

CA Identity Manager

Provisioning Reference Guide

r12.5 SP1



This documentation and any related computer software help programs (hereinafter referred to as the "Documentation") are for your informational purposes only and are subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be used or disclosed by you except as may be permitted in a separate confidentiality agreement between you and CA.

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

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

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

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

The manufacturer of this Documentation is CA.

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

Copyright © 2010 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Product References

This document references the following CA products:

- CA Identity Manager
- CA SiteMinder® Web Access Manager
- CA Directory
- CA Enterprise Log Manager
- CA Role & Compliance Manager

Contact CA

Contact Technical 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 also available on the CA Support website, found at <http://ca.com/docs>.

Contents

Chapter 1: Provisioning Manager	11
User Interface for Provisioning	11
Provisioning Server	11
Administrator Authentication	12
Administrator Login	12
Administrator Authorization	13
 Chapter 2: Advanced Configuration Options	 15
Advanced Configuration Options Overview	15
Global Properties	16
Domain Configuration	17
Provisioning Directory Parameters	20
Authentication Parameters	20
Authorization Parameters	21
Cache Parameters	22
Compatibility Parameters	28
Configuration Setup Parameters	29
Connections Parameters	29
Endpoint Parameters	32
Explore and Correlate Parameters	36
Identity Manager Server Parameters	44
Operation Details Parameters	47
Password Synchronization Parameters	49
Password Parameters	51
Processes Parameters	54
Processor Parameters	56
Search Parameters	57
Servers Parameters	62
Statistics Parameters	62
Synchronization Parameters	64
Transaction Log Parameters	68
 Chapter 3: SPML Service	 73
SPML Overview	73
Benefits of Using SPML	73
When You Would Use the SPML Service	74

SPML Architecture	74
SPML Integration	78
Install SPML	79
SPML Support for FIPS 140-2	80
Uninstall the SPML Service	81
SPML Service Configuration	81
Log On to the SPML Configuration Application	81
Add a New SPML Service	82
Modify an Existing Service	83
Rename an Existing Service	83
Delete an Existing Service	84
Configure SSL Support for Tomcat Servers	84
Configure SPML Client Computer to Support SSL Security	86
CMDRA Commands	87
SPML Feed	90
Using the SPML Manager's Templating Functionality	95
Using Velocity Templates	97
Retrying SPML Requests	102

Chapter 4: Sample SPML Requests 107

Request Execution Types	107
Request Types	108
Add Request	108
Batch Request	108
Cancel Request	110
Delete Request	110
Extended Request	111
Modify Request	115
Propagate Global User Changes	116
Schema Request	117
Search Request	117
Status Request	119
Global Settings	120
Example: Search for Attributes Defined in Global Settings	120
Example: Modify Attributes in Global Settings	121
Account Containers	121
Example: Create an Account Container	122
Example: Create an Account within a Sub-Container	122
Complex Attributes	123
Example: Add a Single-Valued Complex Attribute	124
Example: Add a Multivalued Complex Attribute	124
Request Retries	124

Propagate Global User Changes	125
Example: Modify a Global User and Propagate Changes to Associated Accounts	125
Example: Modify Complex Attribute and Propagate Changes to Accounts	126
Escaping Special Characters in Object Identifiers	126
Escaping Special Characters in Search Filters	127

Chapter 5: etautil Batch Utility **129**

Tasks You Can Perform	129
etautil Syntax	130
etautil Control Statements	132
Multivalued Attributes	136
Use DeletePending	137
Common Error Messages	138
Unknown error nnn opening Common Object Repository	138
End of file reached while expecting an operator	138
Object 'XXXX' operation failed: DB operation failed: Target DN not found.	138
Object 'XXXX' operation failed: No server plug-in found for operation	138
Class 'classname' is not a valid class name	139
Could not find keyword xxxxx for class classname	139
Obtain Operation Details	139
DOS Output from etautil	140

Chapter 6: Provisioning Servers on UNIX **141**

No UNIX GUI Clients or Utilities	141
Command Line Examples	142
Libraries and Executables	142
Registry Access	143
Parser Tables	144
UNIX Services for Provisioning	144
Working with Hung or Crashed Servers	144
Scheduling Periodic Actions	145
Passwords on Command Lines	145
Server Event Logging Destinations	145
Program Exit Definitions	146
C++ Connector Server on Solaris	146

Chapter 7: Program Exits **147**

Program Exits Overview	147
Ordering of Program Exit Invocations	148
Basic Structure of Program Exits	150
Define Common Exