

View SNMP Traps on CA Spectrum

You can view the SNMP traps sent by CA Enterprise Log Manager alerts on the CA Spectrum event model you created for receiving these traps. Received traps are displayed on the Events tab. For the example Configuration_Changes_Alert, the results, ca-elm and 2, are displayed in CA Spectrum with the OIDs 1.3.6.1.4.1.791.9845.2.22 and 1.3.6.1.4.1.791.9845.2.2.

To view SNMP traps on CA Spectrum

1. Log in to CA Spectrum with your CA Spectrum credentials.
2. Bring up the Spectrum Control Panel and start Spectroserver.

   Spectroserver starts.
3. Click OneClick Administrator and log in.

   The Spectrum NFM OneClick application appears.
4. Click Start Console.

   The Spectrum OneClick console appears.
5. Expand the folder created for CA Enterprise Log Manager.
6. Under Universe, select the event model you created for receiving traps sent from CA Enterprise Log Manager.
7. In the right-hand panel, select the Events tab to view traps sent from CA Enterprise Log Manager.

   The value, ca-elm, and event_count=2 is the same data that you could view in CA Enterprise Log Manager.

An unrelated example of how an SNMP trap sent by a CA Enterprise Log Manager alert appears in CA Spectrum OneClick follows. The link is the URL you can paste in a browser to display the CA Enterprise Log Manager table with details presented in CEG format.

More information:

Example: Alerting CA Spectrum of Configuration Changes (see page 365)

Preparing CA NSM to Receive SNMP Traps from Alerts

You can send alerts in the form of SNMP traps from CA Enterprise Log Manager to any destination in your network that receives and interprets traps. Each trap receiver product has its own requirements.

Prepare CA NSM to receive traps from alerts by:

- Verifying that the destination CA NSM system meets system requirements for receiving SNMP trap data from CA Enterprise Log Manager.
- Configuring CA NSM to receive SNMP traps, including enabling support for SNMP v3, modifying port assignments in various files, and starting the required services.

Prepare CA NSM to interpret traps received from action alerts by:

- Creating a custom MIB for each alert you plan to send as an SNMP trap to CA NSM.
- Importing and compiling the CA-ELM.MIB and all custom MIBs.

More information:

CA NSM System Requirements (see page 370)
Configure CA NSM to Receive SNMP Traps (see page 371)
Boilerplate Text for a Custom MIB (see page 350)

CA NSM System Requirements

You can send SNMP traps to CA NSM if your system meets the following CA Enterprise Log Manager interface requirements:

- The CA NSM version is CA NSM r12.2 (GA build).
- CA NSM is installed on Windows Server 2003 EE SP1.
- You have applied the patch T5MK056.caz, which updates the aws_snmpex.dll file and enables CA NSM to receive SNMP v3 traps from CA Enterprise Log Manager.

To apply the patch

1. Download the patch from CA Support.
2. Log on to the server with CA NSM.
3. Stop the SNMP Trap service:
   a. From the Start menu, select Programs, Administrative Tools, Services

      The Services list appears.

   b. Select the SNMP Trap Service, right-click and select Stop from the pop-up menu.

4. Stop all CA NSM services:
   a. Access the command prompt.
   b. Enter the following command:

      ```
      Unicntrl stop all
      ```

5. Copy the download patch file, "T5MK056.caz, to the C:\temp folder.
6. Unzip the patch file with cazipxp.

   ```
   Cazipxp.exe –u T5MK056.caz
   ```

7. Back up the existing aws_snmpex.dll before replacing it.
   a. Navigate to C:\Program Files\CA\SC\CCS\AT\SERVICES\BIN.
   b. Right-click aws_snmpex.dll and select copy.

      A Copy of aws_snmpex.dll is added to the folder.
8. Copy the aws_snmpex.dll from the temp folder to bin folder (C:\Program Files\CA\SC\CCS\AT\SERVICES\BIN)

   CA NSM now meets system requirements. You can configure CA NSM to receive SNMP traps from CA Enterprise Log Manager.

Configure CA NSM to Receive SNMP Traps

Before you can direct alerts to be sent to CA NSM as SNMP traps, you must configure CA NSM to receive traps. You can send both SNMPv2 traps and SNMPv3 traps to CA NSM.

To configure CA NSM to receive SNMP traps from CA Enterprise Log Manager alerts

1. Log on to CA NSM.
2. Enable support for SNMP version3 as follows:
   a. Display the command prompt. From the Start menu, click Run, enter cmd, and click OK.
   b. Type the following:

      ```
      caugui settings
      ```

      The EM Settings window appears.
   c. Click the Event Management tab.
   d. Scroll to display the description: SNMP - Enable SNMP version 3 support.
   e. Select the row and type Y to select YES in the setting column for SNMP Enable SNMP version 3 support.
   f. Click Yes to confirm the change.
   g. Close the window.

3. Change the port used by the SNMP service from the current port, for example 5162, to port 162 as follows:
   a. Open Windows Explorer.
   b. Navigate to the .../System32/drivers/etc folder, typically under C:\WINDOWS.
   c. Back up the Services file. Right click services and select copy.
   d. Open the Services file in a text editor, such as Notepad, and scroll to the entry resembling the following:

      ```
      snmptrap    162/udp    snmp-trap  #SNMP trap
      ```

   e. Edit the snmptrap line to replace the port number 162, with an alternative, for example, 5162. Add the catrapmuxd line where you assign port 162.

      ```
      snmptrap    5162/udp
      catrapmuxd  162/udp    catrapmuxd  #CA Trap Multiplexer
      ```

   f. Save and close the file.
4. Modify the CA Trap Multiplexer configuration file, catrapmux.conf, as follows:
   a. Navigate to C:\Program Files\CA\SC\CCS\WVEM\CAIUSER.
   b. Open CATRAPMUX.CONF in a text editor, such as Notepad.
   c. Scroll to the bottom of the file. Edit the file as needed to include the following entries.

      ```
      CATRAPMUX_CMD:6161
      AWS_SNMP:6162
      catrapd:6163
      snmptrap:5162
      ```

      Note: The first three entries represent default settings.
   d. Save and close the file.
5. Add a line to the snmpv3.dat configuration file to configure SNMP v3 security parameters.
   a. Navigate to the C:\Program Files\CA\SC\CCS\CommonResourcePackages\Misc.
   b. Open snmpv3.dat in a text editor and add the following line at the end of the file.

      ```
      *.*.*.* *.*          test1234:AuthPriv:MD5:test1234:DES:test1234
      ```

      Note: These are the same parameters that you must enter in the V3 Security Parameters dialog in the Alert wizard in order for the SNMP trap to be received by CA NSM. The username and password are what you configure here, the Auth protocol is MD5 and the Encryption protocol is DES.
   c. Save and close the file.

6. Install the CA Trap Multiplexer service:
   a. Access the command prompt.
   b. Run the following command:

      ```
      catrapmuxd uniconfig
      ```

      CA Trap Multiplexer is added to the Services list with a status of Started.
7. Verify that CA Trap Multiplexer is running and start SNMP Trap Service.
   a. From the Start menu, select Programs, Administrative Tools, Services

      The Services list appears.
   b. Examine the status of CA Trap Multiplexer. Verity that the status is Started.
   c. Select the SNMP Trap Service, right-click and select Start from the pop-up menu.
8. Start all services with a Startup Type of Automatic.
   a. Access the command prompt.
   b. Run the following command:

      ```
      Unicntrl start all
      ```

   CA NSM is now configured to receive SNMP v3 traps based on scheduled alerts sent by CA Enterprise Log Manager.

Example: Alerting CA NSM of Configuration Changes

The following example is designed to walk you through a process of alerting CA NSM of configuration changes. This process includes the following procedures:

- Send SNMP traps to CA NSM
- Verify that SNMP traps were sent successfully
- Access the EM Console on CA NSM
- View the SNMP traps received by CA NSM

More information:

Track the Alert Job Progress (see page 376)
Access the EM Console on CA NSM (see page 377)
View SNMP Traps on CA NSM (see page 378)

Send SNMPv3 Traps to CA NSM

When planning what alerts to send to CA NSM, identify query results that would be of interest to the network operations center. For example, consider queries that detect configuration changes. The following example illustrates how to send a scheduled alert based on the Configuration Change Detail query. This alert specifies CA NSM as the SNMP trap destination.

To send SNMPv3 traps to CA NSM

1. Open the Alert Scheduling wizard.
    a. Log on to CA Enterprise Log Manager with the credentials of an Analyst or Administrator.
    b. Click the Alert Management tab and the Alert Scheduling subtab.
    c. Click the Schedule an Action Alert button.
2. Complete the Alert Selection step.
    a. Type a job name. For example, enter Configuration Changes destined for CA NSM.
    b. Verify that the selection type is Queries. Selection of SNMP trap destinations is not allowed for alerts based on tags.
    c. Select the query or queries you identified. For example, select Configuration Change Detail.
3. (Optional) Complete the Alert Filters, Result Conditions, and Schedule Jobs steps as documented in the online help for this wizard.
4. Click the Destination step, and then click the SNMP Trap tab.
5. Examine the destination server and port entries. If not correct, enter the correct IP address for the destination server and port. To add additional destination servers, click add, and enter the additional destination.

6. Specify the SNMP version information. SNMP Version 2 is selected by default.

   a. Click Version 3. CA NSM is configured to accept SNMP v3 traps.

   b. Click V3 Security.

      The SNMP Version 3 Security Parameters dialog appears.

      Important: The entries on this dialog must match the settings in snmpv3.dat that you configured to enable CA NSM to receive SNMP traps from CA Enterprise Log Manager alerts. The recommended setting follows:

      ::*:*:*:* <username>:AuthPriv:MD5:<password>:DES:<password>

   c. Select Authentication. Type the configured user name for username, type the configured password for password, and select MD5 for protocol.

   d. Select Encryption. Type the configured password for password and select DES for protocol.

   e. Click OK.

7. Select the query to send as an SNMP trap.

   In this example, when you select Configuration Change Detail, the fields for that query are displayed as selected. Optionally, you can clear any field you do not want included as a trap.

   Important! When you create a custom MIB for this alert, be sure to define a trap with the fields you select here and in the order shown.

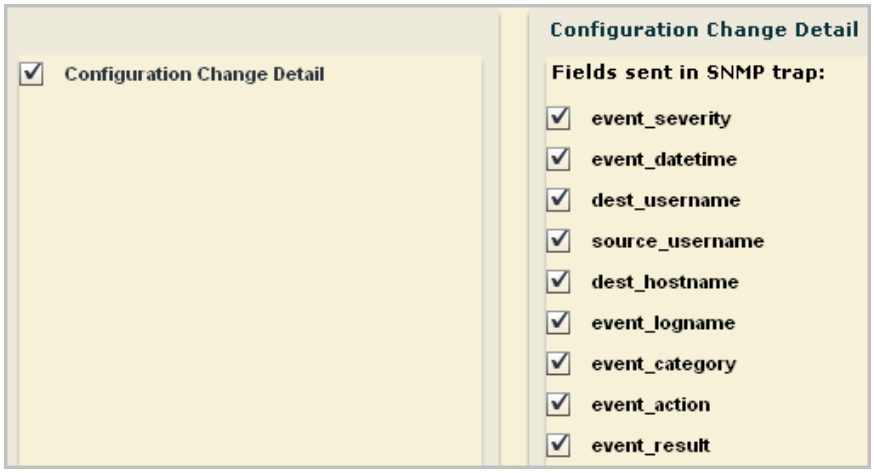

8. Select the number for the final node, x, of the associated elmTrap OID, where all elmTrap OIDs are defined as 1.3.6.1.4.1.791.9845.3.x.

   The initial nodes of the Custom Trap ID are predefined in the CA-ELM.MIB. The final node number is unique to a trap defined in a custom MIB, where the trap reflects a unique set of fields. A custom MIB file defines the traps sent by the CA Enterprise Log Manager alerts that you defined. In the custom trap referenced by the Custom Trap ID, the fields are listed in the same order as the fields sent by the alert. If the OID for the trap in the custom MIB is 1.3.6.1.4.1.791.9845.3.63, select 63 from the number spinner for Custom Trap ID. Or, if you define the alert first, add a trap in your custom MIB for 1.3.6.1.4.1.791.9845.3.63 that defines the query fields you selected.

9. (Optional) Select Servers.

10. Click Save and Close.

    The job appears on the Action Alert Jobs list with the configured job name.

More information:

Access the EM Console on CA NSM (see page 377)

Track the Alert Job Progress

When you schedule an alert, it is a good practice to track the alert job progress the first time it runs. When you track progress, you can verify that the job runs successfully and that the reported results are what you intended to send.

To monitor the alert job progress and preview the results

1. View the alert job you created on the Action Alerts Jobs list. An partial example follows:

Action Alert Jobs

	Job Name	Enabled	Server	Recurrence	Start Time
☐	Configuration Changes destined for CA NSM	true	etr6511l1-sun104	5 minutes	Fri Nov 6 2009 10:57:41 AM

2. (Optional) To track the alert job progress, view System Self Monitoring Events Detail. Double-click any line to display the Event Viewer. Scroll to result_string to view the entire message shown on Result Description.

Action	Result	Result Description
Notification Creation	S	SNMP trap for Action Query [Configuration Change Detail] Alert Name [Configuration Changes destined for CA NSM] on report
Resource Creation	S	Creation of job file while executing action alert for Alert Name [Configuration Changes destined for CA NSM] was Successful.
Alert Creation	S	Run Action Query [Configuration Change Detail] Alert Name [Configuration Changes destined for CA NSM] on reportServer [etr
Resource Modify	S	Update RSSFeed Alert Name [Configuration Changes destined for CA NSM] on reportServer [etr6511l1-sun104] recurrence [5
Resource Execution	S	Query [Configuration Change Detail] run over logDepot [localhost] was successful .

3. Preview the results returned by the queries selected for the alert you created.
   a. Select the Alert Management tab, and the Action Alert subtab.
   b. Click the name of the alert that you scheduled.
   c. View the results for that alert.

   Note: Typically, the data displayed here is the data displayed when browsing the URL sent to the destination server. If a difference exists and you want it to be the same, edit the action alert to reset the dynamic end time for Result Conditions. For example, set it to 'now', '-10 minutes'.

Access the EM Console on CA NSM

You can view the SNMP traps sent by CA Enterprise Log Manager from CA NSM. SNMP traps are displayed as messages on the EM Console.

To access the EM Console on CA NSM

1. Log on to the server where the SNMP trap destination, CA NSM is installed.
2. From the Start menu, select Programs, CA, Unicenter, NSM, Enterprise Management, and EM Classic.

   The EM for Windows window appears.
3. Double-click Windows.

   The <hostname> (Windows) window appears.
4. Double-click Event.

   The Event <hostname> (Windows) window appears.
5. Double-click Console Logs.

   The EM Console (<hostname>) appears.

More information:

View SNMP Traps on CA NSM (see page 378)

View SNMP Traps on CA NSM

Consider the example where an alert is scheduled to run the Configuration Change Detail query. In this example, the Custom Trap ID is set to 1.3.6.1.4.1.791.9845.3.63. Nine fields are sent as an SNMP trap.

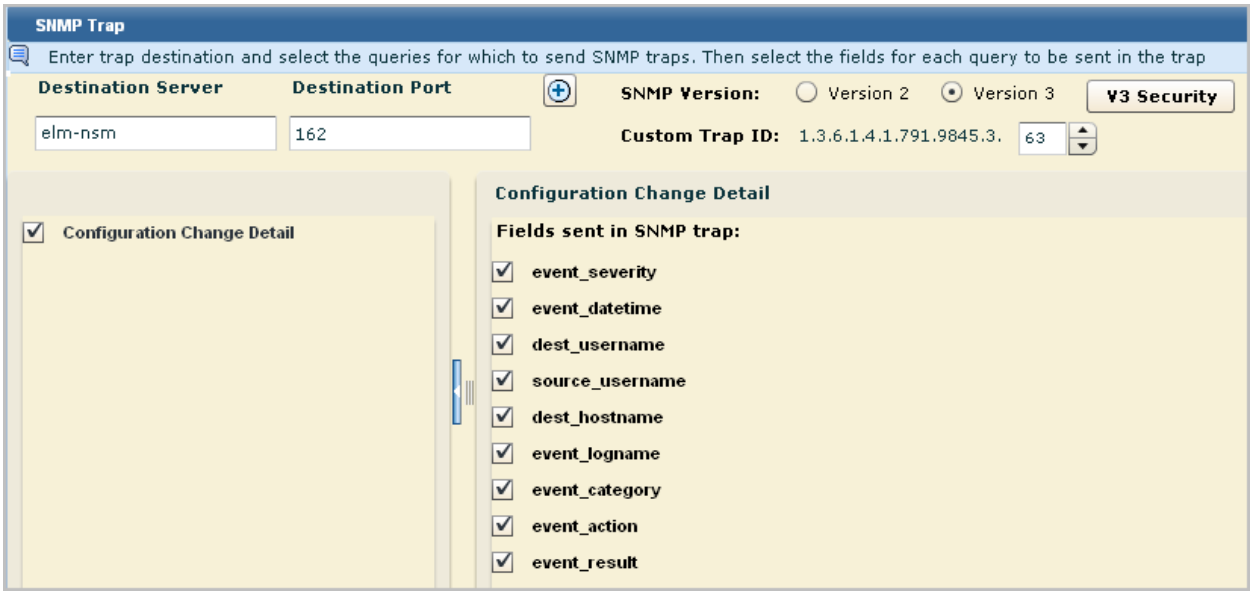

To view the SNMP trap sent by an Alert based on the Configuration Change Detail query

1. When a self monitoring event indicates that an SNMP trap has been successfully sent to CA NSM, access the EM Console on CA NSM.
2. Wait until a log message appears that indicates receipt of an SNMP trap. The message for this trap contains the custom trap ID, 1.3.6.1.4.1.791.9845.3.63.

3. Double-click this message to bring up the message in a format you can copy.

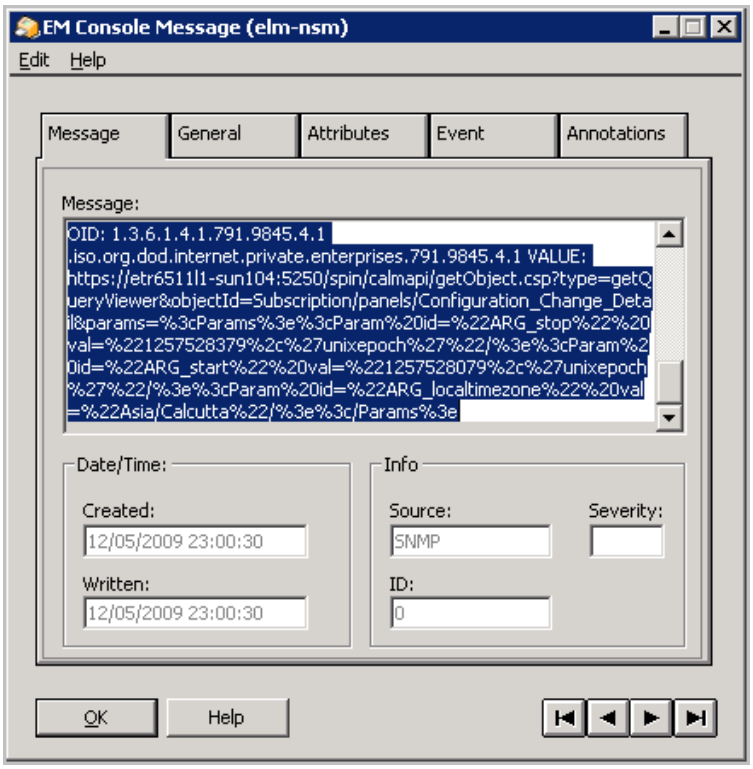

4. Copy the message and paste it into a temporary text file.

   The results resemble the following:

   %CATD_I_060, SNMPTRAP: -u auth user 791 155.35.7.63 etr6511l1-sun104.ca.com 6 63 0:05:00 12

   Specifies that the following data is received as an SNMP trap.

   OID: 1.3.6.1.2.1.1.3.0 system.sysUpTime.0 VALUE: (30000) 0:05:00.00

   Specifies the object ID for uptime in hundredths of a second. This is a known OID through SNMP.

   OID: 1.3.6.1.6.3.1.1.4.1.0 .iso.org.dod.internet.snmpV2.snmpModules.1.1.4.1.0 VALUE: **1.3.6.1.4.1.791.9845.3.63**

   Specifies the object ID for the snmpTrapOID. The value is the custom trap ID you specified when configuring the alert.

   OID: 1.3.6.1.4.1.791.9845.2.80 .iso.org.dod.internet.private.enterprises.791.9845.2.80 VALUE: 2

   Specifies the OID for event_severity and the severity value of 2, which stands for Informational.

OID: 1.3.6.1.4.1.791.9845.2.65 .iso.org.dod.internet.private.enterprises.791.9845.2.65 VALUE: Fri Nov 06 2009 10:53:53 PM

Specifies the OID for event_datetime with the value, the day, date and time when the event with these values occurred.

OID: 1.3.6.1.4.1.791.9845.2.17 .iso.org.dod.internet.private.enterprises.791.9845.2.17 VALUE:

Specifies the object ID for dest_username with no value.

OID: 1.3.6.1.4.1.791.9845.2.1 .iso.org.dod.internet.private.enterprises.791.9845.2.1 VALUE:

Specifies the object ID for source_username with no value.

OID: 1.3.6.1.4.1.791.9845.2.22 .iso.org.dod.internet.private.enterprises.791.9845.2.22 VALUE: etr851l2-elm5

Specifies the object ID for dest_hostname with the hostname of the server where the query results are displayed when you launch the URL.

OID: 1.3.6.1.4.1.791.9845.2.53 .iso.org.dod.internet.private.enterprises.791.9845.2.53 VALUE: EiamSdk

Specifies the object ID for event_logname, EiamSdk, the name of the log file that contains these details.

OID: 1.3.6.1.4.1.791.9845.2.77 .iso.org.dod.internet.private.enterprises.791.9845.2.77 VALUE: Configuration Management

Specifies the object ID for event_category and the value for Category associated with the Configuration Change Detail query.

OID: 1.3.6.1.4.1.791.9845.2.75 .iso.org.dod.internet.private.enterprises.791.9845.2.75 VALUE: Configuration Change

Specifies the object ID for event_action and the value for Action associated with the Configuration Change Detail query.

OID: 1.3.6.1.4.1.791.9845.2.81 .iso.org.dod.internet.private.enterprises.791.9845.2.81 VALUE: S

Specifies the object ID for event_result with the value, S, for Success.

OID: 1.3.6.1.4.1.791.9845.4.1 .iso.org.dod.internet.private.enterprises.791.9845.4.1 VALUE: https://etr6511l1-sun104:5250/spin/calmapi/getObject.csp?type=getQueryViewer&objectId=Subscription/panels/Configuration_Change_Detail&params=%3cParams%3e%3cParam%20id=%22ARG_stop%22%20val=%221257528379%2c%27unixepoch%27%22/%3e%3cParam%20id=%22ARG_start%22%20val=%221257528079%2c%27unixepoch%27%22/%3e%3cParam%20id=%22ARG_localtimezone%22%20val=%22Asia/Calcutta%22/%3e%3c/Params%3e

Specifies the object ID for calmAPIURL under elmDynamicVariables. The value is the URL to the CA Enterprise Log Manager API. After logging in, you can see the query results in the chart or graphic view.

5. Copy the URL at the end of the message, paste it into a browser, and launch the URL.

6. Log on to the CA Enterprise Log Manager API.

   The Configuration Change Detail chart view displays. See the following example:

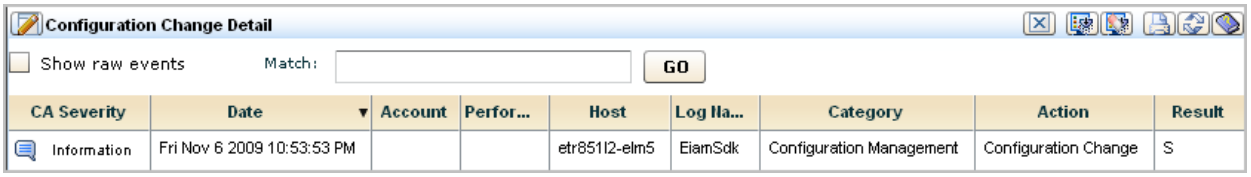

How to Create an Action Alert

The process of creating an Action Alert, using the schedule action alert wizard, has the following main steps:

1. Opening the schedule action alert wizard.
2. Choosing the query or tags on which the alert is based. You can query either the event database, the incident database or both in a single job.
3. (Optional) Setting advanced filters to further define the alert query.
4. (Optional) Setting date range and result conditions
5. (Optional) Defining how often the alert job recurs, and when it is active.
6. (Optional) Configuring automatic alert emails and recipients.
7. (Optional) Selecting whether to run the query on data for the selected server only or to run it for this server and all of its descendants.

More information:

Open Schedule Action Alert Wizard (see page 382)
Create an Advanced Event Filter (see page 482)
How to Set Result Conditions (see page 410)
Set Notification Destinations (see page 384)
Define Alert Job Query Destination (see page 388)

Open Schedule Action Alert Wizard

To create an action alert job, you must use the schedule action alert wizard.

To open the schedule action alert wizard

1. Click the Alert Management tab.

   The Alert Servers list appears.

2. Select the server where you want to schedule an alert job.

   The Server Details pane shows the selected server, displaying the Generated Alerts tab by default.

3. Click the Alert Scheduling tab, and then the Schedule an Alert button.

   The Schedule Action Alerts wizard appears.

   When using the wizard:

   - Click Save and Close to save the action alert and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

More information:

Set Email Notification Destination (see page 385)
Define Alert Job Query Destination (see page 388)

Select an Alert Query

Select tags or queries as the basis for a new action alert job. The query, plus any filters you add, defines the circumstances under which an alert is generated. For example, to create an alert to monitor traffic from a host or port, use the All Events query, add filters to define the source hosts to monitor, and an event threshold.

Note: The Action Alerts query category contains queries designed for various common alert needs.

To select an alert query

1. Open the schedule action alert wizard.
2. Type a job name.
3. Select the time zone you want to schedule the report in from the time zone drop-down menu.
4. Select the Queries or Tags option button to select reports by tag or individually.

   Note: Scheduling alerts by tag lets you add alerts without altering the job itself. If you select the "Identity Management" Tag, any alert with that tag is added to the job at the scheduled run time. You can add a new alert to the job by giving a query the Identity Management tag. This feature also applies to custom tags.

   (Optional) Clear the Enable check box to enable to action alert later rather than as soon as you finish it. The check box is selected by default.

   Note: The ability to create a disabled alert job is designed for use with recurring alerts. If you clear the Enabled check box for a job, and create that job with a single occurrence ("Now" or "Once") it is removed from the Scheduled Alert list.
5. (Optional) Select a tag or tags to narrow the tags and individual reports displayed. This feature matches the behavior of the Report List.
6. Select the tags or individual queries you want, and use the shuttle control to add them to the Selected Queries area. You can select both event and incident queries in a single alert job.
7. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the alert job is scheduled, otherwise the step you select appears.

More information:

Create an Advanced Event Filter (see page 482)
How to Set Result Conditions (see page 410)

Set Alert Job Scheduling Parameters

You can control when your alerts apply by setting their start and end time. You can also control how granular the alert view is by controlling how often the query recurs.

To set alert job scheduling parameters

1. Open the schedule action alert wizard, enter the required information, and advance to the Schedule Jobs step.
2. Set the recurrence interval you want. A lower interval gives a more detailed view, but increases network traffic.

   Before setting a low interval, verify that CA Enterprise Log Manager is synchronized with an NTP server.
3. Set the start and end time you want for the alert job.
4. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the alert job is scheduled, otherwise the step you choose appears.

Set Notification Destinations

You can set one or more of the following destinations for notification of an alert:

E-mail

You can set automatic email notification for an alert to help ensure that the proper personnel are aware of alerts relating to their job role or responsibility. Configure a mail server for your CA Enterprise Log Manager environment before you send alert notification emails.

IT PAM Process

You can run the specified CA IT PAM process if the alert is for a condition that requires notification of the third-party product. Integration with CA IT PAM must be configured under Report Server and IT PAM must have the process defined before you can run the process from alerts.

SNMP Trap

You can send event data captured by an alert to one or more Network Operations Centers (NOCs). You can target management servers such as CA Spectrum or CA NSM using SNMP v2 or SNMP v3 traps. You specify the destinations during the process of scheduling the alert. Integration with SNMP must be configured before you can send alerts using SNMP.

Note: If you do not set a destination, the alert results are published only to the RSS feed.

More information:

Set Email Notification Destination (see page 385)
Set CA IT PAM Information (see page 386)
Example: Alerting CA Spectrum of Configuration Changes (see page 365)
Example: Send an Alert that Runs an IT PAM Process Per Row (see page 334)
Example: Send an Alert that Runs an IT PAM Process Per Query (see page 338)

Set Email Notification Destination

You can set automatic email notification for an alert job, assuring that the proper personnel are aware of alerts relating to their job role or responsibility. This step is optional.

A mail server must be configured for your CA Enterprise Log Manager environment before you can set alert notification emails.

To set alert notification

1. Open the schedule action alert wizard, enter the required information, and advance to the Destination step.
2. Select the Enable email notification check box.
3. Enter at least one recipient email address. You can enter multiple addresses separated by commas.
4. (Optional) Enter From text, a subject line, and a message body for the notification email.

   Note: The message body is constructed in HTML, so all text you enter appears on one line. To create a break after a line, enter <BR/> at the end of the line of text.

More information:

Set CA IT PAM Information (see page 386)

Set CA IT PAM Information

You can set your alert job to run a CA IT PAM process when the alert is generated.

You can run the process once for each query result row, or you can run the configured process once, regardless of the number of rows. If you run it once per row, define summary and description statements using CEG fields to pass the event data to CA IT PAM. Select the fields that are defined to collect data by the query. If you run it once per query, a URL is automatically passed to CA IT PAM that, when launched, displays all rows of event data. In the third-party product that responds to the CA IT PAM process, the URL is appended to the summary text you enter. For example, it appears in the Summary field of CA Service Desk, if it is the third-party product.

To run a CA IT PAM process when the alert is generated

1. Open the schedule action alert wizard, enter the required information, and advance to the Destination step.
2. Click the IT PAM Process tab.

   A check box for each query for this alert job appears in the left pane.
3. Select a query that you want to send to the CA IT PAM process, and do one of the following:
   - Select Run IT PAM Process per row to run the configured process once for each returned row.
   - Clear Run IT PAM Process per row to run the configured process once, regardless of the number of returned rows.
4. Verify the default entries for the process parameters and change if needed. For undefined fields that allows entry of summary or description information, enter a meaningful statement. If you selected Run IT Process per row, use the CEG fields to convey event data. Select the CEG field and click Add next to the target field.
5. If the CA IT PAM process is defined with CEG fields as local parameters in the dataset, select those CEG fields in the Send field values as parameters list.
6. Select another query from the left pane and repeat steps 3 through 6.

Note: When the queries for a scheduled alert job return results, all the information and parameters required to run the configured process are sent to CA IT PAM.

More information:

Set Email Notification Destination (see page 385)

Set SNMP Trap Information

You can set SNMP Trap inform for an alert job, allowing you to send the alert to one or more third-party management systems. When the selected queries return results, a trap that includes returned data for all selected fields from all selected queries is sent to all selected SNMP trap destinations. This step is optional.

To set SNMP Trap information

1. Open the schedule action alert wizard, enter the required information, and advance to the Destination step.
2. Select the SNMP Trap tab.

   The SNMP Trap tab opens, displaying the Destination Server and Destination Port fields, and a list of the queries included in the Action Alert, each with a check box.
3. Examine the default destination server and port entries. If not correct, enter the correct IP address or fully qualified host name and port number.
4. (Optional) Click Add to enter additional Destination Server and Destination Ports.
5. (Optional) To send the alert using SNMP v3, select SNMP Version 3. SNMP Version 2 is the default.
6. If you select SNMP Version 3, click the V3 Security button to set authentication or encryption in the Security Parameters dialog.

   Important: The entries on this dialog must match the settings in snmpv3.dat that you configured to enable CA NSM to receive SNMP traps from CA Enterprise Log Manager alerts. The recommended setting follows:

   ::*:*:* <username>:AuthPriv:MD5:<password>:DES:<password>

   a. Select Authentication. Type the configured user name for username, type the configured password for password, and select MD5 for protocol.

   b. Select Encryption. Type the configured password for password and select DES for protocol.
7. Select the check box next to any query you want to include in the SNMP trap. For example, if you have three queries showing in the list, you could set SNMP to deliver one, two, or all three.

   Selecting a query displays the fields included in each query, each with a check box selected. You can clear any selected field remove that field in the alert.
8. Enter the custom trap ID you want associated with each query. This allows you to send different queries in a single alert to different trap IDs, if required.
9. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the alert job is scheduled, otherwise the step you select appears.

More information:

Set Email Notification Destination (see page 385)
Set CA IT PAM Information (see page 386)

Define Alert Job Query Destination

You can choose which federated event log stores are queried by the alert job.

To choose report destinations

1. Open the schedule action alert wizard, enter the required information, and advance to the Server Selection step.
2. Select any available servers you want to query, and move them to the Selected Servers area using the shuttle control.
3. (Optional) If you want to disable federated queries for this alert job, select "No" from the drop-down menu that appears when you click the Federated Queries entry. Report queries are federated by default.
4. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the alert job is scheduled, otherwise the step you choose appears.

Example: Create an Action Alert for Low Disk Space

Low Available Disk Space is one of the predefined queries with the tag, Action Alerts. Queries with the Action Alerts tag are specifically designed to be used as alerts, but do not become alerts until you schedule them.

The following example shows how to create an action alert from the predefined Low Disk Space query.

1. Click the Queries and Reports tab and the Queries subtab.

   The Query Tag and Query List panes appear.
2. Click the Action Alerts tag.

   The Query List displays the queries tagged with Action Alerts.
3. Click the Low Available Disk Space query in the query list.

   The Low Available Disk Space query appears in the main pane.

4. Click Options and select Schedule Action Alerts.

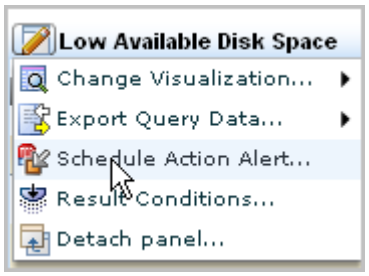

The Schedule Action Alerts wizard appears with the Alert Selection step selected. Low Available Disk Space is preselected under Selected Queries.

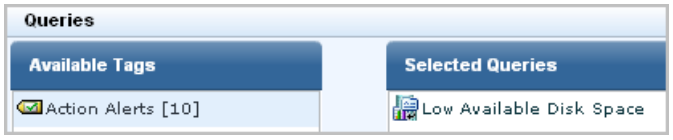

5. Enter a job name, such as Low Disk Space. Clear the Enabled checkbox for now. This lets you save and close the action alert schedule before it is complete without risking an attempt to run it.

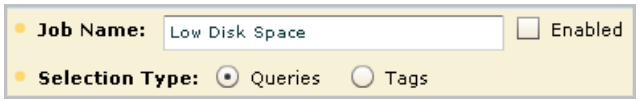

6. You can enter or skip Alert Filters. Filters are additive, that is, when a series of filters are evaluated, they are joined with logical ANDs.

7. Click Result Conditions to override the ones set in the query definition.

   a. To specify the alert should evaluate the disk space for the past hour, enter the date range as 'now' for Dynamic End Time and 'now' '-1 hours' for Dynamic start time.

   b. To specify that you only want to be notified if the query returns a result; and you want to see only the first result returned, select Row Limit and select the value 1. Since the dynamic time range is in hours, select event_hour_datetime as the Time Granularity.

   c. Leave Grouped Events blank since that does not apply to this query.

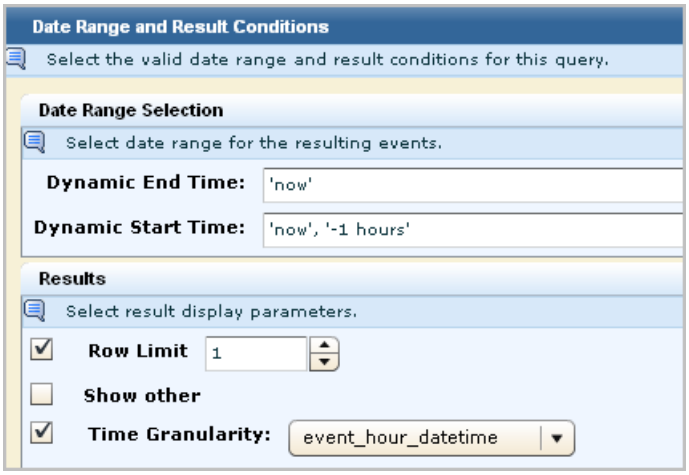

8. Click Schedule Jobs to define the schedule. The default is to start the job immediately with no end date. Set the recurrence interval. For example, set the interval to run the query every hour.

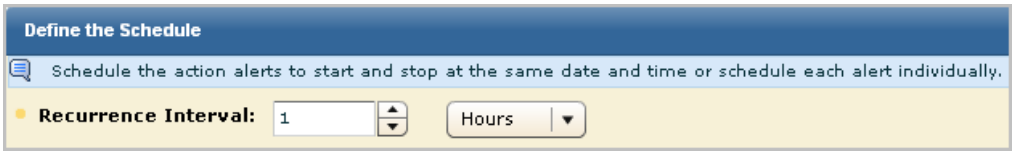

9. Click the Destination step. Select enable-email notification; enter your email address in the Email To field. Optionally, enter a subject and email text. Or, email it to the desired recipients and enter your email address in the From field. If you enter multiple email addresses, separate them with a comma (not a semicolon).

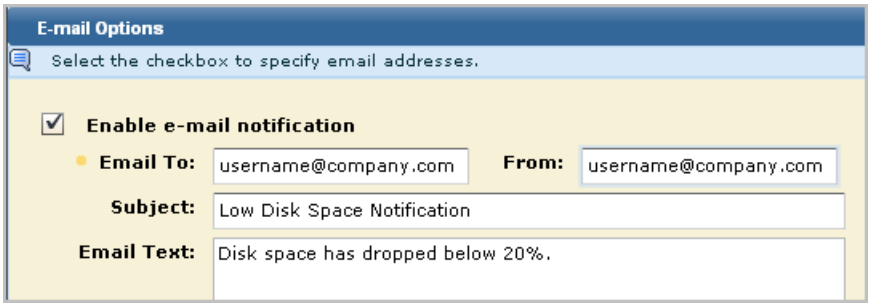

10. Click Server Selection. By default, the query will run on the current CA Enterprise Log Manager server. Select Federated to run the query on this server and all eligible federated queries.
11. Click Alert Selection. Select Enabled.
12. Click Save and Close.

    The action alert job is displayed on the Alert Scheduling subtab.

Action Alert Jobs

Job Name	Enabled	Server	Recurrence	Start Time	End Time	Time Zone	Creator
Low Disk Space	true	calmrhbuildtest01	1 hour	Thu Sep 4 2008 09:52:53 AM		America/New_York	Administrator1

13. Click the Alert Management tab, Action Alerts to view the results of this action alert.

You will receive email notification as requested. An example follows:

Subject: Low disk space Notification

CA Log Manager

RSS Link	https://calmrhbuildtest01:5250/spin/calm/getActionQueryResult.csp?id=J-8424346294223181834-calmrhbuildtest011220531049152-3207977576_actionquery_1220536215721
Alert Name	Low Disk Space
Query Name	Low Available Disk Space
Alert Generated Date	Thu Sep 04 09:50:46 EDT 2008
Tags	[Action Alerts]
Creator	Administrator1
Server	calmrhbuildtest01

Comments

Disk space has dropped below 20%

Alert Data

Hour	Log Manager	Disk Utilization

If you click the RSS Link, a page similar to the following appears:

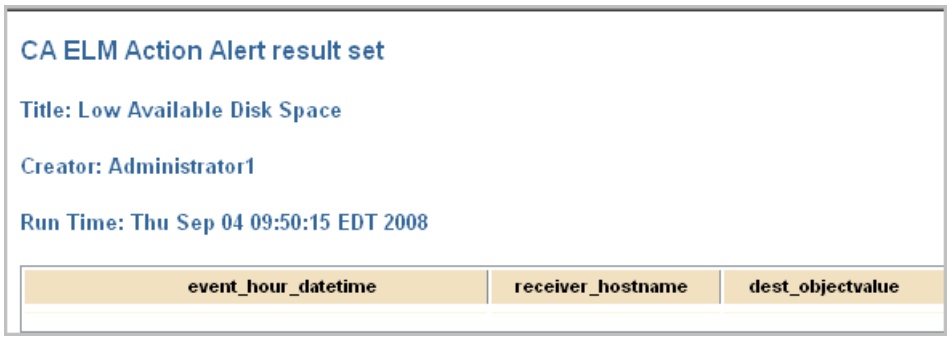

Example: Create an Alert for a Self-Monitoring Event

The predefined query for all self-monitoring events is System All Events Detail. You can copy this query and use it as the basis for defining an alert based on a specific self-monitoring event.

For example, a self-monitoring event is generated when a module requiring you to restart the operating system is downloaded in a subscription update. This self-monitoring event is generated only once. You may want to create an alert as a reminder to restart the operating system, in the event this self-monitoring event is overlooked.

Use the following example as a guide.

1. Create a query based on the query for all self-monitoring events as follows:
   a. Click the Queries and Report tab and the Queries subtab.
   b. Select System All Events Detail in the Query List, expand the Options drop-down list, and select Copy.

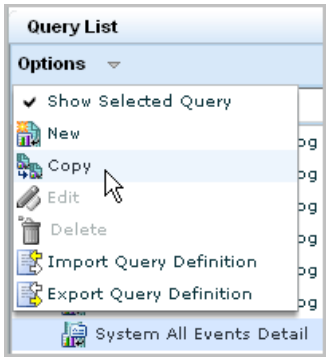

The query design wizard appears with the Details step selected.

c. Replace the name of the copied query with a new name, for example, OS Restart Alert. Optionally, add a short name and new description.

d. Select Action Alerts from Available Tags and move it to Selected Tags.

2. Create query filters as follows:

a. Advance to the Query Filters step. Click the Advanced Filters tab.

b. Click New Event Filter. Select event_logname for Column, leave Equal to for Operator, and select CALM for value.

c. Click New Event Filter. Select receiver_name for Column, leave Equal to for Operator, and enter Subscription.

d. Click New Event Filter. Select result_string for Column, leave Equal to for Operator, and enter the message, OS Updates are installed on this host...Please restart the machine for these updates to have effect !!!

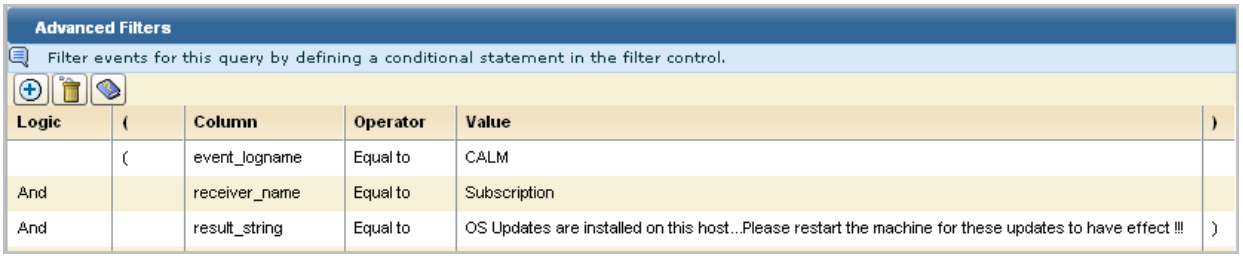

3. Click Save and Close.

The new alert appears in the Query Lists under the User folder.

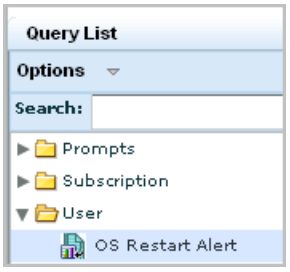

4. Schedule an action alert for the user-defined query as follows:

a. Select the query under the User folder.

b. Click the Edit button in the right pane to display the OS Restart Alert drop-down list, and select Schedule Action Alert.

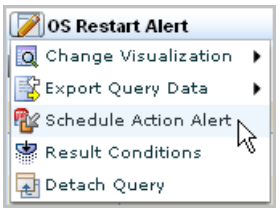

The Schedule Action Alerts wizard appears with the Alert Selection step displayed. OS Restart Alert is preselected under Selected Queries.

c. Enter a job name. For example, enter Restart Operating System Alert.

5. Add an event filter as follows:

   a. Click Alert Filters.

   b. Click New Event Filter.

   c. Select receiver_hostname for Column, leave Equal to for Operator, enter the name of the local CA Enterprise Log Manager for Value.

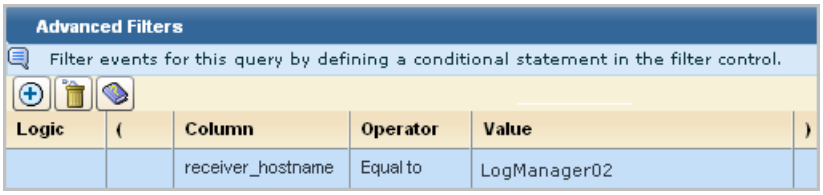

6. Specify the frequency with which to generate the alert when a restart is needed as follows:

   a. Click Schedule Jobs

   b. Set the recurrence interval for the alert generation frequency. For example, select 1 and Days for once a day.

7. Provide your email information as follows to be alerted by email.

   a. Click the Destination step.

   b. Click Enable e-mail notification and provide your email address and any of the other optional information desired.

8. Restrict the notification to when the current server needs to be restarted as follows:

   a. Click Server Selection

   b. Select No for Federated Query.

9. Click Save and Close to save the Alert Job.

   The Action Alert Job appears on the Alert Management tab, Alert Scheduling subtab.

Action Alert Jobs

Job Name	Enabled	Server	Recurrence	Start Time	End Time
Restart Operating System Alert	true	LogManager02	1 day	Tue Oct 28 2008 01:48:23 PM	

Example: Email the Administrator when Event Flow Stops

Administrators need to be notified when any connector on any agent stops collecting events. You can automate this notification when an indicator suggests that this has occurred. You can configure the indicator, which is the elapsed time since a collection server has received events from any connector. You can set the elapsed time to the desired number of minutes, hours, or days. You can extend the query to all collection servers in the federation.

To limit the number of emails sent when a connector goes down, consider only those connectors that have been collecting events up until now. For example, set the alert to return rows only for connectors that did collect events during the hour before this one but did not collect events during the last hour.

To capture this data, select the predefined query, Collection Monitor by Log Manager Agent Connector Down. This query returns the connector name and the agent name when no events are received as defined in Result Conditions in the alert. Use the following example as a guide to generate an alert when no events are received during the last hour from a connector that sent events during the period between one and two hours ago. For the alert destination, specify the email address of the individual to notify. For the schedule to run the query, specify a frequency greater or equal to that of the elapsed time period.

Note: Email Settings must be configured under Administration, Report Server before creating the alert.

To email the Administrator when a connector stops collecting events

1. Select the server from which to run this alert. In a hub and spoke architecture, select a collection server to capture the condition as soon as possible.
2. Select the Alert Management tab and the Alert Scheduling subtab.
3. Click Schedule an Action Alert.
4. Enter a job name, for example, Connector Down.
5. Select from Available Queries, Collection Monitor by Log Manager Agent Connector Down and move it to the Selected Queries list.

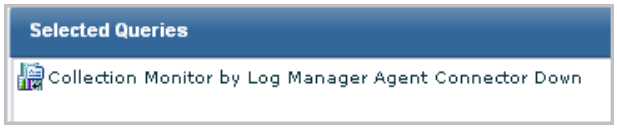

6. Click Result Conditions.
7. Set the time for the last 2 hours.
   a. Select the Predefined Ranges: Last hour.

      This sets the dynamic end time correctly to 'now', '-2 minutes'
   b. Click Edit dynamic time string for Dynamic Start Time.
   c. For Dynamic Time String, replace 62 with 122.
   d. Click OK.

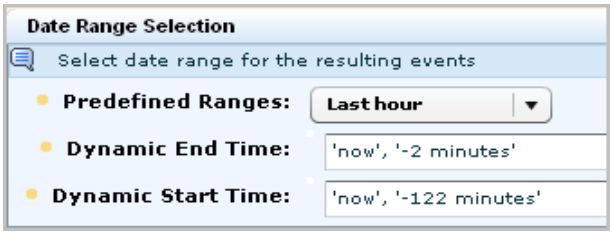

8. Set Result Conditions.
   a. Select Latest grouped event data before and click Edit
   b. Select Now for Reference time and click Add reference time to Dynamic Time string
   c. Click down once on the spinner for Shift time to display -1, select hour from the drop-down list, and click Add time shift to Dynamic Time string.
   d. Click OK.

9. Click the Schedule Jobs step and define the recurrence interval. For example, set the interval for 1 hour.

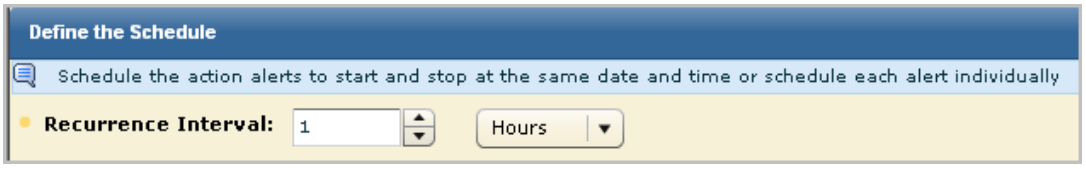

10. Click Destination and complete the E-mail tab.
    a. Select Enable e-mail notification.
    b. Enter the administrator's email address for Email To.
    c. Enter your email address for Email From.
    d. Enter the subject in the Subject field. For example, type Connector may be down.
    e. Enter email text. For example, type: Connector stopped sending events within the last hour.
11. Click the Server Selection step and clear Federated if desired.
12. Click Save and Close.

You could define this alert to query for the date range in days, rather than hours, and then schedule it to run once a day. In this case dynamic end time would be set to 'now', dynamic start time would be set to 'now', '-2 days', and latest grouped event dated before would be set to 'now', '-1 days'.

Configure Action Alert Retention

You can control how many action alerts are saved by the report server, and how long they are retained.

To configure action alert retention

1. Click the Administration tab, and then click the Services subtab.

   The Service List appears.
2. Click Report Server for the global setting or the Report Server host for the local setting.

   The Report Server configuration pane appears.
3. Enter a value in the Maximum Action Alerts entry field. Any Alerts above this threshold are deleted, oldest first.
4. Enter a number of days in the Action Alert Retention entry field, after which alerts are deleted.

   Note: Action Alerts are deleted whenever either threshold is exceeded.
5. Click Save.

Example: Create an Alert for Business_Critical_Sources

You can create a custom query with the Business_Critical_Sources keyed list and schedule an alert based on this query. The keyed list is one that has no default values and no associated predefined query or alert. Use the following end-to-end process as a guide.

1. Install an agent.
2. Configure a connector on that agent to collect events from each business critical source.

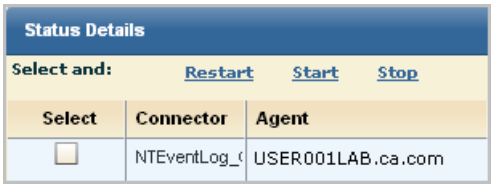

3. Define the hostname values for Business_Critical_Sources user-defined lists (keys).
   a. Click the Administration tab and Services subtab.
   b. Select Report Server from the Service List.
   c. Select Business_Critical_Sources in the User Defined Lists (Keys) area.
   d. Click Add Value in the Values area and enter the hostname of a business critical source.

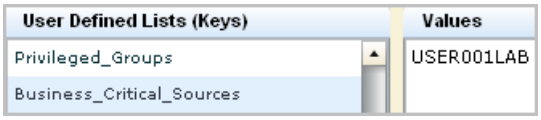

   e. Repeat the last step for each business critical source from which events are collected.
   f. Click Save.

4. Create a query on failed login attempts on business critical sources.
   a. Click Queries and Reports.
   b. Under Query List, enter login in the Search field.
   c. Select Unsuccessful Login Attempt by Host and select Copy from the Options drop-down list.

      The Query Design wizard opens with the name Copy of Unsuccessful Login Attempts by Host.

      Rename to query to Unsuccessful Login Attempts by Business_Critical_Sources.
   d. Select the Query Filters step.
   e. Click the Advanced Filters tab.
   f. Click New Event Filter.

   g. Select source_hostname for the column, select Keyed for the operator, and select Business_Critical_Sources as the value.

Logic	(	Column	Operator	Value	)
		source_hostname	Keyed	Business_Critical_Sources	

   h. Click Save and Close.
5. Schedule an alert based on this custom query.
   a. Click the Queries and Reports tab.
   b. Select Unsuccessful Login Attempts by Business_Critical_Sources under the User folder of the Query List.
   c. Select Schedule Action Alert from the Edit drop-down list.

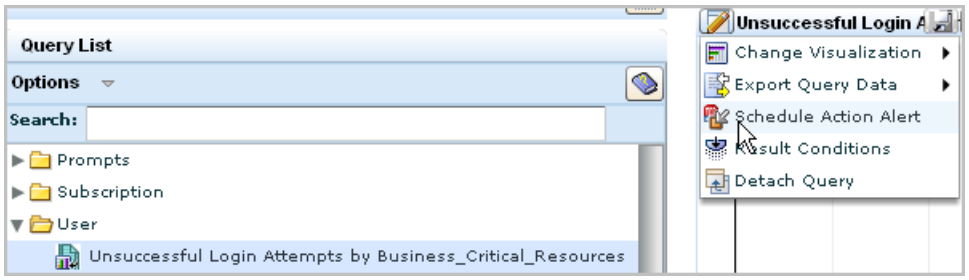

   The Schedule Action Alerts wizard appears.

d. Enter a job name, such as Unsuccessful Login Attempts by Business Critical Resources

e. Click Schedule Jobs and define the schedule.

f. Optionally, specify email options for Destination.

g. Click Save and Close.

6. Verify the job is scheduled.

a. Click the Alert Management tab and the Alert Scheduling subtab.

b. Verify the job name you entered is listed.

Action Alert Jobs	
Job Name	**Enabled**
Unsuccessful Login Attempts by Business Critical Resources	true

7. Check for the generation of the alert.

a. Click the Alert Management tab. The Action Alerts subtab is displayed.

b. View the listed alerts to determine whether the job name you listed appears.

Edit an Action Alert

You can edit an existing Action Alert.

To edit an action alert

1. Click the Alert Management tab.

   The Alert Server list appears.

2. Select the server where the Action Alert you want to edit is scheduled.

   The server details pane appears, showing the Generated Reports tab by default.

3. Click the Scheduled Alerts tab, select the alert you want, and click Edit at the top of the list.

   The Schedule Action Alerts wizard appears

4. Make the changes you want, and click Save and Close.

   The edited Action Alert appears in the Action Alerts list.

Disable or Enable Action Alerts

You can disable one or more action alerts when you no longer want the scheduled queries associated with that action alert to run. You can enable action alerts that were previously disabled, so that they run according to the saved schedule.

To disable or enable an action alert job

1. Click the Alert Management tab, and the Alert Scheduling subtab,

   The Action Alert Jobs list appears, showing the status of each job in the Enabled column. If the job is enabled, the Enabled value is true. If it is disabled, the Enabled value is false.

2. Select the job or jobs you want, and click Enable Selected, or Disable Selected.

   The Action Alert Jobs list displays the new status of all the jobs you enable or disable.

   Note: The ability to disable alert jobs is designed for use with recurring alerts. If you disable an alert job with a single occurrence ("Once") it is removed from the Action Alert Jobs list.

Delete an Action Alert

You can delete an unneeded Action Alert.

To delete an action alert

1. Click the Alert Management tab.

   The Alert Server list appears.

2. Select the server which contains the Action Alert you want to delete.

   The server details pane appears.

3. Click the Scheduled Alerts tab, select the alert you want by clicking on the row, and click Delete at the top of the list. You can select multiple alert jobs for deletion.

   Note: The check boxes beside each alert job are used for enabling or disabling alert jobs.

   A confirmation dialog appears.

4. Click Yes

   A deletion successful message appears

5. Click OK.

   The alert job is removed from the Alert Jobs list.

Chapter 11: Scheduled Reports

This section contains the following topics:

View a Generated Report

You can view a generated report, or save a copy to a location of your choice. Generated reports are stored on the soft appliance with CA Enterprise Log Manager under the following path:

/opt/CA/LogManager/data/reports

To view a generated report

1. Click the Scheduled Reports tab.

   The tab opens, displaying the local CA Enterprise Log Manager host by default.
2. Select the server where the generated reports you want to view are scheduled.

   The server you select is displayed in the details pane.
3. Click the Generated Reports tab if it is not already displayed.

   The Generated Reports List appears.
4. Click the name of the report you want to view.

   The Save dialog appears.
5. Click Save to specify a location to save the report.

More information:

Filter Reports (see page 404)
Annotate a Generated Report (see page 404)
How to Schedule a Report Job (see page 405)
Self-Monitoring Events (see page 423)

Filter Reports

You can set filters to refine the display of available generated reports and scheduled report jobs.

To filter generated or scheduled reports

1. Select the report server where the scheduled or generated reports you want to filter are located, and click the Generated Reports or Scheduled Reports tab.

   The generated reports or scheduled reports list appears.

2. Select the type of recurrence or format by which you want to filter the displayed reports, using the appropriate drop-down menus.

   The list shows the reports that meet your filter qualifications.

More information:

View a Generated Report (see page 403)
Annotate a Generated Report (see page 404)

Annotate a Generated Report

You can add annotations to a generated report, for purposes including report tracking or reviews.

To annotate a generated report

1. Select the report server where the generated reports you want to annotate are located, and click the Generated Reports tab.

   The generated reports list appears

2. Click the Annotations icon beside the report you want to annotate.

   The Report Annotations dialog appears, displaying any previous annotations with the name of their creator, and the time and date of their creation.

3. Type the annotation you want, and click Save.

   The annotation appears in the dialog, which remains open to allow further annotations.

4. (Optional) repeat step 3 to add additional annotations.

5. Click Close when you have no further annotations to add.

   The Report Annotations dialog closes.

More information:

View a Generated Report (see page 403)
Filter Reports (see page 404)

How to Schedule a Report Job

The process of creating a report job, using the report scheduling wizard, has the following main steps:

1. Opening the schedule report wizard.
2. Selecting report templates - To begin scheduling a report job, you must select which report or tag you want to use as a template for the job. You can select a single template or tag, or multiple templates or tags.
3. Creating report filters - You can apply advanced event filters to further customize your report returns, if needed.
4. Setting date range and result conditions - You can set the date range you want the report query to search and other conditions.
5. Scheduling jobs - You must set the day and time that reports are run for both single-occurrence and recurring reports. You can also choose from the available recurrence patterns.
6. Selecting a report format and destination - You can choose the report format you want and email delivery options.
7. Selecting a server - You must select the server to be queried by the report, and whether the server's federated hosts are also to be queried.

Open Schedule Report Wizard

To create a new report job for one or more recurring reports, you must use the schedule report wizard.

To open the schedule report wizard

1. Click the Scheduled Reports tab.

   The Report Servers list appears.

2. Select the server where you want to schedule a report.

   The Report Server Details pane shows the selected server, displaying the Generated Reports tab by default.

3. Click the Report Scheduling tab, and then click Schedule a Report.

   The Schedule Report wizard appears.

   When using the wizard:

   - Click Save and Close to save the scheduled report and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

More information:

Set Scheduling Parameters (see page 414)
Create an Advanced Event Filter (see page 482)
How to Set Result Conditions (see page 410)
Choose Report Query Target (see page 416)

Select a Report Template

The first step in creating a report job is the selection of the report template. If you want to schedule multiple report jobs that share the same filters, scheduling, and destination settings, you can do so by selecting multiple reports or tags as templates.

If you select multiple reports the jobs display separately by report. For example, if you select two individual reports, they share the same scheduling and filter options, but are displayed separately, titled by the report name, in the Generated Reports list.

Users with the Administrator role can create report jobs in a disabled state for later use. User with the Administrator and Analyst roles can enable and disable jobs at a later time. Disabled reports display the value *false* in the Enabled column when viewing the Scheduled Reports tab.

To select a report template

1. Enter a job name.
2. Select the time zone you want to schedule the report in from the time zone drop-down menu.
3. Select the Reports or Tags option button to select reports by tag or individually.

   Note: Scheduling reports by tag lets you add reports without altering the job itself. If you select the "Identity Management" Tag, any report with that tag is added to the job at the scheduled run time. This feature also applies to custom tags.
4. (Optional) Select a tag or tags to narrow the tags and individual reports displayed. This feature matches the behavior of the Report List.
5. (Optional) Clear the Enabled check box to create this report job in a disabled state. The Enabled check box is marked by default.

   Note: The ability to create a disabled report job is designed for use with recurring reports. If you clear the Enabled check box for a job, and create that job with a single occurrence ("Now" or "Once") it is removed from the Scheduled Report list.
6. (Optional) Select the Retain after expiry check box to retain the report configuration after the report is generated.

   Note: After the report is generated, you can edit the report template and reschedule the report.
7. Select the tags or individual reports you want to use as report templates, and use the shuttle control to add them to the Selected Reports area.
8. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the report is scheduled, otherwise the step you select appears.

Using Advanced Filters

You can use SQL-based advanced filters to qualify any function that queries the event log store, including narrowing queries, or adding additional qualifications to simple filters. The Advanced Filters interface helps you create the appropriate filter syntax by providing a form for entering logic columns, operators and values according to your filtering requirements.

Note: This section contains a brief overview of the SQL terms used in advanced filters. To use advanced filters to their full potential you need a thorough understanding of SQL and the Common Event Grammar.

The following SQL terms join multiple filter statements:

And

Displays the event information if *all* the joined terms are true.

Or

Displays the event information if *any* of the joined terms are true.

Having

Refines the terms of the main SQL statement by adding a qualifying statement. For example, you could set an advanced filter for events from specified hosts, and add a "having" statement to return only events of a specified severity level from those hosts.

The following SQL operators are used by advanced filters to create the basic conditions:

Relational Operators

Include the event information if the column bears the appropriate relation to the value you enter. The following relational operators are available:

- Equal to
- Not Equal to
- Less than
- Greater than
- Less than or equal to
- Greater than or equal to

For example, using *Greater than* would include the event information from your chosen column if its value is greater than the value you set.

Like

Includes the event information if the column contains a pattern you enter, using % to set the pattern you want. For example, L% would return any values beginning with L, %L% would return any values with L included as neither first nor last letter.

Not like

Includes the event information if the column does not contain the pattern you specify.

In set

Includes the event information if the column contains one or more of the values in the quote-delineated set you enter. Multiple values in the set must be comma-separated.

Not in set

Includes the event information if the column does not contain one or more of the values in the quote-delineated set you enter. Multiple values in the set must be comma-separated.

Matches

Includes any event information that matches one or more of the characters that you enter, allowing you to search for key words.

Keyed

Includes any event information that is set as a key value during Report Server configuration. You can use key values to set business relevance or other organizational groups.

Not Keyed

Includes any event information that is not set as a key value during Report Server configuration. You can use key values to set business relevance or other organizational groups.

How to Set Result Conditions

You can set a date range and other result conditions for the query, including row limits and base display time period. Result conditions can be altered at any time up to the query's run time, making them a useful way to modify queries without altering the base query or its filters.

You can set the following types of result conditions:

- Date range conditions governing the query's search period
- Display conditions, such as maximum rows
- Grouped Event conditions, such as the most recent grouped events after a given date, or grouped events containing a set number of events.

  Note: If you do not group at least one column when creating a query, users will not be able to edit result conditions from the query display.

Set a Time or Date Range

You can set a time or date range condition for your query. This improves the efficiency of your query by narrowing the portion of the event log store it must search.

You can use a predefined time range, or create a custom time range. For a custom time range to work properly you must set both a beginning and end time. If you only set a single time parameter, it is expressed as a "Where" clause in the query SQL.

To set result conditions

1. Open the result conditions dialog.
2. Select a predefined time range from the drop-down list. For example, if you want to view events received in the last day, select "previous day".

   Note: If you are creating an action alert or scheduled report, the interface displays the following default time ranges:

   - Action Alert: the previous 5 minutes
   - Scheduled Report: the previous 6 hours
3. (Optional) Create a custom time range using the following substeps:
   a. Click Edit beside the 'Dynamic End Time' entry field in the Date Range Selections area. This lets you set the end of the time period you want the query to search.

      The Dynamic Time Specification dialog appears.

   b. Select the reference time you want to base the parameter on, and click Add.

   c. Select the time parameter you want, and click Add. You can add multiple time parameters.

   d. When you are finished adding parameters, click OK.

      The Dynamic Time Specification dialog closes, and the values you choose appear in the 'Dynamic End Time' area. If you use multiple parameters, they form a complete time statement, with each parameter referring to the first. For example, adding the 'Start of the Month,' and 'Day of the Week - Tuesday' values in the 'Dynamic End Time' area will end your query on the first Tuesday of the month.

      Note: When using the 'Number of' values, such as 'Number of days' or Number of hours' you must enter a *negative* number to set a time in the past. Using a positive number will set a future end time, and cause the query to continue sending results as long as at least one qualified event is in the log store.

      For example, adding the 'now,' and 'number of minutes -10' values to the 'Dynamic Start Time' area starts your query 10 minutes before the selected end time.

   e. Repeat step 2 in the 'Dynamic Start Time' area to set the beginning of the time period you want the query to search.

If you do not enter a date range, the query is applied all events in the log store. If you enter an invalid date range, your query might not return any results.

4. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

   If you click Save and Close the new query appears in the Query List, otherwise the Query Design step you choose appears.

More information:

How to Set Result Conditions (see page 410)
Set Display and Group Conditions (see page 413)

Set Display and Group Conditions

You can set conditions that allow you to control the query display and conditions that search for events based on how they are grouped.

To set display and group conditions

1. Open the result conditions dialog.
2. Use the Results check boxes to enable any of the following display qualifications you want:

   Row Limit

   Sets the maximum number of event rows that the query displays, starting with the most recent events.

   Minimum: 1

   Maximum: 5000

   Show Other

   Indicates the presence of other results that are not displayed due to the row limit, allowing you to compare the selected events in the context of all events of that type. For example, if you choose a row limit of 10 for your event viewer display and select show other, events beyond 10 are displayed as a single entry titled Other, showing all remaining events. This setting is only effective when row limit is selected.

   Time Granularity

   Sets the detail level of the time period field used in the query display.
3. Use the Result Conditions to query for various types of grouped event conditions. For example you could set your query to search for the latest grouped event after a selected date, or a certain number of grouped events. A grouped event is a refined event for which you have set a Function and Group Order in the Query Creation step.

   The group conditions use the same time statement system as the time range fields.
4. Click the appropriate arrow to advance to the Query Design step you want to complete next, or click Save and Close.

   If you click Save and Close the new query appears in the Query List, otherwise the Query Design step you choose appears.

More information:

How to Set Result Conditions (see page 410)

Set Scheduling Parameters

You can control when scheduled reports run, whether you want them to recur, and the recurrence interval.

To set scheduling parameters

1. Open the Schedule Report wizard and advance to the Schedule Jobs step.
2. Use the Non Recurring or Recurring radio buttons to select the report generation time and recurrence pattern you want, if any.

   Note: If you use daylight savings time in your environment, do not schedule a report during the changeover time, since it will not be generated. For example, if daylight savings time begins at 2 a.m. March 8, you cannot schedule a report between 2:00:00 and 2:59:59.
3. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the report is scheduled, otherwise the step you choose appears.

More information:

Using Advanced Filters (see page 408)
How to Set Result Conditions (see page 410)
Choose Report Query Target (see page 416)

Select Format and Notification Settings

You can select whether reports are generated in PDF, Excel, or XML format. You can also set up automatic email notification.

To set format and notification

1. Open the Schedule Report wizard and advance to the Destination step.
2. Select the format you want from the Report Format drop-down menu.

   Note: In PDF format, charts are limited to 100 data points, which keeps the chart axis labels clearly legible. If the chart you want to display contains more than 100 points, CA Enterprise Log Manager includes only the first 100 in the published PDF output.
3. Select the Email check box if you want to send a notification when the report is generated.

   The email specification fields appear.
4. Enter email addresses for any users you want to receive the notification. Separate multiple addresses using commas.
5. (Optional) Enter any other specifications you want, including subject, return email address, and body text.
6. (Optional) Select Attach Report to attach a copy of the report in your chosen format to the notification email.
7. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the report is scheduled, otherwise the step you select appears.

More information:

Using Advanced Filters (see page 408)
How to Set Result Conditions (see page 410)
Set Scheduling Parameters (see page 414)
Choose Report Query Target (see page 416)

Choose Report Query Target

You can choose which federated event log stores the report query searches.

To choose report destinations

1. Open the Schedule Report wizard and advance to the Server Selection step.
2. Select any available servers you want to query, and move them to the Selected Servers area using the shuttle control.
3. (Optional) If you want to disable federated queries for this report, select "No" from the drop-down menu that appears when you click the Federated Queries entry. Report queries are federated by default.
4. Advance to the scheduling step you want to complete next, or click Save and Close.

   If you click Save and Close the report is scheduled, otherwise the step you choose appears.

More information:

Using Advanced Filters (see page 408)
How to Set Result Conditions (see page 410)
Set Scheduling Parameters (see page 414)

Example: Schedule Reports with a Common Tag

You can schedule one or more reports to be generated at the specified frequency with the specified end date.

Auditors, Analysts, and Administrators can schedule reports.

To schedule a report

1. Click the Scheduled Reports tab, the Report Scheduling subtab, and then click Schedule a Report.

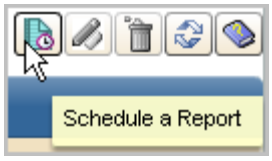

   The Schedule Report wizard appears with Step 1, Report Selection, selected.

2. Enter a job name, select Reports to enable selection of individual reports or select Tags to enable selection of all the reports associated with a selected tag.

   In the following example, selecting the Resource Access tags is an easy way to select the six reports with this tag.

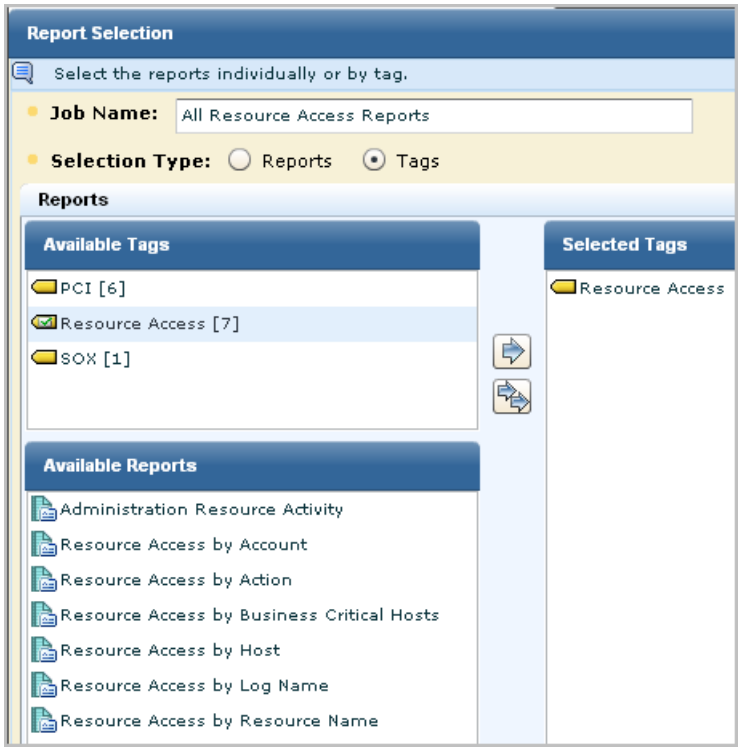

3. (Optional) Click Report Filters and create a new event filter to limit the report to the data you need.

4. (Optional) Click Result Conditions and select a date range and/or result conditions for this query. For example, to search for events that occurred within the last six hours, select 'now' as the dynamic end time and select 'now' and '-6 hours' as the dynamic start time. Select Row Limit and select a number, such as 250.

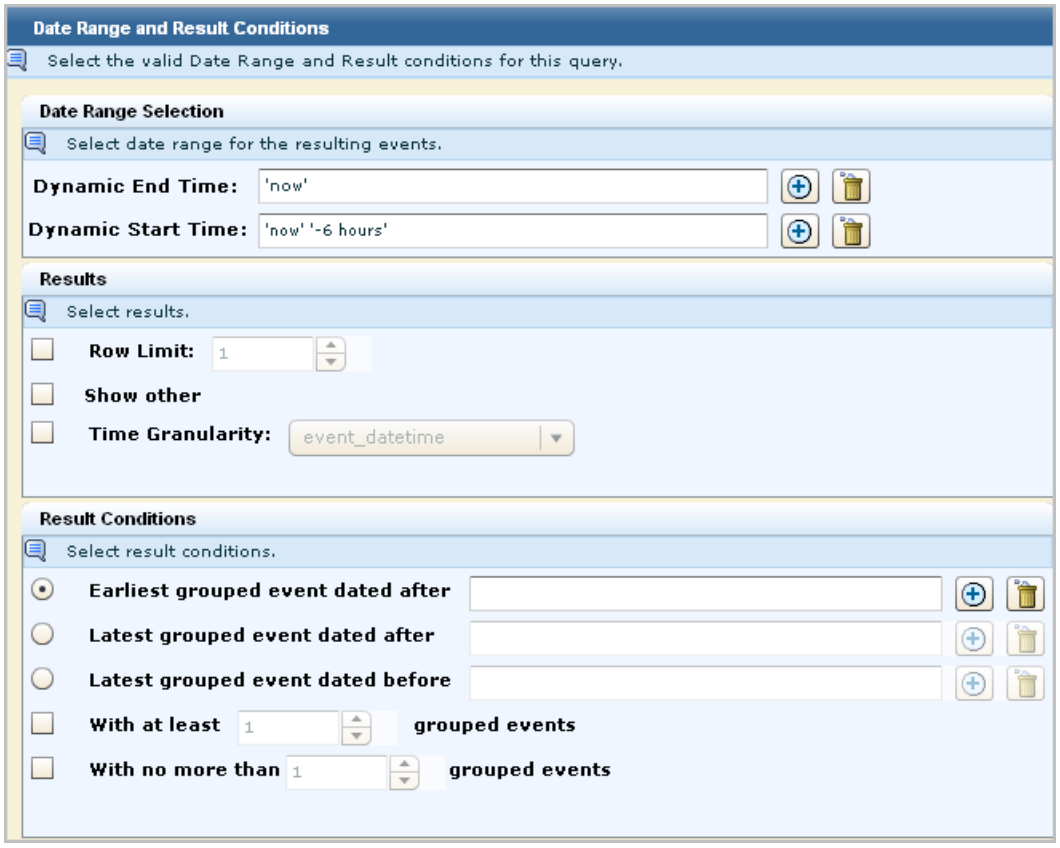

5. Click Scheduled Jobs to schedule the generation for Now or select another option and specify the details.

6. Click Destination, specify whether to format the report as an Excel spreadsheet, a PDF, or XML. A spreadsheet is appropriate for tabular data. A PDF is appropriate for graphics. Optionally, you can send an email notification. (Use the comma as a delimiter between email addresses.) The email can be sent without the report just for confirmation that the scheduled report was generated, or you can send the report as an attachment to the email.

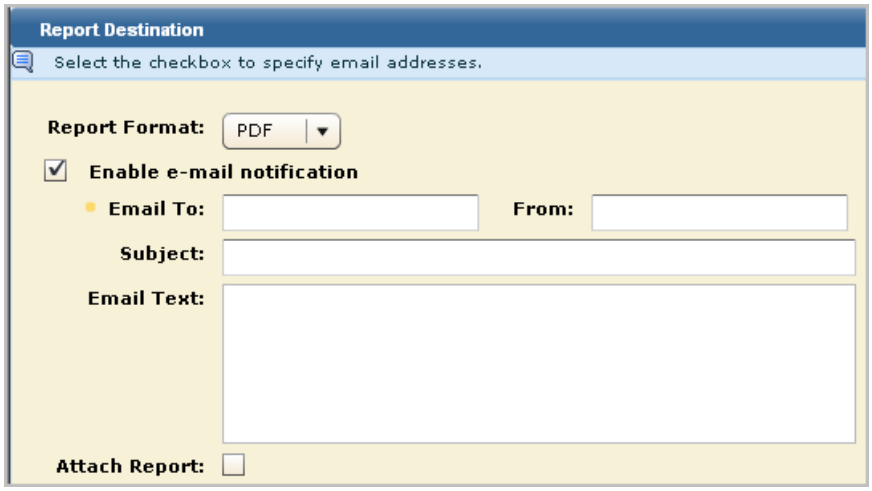

Note: The Administrator can configure reports to be deleted after a specified retention period. Retaining an email copy may be used a backup alternative to manual archiving.

7. Click Server Selection and select one or more severs for the reports and indicate whether to query the server's federation.

8. Click Save and Close.

   The selected reports are scheduled for generation.

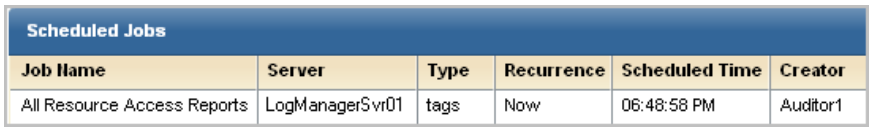

Scheduled Jobs

Job Name	Server	Type	Recurrence	Scheduled Time	Creator
All Resource Access Reports	LogManagerSvr01	tags	Now	06:48:58 PM	Auditor1

Example: Email Daily PCI Reports as PDFs

You can automate the delivery of specified reports in the format you choose to the person you specify at the required frequency.

Before you can specify that scheduled reports be formatted as PDFs and attached to emails, you must configure the following on the Global Service Configuration for the Report Server under the Administration tab and Services subtab.

- Mail Server options:
  - Mail Server
  - SMTP Port (25)
  - Admin eMail
  - SMTP user name and password
- PDF specifications:
  - Company/product name
  - Company/product logo URL
  - Header Font and font size
  - Data font and font size
  - Page orientation, width, and height

Example: Deliver all daily PCI reports, as PDFs, to the auditor's inbox each weekday

1. Click the Scheduled Reports tab and the Report Scheduling subtab.

   The toolbar appears with a Scheduling a Report button.
2. Click Schedule a Report.

   The Report Selection step appears.
3. Select a report as follows:
   a. Type PCI Reports as the job name.
   b. Select Tags as the selection type.
   c. Select PCI from Available Tags and move it to Selected Tags.
4. Schedule the job as follows:
   a. Click the Schedule Jobs step.
   b. Select Daily under recurring.
   c. Select every weekday.

5. Specify the report destination and format as follows:
   a. Click the Destination tab.
   b. Accept the default report format, PDF.
   c. Select Enable e-mail notification.
   d. Type the auditor's email address. Use the following syntax: <email_name>@<company>.com
   e. Select Attach Report.
6. Click Save and Close.

Edit a Scheduled Report Job

You can edit a scheduled report job.

To edit a scheduled report job

1. Click the Scheduled Reports tab.

   The Report Servers list appears.
2. Select the server where the report you want to edit is scheduled.

   The selected server appears in the Report Server Details pane.
3. Select the report job you want, and click Edit at the top of the list.

   The Schedule Report wizard appears.
4. Perform the changes you want, and click Save and Close.

   The edited report appears in the Scheduled Jobs list within 5 minutes, as soon as the list refreshes. Click Refresh to display it immediately.

More information:

How to Schedule a Report Job (see page 405)
Delete a Scheduled Report Job (see page 422)

Enable and Disable Scheduled Report Jobs

You can disable one or more scheduled report jobs when you no longer want the queries associated with that report to run. You can also enable scheduled report jobs that were previously disabled, so that they run according to the saved schedule.

To disable or enable scheduled report jobs

1. Click the Scheduled Reports tab, and the Report Scheduling subtab,

   The Scheduled Jobs list appears, showing the status of each job in the Enabled column. If the job is enabled, the Enabled value is true. If it is disabled, the Enabled value is false.

2. Select the job or jobs you want, and click Enable Selected, or Disable Selected.

   The Scheduled Jobs list displays the new status of all the jobs you enable or disable.

   Note: The ability to disable report jobs is designed for use with recurring reports. If you disable a report job with a single occurrence ("Once") it is removed from the Scheduled Jobs list.

Delete a Scheduled Report Job

You can delete a scheduled report job.

To delete a scheduled report job

1. Click the Scheduled Reports tab.

   The Report Servers list appears.

2. Select the server where the report you want to delete is scheduled.

   The selected server appears in the Report Server Details pane.

3. Click the Report Scheduling tab, select the job you want by clicking on the row, and click Delete at the top of the list. You can select multiple jobs for deletion.

   Note: The check boxes beside each report job are used for enabling or disabling report jobs.

   A confirmation dialog appears.

4. Click Yes

   The report job is removed from the Scheduled Jobs list.

More information:

How to Schedule a Report Job (see page 405)
Edit a Scheduled Report Job (see page 421)

Self-Monitoring Events

Most user actions generate self-monitoring events. These events allow you to track which actions have been taken on or involving the server, and their success or failure. Self-monitoring events are displayed in Event Viewer format for each server on the Scheduled Reports and Alert Management tabs. They can also be accessed as normal or scheduled reports using the Self Monitoring Events report.

More information:

View a Self-Monitoring Event (see page 423)
How to Schedule a Report Job (see page 405)

View a Self-Monitoring Event

You can view relevant self-monitoring events for each server from the Alert Management and Scheduled Reports tabs. The views in each tab are filtered to display relevant alert or report monitoring events. You can remove the filter to display all self-monitoring events.

To view self-monitoring events

1. Click the Scheduled Reports tab, or the Alert Management tab.

   The Report or Alert servers list appears.

2. Select the server whose local self-monitoring events you want to view.

   The server you select is displayed in the details pane.

3. Click the Self-Monitoring Events tab.

   The Self-Monitoring Events viewer pane appears, showing report or alert-related self-monitoring events. You can perform any of the normal reports tasks from the Self-Monitoring Events pane, including:

   - Event Viewer tasks
   - Global or Local filtering
   - Setting Favorites
   - Exporting

Chapter 12: Suppression and Summarization

This section contains the following topics:

Event Refinement Component Versions

CA Enterprise Log Manager retains earlier versions of certain custom event refinement components as you create and edit them. This allows you to refer back to earlier versions. You can view or copy versions of the following components:

- Message Parsing Files
- Data Mapping Files
- Suppression Rules
- Summarization Rules

Each time you create a new custom component, it is designated Version 1.0. When you edit and save a new version of the same object, it is designated Version 2.0. Both versions appear in the appropriate interface area for selection and application.

For example, if you create a custom suppression rule called "NewRule" it appears as NewRule Version 1.0 in the Event Log Store interface list for application. If you then edit that file, it appears as NewRule Version 2.0 in the Event Log Store list.

You can view older versions of event refinement components in the appropriate list. They are read-only and cannot be edited. You can copy an old version and edit it, making it a new version in turn. For example, using the previous example, you could not edit NewRule Version 1.0 once 2.0 exists. You would have to copy Version 1.0 and edit it. Saving those edits creates Version 3.0.

More information:

Edit a Suppression or Summarization Rule (see page 441)

Suppression and Summarization Rules Tasks

Suppression and summarization rules let you control your event flow and manage event log store size by eliminating or combining certain events. Suppression rules prevent native events that match their qualifications from being recorded at all. Summarization rules combine multiple native events into a single refined event, which appears instead of the original component events.

Important! You should create and use suppression and summarization rules cautiously since they can prevent the recording and appearance of certain native events. We recommend testing custom suppression and summarization rules in a test environment before deploying them.

Suppression and summarization tasks can all be carried out from the Log Collection area of the interface. You can create, edit and delete custom suppression and summarization rules.

More information:

How to Create a Suppression Rule (see page 427)
How to Create a Summarization Rule (see page 432)
Apply a Suppression or Summarization Rule (see page 437)
Copy a Suppression or Summarization Rule (see page 440)
Edit a Suppression or Summarization Rule (see page 441)
Delete a Suppression or Summarization Rule (see page 442)
Import a Suppression or Summarization Rule (see page 443)
Export a Suppression or Summarization Rule (see page 444)

Suppression Rule Effects

During planning, you may want to consider the effect of *suppression rules*, which prevent events either from being inserted into the event log store or collected by a connector. Suppression rules are always attached to a connector. You can apply suppression rules at either the agent or group level, or at the CA Enterprise Log Manager server itself. The placement locations have different effects:

- Suppression rules applied at the agent or group levels prevent events from being collected and thus reduce the amount of network traffic *sent* to the CA Enterprise Log Manager server.
- Suppression rules applied at the CA Enterprise Log Manager server prevent events from being *inserted* into the database and thus reduce the amount of information being stored.

There are potential performance considerations in applying suppression rules to events after they arrive at the CA Enterprise Log Manager server, especially if you create multiple suppression rules or the event flow rate is high.

For example, you might want to suppress *some* of the events from a firewall or from some Windows servers that produce duplicate events for the same action. Not collecting these events can speed up the transport of the event logs you do want to keep, and saves processing time on the CA Enterprise Log Manager server. In such cases, you would apply one or more appropriate suppression rules on agent components.

If you want to suppress all events of a certain type from multiple platforms or across your entire environment, you would apply one or more appropriate suppression rules at the CA Enterprise Log Manager server. Evaluation of events with regard to suppression occurs when events arrive at the CA Enterprise Log Manager server. Applying a large number of suppression rules at the server may lead to slower performance as the server must apply suppression rules in addition to inserting events into the event log store.

For smaller implementations, you can perform suppression at the CA Enterprise Log Manager server. You may also choose to apply suppression at the server for deployments where summarization (aggregation) is in use. If you are only inserting a few of the events from an event source that generates large amounts of event information, you may still choose to suppress unwanted events at the agent or agent group level to save processing time on the CA Enterprise Log Manager server.

How to Create a Suppression Rule

You can use suppression rules to prevent large numbers of routine or known and predicted transactions from inflating your event log store and muddling the image of your environment. For example, you might use a suppression rule to eliminate unnecessary syslog information events, particularly in cases where you cannot configure the event source to send only the required set.

The process of creating a suppression rule, using the suppression rule wizard, has the following steps:

1. Opening the suppression rule wizard.
2. Rule Naming - Entering rule name and description information.
3. Event Selection - Identifying an event to suppress, using the CEG normalization attributes and optional advanced filtering.

Note: Once you have created a suppression rule, you must apply it, making it available for use in your environment.

More information:

Open Suppression Wizard (see page 428)
Name a Suppression Rule (see page 428)
Using Advanced Filters (see page 430)
Apply a Suppression or Summarization Rule (see page 437)

Open Suppression Wizard

To create a new suppression rule, or edit an existing one, open the suppression wizard.

To open the suppression wizard

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Suppression and Summarization folder.

   The Suppression and Summarization buttons appear in the details pane.

3. Click New Suppression Rule: 

   The Suppression Wizard opens.

   When using the wizard:

   - Click Save to save the rule file without closing the wizard.
   - Click Save and Close to save the rule file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Name a Suppression Rule

You must name a suppression rule. You can also enter optional description information for reference.

To name a suppression rule

1. Open the suppression wizard.
2. Type a name for the new rule.
3. (Optional) Type description information.
4. Advance to the Filtering step.

Select an Event to Suppress

You must specify the native event that you want the rule to suppress by setting a simple filter for the CEG event normalization fields. These four fields, which are part of the event-specific class, are provided for all events expressed in the CEG, allowing you to identify a native event precisely.

You can specify the combination of event normalization fields you want using the Simple Filters tab. You can also use advanced filters for further detail in event identification. You must specify at least one simple filter for a suppression rule.

To select a suppression rule event

1. Open the suppression wizard, enter the required information, and advance to the Filtering step.
2. Create simple filters to select the event you want by selecting the appropriate check box, and then selecting or entering the value you want. The available fields are as follows:

   Ideal Model

   Describes the broad class of technology involved in the event, for example, Firewall or Network Device.

   Event Category

   Describes broad categories of events within the Ideal Model. For example, all account, user group, and role-related events are recorded under the "Identity Management" Event Category. Each Event Category has one or more classes (sub-categories), so any choice you make changes the available selections in Event Class menu.

   Event Class

   Provides a more detailed classification of events in a specific event category. For example, Identity Management events are divided into one of three classes: account, group or identity. Each Event Class has one or more associated actions, so any choice you make changes the available selections in Event Action menu.

Event Action

Describes common actions for each Event Category and Class. For example, Account Management, a class of the Identity Management category, contains account creation, deletion, and modification actions.

3. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new rule appears in the list, otherwise the step you choose appears.

When you create a new rule, it is saved as version 1.0. If you later edit the rule, a separate copy of the rule is stored as a new version. You can view earlier versions, and apply or copy them as needed.

More information:

Create a Simple Event Filter (see page 479)
Create an Advanced Event Filter (see page 482)
Using Advanced Filters (see page 430)

Using Advanced Filters

You can use advanced filters to qualify any suppression or summarization-related queries of the event log store. The Advanced Filters interface helps you create the appropriate filter syntax by providing a form for entering logic columns, operators and values according to your suppression or summarization rule requirements.

Note: This section contains a brief overview of the terms used in advanced filters for suppression rules and summarization rules. To use advanced filters to their full potential you need a thorough understanding of the filter terms and the Common Event Grammar.

The following terms join multiple filter statements:

And

Displays the event information if *all* the joined terms are true.

Or

Displays the event information if *any* of the joined terms are true.

The following SQL operators are used by advanced filters to create the basic conditions for summarization or suppression:

Match

Includes any event information that matches one or more of the characters in the alphanumeric string that you enter, allowing you to search for key words. This search is case-sensitive.

Match (ignore case)

Includes any event information that matches one or more of the characters in the alphanumeric string that you enter, allowing you to search for key words. This search is not case-sensitive.

Not Match

Includes any event information that does not match one or more of the characters in the alphanumeric string that you enter. This search is case-sensitive.

Not Match (ignore case)

Includes any event information that does not match one or more of the characters in the alphanumeric string that you enter. This search is not case-sensitive.

Regular Expression Match

Includes any event information that matches one or more of the regular expression characters that you enter. This can be used to search in a multibyte environment, and to search using wildcards.

Not Regular Expression Match

Includes any event information that does not match one or more of the regular expression characters that you enter. This can be used to search in a multibyte environment, and to search using wildcards.

Relational Operators

Include the event information if the column bears the appropriate relation to the value you enter. The following relational operators are available:

- Equal to (numeric)
- Not Equal to (numeric)
- Greater than (numeric)
- Greater than or equal to (numeric)
- Less than (numeric)
- Less than or equal to (numeric)

For example, using *Greater than* would include the event information from your chosen column if its value is greater than the value you set.

All of these operators locate only numbers; to search for other characters, select one the "match" operators, as appropriate.

More information

Create an Advanced Event Filter (see page 482)
Name a Suppression Rule (see page 428)

How to Create a Summarization Rule

You can use summarization rules to combine certain native events of a common type into one refined event. This lets you save space in your event log store and simplifies event analysis.

For example, you might create a summarization rule that records a single refined event for every three failed login attempts by a single user. This means that your event log store records only one event rather than three.

The process of creating or editing a summarization rule using the summarization rule wizard has the following main steps:

1. Opening the summarization rule wizard.
2. Summarization Thresholds - Setting the number or frequency of native events that you want to make up a summarized event.
3. Event Selection - Identifying an event to summarize, using the CEG normalization attributes and optional advanced filtering.
4. Summarization - Controlling how the final summarized event will be presented in your reports.

Note: Once you have created a summarization rule, you must apply it to make it available for use in your environment.

More information:

Open Summarization Wizard (see page 433)
Set Summarization Thresholds (see page 433)
Configure a Summarization Display (see page 436)
Using Advanced Filters (see page 430)
Apply a Suppression or Summarization Rule (see page 437)

Open Summarization Wizard

To create a new summarization rule, or edit an existing one, open the summarization wizard.

To open the suppression wizard

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Suppression and Summarization folder.

   The Suppression and Summarization buttons appear in the details pane.

3. Click New Summarization Rule: 

   The Summarization Wizard opens.

   When using the wizard:

   - Click Save to save the rule file without closing the wizard.
   - Click Save and Close to save the rule file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Set Summarization Thresholds

To create or edit a summarization rule, enter general information, and set summarization thresholds. Thresholds are either a number of events, a frequency of occurrence, or a combination of the two, that trigger the creation of a summarized event.

To set summarization thresholds

1. Open the summarization wizard.
2. Enter a name for the new rule. You can also enter optional description information for reference.

3. Define the combination by specifying the number of native events and elapsed time that your rule uses to create a single refined event, using the Event Summarization menus:

   Enable Event Count Threshold

   Controls whether or not the rule uses an event threshold. The event threshold must be greater than one. Selecting this box sets a maximum events value. If this box is cleared, and the event timeout period is enabled, only the time period is considered in summarizing events. If both are enabled, a summarized event is created at every specified time period, as long as at least one qualified raw event occurs.

   Maximum Events

   Defines the number of native events that trigger a summarized event. When the number of native events you specify occurs, a summarized event is created.

   Minimum: 2

   Maximum: 5000

   Enable Event Timeout Period

   Controls whether or not the rule uses a time period threshold. Selecting this box sets a time period value. If this box is cleared, a summarized event occurs only when the event count threshold is reached.

   Time Period

   Defines the time, in seconds, that elapses to trigger a summarized event, if any events of the specified type have occurred. When this threshold is reached, a summarized event is created, as long as at least one qualified native event has occurred. You can set the Time Period to zero, which will result in a summarized event only when the maximum events threshold is reached.

   Minimum: 0

   Maximum: 86400

   For example, in the case of a rule summarizing failed login attempts, selecting 3 in the Maximum Events menu and 10 in the Time Period menu results in a summarized event after three failed login attempts, or every 10 seconds as long as at least 1 failed login occurs.

4. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new rule appears in the list, otherwise the step you choose appears.

More information:

Configure a Summarization Display (see page 436)

Select an Event for Summarization

Specify the native event that you want the rule to summarize by setting a simple filter for the CEG event normalization fields. These four fields, which are part of the event-specific class, are provided for all events expressed in the CEG, allowing you to identify an event.

You can specify the combination of event normalization fields you want using the Simple Filters tab. You can also use advanced filters for further detail in event identification. Specify at least one simple filter for a suppression rule.

To select a summarization rule event

1. Open the summarization wizard and advance to the Filtering step.
2. Create simple filters to select the event you want by selecting the appropriate check box, and then selecting or entering the value you want. The available fields are as follows:

   Ideal Model

   Describes the broad class of technology involved in the event. For example, Firewall and Network Device are idea models.

   Event Category

   Describes broad categories of events. For example, all account, user group, and role-related events are recorded under the "Identity Management" Event Category. Each Event Category has one or more classes (subcategories), so any choice changes the available selections in Event Class menu.

   Event Class

   Provides a more detailed classification of events in a specific event category. For example, Identity Management events are divided into one of three classes: account, group, or identity. Each Event Class has one or more associated actions, so any choice changes the available selections in Event Action menu.

   Event Action

   Describes common actions for each Event Category and Class. For example, Account Management, a class of the Identity Management category, contains account creation, deletion, and modification actions.
3. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new rule appears in the list, otherwise the step you select appears.

More information:

Create a Simple Event Filter (see page 479)
Create an Advanced Event Filter (see page 482)
Configure a Summarization Display (see page 436)
Set Summarization Thresholds (see page 433)

Configure a Summarization Display

Summarization rules control how native events are displayed in the refined event. You configure a summarization display by selecting Summarized by fields and Aggregated fields.

To configure a summarization rule display

1. Open the summarization wizard and advance to the Summarization step.
2. Select the field or fields you want the refined event to be summarized by, using the shuttle control:

   Summarized By

   Controls the field or fields by which the summarized information is grouped. For example, in the case of a rule summarizing failed logins, select source_username to display the number of qualified failed login events for each unique user. You must select one or more Summarized By fields to complete the rule.

3. (Optional) Select the field or fields you want the refined event to be aggregated by:

   Aggregated

   Controls the field or fields by which the summarized information is subdivided, depending on the Summarized By field. For example, in the case of a rule summarizing failed logins, select source_username as a Summarized By field, and dest_hostname as an Aggregated field. This displays the number of qualified failed login events for each unique user, subdivided by the host that the user attempted to log into.

   The aggregated fields' information is retained in the summarized events' raw event field. In the preceding example each unique host on which the user attempted the log on will be stored along with the number of occurrences, in the following format: *hostname1:2,hostname2:5*. This example shows 2 logon attempts from host 1 and 5 attempts from host 2.

   Aggregated fields are optional - you do not have to select an Aggregated field to complete the rule.

4. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new rule appears in the list, otherwise the step you choose appears.

When you create a new rule, it is saved as version 1.0. If you later edit the rule, a separate copy of the rule is stored as a new version. You can view earlier versions, and apply or copy them as needed.

More information:

Set Summarization Thresholds (see page 433)

Apply a Suppression or Summarization Rule

Once you have created a suppression or summarization rule, you must apply it to make it available for use in your environment. This feature helps prevent the application of suppression or summarization rules without proper testing and approval.

To apply a suppression or summarization rule

1. Click the Administration tab, and then the Services subtab.

   The Service List appears.

2. Click the Event Log Store icon.

   The Event Log Store configuration pane appears.

3. Locate and select the suppression or summarization rule you want to apply, using the appropriate shuttle control.

4. Click Save.

   A confirmation message appears on successful application of the rule.

How to Apply Suppression and Summarization on Agent Components

You can assign suppression rules, summarization rules, or both to agent groups, agents, or connectors in your environment. These rules can replace or supplement any suppression or summarization rules applied at the CA Enterprise Log Manager server. Therefore, you streamline the event transmission/reception process by controlling where event refinement takes place.

For example, if you have a Windows agent group, you can associate a suppression rule that eliminates unnecessary Windows events to the agents in the group. You eliminate the need for all incoming events to undergo a Windows-specific check at the CA Enterprise Log Manager server.

You can apply suppression or summarization rules at different levels of the agent folder hierarchy:

- From the Agent Explorer folder, you can apply rules to any agent groups, individual agents, or connectors.
- From a specific agent group folder, you can apply rules to all agents within that group and all the connectors assigned to them.
- From an individual agent, you can apply rules only to that agent and any connectors assigned to it.

The process of applying suppression or summarization rules on the agent components has the following steps:

1. Opening the manage rules wizard.
2. Selecting targets; agent groups, agents, or connectors.
3. Choosing suppression rules to apply.
4. Choosing summarization rules to apply.

You can also remove suppression or summarization rules from multiple agent groups, agents, or connectors using the manage rules wizard.

More information:

Open Manage Summarization Rules Wizard (see page 438)
Select Suppression and Summarization Targets (see page 439)
Choose Suppression Rules to Apply (see page 439)
Choose Summarization Rules to Apply (see page 440)

Open Manage Summarization Rules Wizard

To apply suppression or summarization rules to agent groups, or individual agents or collectors, you can use the manage rules wizard.

To open the manage rules wizard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the Agent Explorer folder, and then click Manage Suppression and Summarization Rules: 

   The manage rules wizard appears.

When using the wizard:

- Click Save to save the file without closing the wizard.
- Click Save and Close to save the file and close the wizard.
- Click Reset to restore the wizard display to the last-saved settings.

Select Suppression and Summarization Targets

To apply suppression or summarization rules to agent components, select targets for the rules.

To select targets

1. Open the manage rules wizard.
2. Select whether you want to apply rules to Agent Groups, Agents, or Connectors.
3. (Optional) Select Delete if you want to remove rules rather than add them.
4. Select the targets you want using the shuttle control.

   Note: You can search for agent or connector names. If no agents or connectors appear in the available list, click Search to display all available agents or connectors.
5. Advance to the rules application step you want.

Choose Suppression Rules to Apply

To finish assigning suppression rules to an agent group, agent, or connector, select which rules to apply.

To choose suppression rules

1. Open the manage suppression rules wizard, and advance to the Apply Suppression Rules step.
2. Choose which of the available rules to apply, using the shuttle control.

   Note: You can search for suppression rules using the Suppression Rules Pattern field.
3. Click Save and Close if you are finished applying rules.

   The rules you select are applied to the chosen targets. If you selected Delete in the Select Targets step, the rules you choose are deleted.

Choose Summarization Rules to Apply

To finish assigning summarization rules to an agent group, agent, or connector, select which rules to apply.

To select summarization rules

1. Open the manage summarization rules wizard, and advance to the Apply Summarization Rules step.
2. Select which of the available rules to apply, using the shuttle control.

   Note: You can search for summarization rules using the Summarization Rules Pattern field.
3. Click Save and Close if you are finished applying rules.
4. The rules you select are applied to the chosen targets. If you selected Delete in the Select Targets step, the rules you select are deleted.

Copy a Suppression or Summarization Rule

You can copy a suppression or summarization rule, allowing you to create a new rule based on an existing one.

To copy a suppression or summarization rule

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.
2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Suppression and Summarization folder.

   The suppression and summarization buttons appear in the details pane.
3. Click the Suppression and Summarization folder that contains the rule you want to copy.

   The folder opens, displaying the rules.
4. Select the rule you want to copy, and click Copy Selected Item: 

   The suppression or summarization wizard opens, displaying the rule.
5. Make any changes you want, and click Save and Close.

   The rule appears in the appropriate list.

Edit a Suppression or Summarization Rule

You can edit a suppression or summarization rule.

To edit a suppression or summarization rule

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Suppression and Summarization folder.

   The suppression and summarization buttons appear in the details pane.

3. Click the Suppression and Summarization folder that contains the rule you want to edit.

4. Select the rule you want to edit, and click the Edit Suppression or Summarization Rule icon.

   The Suppression wizard or the Summarization wizard appears, displaying your selected rule.

5. Make the changes you want, and click Save and Close.

   The rule appears in the appropriate list as new version of the edited rule.

More information:

Event Refinement Component Versions (see page 425)

Delete a Suppression or Summarization Rule

You can delete an unneeded suppression or summarization rule.

To delete a suppression or summarization rule

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Suppression and Summarization folder.

   The suppression and summarization buttons appear in the details pane.

3. Click the Suppression and Summarization folder which contains the rule you want to delete.

4. Select the rule you want to delete and click the Delete Suppression or Summarization Rule icon. The current version is selected by default. You can select an earlier version to delete from the Version pull-down list in the details pane.

   A confirmation dialog appears. If you have applied the rule to an integration, a warning appears. Deleting the rule also removes it from the integration.

5. Click Yes.

   The deleted rule is removed from the appropriate list.

Import a Suppression or Summarization Rule

You can import a suppression or summarization rule, allowing you to move rules from one environment to another. For example you could import rules created in a test environment to your live environment.

To import a suppression or summarization rule

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Suppression and Summarization Rules folder.

   The Import Suppression and Summarization Rule and Export Suppression or Summarization Rule buttons appear in the details pane.

3. Click Import Suppression or Summarization Rule.

   The import file dialog appears.

4. Browse to find the file you want to import, and click OK.

   The Suppression or Summarization Wizard appears, displaying the details of the rule you selected.

5. Make any changes you want, and click Save and Close. If the imported rule shares a name with a rule already in your management database, you are prompted to change the name.

   The imported rule appears in the appropriate suppression or summarization folder.

Export a Suppression or Summarization Rule

You can export a suppression or summarization rule. This lets you share rules between environments. For example, you could export rules created in a test environment to your live environment.

To export a suppression or summarization rule

1. Click the Administration tab, and then click the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Suppression and Summarization Rules folder.

   The Export Suppression or Summarization Rule button appears in the details pane.

3. Click the Suppression Rules or Summarization Rules folder which contains the file you want to export.

   The folder expands, showing the individual files.

4. Select the rule you want to export, and then click Export Suppression or Summarization Rule. The current version is selected by default. You can select an earlier version to export from the Version pull-down list in the details pane.

   An export location dialog appears.

5. Enter or browse to the location where you want to store the exported rule, and click Save.

   An export successful confirmation dialog appears.

6. Click OK.

   The rule is exported.

Create a Windows Event 560 Suppression Rule

Enabling object access auditing on a Windows server creates a significant volume of event traffic, some of which you may wish to eliminate. For example, Windows generates two events each time an administrator opens the Microsoft Management Console (mmc.exe). These events have ID values of 560 and 562.

In this example, you create a new rule that suppresses Windows events with an event_id of 560. Completing the steps in the following procedure gives you an actual suppression rule you can use in your network environment as well as demonstrating how to use the wizard.

To get started with this example, you must log in to a CA Enterprise Log Manager server as a user with the Administrative role and privileges. You cannot create or edit suppression rules while logged in as the EiamAdmin user.

To create a suppression rule for Windows 560 events

1. Open the suppression rule wizard.
2. Type "Windows Event 560 Suppression" in the name entry field, and add the description, "This rule suppresses Window event 560 since the OS also creates Event 562 for the same type of resource access. Its retention is not needed for demonstrating compliance."
3. Advance to the Filtering step and select the following simple filters:
   a. Ideal Model value, Operating System.
   b. Event Category value, Resource Access.
   c. Event Class value, Resource Open.
   d. Event Action value, Resource Activity.
4. Click the Advanced Filters tab, and the New Event Filter button.

   A new filter line appears in the table. You can click a value or the empty space in each table cell to select or enter a new value.

   The Logic operator field defaults to the value, AND. If you have several different types of events that you wanted to suppress, you can enter their event IDs with new lines that use the OR logical operator.
5. Set the advanced field filter values:
   a. Click the value in the Column field and select the field, event_id.
   b. Click the Operator field and select Equal To
   c. Click the Value field and enter the value, 560.
6. Click Save and Close.

   The wizard automatically creates a User folder to contain your suppression rules. You can see this folder by expanding the Suppression Rules folder.

Chapter 13: Mapping and Parsing

This section contains the following topics:

Event States

Information about events in your environment passes through a number of stages, from initial occurrence to possible final display by CA Enterprise Log Manager. Because the term "event" can refer to any one of these stages, we use the following terminology for the possible event states in your environment:

Native Event

Refers to the original occurrence of the state or action that triggers the event, a failed authentication, or firewall violation for example. The appropriate connector or listener service sends native events, parsed and mapped as appropriate, then inserted into the event log store, where it is available for display as raw or refined events.

Raw Event

Refers to the communication sent by the appropriate monitoring agent. Raw events contain information about the native event, often in the form of a syslog string or a name-value pair. This information is stored and searchable unless altered by suppression or summarization rules. Suppressed events are not recorded in the event log store; a set of summarized events is recorded as a single event expressing the outcome of the summarization.

Refined Event

Refers to the event information as mapped and summarized by CA Enterprise Log Manager. This information is stored and searchable.

Recorded Event

Refers to the raw or refined event information in the event log store. Raw events and refined events are always recorded unless suppressed or summarized. Mapped events have both raw and refined information available. This information is stored and searchable.

Consult the following diagram for information about event states:

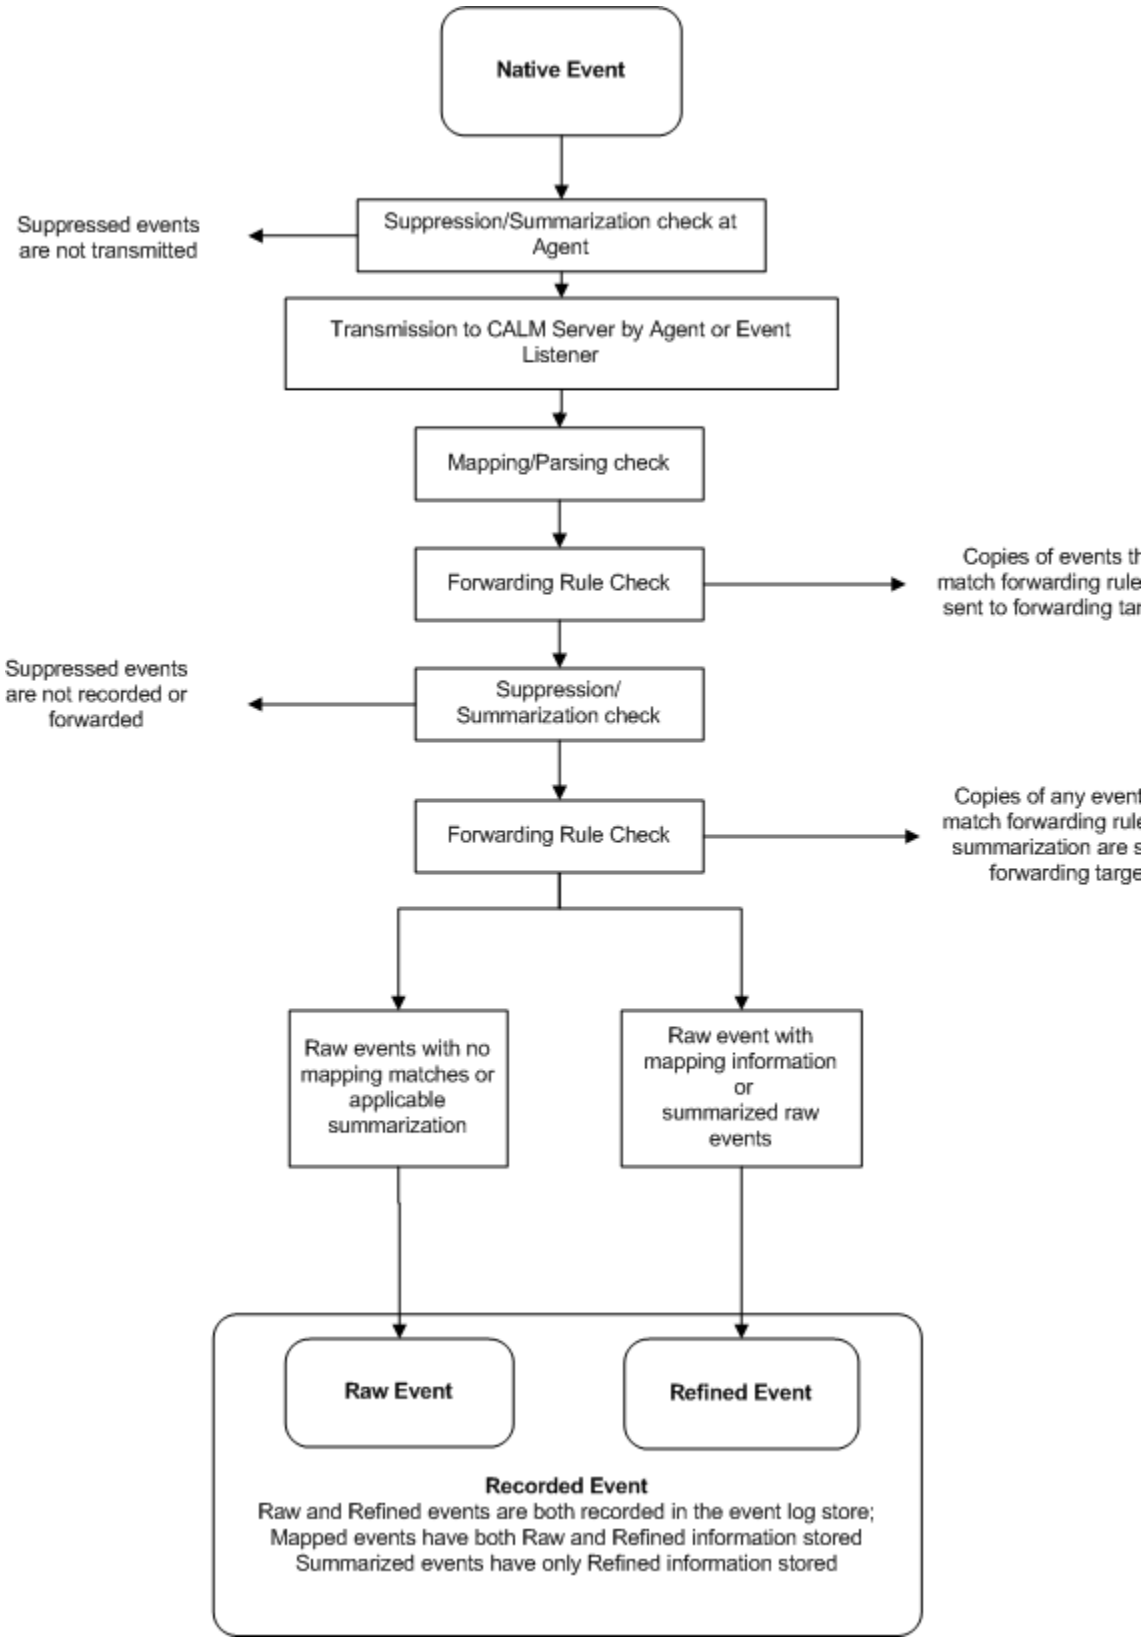
Native Event
Suppression/Summarization check at Agent
Suppressed events are not transmitted
Transmission to CALM Server by Agent or Event Listener
Mapping/Parsing check
Forwarding Rule Check
Copies of events th
match forwarding rule
sent to forwarding tar
Suppression/ Summarization check
Suppressed events are not recorded or forwarded
Forwarding Rule Check
Copies of any event
match forwarding rule
summarization are s
forwarding targe
Raw events with no mapping matches or applicable summarization
Raw event with mapping information or summarized raw events
Raw Event
Refined Event
Recorded Event
Raw and Refined events are both recorded in the event log store;
Mapped events have both Raw and Refined information stored
Summarized events have only Refined information stored

More information:

Suppression and Summarization Rules Tasks (see page 426)
Mapping and Parsing Rules Tasks (see page 450)

Mapping and Parsing Rules Tasks

Message Parsing (XMP) and Data Mapping (DM) file pairs collect and normalize data from specific types of event sources. Most incoming native events pass through the parsing and then the mapping processes to create a reportable event that is inserted into the event log store. Events transmitted through SAPI or iTechnology do not require parsing, and proceed directly to the data mapping stage.

Note: To take full advantage of these advanced features, you need a thorough understanding of the raw and collected events in your environment, the target fields you want to parse, the regular expression syntax, the CEG, and DM and XMP files and how they parse events.

The XML-based XMP files read incoming raw event data and create name-value pairs, according to your specifications. DM files then map the events' name-value pairs assigned by message parsing into the common event grammar. When creating new parsing and mapping files, consider them as part of a process. For example, efficient and complete parsing allows quick and process-effective mapping.

More information:

Event Refinement Component Versions (see page 425)
How to Create a Message Parsing File (see page 450)
How to Create a Data Mapping File (see page 467)

How to Create a Message Parsing File

You can use the parsing file wizard to create, edit, or analyze a Message Parsing (XMP) file. Parsing files read incoming raw event data and create name-value pairs, allowing you to establish mappings even before the data mapping process. This improves overall mapping efficiency.

Note: The Common Event Grammar (CEG) names are not enforced for event parsing, allowing additional flexibility in creating name/value pairs. The CEG fields are available for selection, but the field names and values are not limited to CEG values.

Creating or editing an XMP file has the following steps:

1. Opening the parsing file wizard.
2. Providing file details, including file name, logname, and support information.
3. Locating sample events for file testing and construction.
4. Setting global values that will apply to all events parsed by the file.
5. Creating or editing prematch strings to begin event parsing
6. Selecting prematch filters for parsing filter attachment
7. Creating or editing parsing filters to complete event parsing.
8. Analyzing and saving the new or edited XMP file.

More information:

Define File Details (see page 452)
Load Sample Events (see page 453)
Add Global Fields (see page 454)
Create a Prematch Filter (see page 455)
Analyze the XMP File (see page 466)

Open Parsing File Wizard

To create a message parsing rule, or edit an existing one, you must open the parsing file wizard.

To open the parsing file wizard

1. Click the Administration tab, and then the Library subtab.
2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Mapping and Parsing folder.

   The product integration buttons appear in the details pane.
3. Click New Message Parsing Rule: 

   The Parsing File Wizard opens.

   When using the wizard:

   - Click Save to save without closing the wizard.
   - Click Save and Close to save the file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Define File Details

You can add new parsing file details, including name, source, and reference information. Newly created or edited files are displayed in the User folder in the Mapping and Parsing area.

To add new parsing file details

1. Open the parsing file wizard.
2. Specify Parsing File Information area as outlined in the following substeps:
   a. Type a name for the file. The file name is required, and cannot contain the characters: / \ : * ? “ < > ^;'`,&{}[]. or |.
   b. Type the source logname to identify the logname of the event type you wish the file to parse. The auto-complete feature presents available lognames as you type. The logname you choose will be displayed in the event_logname field of the refined event.
   c. Add a description for reference if needed.
3. (Optional) Add Support Information for reference as outlined in the following substeps.
   a. Click Add Product in the Support Information area.

      A new support information row appears.
   b. Click the New Product or New Version text to enable entry fields, and type the product/version information you want.
4. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Parsing File User folder, otherwise the step you choose appears.

Load Sample Events

You can provide sample events to use in testing the new XMP file by searching through the event log store or by accessing a log file. Sample events provide a template against which you can test the parsing file as you construct it in the other wizard steps. You can also use sample events to test the parsing output in the final step of the wizard.

To provide sample events

1. Open the parsing file wizard and advance to the Load Events step.

   The Load Events screen appears.

2. Select the Log Store or Log File radio button in the Find Sample Events area.

   - If you select Log Store:

     a. Select the sample event source type you want from the Parsing Column drop-down menu. Choose result_string for WMI event sources, or raw_event for syslog event sources.

     b. Select the query you want to use to provide sample events using the Query Tag Filter and Query List.

     The query appears, displaying sample events you can use to test parsing as you advance through the wizard.

     Note: You can use any available or custom query to locate sample events. If you plan to use a custom query, we recommend that you create and test it before beginning the message parsing file design process. We recommend using a sample event file with less than 1500 events for ease of analysis.

   - If you select Log File, browse to find the log file you want, and click Upload.

     Events from the log file appear in the Sample Events pane. You can use the events to test parsing as you advance through the wizard.

     Note: The wizard assumes that each line in the file is an event. Multiple line events are not supported.

3. Click the appropriate arrow to advance to the wizard step you want to complete next.

   If you click Save and Close, the new file appears in the Parsing File User folder, otherwise the step you choose appears.

Add Global Fields

You can add global fields, which are static pairs that match a field name with a specific value. The parsing process adds the global fields to all parsed events, so they are best used for default values such as the ideal model.

To add global fields

1. Open the parsing file wizard and advance to the Global Fields step.

   The Global Fields screen appears.

2. Click Add Global Field in the Global Fields area.

   A new global field row appears in the fields table, displaying New Global Field and New Value entries.

3. Click the New Global Field text to enter the name information you want. The auto-complete feature presents available CEG field names as you type. You can click one to select it, or type a non-CEG field name.

4. Click the New Value text to enter the name information you want.

5. (Optional) Repeat steps 2-4 to add additional global fields as needed.

6. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Parsing File User folder, otherwise the step you select appears.

Create a Prematch Filter

You can create a prematch filter to help the XMP file narrow its search for event information you want to parse. The prematch filter identifies a selected text string to narrow the event selection process, which is then completed by parsing filters. If you consider the parsing file as a funnel, the prematch filter forms the mouth and the parsing filter is the spout.

The more complete your prematch filtering is, the more efficient your parsing process is. This is because narrow prematch categories held reduce the processing effort required to parse events.

For example, if you wanted to parse access attempt events, you might create a prematch filter that searches for the text "login", and add appropriate parsing filters to that prematch filter.

Note: Deleting a prematch filter also removes its associated parsing filter or filters.

To create a prematch filter

1. Open the parsing file wizard and advance to the Match and Parse step.

   The wizard displays any existing prematch filters in the Prematch Filters list. Each one displays the number of prematches to any sample events in parentheses beside it.

2. Click Add a Prematch String at the top of the Prematch Filters list, or select a prematch filter to edit.

   Note: To select a prematch filter, type the first few characters of the prematch string in the Search field. All the prematch strings matching the entered characters are displayed. Within the resulting matching prematch strings, you cannot use the up-down arrows to move a prematch string.

3. Type the text you want the filter to search for in the Prematch String entry field.

   Any sample events that match the text you enter immediately appear, along with the number of matched events found and parsed.

4. (Optional) Click Add prematch based on unmatched events to show all unmatched sample events.

   Any sample events that are currently unmatched appear in the Events area for easy reference in creating a new prematch filter.

5. (Optional) Add or edit additional prematch filters as needed.

6. Set the order in which you want the parsing process to search for prematches, using the up-down arrows beside the Prematch Filters list. Setting prematch filters that match more events higher in the priority list improves the efficiency of your parsing process.

7. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Parsing File User folder, otherwise the step you choose appears.

Create a Parsing Filter

You can create a parsing filter to define how the XMP file parses event data. Each parsing filter is attached to a prematch filter. After the parsing process locates a prematch string, it uses each parsing filter attached to that prematch in turn to locate its specified information. The parsing process returns the first positive match it makes.

When you click the Add a Parsing Filter button in the Match and Parse step of the Message Parsing wizard you start the Parsing File Filter wizard. To create effective parsing filters you need a good understanding of the regular expression syntax.

To create a parsing filter

1. Open the Parsing File Filter wizard, and type a filter name and optional description in the Filter Details page.
2. Click Add new to add a static field value that you want to appear in all events parsed by the filter.

   A static field row appears, displaying New Field and New Value cells.
3. Type an entry in the New Field cell, and type an entry in the New Value cell. The auto-complete feature narrows available CEG field names as you type in the New Field cell, and presents a menu of choices.
4. (Optional) Repeat steps 2-3 to add static field values as needed.
5. Advance to the Regular Expression step.

   The Parsing Expression Testing window opens, displaying any current regular expression. Immediately below the regular expression is the Event pane. This area shows one or more sample events, if you previously loaded sample events. The wizard can test these events against the regular expression as you build it.
6. Click Add or Remove Tokens from Library to display a list of predefined regular expressions you can add for use in the current filter. Select the tokens you want to add and click OK to add them to the Parsing Tokens list.
7. (Optional) Click New Regular Expression Token to create a Parsing Token, and enter its regular expression syntax in the Token Details pane. You can now create custom expressions for your environment. You can add a custom token to your local library by clicking Add Selected Token to the Library at the top of the Parsing Tokens pane.

   Note: When you create a new datetime token, select the 'Treat as a datetime value' check box to enter a format for parsing the time value. This value does not affect the display format.
8. Add regular expression statements for the filter in the Regular Expression entry field. You can drag and drop expressions from the Parsing Tokens list. You can also type or edit the expression directly in the Regular Expression entry field.

   Note: Selecting a token in the Parsing Tokens list displays its regular expression syntax in the Token Details pane. You can view the parsing token mapping in a given rule to repeat it in other parsing rules.

9. (Optional) Select the Dynamic Name/Value Pairs checkbox if your target events include key pairs you want to display. See "Dynamic Parsing" for more information.
10. (Optional) If you want to use dynamic parsing, enter a dynamic parsing expression in the dynamic pairs entry field. For example, enter:

    (_PAIR_KEY_)=(_PAIR_VALUE_);

    Any pairs separated by an equal sign and spaced by a semicolon appear. You can enter more expressions to find pairs displayed in other formats. See Dynamic Parsing for more information.
11. Preview how the file parses the sample events using the Event and Parsed Event panes. As you modify the parsing filter regular expression, parsed portions of the sample event are highlighted in blue text and dynamically parsed pairs appear in green. You can verify the effectiveness of the parsing.
12. (Optional) Change the sample event for additional testing by using the back and forward arrows under the Event pane to move through the available sample events.
13. Click Save and Close when you are satisfied with the regular expression. You can use Reset to return the regular expression to its initial state.

    The Parsing File Filter wizard closes, returning you to the Match and Parse step of the Parsing File wizard.

More information:

Dynamic Parsing (see page 459)
Parsing Tokens (see page 459)
Add a Custom Token to the Library (see page 463)

Dynamic Parsing

You can use dynamic parsing, which allows the display of multiple, unaltered name-value pairs that already exist in the raw event. Unlike normal parsing where each parsed token can be allotted to a CEG field or a user-defined field, the name part of the name/value pair becomes the field and cannot be assigned to any CEG field or user defined field. Dynamic parsing is useful where applications or formats record event data in key pairs that you wish to protect from change, not parsed into CEG names or other values. It also improves parsing performance in the cases where it is applicable.

The regular expression which allows dynamic parsing contains four elements:

1. A pair key indicator "(_PAIR_KEY_)"
2. A pair value indicator "(_PAIR_VALUE_)"
3. A key-value separator between the pair and key value
4. A pair separator between the whole expression and the next expression.

The separators you use must match the structure of the event source you are parsing. If your event source uses a comma as a separator, your regular expression must do as well.

Example

```
(dest_objectclass)=(ServerE);
```

In this example the key-value separator is "=" and the pair separator is ";"

Using this expression after other regular expressions allows the XMP file to locate and display any key pairs that appear in parsed events.

Parsing Tokens

A parsing token is a regular expression template that you can use to build parsing filters. CA Enterprise Log Manager includes a parsing token library that contains predefined parsing tokens. For example, the _IP_ token sets the regular expression that parses the typical IP address format. When you want a parsing filter to extract an IP address you can insert the _IP_ token into the filter rather than constructing the full regular expression language each time.

You can also create your own custom parsing tokens, and add them to the local library, or export them for use in another CA Enterprise Log Manager environment. If you want to export a custom token, add it to the library first. You can also import custom tokens from another CA Enterprise Log Manager environment to create parsing tokens in a test environment and move them to a live environment.

More information:

Datetime Token Values (see page 461)
Add a Custom Token to the Library (see page 463)
Remove a Custom Token from the Library (see page 464)
Import Parsing Tokens (see page 465)
Export Parsing Tokens (see page 466)

Datetime Token Values

CA Enterprise Log Manager supports various syntax options for datetime parsing tokens. You can use these options, in the parsing file datetime format, to customize your datetime stamp appearance.

Each datetime token is composed of one of the following:

- An ordinary character (neither '%' nor a white-space character), which is displayed as entered: a colon to separate time values, for example.

  or

- A conversion specification. Each conversion specification is composed of a '%' character followed by a conversion character which defines the display output: %m to display the month, for example.

CA Enterprise Log Manager supports the following conversion specifications:

%a or %A

Displays the local weekday name, in full or abbreviated form. On Windows, this specification is available in US English only.

%b or %B or %h

Displays the local month name, in full or abbreviated form. On Windows, this specification is available in US English only.

%c

Displays the local date and time.

%C

Displays the century number (0-99).

%d or %e

Displays the day of the month (1-31).

%D

Displays the American style date: Month/Day/Year - the equivalent of entering %m/%d/%y.

Note: The syntax %d/%m/%y is used in Europe. The ISO 8601 standard format is %Y-%m-%d.

%H

Displays the hour on a 24-hour clock (0-23).

%I

Displays the hour on a 12-hour clock (1-12).

%j

Displays the day number of the year (1-366).

%m

Displays the month number (1-12).

%M

Displays the minute (0-59).

%n

Inserts an arbitrary whitespace.

%p

Displays the local equivalent of AM or PM, if any.

%r

Displays the 12-hour clock time: Hour:Minute:Second AM/PM - the equivalent of entering %I:%M:%S %p. If t_fmt_ampm is empty in the local LC_TIME section then the behavior is undefined.

%R

Displays the 24-hour clock time: Hour:Minute - the equivalent of entering %H:%M.

%S

Displays the second (0-60 - 60 can occur for leap seconds).

%t

Displays an arbitrary whitespace.

%T

Displays the 24-hour clock time: Hour:Minute:Second - the equivalent of entering %H:%M:%S.

%U

Displays the week number. Sunday is the first day of the week (0-53). The first Sunday of January is the first day of week 1.

%w

Displays the weekday number (0-6) with Sunday = 0.

%W

Displays the week number with Monday the first day of the week (0-53). The first Monday of January is the first day of week 1.

%x

Displays the date, using the local date format.

%X

Displays the time, using the local time format.

%y

Displays the year in the current century (0-99). When a century is not specified, values in the range 69-99 refer to years in the twentieth century (1969-1999); values in the range 00-68 refer to years in the twenty-first century (2000-2068).

%Y

Displays the year, including century (for example, 1991).

%z

Displays an RFC-822/ISO 8601 standard time zone specification. This specification is not available on Windows.

The default CA Enterprise Log Manager datetime token format is:

%d/%b/%Y:%H:%M:%S %z

Add a Custom Token to the Library

You can add custom parsing tokens to the token library, making them available for other users. For example, if you create a custom token during the message parsing file creation process, and it would be useful for other parsing, you can add it to the library for re-use.

The following procedure assumes that you add tokens during the creation of parsing files or filters.

To add a custom parsing token to the library

1. Open the Message Parsing wizard, and advance to the Match and Parse step.
2. Open the Parsing File Filter wizard, and advance to the Regular Expression step.
3. Click New Regular Expression to create a Parsing Token, and enter its regular expression syntax in the Token Details pane.
4. Select the new parsing token, and click Add selected token to the library.

   A confirmation dialog appears
5. Click Yes.
6. (Optional) Click Add or Remove Tokens from Library to view the new token.

   The parsing tokens library dialog appears, showing custom tokens in black and predefined tokens in green.

Remove a Custom Token from the Library

You can remove unneeded or obsolete custom tokens from the token library. Predefined tokens cannot be removed.

To remove custom tokens from the library

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow next to the Event Refinement Library folder to expand it, and then select the Mapping and Parsing folder.

   The product integration buttons appear in the details pane.

3. Click New Message Parsing Rule: 

   The Parsing File Wizard opens.

4. Advance to the Match and Parse step.

5. Select any prematch filter, and click Edit or click Add a Parsing Filter at the top of the Parsing Filters list.

   The Parsing File Filter wizard appears.

6. Advance to the Regular Expression step.

7. Click Add or Remove Tokens from Library.

   The parsing tokens library dialog appears, showing custom tokens in black and predefined tokens in green.

8. Select the custom token or tokens you want to remove, and click Remove Selected Token from the Library.

   A confirmation dialog appears.

9. Click yes, then Click OK.

Import Parsing Tokens

You can import parsing tokens to add custom parsing tokens created on another management server to your current server, from a test environment to your live environment, for example.

To import parsing tokens

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow next to the Event Refinement Library folder to expand it, and then select the Mapping and Parsing folder.

   The product integration buttons appear in the details pane.

3. Click New Message Parsing Rule: 

   The Parsing File Wizard opens.

4. Advance to the Match and Parse step.

5. Select any prematch filter, and click Edit or click Add a Parsing Filter at the top of the Parsing Filters list.

   The Parsing File Filter wizard appears.

6. Advance to the Regular Expression step.

7. Click Import User Tokens at the top of the Parsing Tokens pane.

   The Import File dialog appears.

8. Browse to find the tokens (.tok) file you want to import, and click OK.

   A confirmation dialog appears.

9. Click Yes if you want to import the file, overwriting any other user tokens in the library.

Export Parsing Tokens

You can export parsing tokens that you have added to the token library to move custom parsing tokens created on the current management server to another server. For example, you could move your custom tokens from a test environment to your live environment.

To export parsing tokens

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow next to the Event Refinement Library folder to expand it, and then select the Mapping and Parsing folder.

   The product integration buttons appear in the details pane.

3. Click New Message Parsing Rule: 

   The Parsing File Wizard opens.

4. Advance to the Match and Parse step.

5. Select any prematch filter, and click Edit or click Add a Parsing Filter at the top of the Parsing Filters list.

   The Parsing File Filter wizard appears.

6. Advance to the Regular Expression step.

7. Click Export User Tokens at the top of the Parsing Tokens pane.

   A download location dialog appears.

8. Select the location where you want to save the exported file, and click Save.

   The exported file is saved in your chosen location.

Analyze the XMP File

You can use the Message Parsing utility to analyze your new or edited file and determine how effective the parsing file is against the sample events. Analysis lets you make modifications to improve the efficacy of the file before saving it.

The utility analyzes an XMP file against your selected sample event set using the following process:

1. Locating all events containing the prematch strings defined in the XMP file. The utility runs a separate search for each prematch string, finding all events containing that string.
2. Finding the first parsing filter for each of the prematched events that can parse the event into tokens.

To analyze the XMP file

Open the parsing wizard and advance to the Parsing Analysis step. The wizard displays the number of matches for the prematch strings and filters. The more matches you have, the more efficient the new or edited XMP file will be. This also allows you to determine if there is any significant information that remains unparsed.

The XMP analysis can take some time to process if the XMP file and the number of sample events are both large. It should not usually take over a minute. You can cancel this process if it is taking too long and then re-analyze using a smaller number of events.

When you create a new rule, it is saved as version 1.0. If you later rule edit the rule, a separate copy of the rule is stored as a new version. You can view earlier versions, and apply or copy them as needed.

How to Create a Data Mapping File

You can use the mapping file wizard to create and edit Data Mapping files, which convert native events into refined events by mapping the parsed text string or field/value pairs to CEG-compatible fields. The mapping file wizard allows you to create and edit various types of mapping to accomplish this.

The process of creating or editing a DM file contains the following steps:

1. Opening the mapping file wizard.
2. Providing file details.
3. Locating and adding sample events using parsing files.
4. Setting direct mappings as needed.
5. Setting function mappings as needed.
6. Setting conditional mappings as needed.
7. Setting block mappings as needed.

   Note: You can set direct or function mappings using block mappings. They are an alternative to setting mappings with steps 4 and 5.
8. Analyzing and saving the completed DM file.

When creating a DM file, you should consider the data mapping priorities of the file itself, as well as the individual mapping types within the file. The completed DM file checks event information in the order of the mapping type screens (steps 4-7 in the wizard). If duplicate mapping types exist, the last value the DM file finds is the one assigned.

For example, if a DM file finds a Direct mapping for a given native event value, and then a different Conditional mapping for the same value, the refined event uses the Conditional mapping result.

Duplicate mappings *within* a given mapping type are handled differently, depending on the type:

- Direct Mappings and Function mappings -- The DM file uses the last matching value found. If a duplicate function mapping is found, the last function will be the one called. For example, you could set a duplicate mapping to call a second function if the first was not found or did not function as expected.
- Conditional mappings and Block mappings -- The DM file applies the first value found, and stops searching. To improve performance, you should place more common conditions earlier in the file for both these mapping types.

More information on the design implications of the mapping order is included in the individual mapping type procedures.

Open Mapping File Wizard

To create a new DM file, or edit an existing one, you must open the mapping file wizard.

To open the mapping file wizard

1. Click the Administration tab, and then the Library subtab.
2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Mapping and Parsing folder.

   The product integration buttons appear in the details pane.
3. Click New Mapping File: 

   The Mapping File Wizard appears.

   When using the wizard:

   - Click Save to save the file without closing the wizard.
   - Click Save and Close to save the file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

More information

Set Concat Function Mapping (see page 473)
Set Block Mappings (see page 476)
Perform Mapping Analysis (see page 477)

Provide File Details

Provide file details for a new DM file. You can save a subscription file as a custom file under a different name.

To provide mapping file details

1. Open the mapping file wizard.
2. Enter a name for the DM file. The file name is required, and cannot contain the characters: / \ : * ? “ < > ^;'`,&{}[]. or |.
3. Select the Parsing File name and version you want to use to parse the sample events from the Parsing File drop-down list.

   The log name field is automatically populated with the name of the parsing file you enter.
4. (Optional) Enter a description.
5. (Optional) Click Add Product in the Support Information area to enter product name and versions for reference.
6. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Mapping File User folder, otherwise the step you select appears.

Provide Sample Events

You can use the Mapping File wizard to search for sample events to use in analyzing the DM file. You can search through the event log store or provide sample events directly from a log file. Sample events provide a template against which to test the mapping output in the final step of the wizard.

To provide sample events

1. Open the mapping file wizard and advance to the Sample Events step.

   The Sample Events screen appears.

2. Select the Log Store or Log File option button in the Find Sample Events area.
3. If you select Log Store:
   a. Select the sample event source type you want from the Parsing Column drop-down menu. Select result_string for WMI event sources, or raw_event for syslog event sources.
   b. Select the query you want to use to provide sample events, using the Query Tag Filter and Query List.

      The query appears, displaying the sample events.

      Note: You can use any available or custom query to locate sample events. If you plan to use a custom query, we recommend that you create and test it before beginning the data mapping file design process.
4. If you select Log File:
   a. Browse to find the log file you want, and click Upload.

      Events from the log file appear in the Sample Events pane.

      Note: The wizard assumes that each line in the file is an event. Multiple line events are not supported.
   b. Click Extract Dynamic Fields, if your sample log file contains dynamic pair values you want to include in the parsed sample.
5. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Mapping File User folder, otherwise the step you select appears.

More information:

Dynamic Parsing (see page 459)

Set Direct Mappings

Direct mappings set 1-1 correspondences between a native event and a single refined event value. Thus, it is best to use direct mappings only for default values, or common values that rarely change, such as the ideal_model field.

A mapping can be set to derive a refined event value in the following ways:

Text value

Sets specific text for a specific CEG field. This value appears each time an appropriate event is mapped. For example, setting the CEG ideal_model field to "Firewall" results in the ideal_model field displaying "Firewall" for all rules that contain that mapping.

Field value

Sets a raw event field whose content is included for a specific CEG or parsed field. A field value is distinguished from a text value by prefacing the value with a dollar sign, $. For example, setting the CEG event_logname field to "$Log" results in any event mapped displaying whatever text appears in the native event Log field.

To set direct mappings

1. Open the mapping file wizard, enter a name and select a Logname for the mapping file, and advance to the Direct Mappings step.

   The Direct Mappings screen appears, displaying current or default mappings. The Name column shows the CEG or parsed field name. The Value column shows either a text value or a field value.

   Note: Select a parsing file in the Provide Sample Events step for parsed field values to appear.

2. Click Add Direct Mapping to add a new mapping entry at the bottom of the table and then select it, or select a current direct mapping to edit.

   The direct mappings for the field, if any, appear in the Mapping Details area.

3. Select a CEG field or parsed event field, if available, to map to from the Field drop down menu. When you begin typing, the auto-complete feature narrows the list of available CEG fields.

4. Enter a new value in the Add Value entry field, and click Add Direct Mapping next to it. Precede the value with "$" to denote a field value rather than a text value.

   The value appears in the Selected Fields area.

5. (Optional) You can enter multiple direct mappings for a single field, using the up and down arrows to set the order in which the DM file considers them. The refined event displays the last direct mapping located by the DM file.

   Note: Adding multiple values decreases performance of the mapping, so you use this feature conservatively.

6. (Optional) Use the shuttle control to move unneeded values to the Available Fields area to prevents them from being considered for the current mapping.
7. When you have added all the direct mappings you want, click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

If you click Save and Close, the new file appears in the Mapping File User folder, otherwise the step you select appears.

Set Function Mappings

Function mappings link a CEG field to a value using a function to retrieve or define the refined event information that appears in the refined event. All function mappings consist of a CEG field name, a predefined or class field value, and the function.

For example, a function mapping can concatenate a series of native event values to a single CEG field using the concat function.

If there are duplicate function mappings, the DM file uses the last one it finds. You could set a duplicate mapping to call a second function if the first was not found or did not function as expected.

To set function mappings

1. Open the mapping file wizard, enter a name and select a Logname for the mapping file, and advance to the Function Mappings step.

   The Function Mappings screen appears, displaying current or default mappings. The Name column shows a CEG or parsed field, the Function column the current linking function, and the Value column a text or field value.

   Note: Select a parsing file in the Provide File Details step for parsed field values to appear.
2. Click Add Function Mapping to add a new mapping entry, or select a current mapping to edit.

   The mapping entry appears in the Mapping Details pane.
3. Select a CEG field to map to from the Field drop-down menu. When you begin typing, the auto-complete feature narrows the list of available CEG fields.

4. Select a function to use for the mapping from the Function drop down menu.

   Note: The concatenate (concat) function works differently than the others, because you specify multiple target values. For more information, see Set a Concat Function Mapping.

5. Enter a target value for the mapping in the Add Value entry field and click the Add Value button next to it. You can precede the value with "$" to denote a field value rather than a specific value.

   The value appears in the Selected Fields area.

6. (Optional) You can enter multiple mappings for a single field, using the up and down arrows to set the order in which the DM file considers them.

   Note: Adding multiple values decreases performance of the mapping, so use stand-alone function mappings only if necessary.

7. (Optional) Use the shuttle control to move unneeded values to the Available Fields area to prevent them from being considered for the current mapping.

8. When you have added all the function mappings you want, click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Mapping File User folder, otherwise the step you select appears.

Set Concat Function Mapping

A Concat function mapping is a type of function mapping. Unlike other function mappings, which specify one target field or value, the concat function specifies multiple mapping targets, which it concatenates into one CEG field.

You can use the Data Mapping wizard to create concat function mappings. Because concat mappings are different from other function mappings, the procedure for creating them is somewhat different.

To set a concat function mapping

1. Open the mapping file wizard, enter a name and select a Logname for the mapping file, and advance to the Function Mappings step.

   The Function Mappings screen appears, displaying current or default mappings. The Name column shows a CEG field, the Function column the current linking function, and the Value column a text or field value.

2. Click Add Function Mapping to add a new mapping entry.

   The mapping entry appears in the Mapping Details pane.

3. Select a CEG field to map to from the Field drop-down menu.

4. Select the concat function from the Function drop-down menu.

   The Format and Value fields appear.

   Note: The value for the concat function is displayed as {...} in the Function Mappings pane. This means that there is a set of values instead of one value.

5. (Optional) Enter a specifier in the Format field to control the placement of the target fields. The format specifier, %s, indicates a field position. Anything other than %s is considered static supporting data to be included in the final table collector field. For example, to separate two target fields with a colon, enter "%s:%s" in the Format field.

6. Click Add Concat Value in the Concat Values area to add a target input/value pair.

7. Enter a value in the Add value entry field, and click Add Value.

   The value appears in the Selected Fields area.

8. Repeat steps 6 and 7 to add additional values to concatenate. You must add at least two target values.

9. When you have added all the concat mappings you want, click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Mapping File User folder, otherwise the step you choose appears.

Set Conditional Mappings

Conditional mappings link a CEG field to different possible results, allowing you to set default and conditional values for a given field. For example, you could use conditional mappings to map success or failure values, or to identify event sources by name or group.

Conditional mappings assign a default value and one or more conditional values to a given CEG field. You can set qualifications for each conditional value. If an event matches those qualifications, the appropriate conditional value is assigned to the chosen field. Otherwise the refined event field displays the default value.

If there are duplicate conditional mappings, the DM file uses the first one it finds, and considers no further mappings. To improve performance, place more common conditions first.

Note: Stand-alone conditional mapping is slower than block mapping. We recommend that you used it only when necessary.

To set conditional mappings

1. Open the mapping file wizard, enter a name and select a Logname for the mapping file, and advance to the Conditional Mappings step.

   The Conditional Mappings screen appears, displaying any current default mappings. The Field column shows the CEG or parsed field name, and the Value column shows the current default value.

   Note: Select a parsing file in the Provide File Details step for parsed field values to appear.

2. Click Add Conditional Mapping in the Conditional Field Mappings list, and select the new row.

   The Mapping Details pane appears, displaying the Field drop-down list and Value shuttle control.

3. Select the CEG field you want to map to from the Field menu. When you begin typing, the auto-complete feature narrows the list of available CEG fields.

4. Enter the default mapping you want in the Add Value entry field, and click Add Value to display it in the Selected Fields pane. You can remove unwanted values by moving them to the Available Fields pane.

5. Click Add Conditional Value in the Conditional Values list.

   A new value appears.

6. Select the New Value text to highlight it and change the name.

   The new name appears in the list, and the filters dialog appears in the details pane.

7. Construct a filter to define the conditional value. For example, you could build one or more filters to link the event_source_address field to IP addresses, identifying event sources with a geographical or other business group.

8. When you have added all the conditional mappings you want, click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Mapping File User folder, otherwise the step you select appears.

Set Block Mappings

Block mappings link a selected condition to a defined series of mappings, allowing you to create a cascade of mappings triggered by that condition. A given block mapping can use any combination of direct or function mappings. Both types of internal block mapping work exactly as they would for stand-alone mappings.

You can create as many blocks as you need for a single mapping file. Each one includes a name and a condition.

If there are duplicate mappings in a given block, the DM file will use the first one it finds, and consider no further mappings. To improve performance, you should place more common conditions first.

To set block mappings

1. Open the mapping file wizard, enter a name and select a Logname for the mapping file, and advance to the Block Mappings step.

   The Block Mappings screen appears, displaying any current block mappings.

2. Click Add Block Mapping in the Block Mappings pane.

   A new block appears in the Block Mappings list.

3. Select the New Block text.

   The Block Definition pane opens, displaying Step 1. Define a Condition

4. Enter a block name, and construct a filter to define the condition for this block. For example, you could define event_result to equal "S" which would invoke the block mappings when a success is detected for the event process.

5. Click the Step 2 bar, and enter any direct mappings you want, using the same process as the stand-alone mapping step.

6. Click the Step 3 bar, and enter any function mappings you want, using the same process as the stand-alone mapping step.

7. When you have added all the block mappings you want, click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new file appears in the Mapping File User folder, otherwise the step you choose appears.

More information:

Using Advanced Filters (see page 408)

Perform Mapping Analysis

You can use the mapping wizard to analyze a data mapping file, allowing you to test and make changes to increase the efficiency of your mapping file. The sample events are tested against the DM file and the results are then validated against the CEG.

To perform mapping analysis, click the Mapping Analysis step of the Mapping File Wizard. The wizard displays a table, showing the parsing result of the sample events you entered in the Sample Event step.

The completed DM file saves your mappings and considers event information in the order of the mapping type screens (wizard steps 4-7). If duplicate mappings exist, the last value the DM file finds is the one assigned. For example, if a DM file finds a direct mapping for a given native event value, and then a different conditional mapping for the same value, the refined event displays the conditional mapping result. More information on the design implications of the mapping order is included in the individual mapping type procedures.

When you create a new rule, it is saved as version 1.0. If you later rule edit the rule, a separate copy of the rule is stored as a new version. You can view earlier versions, and apply or copy them as needed.

Event Forwarding Rules Tasks

Event forwarding rules allow you to select CA Enterprise Log Manager events to forward to remote listeners in outside applications or systems. You can use forwarding rules to identify the events you want to forward, set when they are transmitted, and control how they are received. When an incoming event matches a forwarding rule filter, CA Enterprise Log Manager creates a copy of the event and forwards it. The event is still recorded in the event log store.

Event forwarding rules tasks are carried out from the Log Collection area of the CA Enterprise Log Manager interface. You can create, edit, and delete event forwarding rules. You can also import or export event forwarding rules.

More information:

How to Create Event Forwarding Rules (see page 478)

How to Create Event Forwarding Rules

You can use event forwarding rules to send CA Enterprise Log Manager events to outside applications. For example, you could send events to CA NSM using syslog. Event forwarding rules allow you to set criteria for the events you want to forward, and set one or more receivers.

The process of creating event forwarding rules, using the forwarding rule wizard, has the following steps:

1. Opening the forwarding rule wizard.
2. Setting a name and optional description for the rule.
3. Creating simple and advanced filters to identify events to forward.
4. Setting rule attributes including forwarding destination and CEG fields to include in the forwarded event.

More information:

Name Forwarding Rule (see page 479)
Create a Simple Event Filter (see page 479)
Create an Advanced Event Filter (see page 482)
Using Advanced Filters (see page 430)
Set Forwarding Rule Attributes (see page 483)

Open Forwarding Rule Wizard

To create a forwarding rule, or edit an existing one, open the forwarding rule wizard.

To open the forwarding rule wizard

1. Click the Administration tab, and then click the Library subtab.
2. Select the Forwarding Rules folder.

   The forwarding rule buttons appear in the details pane.
3. Click New Forwarding Rule: 

   The Forwarding Rule Wizard opens.

   When using the wizard:

   - Click Save to save the rule file without closing the wizard.
   - Click Save and Close to save the rule file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Name Forwarding Rule

You must name a forwarding rule. You can also enter description information for reference.

To name a forwarding rule

1. Open the forwarding rule wizard.
2. Enter a name for the rule. The name is required, and cannot contain the characters: / \ : * ? < >;'`,&{}[]. or |.
3. (Optional) Enter description information about the rule for reference.
4. Advance to the Filtering Step.

Create a Simple Event Filter

You can create simple filters to set search parameters for common CEG fields. For example, you could set the Ideal Model field to "Content Management" to identify all events with that value in the Ideal Model CEG field. Simple filters are used by many features, including queries, suppression and summarization rules, and event forwarding rules.

To create a simple filter

1. Select the check box for Ideal Model, or any of the Event fields you want to define, and select a value from the drop-down list, or enter the value you want in the text entry field.
2. (Optional) If you are creating a query filter, select any of the Source, Destination or Agent field check boxes, and enter the value you want in the text entry field.
3. Repeat steps 1-2 to add additional simple filters.
4. Click Save when you have added all the simple filters you want.

More information

Using Advanced Filters (see page 408)
Create an Advanced Event Filter (see page 482)

Using Advanced Filters

You can use SQL-based advanced filters to qualify any function that queries the event log store, including narrowing queries, or adding additional qualifications to simple filters. The Advanced Filters interface helps you create the appropriate filter syntax by providing a form for entering logic columns, operators and values according to your filtering requirements.

Note: This section contains a brief overview of the SQL terms used in advanced filters. To use advanced filters to their full potential you need a thorough understanding of SQL and the Common Event Grammar.

The following SQL terms join multiple filter statements:

And

Displays the event information if *all* the joined terms are true.

Or

Displays the event information if *any* of the joined terms are true.

Having

Refines the terms of the main SQL statement by adding a qualifying statement. For example, you could set an advanced filter for events from specified hosts, and add a "having" statement to return only events of a specified severity level from those hosts.

The following SQL operators are used by advanced filters to create the basic conditions:

Relational Operators

Include the event information if the column bears the appropriate relation to the value you enter. The following relational operators are available:

- Equal to
- Not Equal to
- Less than
- Greater than
- Less than or equal to
- Greater than or equal to

For example, using *Greater than* would include the event information from your chosen column if its value is greater than the value you set.

Like

Includes the event information if the column contains a pattern you enter, using % to set the pattern you want. For example, L% would return any values beginning with L, %L% would return any values with L included as neither first nor last letter.

Not like

Includes the event information if the column does not contain the pattern you specify.

In set

Includes the event information if the column contains one or more of the values in the quote-delineated set you enter. Multiple values in the set must be comma-separated.

Not in set

Includes the event information if the column does not contain one or more of the values in the quote-delineated set you enter. Multiple values in the set must be comma-separated.

Matches

Includes any event information that matches one or more of the characters that you enter, allowing you to search for key words.

Keyed

Includes any event information that is set as a key value during Report Server configuration. You can use key values to set business relevance or other organizational groups.

Not Keyed

Includes any event information that is not set as a key value during Report Server configuration. You can use key values to set business relevance or other organizational groups.

Create an Advanced Event Filter

Advanced filters are used by many features, including query creation, report scheduling, and local and global filters.

To create an advanced filter

1. Click New Event Filter.

   The first row of the event filter table becomes active, and its Logic and Operator columns are populated with the default values "And" and "Equal to" respectively.

2. (Optional) Click the Logic cell and change the logic value as needed.
3. Click the Column cell, and select the event information column you want from the drop-down menu.
4. Click the Operator cell, and select the operator you want from the drop-down menu.
5. Click the Value cell, and enter the value you want.
6. (Optional) Click the open and closed parentheses cells and enter the number of parentheses you need.
7. (Optional) Repeat steps 1 through 6 as needed to add additional filter statements.
8. Click Save when you have entered all the filter statements you want.

More information:

Create a Simple Event Filter (see page 479)
Using Advanced Filters (see page 408)

Set Forwarding Rule Attributes

Set required attributes for a forwarding rule, including forwarding exit points, CEG fields included in the forwarded event, and destination settings.

To set rule attributes

1. Open the forwarding rule wizard and advance to the Policy Attributes step.
2. Set forwarding rule actions in the Actions area:
   a. Select a syslog Facility and a syslog Severity in the appropriate drop-down lists. Any events forwarded by the rule include the syslog attributes you set.
3. Set information about CA Enterprise Log Manager event transmission in the General Information area:
   a. Select whether you want to sent the events identified by the rule before or after suppression and summarization:
   - If you select before, all incoming events are verified against the forwarding rule filters.
   - If you select after, only refined events are verified against the forwarding rule filters; Suppressed events are not forwarded, and summarized events are forwarded only as summarized, not in presummarization detail.

     Note: Choosing before has a larger effect on system performance, because the events are unrefined.

   a. Select the CEG fields you want to be displayed in the transmitted event. If you do not select a CEG field, only the raw event value is sent. If you select *any* CEG field, also select raw_event to forward the raw event.
4. Set forwarding destination information in the Destination area:
   a. Click Add Destination to create a destination row.
   b. Click the text in the Host column to add a destination hostname, or IP address. The IP address can be IPv4 or IPv6.
   c. Click the Port column cell to add the port number that the target application listens on.
   d. Click the text in the Protocol column to select TCP or UDP to set the transmission protocol you want to use.
   e. Repeat steps a-d to add more destinations as next.
5. Click Save or Save and Close.

   The new rule appears in the User subfolder of the Forwarding Rules folder.

About Forwarded syslog Events

The maximum syslog packet size (including PRI, Header, Tag and Content fields) is 1024 bytes, so the forwarded event may not be able to include all of the CEG name-value pairs the user has specified.

When necessary,CA Enterprise Log Manager truncates the message value to keep the length under 1024 bytes. If the forwarding rule specifies CEG fields to include in the generated syslog event, then the generated syslog event's Content field contains the specified CEG name-value pairs.

The name-value pairs have the format *CEG_field_name*:*field_value* from the event that matched the simple filter rule. The string “null” designates a null CEG field value. These CEG fields are in the order specified in the forwarding rule.

The CEG field order specified in the forwarding rule is significant. CA Enterprise Log Manager may truncate the value portion specified, but it will not truncate any CEG field names. If CA Enterprise Log Manager cannot fit the next full CEG field name and the colon and at least one byte of the associated value, then it terminates the syslog content field with the prior CEG name-value pair.

Edit a Forwarding Rule

You can edit a forwarding rule.

To edit a forwarding rule

1. Click the Administration tab, and then click the Library subtab.
2. Expand the Forwarding Rules folder, and click the folder which contains the file you want to edit.
3. Select the rule you want to edit, and click the Edit Forwarding Rule icon.

   The forwarding rule wizard appears, displaying your selected rule.
4. Change the rule as you like, and click Save and Close.

   The rule appears in the appropriate list as a new version of the edited rule.

Delete a Forwarding Rule

You can delete an unneeded forwarding rule.

To delete a forwarding rule

1. Click the Administration tab, and then click the Library subtab.
2. Expand the Forwarding Rules folder, and click the folder which contains the file you want to delete.
3. Select the rule you want to delete and click the Delete Forwarding Rule icon. The current version is selected by default. You can select an earlier version to delete from the Version pull-down list in the details pane.

   A confirmation dialog appears.
4. Click Yes.

   The deleted rule is removed from the appropriate list.

Import a Forwarding Rule

You can import a forwarding rule, allowing you to move rules from one environment to another. For example, import rules created in a test environment to your live environment.

To import a forwarding rule

1. Click the Administration tab, and then click the Library subtab.
2. Select the Forwarding Rules folder.

   The forwarding rule buttons appear in the details pane.
3. Click Import Forwarding Rule.

   The import file dialog appears.
4. Browse to find the rule you want to import, and click OK.

   The Forwarding Rule Wizard appears, displaying the details of the rule you selected.
5. Change the rule as you like, and click Save and Close. If the imported rule shares a name with a rule already in your management database, you are prompted to change the name.

   The imported rule appears in the Event Forwarding Rules user folder.

Export a Forwarding Rule

You can export a forwarding rule, allowing you to move rules from one environment to another. For example, export rules created in a test environment to your live environment.

To export a forwarding rule

1. Click the Administration tab, and then click the Library subtab.
2. Expand the Forwarding Rules folder, and click the folder which contains the file you want to export.
3. Select the rule you want to export, and then click Export Forwarding Rule. The current version is selected by default. You can select an earlier version to export from the Version pull-down list in the details pane.

   An export location dialog appears.
4. Enter or browse to the location where you want to store the exported rule, and click Save.

   An export successful confirmation dialog appears.
5. Click OK.

   The rule is exported.

   Note: If you examine the exported rule, the values for Facility and Severity are shown only numerically. You can use the wizard interface to determine the text descriptions associated with these values.

Chapter 14: Integrations and Connectors

This section contains the following topics:

Integration and Connector Tasks

An integration is a template for connectors. It includes all the components necessary for gathering event information from a specific type of source: a log sensor, XMP and DM files, and optional suppression rules. Integrations are supplied by CA. Users can also create their own integrations.

You can create a custom integration or modify a copy of a predefined integration. You can also create your own XMP or DM files for use in custom integrations, as well as saved integrations containing specific data access information.

After you analyze an event and create the required integration, you can create a connector, using saved configurations, and apply it to an agent, as shown in the following illustration:

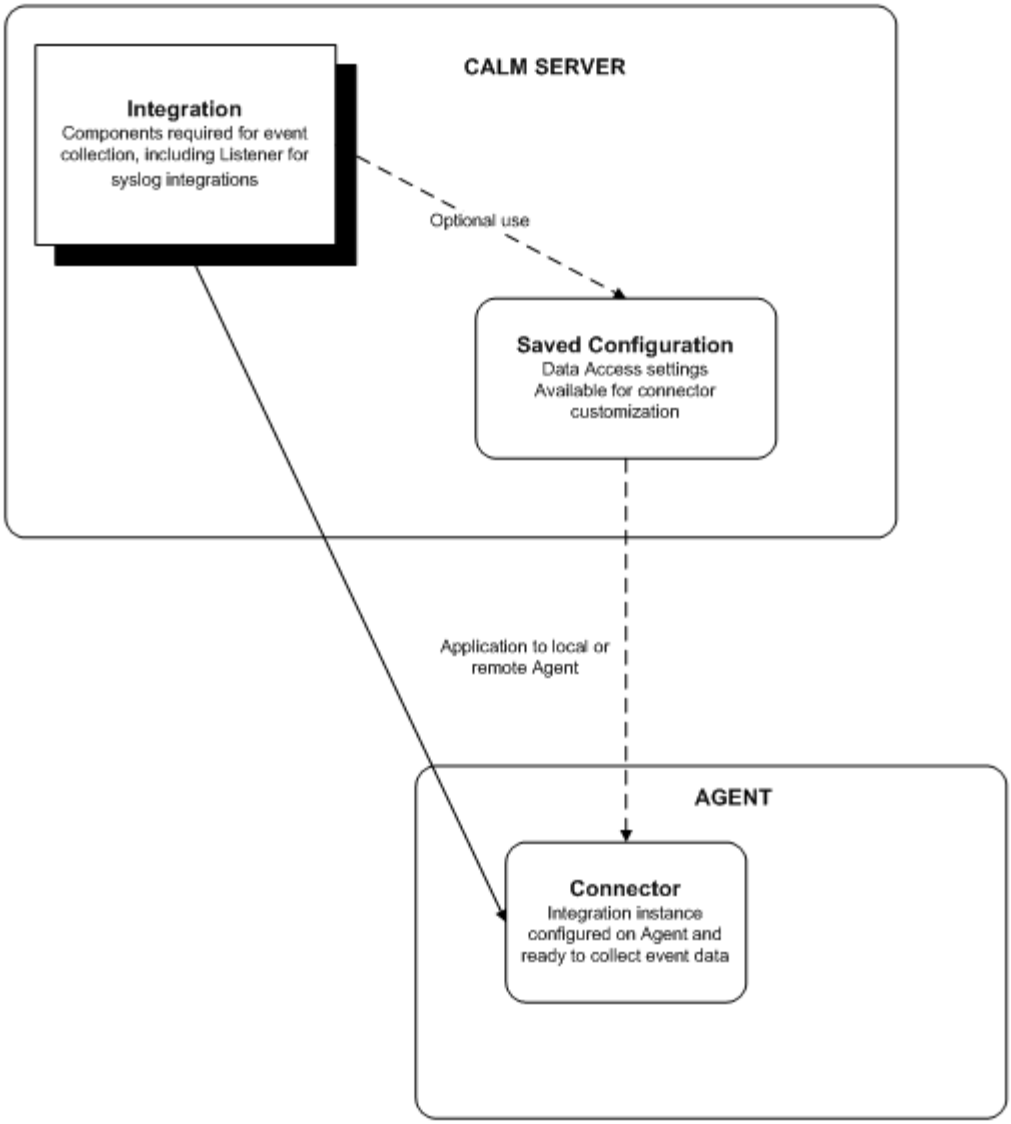

More information:

Exporting and Importing Integration Definitions (see page 502)
About Saved Configurations (see page 508)
View a Connector (see page 506)
Edit a Connector (see page 507)

How to Create an Integration

You can use the integration wizard to create or edit integrations, which serve as templates for the configured connectors that gather or receive events from your environment.

You can create integrations of several types, including WMI and ODBC integrations, which actively gather events of their specified type. You can also create syslog integrations, which receive events passively. Syslog integrations can receive events from more than one source. Therefore the process of creating a syslog integration and connector is slightly different.

To take full advantage of this advanced feature, you need a thorough understanding of the event sources in your environment and their communication types. In addition, you need a thorough understanding of regular expression syntax, the CEG, DM and XMP files, and how they parse events.

Creating an integration includes the following steps:

1. Opening the integration wizard.
2. Adding integration components.
3. Selecting suppression rules.
4. Selecting summarization rules.
5. Setting default configurations. This step does not apply for syslog integrations.

You can also create a custom user integration by copying a subscription integration.

More information:

Add Integration Components (see page 491)
Apply Suppression and Summarization Rules (see page 492)
Set Default Configurations (see page 493)

Open Integration Wizard

To create a new integration or edit an existing one, open the integration wizard.

To open the integration wizard

1. Click the Administration tab, and then the Library subtab.
2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Integrations folder.

   Integration buttons appear in the details pane.
3. Click New Integration: 

   The integration wizard appears.

   When using the wizard:

   - Click Save to save the file without closing the wizard.
   - Click Save and Close to save the file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

More information:

Add Integration Components (see page 491)

Add Integration Components

When you can create an integration, you set key integration details, such as the log sensors, XMP files, and DM files that are used to collect events.

To add integration components

1. Open the integration wizard.
2. Enter a name for the new integration.
3. Select the following required integration components from the drop-down lists:

   Sensor

   Defines the log sensor the integration uses to read events from the log source.

   Configuration Helper

   Defines the helper binary the integration uses to connect to the selected log store. Most integrations do not require a configuration helper.

   Platform

   Refers to the operating system the integration agent is able to run on, *not* the operating system of the application the integration is designed to monitor. The wizard automatically selects the operating system based on your sensor and configuration helper settings.
4. Type a description for the Integration.
5. Select the XMP and DM files you want the integration to use to refine events, using the shuttle controls.
6. If needed, type the name of the native field containing the raw event information you want the integration to parse in the 'target fields' entry field. Some event types contain their raw event information, in one particular field, requiring that the integration is targeted to that field. For example, for NT event log events, this field is named "Message".
7. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new integration appears in the user folder list, otherwise the step you choose appears.

Apply Suppression and Summarization Rules

You can apply both suppression and summarization rules to an integration to streamline event refinement. When the integration is configured as a connector, suppression and summarization rules are applied before being sent to the event log store. The suppression and summarization check is in addition to the suppression and summarization check made at the event log store.

For example, you can apply a suppression rule so that unwanted Windows events are not sent to a WMI agent. Network traffic is reduced and these events never reach the event log store.

Important! Create and use suppression rules cautiously because they prevent the logging and the appearance of certain native events entirely. We recommend testing suppression rules in a test environment before deploying them.

To apply suppression and summarization rules

1. Open the integration wizard and advance to the Suppression Rules step, or the Summarization Rules step.
2. (Optional) Type in the rules pattern entry field to search the available rules. As you type, the rules that match your entry are displayed.
3. Select the rules you want, using the shuttle control.
4. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new integration appears in the user folder list, otherwise the step you select appears.

More information:

Suppression and Summarization Rules Tasks (see page 426)
Add Integration Components (see page 491)
Set Default Configurations (see page 493)
Set File Log Configurations (see page 494)

Set Default Configurations

You can control integration data access settings using default configurations. For example, you can set the domain controller to connect to for WMI communications.

This step does not apply when creating a syslog integration, because syslog integrations inherit their configuration values from the syslog listener.

To set default configurations

1. Open the integration wizard and advance to the Default Configurations step.
2. Complete the required fields.
3. (Optional) Click the Hide button next to any default configuration to conceal it during the creation of a connector. Hidden configurations are not visible to a user creating a connector based on this integration. Therefore, you can set default configurations that cannot be changed when the integration is used to deploy a connector.
4. Click the appropriate arrow to advance to the wizard step to complete next, or click Save and Close.

   If you click Save and Close, the new integration appears in the user folder list, otherwise the step you select appears.

More information:

Add Integration Components (see page 491)

Set File Log Configurations

You can control data access settings for integrations using the file log sensor. You can use the CA-provided default settings for most event collection purposes, but you may want to alter these settings for custom integrations.

To set file log configurations

1. Open the integration wizard, select the File Log sensor type, and advance to the Default Configurations step.
2. Set or edit the anchor rate for the integration:

   UpdateAnchorRate

   Defines the threshold, in events, at which an anchor value is created. If event processing is interrupted, the agent refers to the latest anchor to begin reprocessing. Setting a lower anchor rate reduces the chance of lost events, but affects performance since the anchor value is created more often. Setting a very high anchor rate increases workload, since many events would be reprocessed in the event of a processing interruption.

   Default: 4

   Read from beginning

   Controls whether the agent will begin reading the file from the beginning if event processing is interrupted. If the check box is not selected, the agent will resume reading events using the anchor rate. If the check box is selected the sensor reads the log file from the beginning when you deploy a connector for the first time. Depending on the size of the database and the rate of event generation, the CA Enterprise Log Manager log sensor may take some time to synchronize with real-time events.
3. Set or edit the following configuration values for the targeted event source:

   File archive directory

   Defines the path where the log file is saved after rotation. The archive directory and the directory name can be the same.

   File mask

   Sets a text string used to identify the event source log file. The file mask can use wildcards. For example, to identify a log file named "messages.txt", you could enter the mask *messages** .

File rotation type

Sets the integration to correspond with the file rotation type used by the product from which it receives events. The actual rotation type is set by that product. The following settings are supported by CA Enterprise Log Manager integrations:

- NewFile - used when the integration target is rotated by a utility such as logrotate.
- FileSize - used when the integration target is based on a preset size threshold.
- FileAge - used when the integration target is based on a preset time period. The update generally takes place at or near midnight.

Directory Name

Defines the path for the event source log file.

Event Delimiter

Defines the regular expression that separates individual log entries in a multi-line log file. Each time the log sensor locates the specified delimiter, it begins reading for new events. This allows CA Enterprise Log Manager to receive multiple event entries from a single log file. For example, if each log file entry contains a unique time/date stamp, you could use the regular expression for that timestamp format as the delimiter.

4. (Optional) To add additional event source values, click Repeat: 

   An additional set of configuration value fields appear, allowing you to enter values for event collection from a different event source.

5. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new integration appears in the user folder list, otherwise the step you choose appears.

How to Create a Syslog Listener

You can use the listener wizard to create or edit syslog listeners. The listener controls how syslog events are routed to the CA Enterprise Log Manager server.

Note: You can use the subscription (pre-defined) syslog listener for almost all purposes. Certain users may want to adjust their syslog reception by using custom listeners, and these instructions are included for their purposes.

To take full advantage of this advanced feature, you need a thorough understanding of the syslog event sources in your environment.

Creating an integration includes the following steps:

1. Opening the listener wizard.
2. Adding components.
3. Selecting suppression rules.
4. Selecting summarization rules.
5. Setting default configurations.

More information:

Open Listener Wizard (see page 497)
Add Listener Components (see page 497)
Set Default Configurations (see page 498)
Apply Suppression and Summarization Rules (see page 498)
Add a syslog Time Zone (see page 500)

Open Listener Wizard

To create a new syslog listener or edit an existing one, open the listener wizard.

To open the listener wizard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the arrow beside the Event Refinement Library folder to expand it, and then select the Listeners folder.

   Integration buttons appear in the details pane.

3. Click New Listener: 

   The listener wizard appears.

   When using the wizard:

   - Click Save to save the file without closing the wizard.
   - Click Save and Close to save the file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Add Listener Components

To create a syslog listener, set details such as a name and configuration helper.

To add listener components

1. Open the listener wizard.
2. Type a name for the new listener.
3. (Optional) Select the following component from the drop-down list:

   Configuration Helper

   Defines the helper binary the integration uses to connect to the selected log store. Most integrations do not require a configuration helper.

   Note: The sensor type for a listener is always syslog.

4. (Optional) Type a description for the listener.
5. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new integration appears in the user folder list, otherwise the step you select appears.

Apply Suppression and Summarization Rules

You can apply both suppression and summarization rules to a syslog listener to streamline event refinement. When the listener is used with a connector, incoming events are verified against any applied suppression and summarization rules before being sent to CA Enterprise Log Manager.

For example, if you wanted to create a listener to receive CA Access Control events only, you could apply the CA Access Control successful file access rule. You avoid excess processing because only needed rules are used to verify incoming events.

Important! Create and use suppression rules cautiously because they prevent the logging and the appearance of certain native events entirely. We recommend testing suppression rules in a test environment before deploying them.

To apply suppression or summarization rules

1. Open the listener wizard and advance to the Suppression Rules step, or the Summarization Rules step.
2. (Optional) Type in the rules pattern entry field to search the available rules. As you type, the rules that match your entry are displayed.
3. Select the rules you want, using the shuttle control.
4. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new listener appears in the user folder list, otherwise the step you select appears.

Set Default Configurations

You can control syslog listener data access settings using default configurations. For example, you can set trusted hosts or default communication ports.

To set default configurations

1. Open the listener wizard and advance to the Default Configurations step.
2. Change or add the values you want, including:

   Event Ordering

   Helps ensure that events are sent to the event log store in the same order in which they are received. If event ordering is disabled, the order can be changed if some events are parsed and sent onward more quickly than others. Enabling event ordering can affect performance by slowing event processing and submission.

Thread Count Per Queue

Defines the number of processing threads for each protocol. Using many processing threads speeds processing if event ordering is disabled. If event ordering is enabled, the thread count has no effect. Using many threads can affect performance.

Queue Size

Sets the size of the queue, in number of events, for incoming event information. The queue is used to process and submit events. If the buffer is filled no further events can be received until processed events make room.

Ports

Sets the ports the listener uses to receive events through UDP or TCP. If you specify multiple ports, the service tries to bind to each in turn. The syslog default ports are already set. If you have routed syslog events to other ports, set your CA Enterprise Log Manager reception ports accordingly.

Important! If the agent is running as a non-root user on a UNIX system, change the syslog listener ports to port numbers above 1024. In this case, UDP port 514, the default, is not opened and no syslog events are collected.

Trusted Host

Defines trusted IP addresses for IPv4 or IPv6 - only communications from a trusted host are accepted. If you specify no trusted host, events from all available syslog event sources are accepted. Enter the exact IP address, as recorded in the event_source_address field for trusted hosts. You cannot use wildcards or subnet addresses.

Time Zones

Lets you add time zones for syslog event source computers. syslog does not typically record time. Identify the source systems by full IP address and time zone to receive and adjust events from syslog sources that are in a different time zone than the CA Enterprise Log Manager server. Do not list syslog sources in the same time zone as the server.

3. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new listener appears in the user folder list, otherwise the step you select appears.

Add a syslog Time Zone

Add a time zone for one or more syslog event source computers to receive and correctly adjust events from syslog sources that are in a different time zone than the CA Enterprise Log Manager server.

You can add a syslog time zone when creating an integration, when configuring a connector, or when creating a saved configuration.

Note: When adding a time zone to an environment where daylight savings time applies, be sure that a matching time zone entry exists on the agent host system. Without such an entry, the syslog time zone is unable to process the daylight savings time change, and events show an incorrect time stamp during the daylight savings period.

To add a syslog time zone

1. Access the syslog time zone interface in one of the following ways:
   - Open the syslog integration to which you want to add time zones, and advance to the Default Configuration step.
   - Open the syslog connector to which you want to add time zones, and advance to the Connector Configuration step.
   - Open the saved configuration to which you want to add time zones.

   The syslog time zone interface appears.
2. Click Create Folder at the top of the Time Zones area.

   A new time zone folder appears in the list area, and a time zone drop-down list appears in the right pane.
3. Select a time zone from the drop-down list.

   The zone you select appears next to the folder.
4. Click the arrow next to the folder.

   The folder expands, showing a single untitled event source computer for that time zone.
5. Select the computer icon.

   The IP address entry field appears.
6. Enter an IP Address.

   The address appears next to the computer icon as you type.
7. (Optional) To add additional event source computers, select an existing event source, and click Add Item.

   The folder closes. Open it to display a new untitled event source computer. Go to step 6.
8. (Optional) To add additional time zones, click Create Folder.

A new untitled time zone folder appears. Go to step 3.

9. When you have created all the time zone folders and event source address items you want, click Save.

More information:

Set Connector Configuration (see page 505)
Create a Saved Configuration (see page 508)

Create a New Integration Version

You can create a new version of an existing user-created (custom) integration.

To create a new integration version

1. Click the Administration tab, and then the Library subtab.
2. Expand the Event Refinement Library and Integrations folders, and navigate to the User folder that contains the integration you want.
3. Select the user integration, and click Create New Version.
4. The New Integration wizard appears, displaying the details of the integration you selected.
5. Make the changes you want, and click Save and Close.

   The new integration version appears in the list.

Delete an Integration

You can delete a custom integration. You cannot delete a subscription integration.

To delete a custom integration

1. Click the Administration tab, and then the Library subtab.
2. Expand the Event Refinement Library and Integrations folders, and select the user folder that contains the integration you want to delete.
3. Select the integration you want to delete from the list.
4. Click Delete at the top of the list.

   A confirmation dialog appears.
5. Click Yes.

   The integration is removed from the list.

Exporting and Importing Integration Definitions

You can export and import integration details for use in other management servers. This lets you transfer successful custom integrations between CA Enterprise Log Manager environments, or from a test to a live environment.

More information:

Import Integration Definitions (see page 502)
Export Integration Definitions (see page 503)

Import Integration Definitions

You can import integration definition XML files for use in the local management server.

To import integration details

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder tree appears.

2. Expand the Integrations folder, and navigate to the subfolder where you want to import an integration.

3. Click Import Integration.

   An Import File dialog opens.

4. Enter or browse for the location of the file you want to import, and click OK.

   The query files are imported to the current folder, and a confirmation dialog appears.

5. Click OK.

Export Integration Definitions

You can export integration details for use in other management servers. The export is saved as an XML file.

To export integration details

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder tree appears.

2. Expand the Integrations folder, and navigate to the subfolder containing the integration you want to export.

3. Click Export Integrations

   A download dialog appears.

4. Enter or browse for the location you want to save the XML export files, and click Save.

   The query files are saved to your chosen location, and a confirmation dialog appears.

5. Click OK.

How to Create a Connector

You can create a connector to gather events from a specific operating system or device in your environment. You use an integration or a listener as a template to create a connector, using the new connector wizard. Each new connector is applied to an agent in your environment.

You can create connectors of several types, including WMI and ODBC integrations, which actively gather events of their specified type. You can also create syslog connectors which receive events passively. Syslog connectors can receive events from more than one source, unlike the other types. Therefore the process of creating a syslog connector is slightly different.

The process of creating a connector has the following steps:

1. Opening the connector wizard.
2. Adding connector details, including selecting a listener for syslog connectors.
3. Applying suppression rules.
4. Applying summarization rules.
5. Setting connector configurations.

More information:

Open Connector Wizard (see page 504)
Add Connector Details (see page 504)
Apply Suppression and Summarization Rules (see page 505)
Set Connector Configuration (see page 505)

Open Connector Wizard

To create a new connector or edit an existing one, you must open the connector wizard.

To open the connector wizard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Expand the Agent Explorer folder, and select the agent group where you want to add or edit a connector.

   The agents belonging to the group you selected appear.

3. Select the agent where you want to add or edit a connector.

   Agent management buttons appear in the details pane.

4. Click New Connector: 

   The connector wizard appears.

   When using the wizard:

   - Click Save to save the file without closing the wizard.
   - Click Save and Close to save the file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Add Connector Details

You can add a name and description to identify your connector. You must also choose the integration you want to use as a template for the connector.

To add connector details

1. Open the connector design wizard.

   The wizard opens displaying the Platform and Platform Version for the current agent at the top of the screen.

2. Type a name for the connector.
3. Select the Listener radio button if you want to create a syslog connector, or the Integration radio button for any other type.
4. Select the integration you want to use as a template. The Integration drop-down list shows all available integrations for the current platform version and event source type.
5. (Optional) Select Bypass Platform Version Check to make integrations for *all* versions of the agent platform available in the Integration drop-down list.
6. Type a description for the connector.
7. Advance to the step you want to complete next, or click Save and Close.

   If you click Save and Close, the connector appears in the connectors list.

Apply Suppression and Summarization Rules

When creating or editing a connector, you can select suppression and summarization rules to apply to events handled by the connector. Any suppression or summarization rules you add are applied before the events are transmitted to the CA Enterprise Log Manager server.

To apply suppression or summarization rules

1. Open the connector design wizard and advance to the Apply Suppression Rules step, or the Summarization Rules step.

   A list of available suppression rules appears.
2. (Optional) Type in the rules pattern entry field to search the available rules. As you type, the rules that match your entry are displayed.
3. Select the rule or rules you want to apply, using the shuttle control.
4. Advance to the step you want to complete next, or click Save and Close.

   If you click Save and Close, the connector appears in the connectors list.

Set Connector Configuration

When creating or editing a connector, you can set individual configurations, which determine how the connector receives and transmits events. You can either set the configurations for each connector, or use saved configurations.

Saved configurations are collections of data access settings that you can reuse. You can apply saved configurations to multiple connectors.

To set connector configurations

1. Open the connector design wizard and advance to the Connector Configuration step.
2. If you selected the syslog listener/log sensor, select the integration or integrations you want the connector to use.
3. Select the saved configuration you want from the drop-down list, or alter the displayed configuration values. Connectors inherit their configuration settings from their integration, or the listener in the case of syslog connectors.
4. (Optional) Click the Help link to view the connector guide for the selected integration. The displayed guide provides details.
5. Click Save and Close.

   The connector appears in the connectors list.

View a Connector

You can open the connector list for each agent to view and edit connectors attached to that agent.

To view a connector

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.
2. Expand the Agent Explorer and agent group folders to expose the individual agents
3. Select the agent where the connector you want to view is deployed.
4. Click View Connectors: 

   The Agent Connectors list appears, displaying the connectors deployed on the selected agent.

View a Connector Guide

You can view a guide containing setup and configuration information for each type of CA Enterprise Log Manager connector. The guide contains instructions on how to configure the target product and the connector itself to receive events.

It also contains reference information such as connector log names, and what types of events the connector transmits to CA Enterprise Log Manager.

To view a connector guide

1. Click the Administration tab, and then the Library subtab.
2. Expand the Event Refinement Library, Integrations, and Subscription folders to expose the individual integrations.
3. Select the integration you want to use to create a connector.

   Integration details appear in the right pane.
4. Click the blue Help link just above the Integration Name.

   The connector guide for that integration appears in a new browser window.

Edit a Connector

You can edit an existing connector. Editing a connector creates a new version.

To edit a connector

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.
2. Expand the Agent Explorer and agent group folders to expose the individual agents
3. Select the agent where the connector you want to view is deployed.
4. Click View Connectors: 

   The Agent Connectors list appears, displaying the connectors deployed on the selected agent.
5. Click Edit beside the connector you want to edit.

   The connector wizard opens, displaying the selected connector.
6. Make the changes you want, and click Save and Close.

   The edited connector appears in the list.

About Saved Configurations

A saved configuration is a re-usable collection of settings that allows a connector to collect events from a device or log source. You can use saved configurations to allow a degree of customization without requiring the creation of an entirely new integration.

Configurations differ by integration type. For example, you can save trusted hosts for a syslog connector, or WMI server contact information for a WMI connector.

Saved configurations let you retain this grouped information and apply it to multiple connectors. Since each saved configuration is associated with a particular integration, you can only use a saved configuration on connectors that use that integration.

More information:

Create a Saved Configuration (see page 508)
Integration and Connector Tasks (see page 487)

Create a Saved Configuration

You can create a saved configuration, associating it with a specific integration.

To create a saved configuration

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Open the Event Refinement Library folder, and navigate to the integration where you want to create a saved configuration.

   The integration details appear in the details pane.

3. Click Saved Configurations: 

   The Saved Configurations List appears.

4. Click New.

   The Saved Configuration dialog appears, displaying the default configuration values for the selected integration.

5. Enter the configuration values you want, and click Save and Close.

   A confirmation message appears

6. Click OK.

   The saved configuration appears in the list.

How to Configure Connectors in Bulk

You can configure event collection sources by creating multiple connectors in bulk. You can create multiple connectors at the same time, using the same integrations, and deploy them on various agents in your environment.

The configuration process includes selecting event sources, applying suppression rules, applying summarization rules, and setting connector configurations. Before you can take advantage of this feature, create a list of identification information such as hostnames and IP addresses for the event sources you want to configure. The list must be in comma-separated value (.csv) format.

The process of configuring collection sources, using the bulk connector deployment wizard, has the following steps:

1. Opening the bulk connector deployment wizard
2. Selecting source details
3. Applying suppression rules
4. Applying summarization rules
5. Configuring connector settings
6. Selecting agents and mapping sources

More information:

Open the Configure Collection Sources Wizard (see page 509)
Select Source Details (see page 510)
Apply Suppression Rules (see page 511)
Apply Summarization Rules (see page 511)
Connector Configuration (see page 512)
Select Agents and Map Sources (see page 512)

Open the Configure Collection Sources Wizard

To create connectors on agents, you can use the bulk connector deployment wizard.

To open the bulk connector deployment wizard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the Agent Explorer folder, and then click Configure Collection Sources: 

   The configure collection sources wizard appears.

When using the wizard:

- Click Save to apply the updates without closing the wizard. A confirmation message appears.
- Click Save and Close to apply updates and close the wizard. No confirmation message appears.

Select Source Details

Select source details, and identify which integration to connect to which event sources. You must have a list of required event source details in a .csv file to complete this step.

Note: The .csv file contains the information required to create the connectors. Each column in the .csv file identifies a connector configuration field, and contains values for that field. For example, you can have an IP Address column that lists the IP addresses of hosts from which you want to receive events.

The How to Create an Integration (see page 489) section contains specific configuration fields by log sensor type

To select source details

1. Open the bulk connector deployment wizard.
2. Select the integration your sources use from the Integration drop-down list.
3. Select the integration version from the Version drop-down list.
4. Browse to the location where you have saved the collection source file you want to use. The collection source must be a .csv file.

   The first 100 rows of the collection source file you select appears in the Source File Contents area for review. The first row is set as the column header row, and remains as the header even if you adjust the sample size in step 5.
5. Use the From Row and To Row drop-down lists to narrow which part of the collection source file you want to use.

   The portion of the collection source file you select appears in the Source File Contents area for review. Column headers are not affected by changing the From Row to a value greater than 1.
6. Advance to the next step.

More information:

How to Create an Integration (see page 489)

Apply Suppression Rules

You can select which suppression rules to apply to your bulk configuration change.

To apply suppression rules

1. Open the bulk connector deployment wizard, and advance to the Apply Suppression Rules step.
2. Select which of the available rules to apply, using the shuttle control.

   Note: You can search for suppression rules using the Suppression Rules Pattern field.
3. Advance to the next step.

Apply Summarization Rules

You can select which summarization rules to apply to your bulk configuration change.

To apply summarization rules

1. Open the bulk connector deployment wizard, and advance to the Apply Summarization Rules step.
2. Select which of the available rules to apply, using the shuttle control.

   Note: You can search for suppression rules using the Summarization Rules Pattern field.
3. Advance to the next step.

Connector Configuration

You can set the connector configurations for your bulk connection creation. Each connector you create shares the configurations you set in this step, using either the sources that you collected from the .csv file in Step 1, or Saved Configurations.

To set connector configurations

1. Open the bulk connector deployment wizard, and advance to the Connector Configuration step.

   The page displays the Source Fields you set in step 1. Each header column from the source file appears as a source field. The page also displays default sensor configurations for your chosen integration in the Sensor Configuration area.

2. Set connector configuration in one of the two following ways:

   - Select a Saved Configuration from the appropriate drop-down list.
   - Set individual configurations in the Sensor Configuration area by dragging and dropping source field entries into the configuration entry fields you want. Manually set any required fields for which there are no source field values.

     For example, your source list contains a UserName column. Drag it into the User Name Sensor Configuration field, if one exists for the sensor type you are configuring.

3. (Optional) Click Repeat to add more Sensor Configuration fields as needed.
4. Advance to the next step.

Select Agents and Map Sources

You can select the agents where you want to create the connectors you have configured. Map the event sources you selected in Step 1 to the agents you want to target for connector deployment.

To select agents and map sources

1. Open the bulk connector deployment wizard, and advance to the Select Agents and Map Sources step.

   The page displays a list of sources based on the sources you uploaded in Step 1. Each source is numbered in row order, so Source 1 represents the first row you specified in your source list.

2. Search for the agents you want to target by Agent Group, Platform, or Agent Name.
3. Drag the desired source or sources to each target agent folder, and click to save that connector mapping.
4. Click Save, or Save and Close.

   Connectors based on the sources you select are configured on the selected agents.

Update Multiple Connector Configurations

You can update multiple connectors that use the same log sensor by changing one or more of the default configurations. For example, you could change the log file rotation type on multiple connectors using the log file sensor.

To update multiple connector configurations

1. Click the Administration tab, and then the Library subtab.
2. Expand the Event Refinement Library, Integrations, and Subscription folders.
3. Select an integration which uses the log sensor of the type to which you want to apply configuration changes.
4. Click Apply batch update to connectors: 

   The Update Connectors wizard appears, displaying the Select Connectors page.
5. Select the connectors to which you want to apply updates, and advance to the Default Configurations page.
6. Enter the value you want in each field you want to update, and check the box beside it.
7. Click Run.

   A confirmation message appears.

Chapter 15: Event Correlation and Incident Management

This section contains the following topics:

Chapter 16: Correlation Rule Tasks

Correlation rules can report patterns of events, helping you to identify suspicious activities or dangerous conditions in your environment. Each time events match a correlation rule's criteria, CA Enterprise Log Manager creates an incident.

You can perform the following correlation rule tasks if you have the Administrator role:

- Create, edit, or delete a correlation rule
- Create a correlation rule group
- Import or export a correlation rule.
- Apply correlation rules and associate notification destinations in your environment.

This section contains the following topics:

About Correlation Rules

You can apply predefined correlation rules, use the correlation rule wizard to create custom correlation rules for your environment, or modify existing rules. Correlation rules allow you to identify groups of events that may indicate attacks or other security risks. You must have the Administrator role to create or edit correlation rules.

When you create a correlation rule, you must select which of the three types to create. The rule template controls what event or events are considered an incident. The following templates are available:

- Simple Filter - allows you to search for a single event or state. This templates creates an incident from a single event.
- Counting Template - allows you to search for a set of identical events. You can control how many events of the same type the rule searches for. Each time the rule detects the number of events you set, it triggers an incident.
- State Transition Template - allows you to search for a related series of events. When one specific event or state occurs, followed by one or more others, the rule creates an incident. You can define the states that the rule searches for, and set the number of states.

Note: Effective correlation requires a full view of incoming events. For this reason you should consider avoiding applying suppression or summarization rules at the agent level. Any events that are suppressed or summarized at the agent are not considered for correlation and incident creation.

Event correlation can result in significant network traffic. For this reason you may wish to consider assigning a dedicated Correlation Server. See the *CA Enterprise Log Manager Implemention Guide* for more information about server roles.

If there too many incident messages for the correlation service to process, the correlation service maintains a queue of up to 10,000 messages. Any further messages are lost. CA Enterprise Log Manager generates a self-monitoring event if this occurs.

More information:

About Incident Notifications (see page 523)

Using Pre-Defined Correlation Rules

CA Enterprise Log Manager provides a large number of pre-defined correlation rules for use in your environment, organized by type or regulatory requirement. For example, in the Correlation rules folder of the Library interface, you can see a folder titled PCI, containing rules for various PCI requirements. You can also see a folder titled Identity, which contains general-purpose rules on authorization and authentication.

There are three main types of rules, any or all of which may be included in each category. This topic gives an example of choosing and applying one of each type.

Example - Select and Apply a Simple Rule

Simple correlation rules detect the presence of one state or occurrence. For example, you can apply a rule that alerts you to account creation activity outside normal office hours. Before applying any rule, you should ensure that you have created the Notification Destinations that you want for your environment.

To select and apply the Account Creation Outside Normal Office Hours rule

1. Click the Administration tab, then the Library subtab, and expand the Correlation Rules folder.
2. Expand the PCI folder, then the Requirement 8 folder, and select the Account Creation Outside Normal Office Hours rule.

   The rule details appear in the right pane.
3. Review the rule details to ensure that the rule is appropriate for your environment. In this case, the filters define the account creation action, and set the normal business hours by time and day of the week.
4. (Optional) Click Edit at the top on the pane to modify the filter settings, if required. For example, you could change the normal work hours to fit your local specifications.

   The Manage Rule wizard opens, populated with the rule details.
5. Add any notification details you want in the Manage Rule wizard. Notification details provide the message content that is delivered as specified in Notification Destinations.
6. Once you have finished preparing the rule, click Save and Close in the wizard. When you edit and save a pre-defined correlation rule, CA Enterprise Log Manager automatically creates a new version, preserving the original version.
7. Click the Services subtab, and expand the Correlation Service node.
8. Select the server you want to apply the rule on. If you have identified a Correlation Server you should select that server.
9. Click Apply in the Rule Configuration area, and select the new version of the Account Creation Outside Normal Business Hours rule, along with the Notification Destination you want associated with it.

10. Click OK to close the dialog and activate the rule.

Example - Select and Apply a Counting Rule

Counting correlation rules identify a series of identical states or occurrences. For example, you can apply a rule that alerts you to five or more failed logins by an Administrator account. Before applying any rule, you should ensure that you have created the Notification Destinations that you want for your environment.

To select and apply the 5 Failed Logins by Administrator Account rule

1. Click the Administration tab, then the Library subtab, and expand the Correlation Rules folder.
2. Expand the Threat Management folder, then the Suspicious Account and Login Activity folder, and select the 5 Failed Logins by Administrator Account rule.

   The rule details appear in the right pane.
3. Review the rule details to ensure that the rule is appropriate for your environment. In this case, the filters define an Administrator account as a username belonging to the 'Administrators' keyed list, and sets the count threshold to 5 events in 60 minutes.
4. (Optional) Click Edit at the top on the pane to modify the filter settings, if required. For example, you could change the time threshold to 3 events in 30 minutes.

   The Manage Rule wizard opens, populated with the rule details.
5. Add any notification details you want in the Manage Rule wizard. Notification details provide the message content that is delivered as specified in Notification Destinations.
6. Once you have finished preparing the rule, click Save and Close in the wizard. When you edit and save a pre-defined correlation rule, CA Enterprise Log Manager automatically creates a new version, preserving the original version.
7. Click the Services subtab, and expand the Correlation Service node.
8. Select the server you want to apply the rule on. If you have identified a Correlation Server you should select that server.
9. Click Apply in the Rule Configuration area, and select the new version of the 5 Failed Logins by Administrator Account rule, along with the Notification Destination you want associated with it.
10. Click OK to close the dialog and activate the rule.

Example - Select and Apply a State Transition Rule

State transition correlation rules identify a series of states or occurrences in turn. For example, you can apply a rule that alerts you to failed logins followed by a successful login from the same user account. Before applying any rule, you should ensure that you have created the Notification Destinations that you want for your environment.

1. Click the Administration tab, then the Library subtab, and expand the Correlation Rules folder.
2. Expand the Identity folder, then the Authentication folder, and select the Failed Logins Followed by Success rule.

   The rule details appear in the right pane.
3. Review the rule details to ensure that the rule is appropriate for your environment. In this case, the details pane displays the two states that the rule tracks. The first is five or more failed logins by the same user account or identity. The second is a successful login by that same user or identity.
4. (Optional) Click Edit at the top on the pane to modify the state settings, if required.

   The Manage Rule wizard opens, displaying the two states that make up the rule.
5. Double-click any state you want to change.

   The State Definition wizard appears, displaying the details of the state.
6. Make any state changes you want to the state you selected., and click Save and Close to return to the Manage Rule wizard. For example, the first state checks for 5 failed logins in 10 minutes. You could change the failed login threshold, or the time, or both.
7. Add any notification details you want in the Manage Rule wizard. Notification details provide the message content that is delivered as specified in Notification Destinations.
8. Once you have finished preparing the rule, click Save and Close in the wizard. When you edit and save a pre-defined correlation rule, CA Enterprise Log Manager automatically creates a new version, preserving the original version.
9. Click the Services subtab, and expand the Correlation Service node.
10. Select the server you want to apply the rule on. If you have identified a Correlation Server you should select that server.
11. Click Apply in the Rule Configuration area, and select the new version of the Failed Logins Followed by Success rule, along with the Notification Destination you want associated with it.
12. Click OK to close the dialog and activate the rule.

More information:

About Incident Notifications (see page 523)
About Correlation Rules (see page 518)
Set Notification Defaults (see page 525)

Using Keyed Lists with Correlation Rules

All correlation rules are built from one or more filters. Some predefined rule filters are designed to select all values from a given table where a certain attribute field contains a value used as criteria for compiling the list of key values.

You can use keyed lists in the creation or application or correlation rules to provide predefined or custom values to rule filters. You can also update the keyed value lists manually or automatically to keep the lists current. You can use keyed lists with correlation rules just as you do with reports.

For more information on using keyed lists, see the Queries and Reports chapter of this guide.

More information:

Preparing to Use Reports with Keyed Lists (see page 275)
Approaches to Maintaining Keyed Lists (see page 279)
Creating Keyed Values for Predefined Reports (see page 285)

Example: Creating a CSV File for Testing

This example illustrates the creation of a CSV file for correlation rule testing. It is intended to test a rule that searches for 5 failed logins followed by a successful login from a single user.

To create a CSV file to test a failed login followed by success rule

1. Log in to CA Enterprise Log Manager as an Administrator, and click the Queries and Reports tab.
2. Search for the "Five Failed Logins by in Last 1 Hour by Performer" query.
3. Run it and view the results. If there are results, proceed to the next step. If not, create a dummy user, log out, and create failed logins using the new dummy user.
4. Export the query to CSV, and open the CSV file in Excel.
5. Add other user details as needed. For example, add information to reflect the sucessful login.
6. Save the CSV file when it has all the event information you need.
7. Open the rule you want to test in the Library Explorer, and click the Rule Test tab in the details pane.
8. Load the CSV file, and confirm that the proper incidents are created.

About Incident Notifications

You can set notifications, which pass information about an incident, to be triggered automatically when an incident is created, or launch them manually after viewing the incident. In either case you must first define the notification destinations you want to use in your environment.

You must create notifications in two parts:

1. Notification Destination, which can contain any combination of the available destination types. For example, a destination might contain email addresses, SNMP server credentials, and an IT PAM process name. Destinations can be assigned to multiple rules.
2. Notification Details, which are added to individual rules, and contain the information delivered by the notification; email subjects and text, SNMP data, IT PAM process parameters, for example.

Automatic notifications require a correlation rule with notification details, and an associated notification destination. If both components are present, each time the rule creates an incident, an automatic notification is sent to the specified destination or destinations. The combination of destinations and details allows you to set up modular notification. For example, you could route the same notification information to different regional service desks or IT personnel.

You can also assign destinations from existing incidents. When you open an incident and assign a Notification Destination, the notification details specified in the rule are sent immediately. The rule must include notifications in order to send manual notifications.

More information:

Set Notification Defaults (see page 525)
About Correlation Rules (see page 518)

How to Design and Apply Incident Notifications

You can set up notifications for your correlation rules. Notifications allow you to pass key information on detected incidents to the staff you specify, or create CA IT PAM service desk tickets automatically.

Use the following process to design and set up notifications in your environment:

1. Plan and create notification destinations.
2. Select the pre-defined correlation rules, or create custom rules you want to use in your environment.
3. Add notification details to the rules for which you want to set notifications.
4. Apply correlation rules to CA Enterprise Log Manager servers, and assign notification destinations.

More information:

About Correlation Rules (see page 518)
Set Notification Defaults (see page 525)
Correlation Service Considerations (see page 129)

Set Notification Defaults

You can set notification details in a rule, which specify notification content but not destinations. For example, you can set email subject line and content text, but not the delivery addresses, which are controlled by notification destinations. This system allows you to set up standard content (using details), which can be delivered to various recipients (using destinations).

You can include any combination of the available notification types in a single rule's notification details.

To set notification details

1. Open the correlation rule wizard, enter the required rule definitions, and advance to the Notification Details step.
2. Select the email tab and use the following steps to add email notification information:
   a. Enter a subject line for the notification email.
   b. (Optional) When entering text in either field of the email tab, you can use the Data Fields drop down list and the Add button to insert data field variables. For example you could choose agent_address and click Add.

      "%agent_address%" appears in the text field. When a rule generates an email, the value of the agent_address field is displayed in place of the variable.
   c. Enter message body for the notification email.

      Note: The message body is constructed in HTML, so all text you enter appears on one line. To create a break after a line, enter <BR/> at the end of the line of text.
3. Select the Process tab and use the following steps to add CA IT PAM process parameters:
   a. Enter the name of an IT PAM process to which you want to pass incident information, such as:
      /CA_ELM/EventAlertOutput
   b. Click Add Parameter to specify a parameter and its value.

      The Add Process Parameter dialog appears.
   c. Type a parameter name in the Name field, 'Severity' for example.
   d. Define a value by typing in the value area, or selecting a CEG field from the drop-down list and clicking Add data field. Event information from the CEG field you specify is passed to the named parameter. Continuing the example from the previous step, you could select 'event_severity' to present the value of the event_severity field as the IT PAM Severity parameter.
   e. Repeat Steps a-c to add additional parameters and values as needed.

f. When you have added all the CEG fields you want for the current parameter, Click OK.

Note: You can type, and add multiple CEG fields as needed to define a parameter. For example, if you want to define the Description parameter for a notification used with an account guessing rule, you could enter:

This incident reports four failed logins by by %dest_identity_unique_name% on %dest_hostname% occurred within 10 minutes.

The *%value%* structure is the result of selecting a CEG field and using the Add data field button as described in step b.

4. Select the SNMP tab and use the following steps to add SNMP trap settings:
   a. Adjust the Custom Trap ID as required by your SNMP transmission target.
   b. Enter the name of a CEG field that you want to send in the entry area. Typing in the field narrows available choices in the drop-down list as you type.
   c. Click Add.
   d. The CEG field name appears in the selected fields area. Any event information in that field is sent for rules using this notification template. You must specify at least one CEG field.
   e. Repeat steps b-c to send additional CEG fields.
5. Click the appropriate arrow to advance to the wizard step you want to complete next, or click Save and Close.

   If you click Save and Close, the new rule appears in the appropriate folder, otherwise the step you choose appears.

More information:

About Correlation Rules (see page 518)
About Incident Notifications (see page 523)

Chapter 17: Incident Management Tasks

A CA Enterprise Log Manager incident is composed of one or more events, as identified and linked by a correlation rule. Each time a correlation rule detects an event or events that satisfy its criteria, CA Enterprise Log Manager creates an incident.

You can perform the following incident management tasks if you have the Administrator role:

- View the details of incidents created by the correlation rules in your environment.
- Filter the incident list, or set result conditions to locate particular incidents or types of incidents, or narrow your incident view.
- Apply notification destinations to existing incidents, controlling responses such as email notification.
- Export incident information.
- Schedule action alerts based on an incident.
- Merge existing incidents into one new incident.

Note: For detailed information on Incident Management tasks, see the *CA Enterprise Log Manager Online Help.*

View Incident Details

You can view details of incidents in your environment, including status, priority and history information. You can view only those incidents that are routed to the correlation server that you are logged in to. You can control how CA Enterprise Log Manager servers route events by configuring the Correlation Service.

To view incident details

1. Click the Incidents tab, select the incident you want to investigate, and double-click in the incident row.

   The Details dialog appears, displaying the basic details of the incident, including name, date and severity.

2. Change the Priority or Status settings using the appropriate drop-down menus.
3. (Optional) Click the History tab to view information such as number and and time of events added to the incident, or automatic notifications triggered.
4. Click OK or Apply.

More information:

Correlation Service Considerations (see page 129)

Chapter 18: Agents

This section contains the following topics:

Plan Agent Installation

When planning agent installation, the planner needs to determine how many agents are needed and where to install them. The individual who installs the agents may do this planning or it may be performed by a network administrator or systems architect.

To plan agent installations

1. Create an electronic version of an agent installation planning table that is suitable for recording the information you need. Consider using the following sample table and column headings.

Event source platform	Host name or IP address where event source is running	One of the following: - Agentless-Direct - Agentless-Collection Point - Agent on End-point	Host name or IP address where agent is to be installed

2. Identify each event source to target for log collection and record the location and platform in your agent planning table.

3. Consider the following the costs and benefits of each solution type.

	Benefit	Cost or Limitation
Agentless - Direct from CA Enterprise Log Manager - no installed agent	No agent installation is required	Can accommodate collection of only those event sources compatible with the soft appliance platform. Costs of Agentless-Collection point also apply.
Agentless - agent on collection point	No agent needs to be installed on the host where the event source is running. Consolidating collection to a common point reduces the number of agents that need to be installed compared with agent-based collection.	Suppression rules can be applied only at the CA Enterprise Log Manager server. This lacks the advantage of reducing network traffic. The communication of events between the source and the CA Enterprise Log Manager server is not encrypted. The event source must be able to be accessed remotely.
Agent-based - agent on end-point	You can apply suppression rules at the source rather than at the CA Enterprise Log Manager server. This reduces the network traffic between the point of collection and the CA Enterprise Log Manager server. The communication of logs between the source and the CA Enterprise Log Manager server is encrypted. Can accommodate the highest event volume of the three solutions.	An agent must be installed where the event source is running.

4. Record your preferred solution for each event source.
5. Sort your agent planning table by column 3 and then by column 2. This displays all agent-based event sources that are running on the same host in blocks.
6. For event sources you mapped to agent-based, find the first occurrence in a block of the same name, and copy that data from column 2 to column 4. This results in just one entry for each host on which one or more event sources are running. (You never need more than one agent on a given host, regardless of the number of event sources.)

7. For event sources you mapped to agentless-collection point, plan where to install agents as follows:
    a. Identify the number of collection points needed to accommodate the event sources identified as candidates for remote agents.
    b. Plan groupings of event sources identified as candidates for remote agents by common network location area.
    c. For each group of event sources, identify the server to use as the collection point. This can be a dedicated server. Record your decision in column 4.
8. If you have recorded event sources for which there is no mapped solution and you are using a legacy CA adapter, see the *CA Enterprise Log Manager Implementation Guide* for further details.
9. Hand off the data you recorded in the fourth column of your agent planning table to the user who is to install the agents.

Planning Agent Configuration

The EiamAdmin user installs agents based on the determination of the best collection method. Methods evaluated include the following:

- Agentless log collection directly from the CA Enterprise Log Manager server, sometimes referred to as Direct collection.
- Agentless log collection from a collection point
- Agent-based log collection from the host where the event source is running

The analysis that precedes installation can uncover some of the information needed by the Administrator who configures the agents and connectors.

The first step in configuring agents is to get the agent planning spreadsheet from the EiamAdmin or the alternative used to document where agents are installed. After configuring the first Administrator, the EiamAdmin user provides the Administrator with the annotated agent installation planning worksheet. The first Administrator, in turn, plans the needed connectors for each agent before beginning configuration.

The Administrator configures each agent installed by the EiamAdmin. In addition, the Administrator configures a connector for each event source, regardless of the collection method (agentless-direct, agentless-collection point, or agent-based). The Administrator configures connectors on each agent while logged on to the CA Enterprise Log Manager server which is to receive events from collected by that agent.

Note: The fewer connectors configured on an agent, the better the performance.

An exception to this process is when agent installation is performed silently. In this case, it is the installer who configures the connectors. The connectors configured on an agent enable the agent to collect raw events from specific event sources. The connectors translate raw events to refined events and pass the refined events to CA Enterprise Log Manager.

Creating agent groups is optional. If no custom agent groups are created, agents are assigned to the Default Agent Group. Administrators create agent groups for the following reasons:

- To enable reporting of events collected by agents in the same agent group.
- To enable the assignment of a different administrative user to different groups of agents. (User access can be limited to specified agent groups with access policies.)

Collected event logs are sent to a CA Enterprise Log Manager server for processing and initial storage. Administrators must configure the server that is to receive logs for each agent or agent group. Assigning a server to an agent group is a quick way to assign the server to all agents in the agent group.

Planning Direct Log Collection

The CA Enterprise Log Manager is installed with a default agent that can be used for direct log collection. It is called direct collection because the use of the default agent requires no agent installation. The default agent can collect events from almost any event source, with the following limitations.

- The log sensor must be able to run on the soft appliance; some log sensors, such as the WMI log sensor, are tied to a specific platform.
- The event source must be able to be accessed remotely.

You configure the default agent just like you configure a separately installed agent. Direct log collection by the default agent is ideal for a very small system.

Event Sources for Direct Log Collection

CA Enterprise Log Manager provides log sensors that can run on the CA Enterprise Log Manager server to facilitate agentless direct log collection. As of the release of this document, the following are supported:

- syslog
- WinRM
- ODBC
- TIBCO

To determine the integrations supported by the default agent

1. Select a CA Enterprise Log Manager server from the Agent Explorer under the Administration tab, Log Collection subtab
2. Click Create New Connector.

The Integration drop-down includes the integrations from which you can create a connector for deployment on the default agent. Each integration, on which connectors are based, is designed to retrieve events from a specific event source.

For a complete list of supported log sensors and integrations, see the CA Enterprise Log Manager product page at Customer Support.

Note: A log sensor is an integration component designed to read from a specific log type such as a database, syslog, file, or SNMP.

Planning Agentless Log Collection

Agentless log collection can be implemented by installing an agent on a collection server that collects events from multiple remote event sources.

Consider the following when planning the configuration of agentless log collection from a collection server:

- The fewer connectors deployed to an agent, the better the performance.
- The maximum number of connectors you should configure on a given agent depends on whether the agent is installed on a dedicated server, how powerful this server is, and the kinds of event sources you are targeting. As a rule-of-thumb, you should probably configure no more than forty or fifty connectors on a single agent.
- There is no performance advantage to grouping connectors of the same type configured on different agents to the same collection server. By extension, there is no performance advantage to directing events from same types of event sources to a given CA Enterprise Log Manager server in a federation.

Planning Agent-Based Log Collection

After the installer installs an agent on a server with local event sources, Administrators configure a connector on that agent for each locally running event source.

If there are many target servers with the same types of event sources, consider grouping those target servers into an agent group and performing configuration at the agent group level.

Guaranteed delivery can be a problem for direct collection of syslogs. To counter this, configure a syslog listener on an agent installed with the syslog event source.

Selecting the Level to Configure

The Subscription, Apply Suppression Rules, and Status and Command options can be selected from various levels. For example, Subscription configuration can be initiated from the following levels:

- Agent Explorer
- Default Agent Group or a user-defined agent group
- Agent

To configure an option so that it applies to all agents in all groups, select Agent Explorer and then click the button for the action you want to perform.

To configure an option so that it applies to all agents in a selected group, select the group name and then click the button for the action you want to perform.

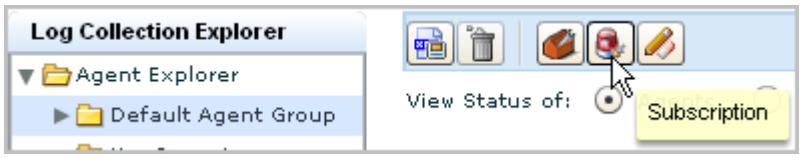

To configure an option so that it applies to just one agent, select the agent and then click the button for the action you want to perform.

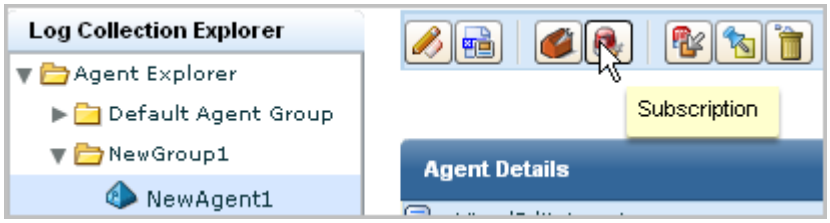

Agent Management Tasks

The Agent Explorer lets you view and manage the event collection agents in your environment. You can use the Agent Explorer interface to perform management tasks in the following areas:

- Agent configuration—Lets you rename agents, and configure the groups to which they belong, and their associated groups.
- Agent groups—Lets you group agents, by geography, business importance, or event source type, for example. You can also assign grouped agents to different CA Enterprise Log Manager servers as needed.
- Subscriptions—Lets you view and apply available updates to agents.
- Agent command and status—Lets you view the current status of agents and stop and start them as needed.

More information:

Download Agent Binaries (see page 538)
Update the Agent Authentication Key (see page 537)
How to Create an Agent Group (see page 544)
How to Apply Subscription Updates (see page 552)
How to Configure Agent Management (see page 546)

Update the Agent Authentication Key

You can view and update the key used by agents to register with the CA Enterprise Log Manager server. Changing this key regularly helps prevent unauthorized agents from being installed in your environment. By default, the key is the same for all CA Enterprise Log Manager servers across application instances. However, you can set this key to be unique by application instance.

The agent installer must enter this agent key as the Authorization Code in the agent installation wizard.

To update the agent authentication key

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the Agent Explorer folder.

   Agent management buttons appear in the details pane.

3. Click Agent Authentication Key: 

   The Agent Authentication Key pane appears.

4. Type a new key in the Enter Key and Confirm Key fields, and click Save.

   A success confirmation message appears.

Download Agent Binaries

You can download agent binaries and install them on your local computer without using other installation media.

For more information about agent installation, see the *CA Enterprise Log Manager Agent Installation Guide*.

To download agent binaries

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the Agent Explorer folder.

   Agent management buttons appear in the details pane.

3. Click Download Agent Binaries: 

   The Agent Binaries List appears, showing the available agents and their current versions.

4. Click the agent you want to download.

   The download dialog appears.

5. Select the location where you want to save the agent binary file, and click Save.

   The file is saved in the chosen location, and a confirmation message appears.

Configure an Agent

You can configure an installed and registered agent after you access it in the Agent Explorer.

To configure an agent

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Click the Agent Explorer folder.

   The folder expands, displaying Agent Group folders.

3. Select the agent you want to configure, and click Edit at the top of the pane.

   Agent details appear in the details pane.

4. Make any changes you want, including:

   User Name

   Defines the user name under which the agent runs.

   Port

   Sets the port the agent uses to communicate with CA Enterprise Log Manager.

   Agent Group

   Defines the group to which the agent belongs.

   Max Number of Files

   Sets the maximum number of files that can be created in the event reception file queue. The Max Number limit is 1000 files.

   Max Size per File

   Sets the maximum size, in MB, for each file in the event reception file queue. When a file reaches the maximum size, CA Enterprise Log Manager creates a new file. The Max Size limit is 2048 MB.

   Event Sending Mode

   Defines which of the following transmission styles the agent uses:

   - Failover - The agent sends events to the first server in the Log Manager Servers list. If communication is lost, it then attempts to communicate with each server in the order listed, until communication is restored.
   - Round Robin - The agent sends events to each server in the Log Manager Servers list in turn, attempting to contact the next server in the list after one hour. This period is not configurable.

Enable Events Encryption

Sets the agent to use AES128 to encrypt the events it transmits. Enabling event encryption will affect performance.

Enable Dispatch Scheduling

Sets the agent to send events only in a certain time span. Selecting the Enable Dispatch Sending check box displays Start Time and End Time fields. Enter the GMT time values you want in 24-hour clock format, with the following qualifications:

- There must be at least one hour between the start and end time values.
- If the start time value is higher than the end time value, the end time is set the day after the start time. For example, if you set the start time at 23, and the end time at 6, the transmission time span runs from 11 pm GMT through 6 am GMT of the following day.

Log Manager Servers

Controls the CA Enterprise Log Manager servers to which the agent routes events, and the order in which they are contacted. You can use the shuttle control to select available servers and the arrow buttons to the right of the selected servers to set communication priority.

Note: Update your CA Enterprise Log Manager servers before you update agents. CA Enterprise Log Manager servers support agents at or below their current version number. To help ensure proper storage of collected events when you configure or update agents, verify that the agent sends events only to CA Enterprise Log Manager servers whose level is the same as the agent or higher.

5. Click Save.

More information

How to Create an Agent Group (see page 544)

Tampered Configuration File Handling

Agents use a configuration file stored in memory when they are running. If someone tampers with a configuration file while the agent is running, the agent does not use the tampered file. When an agent receives a new configuration from the CA Enterprise Log Manager server, the agent replaces the disk file with the received file before restarting. In this way, a tampered file is automatically replaced with the correct file.

If someone restarts the agent from external source after tampering with the file, the Agent detects that the file is tampered and shuts down. The agent does not accept any configuration data, including from the CA Enterprise Log Manager server list from the tampered file.

The Agent Explorer shows the agent as not responding. Use the CA Enterprise Log Manager server status and command tools to reset the agent configuration. The agent resumes working properly after this action.

View Agent Dashboard

You can view the agent dashboard to view the status of agents in your environment. The dashboard also displays details such as the current FIPS mode (FIPS or non-FIPS), and usage details. These include events per second load, CPU percentage use, and most recent update date and time.

To view the agent dashboard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Select the Agent Explorer folder.

   Agent management buttons appear in the details pane.

3. Click Agent Status Monitor and Dashboard: 

   The agent search panel appears, displaying status for all available agents in a details chart. For example:

   Total: 10 Running: 8 Pending: 1 Stopped: 1 Not Responding: 0

4. (Optional) Select agent search criteria to narrow the list of displayed agents. You can select any one or more of the following criteria:
   - Agent Group—returns only agents assigned to the selected group
   - Platform—returns only agents running on the selected platform
   - Status—returns only agents with the Status you select, Running, for example.
   - Agent name pattern—returns only agents containing the specified pattern.
5. Click Show Status.

   A list of agents meeting your search criteria appears, displaying information including:
   - Local connector name and version
   - Current CA Enterprise Log Manager server
   - Agent FIPS Mode (FIPS or non-FIPS)
   - Last recorded event per second load handled by the agent
   - Last recorded CPU usage value
   - Last recorded memory usage value
   - Most recent configuration update
   - Configuration update status

View and Control Agent or Connector Status

You can monitor the status of agents or connectors in your environment, restart agents, and start, stop, and restart connectors as needed.

You can view agents or connectors from different levels of the Agent Explorer folder hierarchy. Each level narrows the available view accordingly:

- From the Agent Explorer folder, you can view all agents or connectors assigned to the current CA Enterprise Log Manager server.
- From a specific agent group folder, you can view agents and connectors assigned to that agent group.
- From an individual agent, you can view only that agent and any connectors assigned to it.

You can determine the FIPS mode (FIPS or non-FIPS) for an agent from all three levels.

To view agent or connector status

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.

2. Select the Agent Explorer folder.

   Agent management buttons appear in the details pane.

3. Click Status and Command: 

   The status panel appears.

4. Select Agents or Connectors.

   The agent or connector search panel appears.

5. (Optional) Select agent or connector update search criteria. If you enter no search terms, all available updates appear. You can select any one or more of the following criteria to narrow your search:

   - Agent Group—returns only agents and connectors assigned to the selected group.
   - Platform—returns only agents and connectors running on the selected operating system.
   - Agent name pattern—returns only agents and connectors containing the specified pattern.
   - (Connectors only) Integration—returns only connectors using the selected integration.

6. Click Show Status.

   A details chart appears, displaying status for agents or connectors that match your search. For example:

   Total: 10 Running: 8 Pending: 1 Stopped: 1 Not Responding: 0

   Note: If you update the configuration of an agent, CA Enterprise Log Manager requires a maximum of five minutes to synchronize the updated status of this agent with the other agents in a federation.

7. (Optional) Click the status display to view details in the Status pane at the bottom of the chart.

   Note: You can click the On Demand button for an agent or connector to refresh the status display.

8. (Optional) If you are viewing connectors, select any connector and click Restart, Start, or Stop. If you are viewing agents, select any agent and click Restart.

How to Create an Agent Group

You can create an agent group to organize your agents by location, operating system, or any other convenient category. The process of creating an agent group using the Agent Group wizard has the following steps:

1. Opening the agent group wizard.
2. Entering group details.
3. Adding agents.

More information

Open Agent Group Wizard (see page 544)
Add Agent Group Details (see page 545)

Open Agent Group Wizard

To create an agent group or edit an existing one, open the agent group wizard.

To open the agent group wizard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.
2. Click the Agent Explorer folder.

   Agent management buttons appear in the details pane.
3. Click New Agent Group: 

   The agent group wizard appears.

   When using the wizard:

   - Click Save to save the file without closing the wizard.
   - Click Save and Close to save the file and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Add Agent Group Details

You can add identifying details to your agent group.

To add agent group details

1. Open the Agent Group wizard.
2. Enter a name for the group and an optional description for reference.
3. Select the Event Sending Mode.
4. (Optional) Enter the first few characters of the server you want to add in the Name Pattern field.

   Servers matching your search appear in the Available area.
5. Select the servers you want to add using the shuttle control, and arrange them in the order you want them to appear in the agent group display using the up and down arrows.
6. Advance to the next step, or click Save and Close.

   If you click Save and Close the group is created, otherwise the step you choose appears.

Add Agents to an Agent Group

You can add agents to a group for administrative purposes. For example, you might create groups by geographical region or operating system.

Note: The properties of an agent group apply to all the agents in that agent group.

To add agents to a group

1. Open the Agent Group wizard and advance to the Agents step.
2. (Optional) Select agent search criteria. If you enter no search terms, all agents appear. You can select any one or more of the following criteria to narrow your search:
   - Agent Group—returns only agents assigned to the selected group
   - Platform—returns only agents running on the selected platform
   - Agent name pattern—returns only agents containing the specified pattern
3. Click Search.

   Agents matching your search appear in the Available Agents area.
4. Select the agents you want to add using the shuttle control, and arrange them in the order you want them to appear in the agent group display using the up and down arrows.

   Note: You cannot move an agent into an agent group that does not have configured CA Enterprise Log Manager severs.
5. Click Save and Close.

   The agent group appears in the list.

Note: If you remove an agent from a user-created agent group, the agent is moved into the Default Agent Group, and the agent inherits the properties of the Default Agent Group. However, you cannot delete an agent from the Default Agent Group.

How to Configure Agent Management

You can configure your agents or agent groups to report to different CA Enterprise Log Manager servers in your federated environment. This lets you configure groups or agents to send event information to chosen CA Enterprise Log Manager servers.

The process of configuring agent management using the assign Log Manager Servers wizard has the following steps:

1. Opening the assign Log Manager Servers wizard.
2. Selecting agent or agent group targets to assign.
3. Selecting the CA Enterprise Log Manager servers to assign agents or groups.

More information:

Open Assign Log Manager Servers Wizard (see page 547)
Select Target Agents (see page 548)
Select Log Managers (see page 548)

Open Assign Log Manager Servers Wizard

To configure agent or agent group assignments, open the assign Log Manager Servers wizard.

To open the assign Log Manager Servers wizard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears, displaying agent management buttons in the details pane.

2. Click Log Manager Servers: 

   The assign Log Manager Servers wizard appears.

   When using the wizard:

   - Click Save to save without closing the wizard.
   - Click Save and Close to save the assignment and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

More information:

How to Configure Agent Management (see page 546)
Select Target Agents (see page 548)
Select Log Managers (see page 548)

Select Target Agents

To assign agents to a server for event reception and archiving purposes, you must choose which agent or group to assign to a specific CA Enterprise Log Manager server.

To select target agents

1. Open the assign Log Manager servers wizard.
2. Choose whether you want to assign agents by group, or individually.
3. If you choose Groups, use the shuttle control to select the groups you want to assign. You can use the search entry field to locate groups by typing a partial name. The available groups are filtered as you type.
4. If you choose Agents, use the shuttle control to select the individual agents you want to assign. You can use the Agent Group and Platform drop-down lists, and the search entry field to locate agents.
5. Advance to the Select Log Manager step.

More information:

How to Configure Agent Management (see page 546)
Open Assign Log Manager Servers Wizard (see page 547)
Select Log Managers (see page 548)

Select Log Managers

You must choose which CA Enterprise Log Manager server you want to assign agents or agent groups to.

To select log manager servers

1. Open the assign Log Manager servers wizard, select target agents, and advance to the select log managers step.
2. Use the shuttle control to select the servers you want to assign agents or agent groups to. You can use the Name Pattern entry field to search the available list.
3. Click Save and Close.

   The agents or grouped agents are assigned to the servers you selected.

More information:

How to Configure Agent Management (see page 546)
Open Assign Log Manager Servers Wizard (see page 547)
Select Target Agents (see page 548)

How to Protect Agents from Impact of Server IP Address Changes

When you install an agent, you assign a primary CA Enterprise Log Manager server for the agent to contact first with collected events. When you configure an agent, you add other CA Enterprise Log Manager servers in an ordered list. When an agent that is ready to send collected logs to the primary server cannot reach it, the agent contacts each secondary server in the list until it finds an available one. Configuration of an ordered list of secondary servers guarantees log delivery from agent to server. An agent can send events to only one CA Enterprise Log Manager at a time; that is, there is no event duplication.

When the servers that you select to manage an agent are assigned new IP addresses, the ability of the agent to forward collected events to a server on its list can be impacted. By taking precautionary steps, you can help ensure that servers remain highly available to the agents that use them, even when these servers undergo manual or dynamic reassignment of their IP addresses.

The IP address of an installed CA Enterprise Log Manager could change in the following cases:

- Dynamic reassignment by DHCP

  The CA Enterprise Log Manager server in a single-server system is configured to have DHCP assign its IP address. At some point after this server is selected to manage agents, DHCP assigns it a new IP address. This can occur when the CA Enterprise Log Manager is offline long enough for the IP address lease to elapse. No user notification is required when IP addresses are dynamically changed.

- Manual reassignment

  The CA Enterprise Log Manager servers are configured with static IP addresses. Due to a site process where IP addresses are reassigned as part of deploying a new subnet, new IP addresses are manually assigned to the CA Enterprise Log Manager servers.

Take appropriate action to ensure high availability of servers to agents, when those servers' IP addresses are subject to change.

More information:

Ensure Availability of Servers with Dynamic IP Addresses (see page 550)
Ensure Availability of Servers During Reassignment of Static IP Addresses (see page 550)

Ensure Availability of Servers with Dynamic IP Addresses

If you select DHCP when installing a single-server CA Enterprise Log Manager, specify the hostname (not the IP address) of that CA Enterprise Log Manager when installing each agent. This will ensure that any DHCP-reassignment of the IP address of the CA Enterprise Log Manager server will not impact the agents that use it.

If you specify the IP address of the CA Enterprise Log Manager server when installing the agents and that dynamic IP address changes, you will need to reinstall the agents to restore the availability of the single-system CA Enterprise Log Manager server. To avoid this potential problem, we recommend you install an additional CA Enterprise Log Manager server and add it as a secondary server for all agents. This will ensure high availability of servers.

Ensure Availability of Servers During Reassignment of Static IP Addresses

If you install the CA Enterprise Log Manager servers with static IP addresses and later plan to renumber them, use the following workflow to ensure continuous high availability to the agents with these servers on their ordered lists. It is not necessary to restart the agent after any step since the agent refreshes its configuration information every 5 minutes by default.

Important! If the agent is configured with only one CA Enterprise Log Manager server in a multi-server system, be sure to add a second server to the ordered list before assigning new IP addresses. If this is not done, you may need to reinstall and reconfigure the agent after the server IP address assignment to reinstate the server's availability.

To ensure that agents can reach a CA Enterprise Log Manager server on their ordered list during a reassignment of static IP addresses to servers

1. If planning IP address reassignment to a CA Enterprise Log Manager server where that server is the only CA Enterprise Log Manager, install a temporary CA Enterprise Log Manager that points to the CA EEM on the original CA Enterprise Log Manager.

2. If one or more additional CA Enterprise Log Manager servers have not be configured for any agents in a multi-server system, assign at least one additional server to the ordered list.
   a. Select Agent Explorer and click Log Manager Servers.

      The Assign Log Manager Server wizard appears with the Select Targets step selected.

   b. Select Agents or Groups, depending on how you make such assignments.
   c. Select the affected agents or groups as targets from the Available list and move them to the Selected list.
   d. Click the Select Log Manager Servers step.
   e. Select a CA Enterprise Log Manager from the Available list and move it to the Selected list.
   f. Click Save and Close.
3. Remove half of the ordered list of CA Enterprise Log Manager servers from the agent or agent group configuration using the Assign Log Manager Server wizard.
4. Assign the new static IP addresses to those servers you removed from the selected list.
5. Re-add the servers with the new IP addresses.
6. Wait for the agent to refresh its configuration information.

   Note: You can manually restart the agent to refresh the information immediately.
7. Remove the other half of the ordered list.

   Note: In the case where you added a temporary server, you can retain it for failover purposes or you can uninstall it and then delete it.
8. Reassign those IP addresses.
9. Re-add those servers to restore the original ordered list.

More information:

How to Configure Agent Management (see page 546)
Delete a Service Host (see page 121)

How to Apply Subscription Updates

You can apply CA subscription updates to agents or connectors. The process of applying subscription packages using the Updates List wizard has the following steps:

1. Opening the updates list wizard
2. Selecting one of the following update types and specifying search criteria for available update packages:
   - Agent updates
   - Integration updates for connectors

   Note: If there are agent and connector updates, you must apply the agent updates first in order for the update to be completed properly.
3. Selecting the agents or connectors to update to the latest version.

More information:

Open Updates List Wizard (see page 552)
Select Agents or Connectors for Update (see page 553)
Update Agent or Connector Integration Versions (see page 554)

Open Updates List Wizard

To update agents or connectors to the latest version, open the updates list wizard.

To open the updates list wizard

1. Click the Administration tab, and then the Log Collection subtab.

   The Log Collection folder list appears.
2. Click the Agent Explorer folder

   Agent management buttons appear in the details pane.
3. Click Subscription: 

   The updates list wizard appears.

   When using the wizard:
   - Click Save to save without closing the wizard.
   - Click Save and Close to save your progress and close the wizard.
   - Click Reset to restore the wizard display to the last-saved settings.

Select Agents or Connectors for Update

You can check for available updates by specifying search criteria for candidate agents or connectors.

To select agents or connectors for update

1. Open the updates list wizard.

   The Updates Selection List appears.

2. Select Agent Updates or Connector Updates.

   Note: If there are agent and connector updates, you must apply the agent updates first in order for the update to be completed properly.

3. Complete the agent or connector update search criteria:

   a. Select an Agent Group from the drop-down list.

   b. Select the platform from the drop-down list.

   c. Type the agent name pattern, using wild cards.

   d. (Connector Integration updates only) Select the Integration from the drop-down list.

4. Click Search.

   Update packages matching your search qualifications appear in the next wizard step, Version Selection. You must advance to the Version Selection step to view and apply them.

More information:

Open Updates List Wizard (see page 552)
Update Agent or Connector Integration Versions (see page 554)

Update Agent or Connector Integration Versions

You can compare the version of each listed agent or connector with the downloaded update versions to determine whether an update is needed, and then specify whether to replace the current version with a different version.

To update agents or connectors

1. Open the updates list wizard, and select the agents or connectors you want to consider for update.
2. Advance to the Version Selection step.

   The list of agents or connectors meeting your search criteria appear

   - Each agent is displayed with its current version and a drop-down list showing available update versions.
   - Each connector is displayed with its current integration version and a drop-down list showing available update versions.
3. (Optional) Select the Bypass OS version to see all available updates for your selected operating system, regardless of version.
4. Select the agents or connectors to which you want apply updates, and click Save and Close.

   The agent installs the updates, replacing the current version with the selected agent or integration update.

Note: You can verify that all agents or connectors have the latest version applied by reviewing these details after the selected updates have been applied.

Chapter 19: Custom Certificates

This section contains the following topics:

Implementing Custom Certificates

The installation process generates two certificates and places them in the /opt/CA/SharedComponents/iTechnology directory of the CA Enterprise Log Manager server. You can use the installed certificates as is. These certificates have the following names, where *ApplicationName* is CAELM for the CA Enterprise Log Manager product.

- *ApplicationName*Cert.cer

  This certificate is used by all CA Enterprise Log Manager services to communicate with the management server. The entry for this certificate also exists under the CALM.cnf file.

- *ApplicationName*_AgentCert.cer

  This certificate is used by all the Agents to communicate with the CA Enterprise Log Manager server.

  Important! Replacing the CAELM_AgentCert.cer certificate with a custom certificate in an environment with active agents requires reinstallation of these agents.

To use custom certificates, you must first obtain a trusted root certificate from a Root Certificate Authority (CA). A certificate authority can issue multiple certificates in the form of a tree structure. All certificates below the trusted root certificate inherit the trustworthiness of the root certificate. This process assumes that if both certificates are being replaced, the custom service certificate and the custom agent certificate have the same trusted root.

Only custom certificates with .cer extensions are supported. After you obtain a trusted root certificate, the typical sequence of actions to implement custom certificates follows:

1. Add the Trusted Root certificate to iAuthority.conf on the management CA Enterprise Log Manager server or standalone CA EEM.

2. If you are replacing CAELM_AgentCert.cer, add the Trusted Root certificate to iControl.conf on the management CA Enterprise Log Manager, then repeat this addition on every other CA Enterprise Log Manager.
3. If you are replacing CAELMCert.cer, add this custom certificate's common name to the AdministerObjects scoping policy on the management CA Enterprise Log Manager or standalone CA EEM.
4. Add the custom certificates to the iTechnology folder of each CA Enterprise Log Manager server and add the name and password for each certificate in separate configuration files.

More information:

Add the Trusted Root Certificate to the Management CA Enterprise Log Manager Server (see page 556)
Add the Trusted Root Certificate to All Other CA Enterprise Log Manager Servers (see page 557)
Add the Certificate Common Name to an Access Policy (see page 558)
Deploy the New Certificates (see page 559)

Add the Trusted Root Certificate to the Management CA Enterprise Log Manager Server

First, you obtain a Trusted Root Certificate in PEM format from the Certifying Authority (CA). Then you add this Trusted Root Certificate into the iTechnology SPIN web interface of the management server or standalone CA EEM.

To add the Trusted Root Certificate to the management CA Enterprise Log Manager

1. Browse to the CA iTechnology SPIN web interface of the management CA Enterprise Log Manager server or the standalone CA EEM.

   https://<management_ELM_hostname>:5250/spin

   https://<EEM_hostname>5250/spin

   The CA iTechnology SPIN page appears.
2. Select iTech Administrator from the drop-down list and click Go.

   The iTechnology Administrator page appears with a Login link.
3. Click Login.

   The CA iTechnology logon dialog appears.
4. Enter the EiamAdmin credentials, select iAuthority, and click Log In.

5. Select the iAuthority tab and add the Trusted Root to iAuthority.conf as follows:
   a. Enter a Label for the certificate. Do not enter "myself" as the label.
   b. Browse and select the .cer file.
   c. Click Add Trusted Root.

   The confirmation message indicates that the trust root is added to the iAuthority.conf, a file that exists only on the management server or on the standalone CA EEM.
6. If you use a standalone CA EEM, skip to the last step.
7. If you are replacing the CAELM_AgentCert.cer certificate with a custom certificate, add the Trusted Root to iControl.conf as follows:
   a. Select the Configure tab.
   b. Enter the same Label for the certificate that you entered in the previous step.
   c. Browse and select the same root PEM (.cer) file that you selected in a previous step.
   d. Click Add Trusted Root.

   The confirmation message indicates that the trusted root of the custom certificate is added to the iControl.conf file in the iTechnology directory of the management CA Enterprise Log Manager server.
8. Click Logout and close the iTechnology SPIN.

Add the Trusted Root Certificate to All Other CA Enterprise Log Manager Servers

If you are replacing the CAELM_AgentCert.cer certificate with a custom certificate, you must add the Trusted Root Certificate into the iTechnology SPIN web interface of each additional CA Enterprise Log Manager server. In this procedure, you add the Trusted Root Certificate to CA iControl. This procedure is not needed if you are replacing only the CAELMCert.cer certificate.

To add the Trusted Root Certificate to CA iControl of each non-management CA Enterprise Log Manager server

1. Log into the SPIN UI on the iGateway where a non-management server is running. Use the following URL:

   https://<ELM_hostname>:5250/spin/

   The CA iTechnology SPIN page appears.
2. Select iTech Administrator from the drop-down list and click Go.

   The iTechnology Administrator page appears with a Login link.

3. Click Login.

   The CA iTechnology logon dialog appears.

4. Enter the EiamAdmin credentials, select iAuthority, and click Log In.

5. Select the Configure tab and add the Trusted Root as follows:

   a. Enter the same Label for the certificate that you entered in the previous step.

   b. Browse and select the .cer file.

   c. Click Add Trusted Root.

   The trusted root of the custom certificate is added to the iControl.conf file in the iTechnology directory. A confirmation message appears.

6. Click Logout and close the iTechnology SPIN.

Add the Certificate Common Name to an Access Policy

The CAELMCert.cer certificate is used by all CA Enterprise Log Manager services to communicate with the management CA Enterprise Log Manager server. If you replace CAELMCert.cer with a custom certificate, you must add this custom certificate's common name (cn) to the AdministerObjects policy on the management server or the standalone CA EEM server.

Note: It is not necessary to delete [User] CERT_CAELM identity, the common name of the default certificate, from this policy.

To add the custom certificate's common name to the AdministerObjects policy

1. Browse to the management CA Enterprise Log Manager server or the standalone CA EEM server by entering the appropriate URL.

   https://<management_server_hostname>:5250/spin/calm

   https://<EEM_server_hostname>:5250/spin/eiam

2. Log in with Administrative privileges to the CA Enterprise Log Manager management server. If accessing a standalone CA EEM, log in as the EiamAdmin user.

3. Click the Administration tab, the User and Access Management subtab, and the Access Policy link in the left pane. If logged into a standalone CA EEM, click the Manage Access Policies tab.

4. Click the Scoping Policies link.

   The Policy Table of scoping policies appears in the main pane.

5. Scroll to the Administer Objects policy and select the AdministerObjects link.

   The AdministerObjects policy opens in edit mode.

6. Add the common name (cn) of the custom certificate as follows:
   a. Enter the common name of the custom certificate in the Identity field.
   b. Click the arrow to move your entry.

      [User}<*custom certificate cn*> appears in the Selected Identities list.
7. Click Save.

   The AdministerObjects policy is saved with the addition of the common name of your custom certificate as an identity granted read and write access to the resources listed in this policy.
8. Click Close and log out of the CA Enterprise Log Manager user interface.

Deploy the New Certificates

CA Enterprise Log Manager uses two certificates. You can replace one or both of the predefined certificates with custom certificates. To deploy new certificates, you log on to the soft appliance, stop iGateway, add the new certificates, modify the respective configuration files, and then restart iGateway.

Before you deploy new certificates, verify that:

- The Trusted Root Certificate has been added to the iTechnology iAuthority of the management server or standalone CA EEM your CA Enterprise Log Manager servers use.
- If you are replacing CAELM_AgentCert.cer with a custom certificate, the Trusted Root Certificate has been added to the iTechnology iControl of each CA Enterprise Log Manager server.
- The custom certificate's common name has been added to the AdministerObjects access policy. This refers to the custom certificate that is to replace CAELMCert.cer.

To deploy the new certificates

1. Access the host where the CA Enterprise Log Manager server is installed.
2. Use your **caelmadmin** credentials to log on to the CA Enterprise Log Manager server.
3. At the command prompt, switch users to root, that is:

   ```
   su - root
   ```
4. Change directories to /opt/CA/SharedComponents/iTechnology with the following shortcut:

   ```
   cd $IGW_LOC
   ```

5. Stop iGateway:

   ./S99igateway stop

6. To replace CAELMCert.cer:

   a. Copy the custom *ApplicationName*Cert.cer certificate and the *ApplicationName*Cert.key key file into the iTechnology directory.

   b. Open the CALM.cnf file. Replace the certificate name with the new name.

   c. Replace the existing key file name with the new key file name.

7. To replace CAELM_AgentCert.cer:

   a. Copy the custom *ApplicationName*_AgentCert.cer certificate and the *ApplicationName*_AgentCert.key key file into the iTechnology directory.

   b. Open the AgentManager.conf file. Replace the certificate name with the new name.

   c. Replace the existing key file name with the new key file name.

8. Start iGateway.

   ./S99igateway start

   All agents installed after this deployment automatically use the custom certificate, if CAELM_AgentCert.cer was replaced.

Appendix A: Accessibility Features

CA is committed to ensuring that all customers, regardless of ability, can successfully use its products and supporting documentation to accomplish vital business tasks. This section outlines the accessibility features that are part of CA Enterprise Log Manager.

Accessibility Mode

You can set CA Enterprise Log Manager to use an accessibility mode, which displays all graphic panels in queries and reports as tables instead. To enter accessibility mode, select the Activate Accessibility check box on the login screen.

Accessibility Controls

You can use keyboard controls to navigate through CA Enterprise Log Manager, as shown in the following table:

Tasks	Keyboard Controls
Switch between open applications	CTRL-TAB
Select a file in a open window	CTRL-TAB
Help	F1
Button Click	Space or Enter
Check Box Selection	Space or Enter
Open Menu, Combo Box	CTRL + Down Arrow
List Navigation	CTRL + Down Arrow to set focus Up/Down arrows to navigate Space or Enter to select list item
Radio Button Group	CTRL + Down Arrow to set focus Up/Down arrows to navigate Space or Enter to select list item
Close Active Window	ALT F4
Double-click	CTRL + D

CA Enterprise Log Manager Language Display Settings

You can view the CA Enterprise Log Manager interface in the following languages other than English:

- French
- Italian
- German
- Spanish
- Japanese

Change the language settings in your browser window. For example, if you are using Microsoft Internet Explorer, open the Internet Options dialog, and add or select the primary language you want to view.

If you select one of the five supported languages, when you open the CA Enterprise Log Manager interface it appears in the appropriate language. Interface labels and tabs are translated, but certain elements are not: Tags titles, and data strings in report results, for example, remain in English.

Note: If you are displaying CA Enterprise Log Manager when you change languages, refresh your browser window to see the new language. If you are logged in when you do, you are returned to the login screen in the new language.

More information:

Manual Localization for CA Enterprise Log Manager (see page 563)

Manual Localization for CA Enterprise Log Manager

You can manually localize CA Enterprise Log Manager by creating your own language files. This allows you to display the CA Enterprise Log Manager interface in other languages than those already supported. You can do this by copying existing files to use as templates.

To manually localize CA Enterprise Log Manager

1. Log in to your CA Enterprise Log Manager server host, navigate to opt/CA/LogManager/local, and choose the files you want to use as templates. There are two files for each language:

   - content.properties - contains text describing various content, such as report and query names, and descriptions.
   - ui.properties - contains text strings for interface feature titles, such as tab labels and headers.

   Each file is preceded by a standard language prefix. For example, the German content file is named de_content.properties. The English interface file is named en_ui.properties.

2. Copy one of each file type, and rename them using the standard prefix. For example if you wanted to create a localization file for Portuguese, you would copy files and rename them pt_content.properties, and pt__ui.properties.

   Note: The standard language prefixes can be located in your browser's supported language list.

3. Open the files, and translate the original language strings into the language you want. For example, if you had copied the English files, you would replace each English text string with your desired language.

4. Save the manually translated files in the location given in Step 1, on each CA Enterprise Log Manager server where you want them available.

5. Set your browser to the target language, and log into CA Enterprise Log Manager.

More information:

CA Enterprise Log Manager Language Display Settings (see page 562)

Appendix B: Accessing Collected Events with ODBC and JDBC

This section contains the following topics:

About ODBC/JDBC Access in CA Enterprise Log Manager

CA Enterprise Log Manager provides read-only ODBC and JDBC access to the event log store to allow you to do the following things:

- Configure custom reports with an external reporting utility, like BusinessObjects' Crystal Reports
- Retrieve selected log information using external applications

These features allow you to create and format your own custom reports using log information already retrieved by CA Enterprise Log Manager. In addition, you can retrieve data for use with your existing correlation engines, malware detection packages, and other functions.

After installing the CA-supplied client on the system you plan to use to access the event log store, configure a connection to the data source and then begin retrieving data. The subscription service installs the server-side components.

Creating ODBC and JDBC Queries for Use with CA Enterprise Log Manager

Support for ODBC and JDBC in CA Enterprise Log Manager is limited to read-only queries using SELECT statements on the following tables:

- VIEW_EVENT - contains events from the Event Database
- VIEW_INCIDENT - contains incidents from the Incident Database
- VIEW_INCIDENTEVENT_BYID - contains incident component events from the Incident Event Database

Build your SELECT statements using ANSI SQL format and rules. The server-side components contain a SQL parsing engine. The parser implements a large portion of the entry level SQL as defined in the X3.135-1992, "Database Language SQL" specification. The parser also supports SQL features from ANSI SQL99 and commercial databases like Microsoft SQL Server and Oracle. The parser is also compliant with the ODBC minimal grammar specification.

The common event grammar serves as the schema for this table. See the *CEG Reference Guide* for details on the schema.

SQL Support Limitations

CA Enterprise Log Manager does not support stored procedure calls, data control language (DCL) commands, or data definition language (DDL) commands.

CA Enterprise Log Manager does not support the following data manipulation language (DML) operations and keywords:

- UNION
- JOIN
- Nested SELECT
- INSERT
- UPDATE
- DELETE
- Transaction control operators like COMMIT, ROLLBACK, and so forth

Supported SQL Functions

The following are the supported SQL functions for use in building SELECT statements:

- ABS(*numeric_exp*)
- ROUND(*numeric_exp, integer_exp*)
- LCASE(*string_exp*)
- LOWER(*string_exp*)
- LENGTH(*string_exp*)
- LTRIM(*string_exp*)
- RTRIM(*string_exp*)
- SUBSTRING(*string_exp, start,length*)
- UCASE(*string_exp*)
- UPPER(*string_exp*)
- IFNULL (*expr, default_val*)
- ISNULL (*expr, default_val*)
- NVL (*expr, default_val*)
- CONVERT (*value_exp, data_type*)
- CURDATE()
- CURTIME()
- CURTIMESTAMP()
- DATEADD(*datepart, number, date*)
- TIMESTAMPADD(*datepart, number, date*)
- DATEDIFF(*datepart, startdate, enddate*) For datepart, the values Year, Day, and second are allowed.
- TIMESTAMPDIFF(*datepart, startdate, enddate*) For datepart, the values Year, Day, and second are allowed.
- DAYOFMONTH(*date_exp*)
- DAYOFWEEK(*date_exp*)
- DAYOFYEAR(*date_exp*)
- HOUR(*time_exp*)

- MINUTE(*time_exp*)
- MONTH(*date_exp*)
- NOW( )
- SECOND(*time_exp*)
- WEEK(*date_exp*)
- YEAR(*date_exp*)
- AVG([ALL | DISTINCT]*expression*)
- SUM([ALL | DISTINCT]*expression*)
- COUNT({[ALL | DISTINCT]*expression*] | *})
- MAX([ALL | DISTINCT]*expression*)
- MIN([ALL | DISTINCT]*expression*)

How Queries are Processed

CA Enterprise Log Manager processes an ODBC or JDBC client-initiated query in the following way:

1. A client application sends a SELECT statement through an ODBC connection to the CA Enterprise Log Manager server.
2. The CA Enterprise Log Manager server validates the SELECT statement. If the validation is successful, the CA Enterprise Log Manager server creates a data structure representing the query.

   Any errors encountered are returned directly to the client driver.
3. The CA Enterprise Log Manager server converts the SQL elements into a query that it can use. If the conversion is successful, the CA Enterprise Log Manager server runs the query.

   Any errors encountered are returned to the client driver.
4. The CA Enterprise Log Manager server manages state information, including an expiration timer, for each query so that it can be canceled in the event the session is closed or the query expires.
5. The CA Enterprise Log Manager server translates the query results and sends them back to the ODBC client driver, and the client application then receives the data.

Result Column Alias

As part of its query state management, CA Enterprise Log Manager provides support for aliased result column names. Therefore you can display common event grammar fields in your custom reports using your own labels and headings.

The alias names are persisted and used for proper data mapping when the CA Enterprise Log Manager server transfers a result set back to the client driver.

Result Limits

To manage disk space, CA Enterprise Log Manager limits the number of result rows. CA Enterprise Log Manager uses a subset of the Transact-SQL TOP keyword with only a fixed value. The percentage variant of the keyword is not supported.

The default TOP value used in CA Enterprise Log Manager is 5000 rows, with a maximum value of 50,000 rows.

CA Enterprise Log Manager-specific Error Codes

The following are the ODBC and JDBC error codes that can occur while accessing the CA Enterprise Log Manager event log store. Each error message provides specific error details.

- 88 – do not support

  The accompanying error message provides details about the unsupported functionality.

- 300 – generic error

  The accompanying error message provides details.

- 301 – error to convert multibyte character parameters
- 302 – authentication error
- 304 – invalid expression (column) in OrderBy or GroupBy clause

The following errors are SQL statement execution errors:

- 305 – error to start query
- 306 – error to get query status
- 307 – error to parse query results XML
- 308 – query execution error
- 309 – query executed with errors
- 310 – error to fetch query results
- 311 – query timeout

Example: Use an Access Filter to Limit ODBC Results

You can create an access filter to limit the data returned to an ODBC access request. When members the named application group accesses CA Enterprise Log Manager event data using the ODBC client, they see only the information allowed by the filter.

This example assumes that you have an application group named UNIX_Analysts and that you want to restrict all members of that group to see only UNIX events from the event log store. The filter created in this example limits event data views from within the CA Enterprise Log Manager user interface and external requests through ODBC.

More information about access filters is available in the online help.

To create an access filter

1. Log in to CA Enterprise Log Manager as an Administrator user.
2. Click the Administration tab, then click the User and Access Management subtab.
3. Click New Access Filter  . The Access Filter Design wizard starts.
4. Enter UNIX Analysts for the Name field, and the phrase, UNIX Analysts access filter in the Description field, and then click step 2 Identities at the top of the dialog.
5. Change the Type to Application Group and type UNIX in the Name field and then click Search Identities.

   A list of identities matching your search criteria appear in a shuttle control so that you can select the desired identities.
6. Select the UNIX_Analysts application group from the Available Identities list and click the right arrow shuttle control to move the selection into the Selected Identities list.
7. Click step 3 Access Filters at the top of the dialog.
8. Click New Event Filter  to add a line, and then click the field area under Column.
9. Select event_logname from the drop-down list, and then click the field area under Value.

10. Select Unix from the drop-down list. Your dialog resembles the following:

Equation 1: This illustration shows the completed advanced filter page of the Access Filter Design wizard.

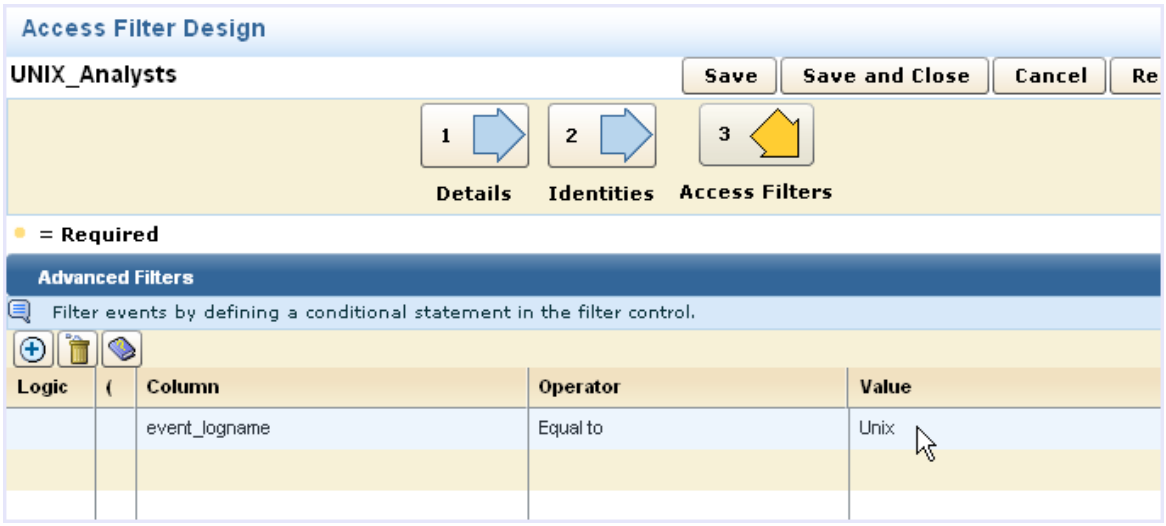

11. Click Save and Close.

Example: Preparing to Use ODBC and JDBC Clients with Crystal Reports

Preparing for ODBC or JDBC client access to CA Enterprise Log Manager events using BusinessObjects Crystal Reports involves the following steps:

1. Create a CA Enterprise Log Manager user to allow access to the database.
2. Verify that the ODBC Service is using SSL encryption and port 17002.
3. Install the ODBC client, or copy the JDBC files to the server where Crystal Reports resides.

   See the *Implementation Guide* for information about these installations.
4. Configure operating system components:
   a. Create and test an ODBC datasource in the Windows Control Panel.
   b. Edit the Crystal Reports configuration file to use the JDBC client.
5. Create events for collection by CA Enterprise Log Manager.

   If you are sure that the event log store already contains events of the type queried, you can omit this step.

Note: This process and the related example assume that you are familiar with creating basic SQL statements and using Crystal Reports. More information about using Crystal Reports is available in the BusinessObjects online help.

Create a CA Enterprise Log Manager User for ODBC or JDBC Access

Use this procedure to create a user account named ELM_Access for use with your JDBC and ODBC clients.

CA Enterprise Log Manager users with the dataaccess permission can use ODBC or JDBC to access event data. Each of the default user roles provided with CA Enterprise Log Manager has this permission. For this example, you could create a user with any of the default CA Enterprise Log Manager roles - Administrator, Analyst, or Auditor.

To create the new user

1. Log in to the CA Enterprise Log Manager server as an Administrator user.
2. Click the Administration tab and the User and Access management subtab.
3. Click the Users button to display the embedded CA EEM user interface.
4. Click New User. The New User dialog opens.
5. Click Add Application User Details and give the account Administrator application privileges.

   Note: In a production environment, you would assign the least permission needed for a user to access the data. You can limit access in many ways including user roles, access policies, and access filters. See the *Online Help* for more information.
6. Complete the user record as needed and assign a password, making a note of it for later use in this example.
7. Save the user and exit the Users window.

Configure the ODBC Service Settings

Use this procedure to configure the CA Enterprise Log Manager ODBC and JDBC service settings.

Note: Changes made to this area cause a restart of the server-side processes that enable ODBC and JDBC communications.

To configure the ODBC Service

1. Log in to the CA Enterprise Log Manager server as an Administrator user.
2. Click the Administration tab and the Services subtab.
3. Click the ODBC Service node.

4. Accept the default settings:
   - Enable Service: True (allows ODBC and JDBC connections on the CA Enterprise Log Manager server)
   - Port: 17002
   - Encrypted (SSL) is selected
   - Session Timeout: 15 minutes
   - Logging Level: accept the default value, NOTSET
5. Click Save.

Create an "elm" ODBC Data Source

Use this procedure to create a data source named "CA-ELM."

Note: Install the ODBC client before you can configure the data source.

To create a data source

1. Access the Windows Control Panel.
2. Open the Administrative Tools folder and start the Data Sources (ODBC) utility.
3. Click Add to display the Create New Data Source dialog.
4. Select the entry, DataDirect OpenAccess SDK 6.0 and click Finish.

   The DataDirect OpenAccess SDK ODBC Driver Setup utility displays a configuration screen.
5. Enter CA-ELM in the Data Source Name field, and provide a text description.
6. Enter the name of your CA Enterprise Log Manager server in the Service Host field. This example uses ca-elm.
7. Enter 17002 in the Service Port field.
8. Select the Encrypted SSL check box.
9. Enter the following Custom Properties:

   querytimeout=600
   (seconds);queryfederated=true;queryfetchrows=1000;offsetmins=0;suppressNoncriticalErrors=false
10. Click Apply and then click Test Connection.

    If your connection parameters are correct, a successful connection message appears.
11. Click OK to return to the ODBC Data Source Administrator dialog and then click OK again to exit the utility.

ODBC Data Source Considerations

The following are the descriptions of the ODBC data source fields as they relate to CA Enterprise Log Manager:

Data Source Name

Create a name for this data source. Client applications that want to use this data use this name to connect to the data source.

Service Host

Specifies the name of the CA Enterprise Log Manager server which the client connects. You can use either a hostname or an IPv4 address.

Service Port

Specifies the TCP service port on which the CA Enterprise Log Manager server listens for ODBC client connections. The default value is 17002. The value you set here must match the setting for the ODBC Server service or the connection fails.

Service Data Source

Leave this field blank, otherwise the connection attempt fails.

Encrypted SSL

Specifies whether to use encryption on the communications between the client and the CA Enterprise Log Manager server. The default value is to have SSL enabled. The value you set here must match the setting for the ODBC Server service or the connection fails.

Custom Properties

Specifies the connection properties for use with the event log store. The delimiter between the properties is a semi-colon with no space. The recommended default values include the following:

querytimeout

Specifies the timeout value in seconds with no data returned after which the query is closed. The following is the syntax for this property:

```
querytimeout=300
```

queryfederated

Specifies whether to perform a federated query. Setting this value to false performs a query only on the CA Enterprise Log Manager server to which the database connection is made. The following is the syntax for this property:

```
queryfederated=true
```

queryfetchrows

Specifies how many rows to retrieve in a single fetch operation, if the query is successful. The minimum value is 1, and the maximum value is 5000. The default value is 1000. The following is the syntax for this property:

queryfetchrows=1000

offsetmins

Specifies the offset for the timezone for this ODBC client. A value of 0 uses GMT. You can use this field to set your own timezone offset from GMT. The following is the syntax for this property:

offsetmins=0

suppressNoncriticalErrors

Indicates the Interface Provider's behavior in case of noncritical errors such as a database not responding or a host not responding.

The following is the syntax for this property:

suppressNoncriticalErrors=false

Edit the Crystal Reports Configuration File

Before you can use Crystal Reports with the CA Enterprise Log Manager JDBC client, you must first provide some configuration settings. After you configure the Crystal Reports XML configuration file, you are ready to prepare and send ANSI SQL standard queries to the CA Enterprise Log Manager event log store.

To configure Crystal Reports settings for JDBC

1. Ensure that you copy the JDBC client JAR files to the Crystal Reports server before you edit the configuration file.

   More information is available in the *Implementation Guide*.

2. Access the server where Crystal Reports resides.
3. Locate the file, CRConfig.xml, and open it for editing.
4. Locate the <DataDriverCommon> tag and the <Classpath> tag section under it.
5. Add the location of the JDBC JAR files for the JDBC client to the classpath.
6. Change the value in the JDBC URL tags to the following:

   jdbc:ca-elm://127.0.0.1:17002;encrypted=1;ServerDataSource=Default;CustomProperties=(querytimeout=600;queryfederated=true;queryfetchrows=1000;offsetmins=0;suppressNoncriticalErrors=false)

   See the section on JDBC URL considerations for more explanation of these parameters.

   Refer to the documentation supplied with Crystal Reports for additional information on setting connection parameters in that product.

7. Change the value in the JDBC Classname tags to be the following:

   com.ca.jdbc.openaccess.OpenAccessDriver

8. Save the file and exit.

More information:

JDBC URL Considerations (see page 576)

JDBC URL Considerations

When using the JDBC client to access event data stored in CA Enterprise Log Manager, you need both the JDBC Classpath and a JDBC URL. The JDBC Classpath names the driver JAR file locations. The JDBC URL defines the parameters the classes in the JARs use when they load.

The following is a complete, sample JDBC URL:

jdbc:ca-elm://127.0.0.1:17002;encrypted=1;ServerDataSource=Default;CustomProperties=(querytimeout=600;queryfederated=true;queryfetchrows=1000;offsetmins=0;suppressNoncriticalErrors=false)

The following descriptions explain the URL components:

jdbc.ca-elm:

Defines the protocol:subprotocol string that designates the JDBC driver provided with CA Enterprise Log Manager.

//IP Address:Port;

Names the IP address that represents the CA Enterprise Log Manager server whose data you want to access. The port number is the port to use for the communications, and must match the setting in the CA Enterprise Log Manager ODBC Service configuration panel. If the ports do not match, the connection attempt fails.

encrypted=0|1;

Determines whether SSL encryption is used for the communications between the JDBC client and the CA Enterprise Log Manager server. The default value is 0, not encrypted, and does not require specification in the URL. Setting encrypted=1 turns encryption on. Set the connection to encryption explicitly. In addition, this setting must match what you configure in the CA Enterprise Log Manager ODBC Service dialog or the connection attempt fails.

ServerDataSource=Default

Specifies the name of the data source. Set this value to *Default* for access to the CA Enterprise Log Manager event log store.

CustomProperties=(x;y;z)

These properties are the same as the ODBC custom properties. If you do not specify them explicitly, the default values shown in the example URL apply.

More information

ODBC Data Source Considerations (see page 574)

Create Events for the ODBC Example

For the example query to display, create some relevant events by causing failed activities like the following:

- Log in incorrectly to the CA Enterprise Log Manager server several times.
- Log in incorrectly to an Agent host whose events go to a specific CA Enterprise Log Manager server.
- Access a network or system resource with incorrect credentials.

Use Crystal Reports to Access the Event Log Store with ODBC

You can use the ODBC access feature to query CA Enterprise Log Manager event data from a third party reporting tool like BusinessObjects Crystal Reports. After you complete the required installation and configurations, you are ready to prepare and send ANSI SQL standard queries to the CA Enterprise Log Manager event log store.

The database schema for the event log store is the common event grammar (CEG). The CA Enterprise Log Manager online help contains a CEG reference component to help you create queries. You can also review the underlying SQL statements for the out-of-the-box queries, but use ANSI SQL to access the database from outside CA Enterprise Log Manager.

To access event data from Crystal Reports

1. Complete the prerequisite installation and configuration tasks.
2. Start Crystal Reports and access the Standard Report Wizard.

3. Create an ODBC connection in the Data dialog, and select the ODBC datasource you created in the Windows Control Panel.

Equation 2: This illustration shows the Standard Report Creation Wizard from BusinessObjects Crystal Reports tool.

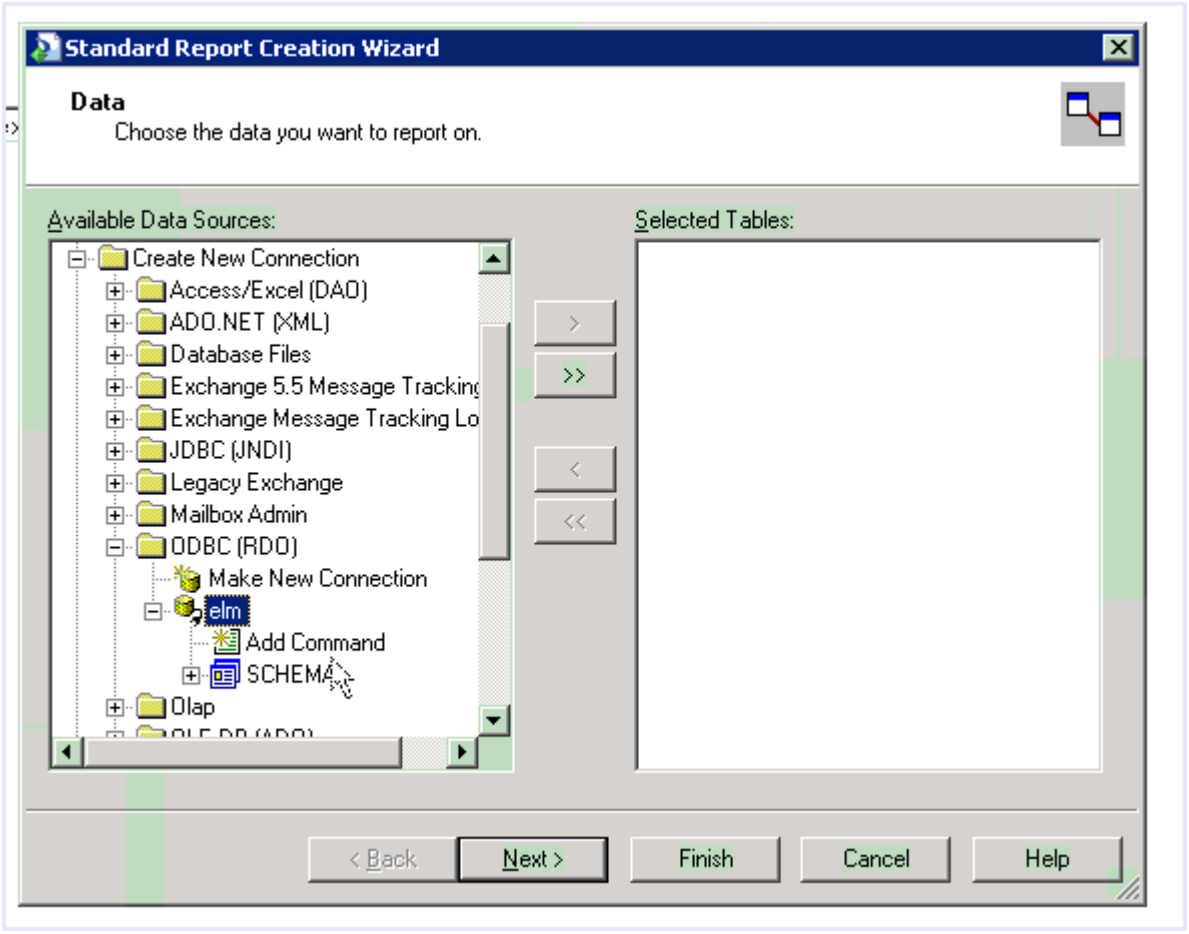

4. Use the Add Command functionality to create a query in the SQL entry area.

   For example, you could create the following query:

   SELECT source_username as source_username , SUM(event_count) AS FUNC_SUM_event_count FROM view_event WHERE event_result = 'F' GROUP BY source_username ORDER BY FUNC_SUM_event_count DESC;

5. Click OK to complete the query entry.

   A report template appears in which to place the data columns returned by the query.

6. Drag and drop the fields from the Field Explorer, in the upper right corner, into the report template as columns.

   Running the query displays the values associated with the fields. You can use Crystal Reports to create any visualization or customization you need.

7. (Optional) Compare report results against the out-of-the-box report, "Failed Activity by Performer."

Accessing Events from Crystal Reports with JDBC

The following tasks enable you to use JDBC to access the event log store:

1. Copy the JDBC JAR files for the client to the server on which Crystal Reports resides.
2. Edit the Crystal Reports configuration file.
3. Use Crystal Reports to send a query.

Note: This process and the related example assume that you are familiar with creating basic SQL statements and about how to use Crystal Reports. More information about using Crystal Reports is available in the BusinessObjects online help.

Copy the JDBC Driver JAR Files

Before you can use the JDBC driver to access events from a CA Enterprise Log Manager server, copy the related JAR files to the server you want to use for access.

To copy the files

1. Open the ISO image or access the Application installation DVD.
2. Navigate to the directory, \CA\ELM\JDBC.
3. Copy the JAR files to the server on which Crystal Reports resides.

   The report package you are using requires a specific location for these files. Consult the documentation supplied with your application.

4. Make a note of the directory where you place these files for reference path when configuring connections.

Use Crystal Reports to Access the Event Log Store with JDBC

You can use the JDBC access feature to query CA Enterprise Log Manager event data from a third party reporting tool like BusinessObjects Crystal Reports.

The database schema for the event log store is the common event grammar (CEG). The CA Enterprise Log Manager online help contains a CEG reference component to help you create queries. You can also review the underlying SQL statements for the out-of-the-box queries, but use ANSI SQL to access the database from outside CA Enterprise Log Manager.

To access event data from Crystal Reports

1. Start Crystal Reports and access the Standard Report Wizard.
2. Create a JDBC connection in the Data dialog.

   Note: Use the value, *Default*, for the Database name in the Connection Information dialog when it appears.
3. Use the Add Command functionality to create and run the following query in the SQL entry area:

   ```
   SELECT source_username as source_username , SUM(event_count) AS FUNC_SUM_event_count FROM view_event WHERE  event_result = 'F' GROUP BY source_username ORDER BY FUNC_SUM_event_count DESC;
   ```
4. Drag and drop the fields from the Field Explorer on the right into the report template as columns.

   Running the query displays the values associated with the fields. You can use the Crystal Reports tools to create any visualization or customization you need.
5. (Optional) Compare report results against the out-of-the-box report, "Failed Activity by Performer."

Remove the ODBC Client on Windows Systems

On all Windows platforms, the Remove option of the client installation package deletes product files and entries in the system information.

Important! If you created any IP source files under the Local ODBC Client installation directory, back the files up to a different location before removing the ODBC client on a Windows system.

If you have the Local ODBC Client installed and want to install it in a different location, use the Remove option. Remove the installed the Local ODBC Client, then re-install it in the new location.

To remove the ODBC client

1. Access the Add or Remove Programs utility on the Windows Control Panel.
2. Locate and select the entry, CA Enterprise Log Manager ODBC Driver.
3. Click Remove.

Remove the JDBC Client

To uninstall the JDBC client, remove the installation directory.

Glossary

access filter

An *access filter* is a filter that the Administrator can set to control what event data non-Administrator users or groups can view. For example, an access filter can restrict the data specified identities can view in a report. Access filters are automatically converted into obligation policies.

access policy

An *access policy* is a rule that grants or denies an identity (user or user group) access rights to an application resource. CA Enterprise Log Manager determines whether policies apply to the particular user by matching identities, resources, resource classes, and evaluating the filters.

account

An *account* is a global user who is also a CALM application user. A single person could have more than one account, each with a different user-defined role.

action alert

An *action alert* is a scheduled query job, which can be used to detect policy violations, usage trends, login patterns, and other event actions that require near-term attention. By default, when the alert queries return results, the results are displayed on the CA Enterprise Log Manager Alerts page and are also added to an RSS Feed. When you schedule an alert, you can specify additional destinations, including email, a CA IT PAM event/alert output process, and SNMP traps.

action query

An *action query* is a query that supports an Action Alert. It is run on a recurring schedule to test for the conditions outlined by the Action Alert to which it is attached.

Administrator role

The *Administrator role* grants users the ability to perform all valid actions on all CA Enterprise Log Manager resources. Only Administrators are permitted to configure log collection and services or manage users, access policies, and access filters.

agent

An *agent* is a generic service configured with connectors, each of which collects raw events from a single event source and then sends them to a CA Enterprise Log Manager for processing. Each CA Enterprise Log Manager has an onboard agent. Additionally, you can install an agent on a remote collection point and collect events on hosts where agents cannot be installed. You can also install an agent on the host where event sources are running and benefit from the ability to apply suppression rules and encrypt transmission to the CA Enterprise Log Manager.

agent explorer

The *agent explorer* is the store for agent configuration settings. (Agents can be installed on a collection point or on the endpoints where the event sources exist.)

agent group

An *agent group* is a tag that users can apply to selected agents that lets user apply an agent configuration to multiple agents at once and retrieve reports based on the groups. A given agent can belong to only one group at a time. Agent groups are based on user-defined criteria such as geographical region or importance.

agent management

Agent management is the software process that controls all agents associated with all federated CA Enterprise Log Managers. It authenticates agents that communicate with it.

alert server

The *alert server* is the store for action alerts and action alert jobs.

Analyst role

The *Analyst role* grants users the ability to create and edit custom reports and queries, edit and annotate reports, create tags, and schedule reports and action alerts. Analysts can also perform all Auditor tasks.

application group

An *application group* is a product-specific group that can be assigned to a global user. Predefined application groups for CA Enterprise Log Manager, or roles, are Administrator, Analyst and Auditor. These application groups are only available for CA Enterprise Log Manager users; they are not available for assignment to users of other products registered to the same CA EEM server. User-defined application groups must be added to the CALM Application Access default policy so that its users can access the CA Enterprise Log Manager.

application instance

An *application instance* is a common space in the CA EEM repository where all the authorization policies, users, groups, content, and configurations are stored. Typically, all CA Enterprise Log Manager servers in an enterprise use the same application instance (CAELM, by default). You can install CA Enterprise Log Manager servers with different application instances, but only severs that share the same application instance can be federated. Servers configured to use the same CA EEM server but with different application instances share only the user store, password policies, and global groups. Different CA products have different default application instances.

application resource

An *application resource* is any of the CA Enterprise Log Manager-specific resources to which CALM access policies grant or deny specified identities the ability to perform application-specific actions such as create, schedule and edit. Examples include report, alert, and integration. See also global resource.

application user

An *application user* is a global user that has been assigned application-level details. CA Enterprise Log Manager application user details include the user group and any restrictions on access. If the user store is the local repository, application user details also include the logon credentials and password policies.

AppObjects

The *AppObjects,* or Application Objects, are product-specific resources stored in CA EEM under the application instance for a given product. For the CAELM application instance, these resources include report and query content, scheduled jobs for reports and alerts, agent content and configurations, service, adapter, and integration configurations, data mapping and message parsing files, and suppression and summarization rules.

archive catalog

See catalog.

archive query

An *archive query* is a query of the catalog that is used to identify the cold databases that need to be restored and defrosted for querying. An archive query is different from a normal query in that it targets cold databases, whereas a normal query targets hot, warm, and defrosted databases. Administrators can issue an archive query from the Administration tab, Log Collection subtab, Archive Catalog Query option.

archived databases

The *archived databases* on a given CA Enterprise Log Manager server include all warm databases that are available for querying but need to be manually backed up before they expire, all cold databases that have been recorded as backed up, and all databases that have been recorded as restored from backup.

audit records

Audit records contain security events such as authentication attempts, file accesses, and changes to security policies, user accounts, or privileges. Administrators specify which types of events should be audited and what should be logged.

Auditor role

An *Auditor role* grants users access to reports and the data they contain. Auditors can view reports, the report template list, the scheduled report job list, the generated report list. Auditors can schedule and annotate reports. Auditors do not have access to the RSS (Rich Site Summary) feeds unless the configuration is set to require no authentication for viewing action alerts.

auto-archive

Auto-archive is a configurable process that automates the moving of archive databases from one server to another. In the first auto-archive phase, the collection server sends newly archived databases to the reporting server at the frequency you specify. In the second phase, the reporting server sends aging databases to the remote storage server for long-term storage, eliminating the need for a manual backup and move procedure. Auto-archiving requires you configure passwordless authentication from the source to the destination server.

CA adapters

The *CA Adapters* are a group of listeners that receive events from CA Audit components such as CA Audit clients, iRecorders, and SAPI recorders as well as sources that send events natively through iTechnology.

CA Enterprise Log Manager

CA Enterprise Log Manager is a solution that helps you collect logs from widely dispersed event sources of different types, check for compliance with queries and reports, and keep records of databases of compressed logs you have moved to external, long-term storage.

CA IT PAM

CA IT PAM is the short form for CA IT Process Automation Manager. This CA product automates processes you define. CA Enterprise Log Manager uses two processes--the process of creating an event/alert output process for a local product, such as CA Service Desk, and the process of dynamically generating lists that can be imported as keyed values. Integration requires CA IT PAM r2.1.

CA Spectrum

CA Spectrum is a network fault management product that can be integrated with CA Enterprise Log Manager for use as a destination for alerts sent in the form of SNMP traps.

CA Subscription Server

The *CA Subscription Server* is the source for subscription updates from CA.

CAELM

CAELM is the application instance name that CA EEM uses for CA Enterprise Log Manager. To access CA Enterprise Log Manager functionality in CA Embedded Entitlements Manager, enter the URL, https://<ip_address>:5250/spin/eiam/eiam.csp, select CAELM as the application name and enter the password of the EiamAdmin user.

caelmadmin

The *caelmadmin* user name and password are credentials required to access the operating system of the soft appliance. The caelmadmin user ID is created during the installation of this operating system. During installation of the software component, the installer must specify the password for the CA EEM superuser account, EiamAdmin. The caelmadmin account is assigned this same password. We recommend that the server administrator ssh in as the caelmadmin user and change this default password. Although the administrator cannot ssh in as root, the administrator can switch users to root (su root) if needed.

caelmservice

The *caelmservice* is a service account that allows iGateway and the local CA EEM services to run as a non-root user. The caelmservice account is used for installing operating system updates downloaded with subscription updates.

calendar

A *calendar* is a means of limiting the times that an access policy is effective. A policy allows specified identities to perform specified actions against a specified resource during a specified time.

CALM

CALM is a predefined resource class that includes the following CA Enterprise Log Manager resources: Alert, ArchiveQuery, calmTag, Data, EventGrouping, Integration, and Report. Actions permitted on this resource class are Annotate (Reports), Create (Alert, calmTag, EventGrouping, Integration, and Report), Dataaccess (Data), Run (ArchiveQuery), and Schedule (Alert, Report).

CALM Application Access policy

The *CALM Application Access policy* is an access control list type of scoping policy that defines who can log into the CA Enterprise Log Manager. By default, the [Group] Administrator, [Group] Analyst and [Group] Auditor are granted logon access.

calmTag

The *calmTag* is a named attribute on the AppObject used when creating a scoping policy to limit the users to reports and queries belonging to certain Tags. All reports and queries are AppObjects and have calmTag as an attribute. (This is not to be confused with the resource Tag.)

catalog

The *catalog* is the database on each CA Enterprise Log Manager that maintains the state of archived databases as well as acting like a high level index across all databases. State information (warm, cold, or defrosted) is maintained for all databases that have ever been on this CA Enterprise Log Manager and any database that has been restored to this CA Enterprise Log Manager as a defrosted database. Indexing ability extends to all hot and warm databases in the event log store on this CA Enterprise Log Manager.

CEG fields

CEG fields are labels used to standardize the presentation of raw event fields from disparate event sources. During event refinement, CA Enterprise Log Manager parses raw event messages into a series of name/value pairs, then maps the raw event names to standard CEG fields. This refinement creates name/value pairs consisting of CEG fields and values from the raw event. That is, different labels used in raw events for the same data object or network element are converted to the same CEG field name when raw events are refined. CEG fields are mapped to OIDs in the MIB used for SNMP traps.

certificates

The predefined *certificates* used by CA Enterprise Log Manager are CAELMCert.cer and CAELM_AgentCert.cer. All CA Enterprise Log Manager services use CAELMCert.cer to communicate with the management server. All agents use CAELM_AgentCert.cer to communicate with their collection server.

cold database state

A *cold database state* is applied to a warm database when an Administrator runs the LMArchive utility to notify CA Enterprise Log Manager that the database has been backed up. Administrators must back up warm databases and run this utility before they are deleted. A warm database is automatically deleted when its age exceeds the Max Archive Days or when the configured Archive Disk Space threshold is reached, whichever comes first. You can query the archive database to identify databases in the warm and cold states.

collection point

A *collection point* is a server on which an agent is installed, where the server has network proximity to all of the servers with event sources associated with its agent's connectors.

collection server

A *collection server* is a role performed by a CA Enterprise Log Manager server. A collection server refines incoming event logs, inserts them into the hot database, compresses the hot database, and auto-archives, or copies, it to the related reporting server. The collection server compresses the hot database when it reaches the configured size and auto-archives it on the configured schedule.

Common Event Grammar (CEG)

Common Event Grammar (CEG) is the schema that provides a standard format to which CA Enterprise Log Manager converts events using parsing and mapping files, before storing them in the Event Log Store. The CEG uses common, normalized fields to define security events from different platforms and products. Events that cannot be parsed or mapped are stored as raw events.

computer security log management

Computer Security Log Management is defined by NIST as "the process for generating, transmitting, storing, analyzing, and disposing of computer security log data."

connector

A *connector* is an integration for a particular event source that is configured on a given agent. An agent can load multiple connectors of similar or dissimilar types into memory. The connector enables raw event collection from an event source and rule-based transmission of converted events to an event log store, where they are inserted into the hot database. Out-of-the-box integrations provide optimized collection from a wide range of event sources, including operating systems, databases, web servers, firewalls, and many types of security applications. You can define a connector for a homegrown event source from scratch or using an integration as a template.

content updates

Content updates are the non-binary portion of subscription updates that are saved in the CA Enterprise Log Manager management server. Content updates include content such as XMP files, DM files, configuration updates for CA Enterprise Log Manager modules, and public key updates.

custom MIB

A *custom MIB* is a MIB you create for an action alert sent to an SNMP trap destination, such as CA NSM. The custom trap ID specified in the action alert assumes the existence of an associated custom MIB that defines the selected CEG fields sent as a trap.

data access

Data access is a type of authorization granted to all CA Enterprise Log Managers through the Default Data Access policy on the CALM resource class. All users have access to all of the data except where restricted by data access filters.

data mapping (DM)

Data mapping is the process of mapping the key value pairs into the CEG. Data mapping is driven by a DM file.

data mapping (DM) files

Data mapping (DM) files are XML files that use the CA Common Event Grammar (CEG) to transform events from the source format into a CEG-compliant form that can be stored for reporting and analysis in the Event Log Store. One DM file is required for each log name before the event data can be stored. Users can modify a copy of a DM file and apply it to a specified connector.

database states

The *database states* include hot for the uncompressed database of new events, warm for a database of compressed events, cold for a backed up database, and defrosted for a database restored to the event log store where it was backed up. You can query hot, warm, and defrosted databases. An archive query displays information on cold databases.

default agent

The *default agent* is the onboard agent that is installed with the CA Enterprise Log Manager server. It can be configured for direct collection of syslog events as well as events from various non-syslog event sources such as CA Access Control r12 SP1, Microsoft Active Directory Certificate Service, and Oracle9i databases.

defrosted database state

A *defrosted database state* is the state applied to a database that has been restored to the archive directory after the Administrator runs the LMArchive utility to notify CA Enterprise Log Manager that it has been restored. Defrosted databases are retained for the number of hours configured for the Export Policy. You can query for event logs in databases that are in the hot, warm, and defrosted states.

defrosting

Defrosting is the process of changing the state of a database from cold to defrosted. This process is performed by CA Enterprise Log Manager when notified by the LMArchive utility that a known cold database has been restored. (If the cold database is not restored to its original CA Enterprise Log Manager, the LMArchive utility is not used and defrosting is not required; recataloging adds the restored database as a warm database.)

delegation policy

A *delegation policy* is an access policy that lets a user delegate their authority to another user, application group, global group, or dynamic group. You must explicitly delete the delegation policies created by the deleted or disabled user.

direct log collection

Direct log collection is the log collection technique where there is no intermediate agent between the event source and the CA Enterprise Log Manager software.

dynamic user group

A *dynamic user group* is composed of global users that share one or more common attributes. A dynamic user group is created through a special dynamic user group policy where the resource name is the dynamic user group name and membership is based on a set of filters configured on user and group attributes.

dynamic values process

A *dynamic values process* is a CA IT PAM process that you can invoke to populate or update the values list for a selected key that is used in reports or alerts. You provide the path to the Dynamic Values Process as part of IT PAM configuration on the Report Server Service List under the Administration tab. You click Import Dynamic Values list on the Values section associated with Key Values on this same UI page. Invoking the dynamic values process is one of three ways you can add values to your keys.

EEM User

The *EEM User*, configured in the Auto-Archiving section of the Event Log Store, specifies the user who can perform an archive query, recatalog the archive database, run the LMArchive utility, and run the restore-ca-elm shell script to restore archive databases for examination. This user must be assigned the predefined role of Administrator or a custom role associated with a custom policy that permits the edit action on the Database resource.

EiamAdmin user name

EiamAdmin is the default superuser name assigned to the installer of the CA Enterprise Log Manager servers. While installing the first CA Enterprise Log Manager software, the installer creates a password for this superuser account, unless a remote CA EEM server already exists. In that case, the installer must enter the existing password. After installing the soft appliance, the installer opens a browser from a workstation, enters the URL for CA Enterprise Log Manager and logs in as EiamAdmin with the associated password. This first user sets the user store, creates password policies, and creates the first user account with an Administrator role. Optionally, the EiamAdmin user can perform any operation controlled by the CA EEM.

entitlement management

Entitlement management is the means of controlling what users are allowed to do once they are authenticated and logged on to the CA Enterprise Log Manager interface. This is achieved with access policies associated with roles assigned to users. Roles, or application user groups, and access policies can be predefined or user-defined. Entitlement management is handled by the CA Enterprise Log Manager internal user store.

EPHI-related reports

The *EPHI-related reports*, are reports that focus on HIPAA security, where EPHI stands for Electronic Protected Health Information. These reports can help you demonstrate that all individually identifiable health information related to patients this is created, maintained, or transmitted electronically is protected.

event aggregation

Event aggregation is the process by which similar log entries are consolidated into a single entry containing a count of the number of occurrences of the event. Summarization rules define how events are aggregated.

event categories

Event categories are the tags used by the CA Enterprise Log Manager to classify events by their function before inserting them into the event store.

event collection

Event collection is the process of reading the raw event string from an event source and sending it to the configured CA Enterprise Log Manager. Event collection is followed by event refinement.

event filtering

Event filtering is the process of dropping events based on CEG filters.

event forwarding rules

Event forwarding rules specify that selected events are to be forwarded to third-party products, such as those that correlate events, after being saved in the event log store.

event log storage

Event log storage is the result of the archiving process, where the user backs up a warm database, notifies CA Enterprise Log Manager by running the LMArchive utility, and moves the backed up database from the event log store to long term storage.

event log store

The *event log store* is a component on the CA Enterprise Log Manager server where incoming events are stored in databases. The databases in the event log store must be manually backed up and moved to a remote log storage solution before the time configured for deletion. Archived databases can be restored to an event log store.

event refinement

Event refinement is the process where a collected raw event string is parsed into constituent event fields and mapped to CEG fields. Users can run queries to display the resulting refined event data. Event refinement follows event collection and precedes event storage.

event refinement library

The *event refinement library* is the store for predefined and user-defined integrations, mapping and parsing files, as well as suppression and summarization rules.

event source

An *event source* is the host from which a connector collects raw events. An event source can contain multiple log stores, each accessed by a separate connector. Deploying a new connector typically involves configuring the event source so that the agent can access it and read raw events from one of its log stores. Raw events for the operating system, different databases, and various security applications are stored separately on the event source.

event/alert output process

The *event/alert output process* is the CA IT PAM process that invokes a third-party product to respond to alert data configured in CA Enterprise Log Manager. You can select CA IT PAM Process as a destination when you schedule an alert job. When an alert runs the CA IT PAM process, CA Enterprise Log Manager sends CA IT PAM alert data and CA IT PAM forwards it along with its own processing parameters to the third party product as part of the event/alert output process.

event_action

The *event_action* is the fourth-level event-specific field in event normalization used by the CEG. It describes common actions. Examples of types of event actions include Process Start, Process Stop, and Application Error.

event_category

The *event_category* is the second-level event-specific field in event normalization used by the CEG. It provides a further classification of events with a specific ideal_model. Event category types include Operational Security, Identity Management, Configuration Management, Resource Access, and System Access.

event_class

The *event_class* is the third-level event-specific field in event normalization used by the CEG. It provides a further classification of events within a specific event_category.

events

Events in CA Enterprise Log Manager are the log records generated by each specified event source.

federation servers

Federation servers are CA Enterprise Log Manager servers connected to one another in a network for the purpose of distributing the collection of log data but aggregating the collected data for reporting. Federation servers can be connected in a hierarchical or meshed topology. Reports of federated data include that from the target server as well as that from children or peers of that server, if any.

filter

A *filter* is a means by which you can restrict an event log store query.

FIPS 140-2

FIPS 140-2 is the Federal Information Processing Standard. This federal standard specifies the security requirements for cryptographic modules used within a security system that protects sensitive but unclassified information. The standard provides four increasing, qualititative levels of security intended to cover a wide range of potential applications and environments.

FIPS 140-2 compatible

FIPS 140-2 compatible is a designation for a product that can *optionally* use FIPS-compliant cryptographic libraries and algorithms to encrypt and decrypt sensitive data. CA Enterprise Log Manager is a FIPS-compatible log collection product because you can select whether to run in FIPS mode or non-FIPS mode.

FIPS 140-2 compliant

FIPS 140-2 compliant is a designation for a product that by default uses *only* encryption algorithms certified by an accredited Cryptographic Module Testing (CMT) laboratory. CA Enterprise Log Manager can use cryptographic modules based on the certified RSA BSAFE Crypto-C ME and Crypto-J libraries in FIPS mode, but may not do so by default.

FIPS mode

FIPS mode is the setting that requires CA Enterprise Log Manager servers and agents to use FIPS-certified cryptographic modules from RSA for encryption. The alternative setting is non-FIPS mode.

folder

A *folder* is a directory path location that CA Enterprise Log Manager management server uses to store the CA Enterprise Log Manager object types. You reference folders in scoping policies to grant or deny users the right to access a specified object type.

function mappings

Function mappings are an optional part of a Data Mapping file for a product integration. A function mapping is used to populate a CEG field when the needed value cannot be retrieved directly from the source event. All function mappings consist of a CEG field name, a pre-defined or class field value and the function used to obtain or calculate the value.

global configuration

The *global configuration* is a series of settings that apply to all CA Enterprise Log Manager servers that use the same management server.

global filter

A *global filter* is a set of criteria you can specify that limits what is presented in all reports. For example, a global filter of the last 7 days reports events generated in the last seven days.

global group

A *global group* is a group that is shared across application instances registered to the same CA Enterprise Log Manager management server. Any user can be assigned to one of more global groups. Access policies can be defined with global groups as Identities granted or denied the ability to perform selected actions on selected resources.

global resource

A *global resource* for the CA Enterprise Log Manager product is a resource shared with other CA applications. You can create scoping policies with global resources. Examples include user, policy, and calendar. See also application resource.

global user

A *global user* is the user account information that excludes application-specific details. The global user details and global group memberships are shared across all CA applications that integrate with the default user store. Global user details can be stored in the embedded repository or in an external directory.

hierarchical federation

A *hierarchical federation* of CA Enterprise Log Manager servers is a topology that establishes a hierarchical relationship between servers. In its simplest form, server 2 is a child of server 1 but server 1 is not a child of server 2. That is, the relationship is one-way only. A hierarchical federation can have multiple levels of parent-child relationships and a single parent server can have many child servers. A federated query return results from the selected server and its children.

hot database state

A *hot database state* is the state of the database in the event log store where new events are inserted. When the hot database reaches a configurable size on the collection server, the database is compressed, cataloged, and moved to warm storage on the reporting server. Additionally, all servers store new self-monitoring events in a hot database.

HTTP proxy server

An *HTTP proxy server* is a proxy server that acts like a firewall and prevents Internet traffic from entering or leaving the enterprise except through the proxy. Outgoing traffic can specify an ID and password to bypass the proxy server. The use of a local HTTP proxy server in subscription management is configurable.

ideal_model

ideal_model represents the technology expressing the event. This is the first CEG field in a hierarchy of fields used for event classification and normalization. Examples of an ideal model include antivirus, DBMS, firewall, operating system, and web server. Check Point, Cisco PIX and Netscreen/Juniper firewall products could be normalized with a value of "Firewall" in the field ideal_model.

identity

An *identity* in CA Enterprise Log Manager is a user or group that is allowed access to the CAELM application instance and its resources. An identity for any CA product can be a global user, an application user, a global group, an application group, or a dynamic group.

identity access control list

An *identity access control list* lets you specify different actions each selected identity can take on the selected resources. For example, with an identity access control list, you can specify that one identity can create reports and another can schedule and annotate reports. An identity access control list differs from an access control list in that it is identity-centric rather than resource-centric.

installer

The *installer* is the individual who installs the soft appliance and the agents. During the installation process, the caelmadmin and EiamAdmin user names are created and the password specified for EiamAdmin is assigned to caelmadmin. These caelmadmin credentials are required for the first access to the operating system; the EiamAdmin credentials are required for the first access to the CA Enterprise Log Manager software and for installing agents.

integration

Integration is the means by which unclassified events are processed into refined events so that they can be displayed in queries and reports. Integration is implemented with a set of elements that enables a given agent and connector to collect events from one of more types of event sources and send them to CA Enterprise Log Manager. The set of elements includes the log sensor and the XMP and DM files that are designed to read from a specific product. Examples of predefined integrations include those for processing syslog events and WMI events. You can create custom integrations to enable the processing of unclassified events.

integration elements

Integration elements include a sensor, a configuration helper, a data access file, one or more XMP message parsing (XMP) files, and one or more data mapping files.

iTech event plugin

The *iTech event plugin* is a CA adapter that an Administrator can configure with selected mapping files. It receives events from remote iRecorders, CA EEM, iTechnology itself, or any product that sends events through iTechnology.

key values

Key values are user-defined values assigned to a user-defined list (key group). When a query uses a key group, the search results include matches to any of the key values in the key group. There are several predefined key groups, some of which contain predefined key values, which are used in predefined queries and reports.

LMArchive utility

The *LMArchive utility* is the command line utility that tracks the backup and restoration of archive databases to the event log store on a CA Enterprise Log Manager server. Use LMArchive to query for the list of warm database files that are ready for archiving. After backing up the listed database and moving it to long-term (cold) storage, use LMArchive to create a record on CA Enterprise Log Manager that this database was backed up. After restoring a cold database to its original CA Enterprise Log Manager, use LMArchive to notify CA Enterprise Log Manager, which in turn changes the database files to a defrosted state that can be queried.

LMSEOSImport utility

The *LMSEOSImport* utility is a command line utility used to import SEOSDATA, or existing events, into CA Enterprise Log Manager as part of the migration from Audit Reporter, Viewer, or Audit Collector. The utility is supported only on Microsoft Windows and Sun Solaris Sparc.

local event

A *local event* is an event that involves a single entity, where the source and the destination of the event is the same host machine. A local event is type 1 of the four event types used in the Common Event Grammar (CEG).

local filter

A *local filter* is a set of criteria you can establish while viewing a report to limit the displayed data for the current report.

log

A *log* is an audit record, or recorded message, of an event or a collection of events. A log may be an audit log, a transaction log, an intrusion log, a connection log, a system performance record, a user activity log, or an alert.

log analysis

Log analysis is the study of log entries to identify events of interest. If logs are not analyzed in a timely manner, their value is significantly reduced.

log archiving

Log archiving is the process of that occurs when the hot database reaches its maximum size, where row-level compression is done and the state is changed from hot to warm. Administrators must manually back up the warm databases before the threshold for deletion is reached and run the LMArchive utility to record the name of the backups. This information then becomes available for viewing through the Archive Query.

log entry

A *log entry* is an entry in a log that contains information on a specific event that occurred on a system or within a network.

log parsing

Log parsing is the process of extracting data from a log so that the parsed values can be used in a subsequent stage of log management.

log record

A *log record* is an individual audit record.

log sensor

A *log sensor* is an integration component designed to read from a specific log type such as a database, syslog, file, or SNMP. Log sensors are reused. Typically, users do not create custom log sensors.

management server

The *management server* is a role assigned to the first CA Enterprise Log Manager server installed. This CA Enterprise Log Manager server contains the repository that stores shared content, such as policies, for all its CA Enterprise Log Managers. This server is typically the default subscription proxy. While not recommended for most production environments, the management server can perform all roles.

mapping analysis

A *mapping analysis* is a step in the Mapping File wizard that lets you test and make changes to a data mapping (DM) file. Sample events are tested against the DM file and results are validated with the CEG.

meshed federation

A *meshed federation* of CA Enterprise Log Manager servers is a topology that establishes a peer relationship between servers. In its simplest form, server 2 is a child of server 1 and server 1 is a child of server 2. A meshed pair of servers has a two-way relationship. A meshed federation can be defined such that many servers are all peers of one another. A federated query returns results from the selected server and all its peers.

message parsing

Message parsing is the process of applying rules to the analysis of a raw event log to get relevant information such as timestamp, IP address, and user name. Parsing rules use character matching to locate specific event text and link it with selected values.

message parsing file (XMP)

A *message parsing file (XMP)* is an XML file associated with a specific event source type that applies parsing rules. Parsing rules break out relevant data in a collected raw event into name/value pairs, which are passed to the data mapping file for further processing. This file type is used in all integrations, and in connectors, which are based on integrations. In the case of CA Adapters, XMP files can also be applied at the CA Enterprise Log Manager server.

message parsing library

The *message parsing library* is a library that accepts events from the listener queues and uses regular expressions to tokenize strings into name/value pairs.

message parsing token (ELM)

A *message parsing token* is a re-usable template for building the regular expression syntax used in CA Enterprise Log Manager message parsing. A token has a name, a type, and a corresponding regular expression string.

MIB (management information base)

The *MIB (management information base)* for CA Enterprise Log Manager, CA-ELM.MIB, must be imported and compiled by each product that is to receive alerts in the form of SNMP traps from CA Enterprise Log Manager. The MIB shows the origin of each numeric object identifier (OID) used in an SNMP trap message with a description of that data object or network element. In the MIB for SNMP traps sent by CA Enterprise Log Manager, the textual description of each data object is for the associated CEG field. The MIB helps ensure that all name/value pairs sent in an SNMP trap are correctly interpreted at the destination.

module (to download)

A *module* is a logical grouping of component updates that is made available for download through subscription. A module can contain binary updates, content updates, or both. For example, all reports make up one module, all sponsor binary updates make up another module. CA defines what makes up each module.

native event

A *native event* is the state or action that triggers a raw event. Native events are received and parsed/mapped as appropriate, then transmitted as raw or refined events. A failed authentication is a native event.

NIST

The *National Institute of Standards and Technology (NIST)* is the federal technology agency that provides recommendations in its Special Publication 800-92 *Guide to Computer Security Log Management* that were used as the basis for the CA Enterprise Log Manager.

non-FIPS mode

Non-FIPS mode is the default setting that permits CA Enterprise Log Manager servers and agents to use a combination of encryption techniques, some of which are not FIPS-compliant. The alternative setting is FIPS mode.

non-interactive ssh authentication

Non-interactive authentication enables files to move from one server to another without requiring the entry of a passphrase for authentication. Set non-interactive authentication from the source server to the destination server before configuring auto archiving or using the restore-ca-elm.sh script.

obligation policy

An *obligation policy* is a policy that is created automatically when you create an access filter. You should not attempt to create, edit, or delete an obligation policy directly. Instead, create, edit or delete the access filter.

observed event

An *observed event* is an event that involves a source, a destination, and an agent, where the event is observed and recorded by an event-collection agent.

ODBC and JDBC access

ODBC and JDBC access to CA Enterprise Log Manager event log stores supports your use of event data with a variety of third-party products, including custom event reporting with third-party reporting tools, event correlation with correlation engines, and event evaluation by intrusion and malware detections products. Systems with Windows operating systems use ODBC access; those with UNIX and Linux operating systems use JDBC access.

ODBC server

The *ODBC server* is the configured service that sets the port used for communications between the ODBC or JDBC client and the CA Enterprise Log Manager server and specifies whether to use SSL encryption.

OID (object identifier)

An *OID (object identifier)* is a unique numeric identifier for a data object that is paired with a value in an SNMP trap message. Each OID used in an SNMP trap sent by CA Enterprise Log Manager is mapped to a textual CEG field in the MIB. Each OID that is mapped to a CEG field has this syntax: 1.3.6.1.4.1.791.9845.x.x.x, where 791 is the enterprise number for CA and 9845 is the product identifier for CA Enterprise Log Manager.

parsing

Parsing, also called message parsing (MP), is the process of taking raw device data and turning it into key-value pairs. Parsing is driven by an XMP file. Parsing, which precedes data mapping, is one step of the integration process that turns the raw event collected from an event source into a refined event you can view.

parsing file wizard

The *parsing file wizard* is a CA Enterprise Log Manager feature that Administrators use to create, edit, and analyze eXtensible Message Parsing (XMP) files stored in the CA Enterprise Log Manager management server. Customizing the parsing of incoming event data involves editing the pre-matched strings and filters. New and edited files are displayed in the Log Collection Explorer, Event Refinement Library, Parsing Files, User folder.

pozFolder

The *pozFolder* is an attribute of the AppObject, where the value is the parent path of the AppObject. The pozFolder attribute and value is used in the filters for access policies that restrict access to resources such as reports, queries, and configurations.

profile

A *profile* is an optional, configurable, set of tag and data filters that can be product-specific, technology-specific or confined to a selected category. A tag filter for a product, for example, limits the listed tags to the selected product tag. Data filters for a product display only data for the specified product in the reports you generate, the alerts you schedule, and the query results you view. After you create the profile you need, you can set that profile to be in effect whenever you log in. If you create several profiles, you can apply different profiles, one at a time, to your activities during a session. Predefined filters are delivered with subscription updates.

prompt

A *prompt* is a special type of query that displays results based on the value you enter and the CEG fields you select. Rows are returned only for events where the value you enter appears in one or more of the selected CEG fields.

query

A *query* is a set of criteria used to search the Event Log Stores of the active CA Enterprise Log Manager server and, if specified, its federated servers. A query targets hot, warm, or defrosted databases specified in the where clause of the query. For example, if the where clause limits the query to events with source_username="myname" in a certain time frame and only ten of the 1000 databases contain records meeting this criteria based on information contained in the catalog database, the query will run against only those ten databases. A query can return a maximum of 5000 rows of data. Any user with a predefined role can run a query. Only Analysts and Administrators can schedule a query to distribute an action alert, create a report by selecting the queries to include, or create a custom query using the Query Design wizard. See also archive query.

query library

The *query library* is the library that stores all predefined and user-defined queries, query tags, and prompt filters.

raw event

A *raw event* is the information triggered by a native event that is sent by a monitoring agent to the Log Manager collector. The raw event is often formatted as a syslog string or name-value pair. It is possible to review an event in its raw form in CA Enterprise Log Manager.

recataloging

A *recataloging* is a forced rebuild of the catalog. A recatalog is required only when restoring data to an event log store on a different server than the one on which it was generated. For example, if you designated one CA Enterprise Log Manager to act as a restore point for investigations on cold data, you would then need to force a recatalog of the database after restoring it to the designated restore point. A recatalog is automatically performed when iGateway is restarted, if needed. Recataloging a single database file can take several hours.

recorded event

A *recorded event* is the raw or refined event information after it is inserted into the database. Raw events are always recorded unless suppressed or summarized, as are refined events. This information is stored and searchable.

refined event

A *refined event* is mapped or parsed event information derived from raw or summarized events. CA Enterprise Log Manager performs the mapping and parsing so that the stored information is searchable.

remote event

A *remote event* is an event that involves two different host machines, the source and the destination. A remote event is type 2 of the four event types used in the Common Event Grammar (CEG).

remote storage server

A *remote storage server* is a role assigned to a server that receives auto-archived databases from one or more reporting servers. A remote storage server stores cold databases for the required number of years. The remote host used for storage typically does not have CA Enterprise Log Manager or any other product installed. For auto-archiving, configure non-interactive authentication.

report

A *report* is a graphical or tabular display of event log data that is generated by executing predefined or custom queries with filters. The data can be from hot, warm, and defrosted databases in the event log store of the selected server and, if requested, its federated servers.

report library

The *report library* is the library that stores all predefined and user-defined reports, report tags, generated reports and scheduled report jobs.

report server

The *report server* is the service that stores configuration information such as the email server to use when emailing alerts, the appearance of reports that are saved to PDF format, and the retention of policies for reports saved to the Report Server and alerts sent to the RSS feed.

reporting server

A *reporting server* is a role performed by a CA Enterprise Log Manager server. A reporting server receives auto-archived warm databases from one or more collection servers. A reporting server handles queries, reports, scheduled alerts, and scheduled reports.

restore point server

A *restore point server* is a role performed by a CA Enterprise Log Manager server. To investigate "cold" events, you can move databases from the remote storage server to the restore point server with a utility, add the databases to the catalog, and then conduct queries. Moving cold databases to a dedicated restore point is an alternative to moving them back to their original reporting server for investigation.

RSS event

An *RSS event* is an event generated by CA Enterprise Log Manager to convey an Action Alert to third-party products and users. The event is a summary of each Action Alert result and a link to the result file. The duration for a given RSS feed item is configurable.

RSS feed URL for action alerts

The *RSS feed URL for action alerts* is: https://{elmhostname}:5250/spin/calm/getActionQueryRssFeeds.csp. From this URL, you can view action alerts subject to the configuration for maximum age and quantity.

RSS feed URL for subscription

The *RSS feed URL for subscription* is a preconfigured link used by online subscription proxy servers in the process of retrieving subscription updates. This URL is for the CA Subscription Server.

SafeObject

SafeObject is a predefined resource class in CA EEM. It is the resource class to which AppObjects, stored under the scope of Application, belong. Users who define policies and filters for granting access to AppObjects refer to this resource class.

SAPI collector

The *SAPI collector* is a CA adapter that receives events from CA Audit Clients. CA Audit Clients send with the Collector action that provides build-in failover. Administrators configure the CA Audit SAPI Collector with, for example, selected ciphers and DM files.

SAPI recorder

A *SAPI recorder* was the technology used to send information to CA Audit before iTechnology. SAPI stands for Submit API (Application Programming Interface). CA Audit recorders for CA ACF2, CA Top Secret, RACF, Oracle, Sybase, and DB2 are examples of SAPI recorders.

SAPI router

The *SAPI router* is a CA adapter that receives events from integrations, such as Mainframe, and sends them to a CA Audit router.

saved configuration

A *saved configuration* is a stored configuration with the values for the data access attributes of an integration that can be used as a template when creating a new integration.

scoping policy

A *scoping policy* is a type of access policy that grants or denies access to resources stored in the management server, such as AppObjects, users, groups, folders, and policies. A scoping policy defines the identities that can access the specified resources.

scp utility

The *scp* secure copy (remote file copy program) is a UNIX utility that transfers files between UNIX computers on a network. This utility is made available at CA Enterprise Log Manager installation for you to use to transfer subscription update files from the online subscription proxy to the offline subscription proxy.

self-monitoring event

A *self-monitoring event* is an event that is logged by CA Enterprise Log Manager. Such events are automatically generated by acts performed by logged in users and by functions performed by various modules such as services and listeners. The SIM Operations Self Monitoring Events Details report can be viewed by selecting a report server and opening the Self Monitoring events tab.

services

The CA Enterprise Log Manager *services* are event log store, report server, and subscription. Administrators configure these services at a global level, where all settings apply to all CA Enterprise Log Managers by default. Most global settings for services can be overridden at the local level, that is, for any specified CA Enterprise Log Manager.

SNMP

SNMP is the acronym for Simple Network Management Protocol, an open standard for sending alert messages in the form of SNMP traps from an agent system to one or more management systems.

SNMP trap contents

An *SNMP trap* consists of name/value pairs, where each name is an OID (object identifier) and each value is one returned from the scheduled alert. Query results returned by an action alert consist of CEG fields and their values. The SNMP trap is populated by substituting an OID for each CEG field used for the name of the name/value pair. The mapping of each CEG field to an OID is stored in the MIB. The SNMP trap only includes name/value pairs for the fields you select when you configure the alert.

SNMP trap destinations

One or more *SNMP trap destinations* can be added when you schedule an action alert. Each SNMP trap destination is configured with an IP address and port. The destination is typically a NOC or a management server such as CA Spectrum or CA NSM. An SNMP trap is sent to configured destinations when queries for a scheduled alert job returns results.

soft appliance

A *soft appliance* is a fully functional software package that contains the software as well as the underlying operating system and all dependant packages. It is installed onto end-user provided hardware by booting from the soft appliance installation media.

subscription client

A *subscription client* is a CA Enterprise Log Manager server that gets content updates from another CA Enterprise Log Manager server called a subscription proxy server. Subscription clients poll the configured subscription proxy server on a regular schedule and retrieve new updates when available. After retrieving updates, the client installs the downloaded components.

subscription module

The *subscription module* is the service that enables subscription updates from the CA Subscription Server to be automatically downloaded and distributed to all CA Enterprise Log Manager servers, and all agents. Global settings apply to local CA Enterprise Log Manager servers; local settings include whether the server is an offline proxy, an online proxy, or a subscription client.

subscription proxies (for client)

The *subscription proxies for client* make up the subscription proxy list that the client contacts in a round robin fashion to get CA Enterprise Log Manager software and operating system updates. If one proxy is busy, the next one in the list is contacted. If all are unavailable and the client is online, the default subscription proxy is used.

subscription proxies (for content updates)

S*ubscription proxies for content updates* are the subscription proxies selected to update the CA Enterprise Log Manager management server with content updates that are downloaded from the CA Subscription Server. Configuring multiple proxies for redundancy is a good practice.

subscription proxy (default)

The *default subscription proxy* is typically the CA Enterprise Log Manager server that is installed first and may also be the Primary CA Enterprise Log Manager. The default subscription proxy is also an online subscription proxy and, therefore, must have Internet access. If no other online subscription proxies are defined, this server gets subscription updates from the CA Subscription server, downloads binary updates to all clients, and pushes content updates to CA EEM. If other proxies are defined, this server still gets subscription updates, but is contacted by clients for updates only when no subscription proxy list is configured or when the configured list is exhausted.

subscription proxy (offline)

An *offline subscription proxy* is a CA Enterprise Log Manager server that gets subscription updates through a manual directory copy (using scp) from an online subscription proxy. Offline subscription proxies can be configured to download binary updates to clients that request them and to push the latest version of content updates to the management server if it has not yet received them. Offline subscription proxies do not need Internet access.

subscription proxy (online)

An *online subscription proxy* is a CA Enterprise Log Manager with Internet access that gets subscription updates from the CA Subscription server on a recurring schedule. A given online subscription proxy can be included in the proxy list for one or more clients, who contact listed proxies in round-robin fashion to request the binary updates. A given online proxy, if so configured, pushes new content and configuration updates to management server unless already pushed by another proxy. The subscription update directory of a selected online proxy is used as the source for copying updates to offline subscription proxies.

subscription updates

*Subscription update*s refer to the binary and non-binary files that are made available by CA Subscription server. Binary files are product module updates that are typically installed on the CA Enterprise Log Managers. Non-binary files, or content updates, are saved to the management server.

summarization rules

Summarization rules are rules that combine certain native events of a common type into one refined event. For example, a summarization rule can be configured to replace up to 1000 duplicate events with the same source and destination IP addresses and ports with a single summarization event. Such rules simplify event analysis and reduce log traffic.

suppression

Suppression is the process of dropping events based on CEG filters. Suppression is driven by SUP files.

suppression rules

Suppression rules are rules you configure to prevent certain refined events from appearing in your reports. You can create permanent suppression rules to suppress routine events of no security concern and you can create temporary rules to suppress the logging of planned events such as the creation of many new users.

tag

A *tag* is a term or key phrase that is used to identify queries or reports that belong to the same business-relevant grouping. Tags enable searches based on business-relevant groupings. Tag is also the resource name used in any policy that grants users the ability to create a tag.

URL for CA Embedded Entitlements Manager

The *URL for CA Embedded Entitlements Manager* (CA EEM) is: https://<ip_address>:5250/spin/eiam. To log in, select CAELM as the application and enter the password associated with the EiamAdmin user name.

URL for CA Enterprise Log Manager

The *URL for CA Enterprise Log Manager* is: https://<ip_address>:5250/spin/calm. To log in, enter the user name defined for your account by the Administrator and the associated password. Or, enter the EiamAdmin, the default superuser name, with the associated password.

user group

A *user group* can be an application group, a global group, or a dynamic group. Predefined CA Enterprise Log Manager application groups are Administrator, Analyst, and Auditor. CA Enterprise Log Manager users may belong to global groups through memberships apart from CA Enterprise Log Manager. Dynamic groups are user-defined and created through a dynamic group policy.

user role

A *user role* can be a predefined application user group or a user-defined application group. Custom user roles are needed when the predefined application groups (Administrator, Analyst, and Auditor) are not sufficiently fine-grained to reflect work assignments. Custom user roles require custom access policies and modification of predefined policies to include the new role.

user store

A *user store* is the repository for global user information and password policies. The CA Enterprise Log Manager user store is the local repository, by default, but can be configured to reference CA SiteMinder or a supported LDAP directory such as Microsoft Active Directory, Sun One, or Novell eDirectory. No matter how the user store is configured, the local repository on the management server contains application-specific information about users, such as their user role and associated access policies.

varbind

A *varbind* is an SNMP variable binding. Each varbind is made up of an OID, a type, and a value. You add varbinds to a custom MIB.

visualization components

Visualization components are available options for displaying report data including a table, a chart (line graph, bar graph, column graph, pie chart), or an event viewer.

warm database state

The *warm database state* is the state that a hot database of event logs is moved into when the size (Maximum Rows) of the hot database is exceeded or when a recatalog is performed after restoring a cold database to a new event log store. Warm databases are compressed and retained in the event log store until their age in days exceeds the configured value for Max Archive Days. You can query for event logs in databases that are in the hot, warm, and defrosted states.

XMP file analysis

XMP file analysis is the process performed by the Message Parsing utility to find all events containing each pre-match string and, for each matched event, parse the event into tokens using the first filter found that uses the same pre-match string.

Index

A

C

D

E

F

G

I

K

L

M

P

Q

R

S

T

U

V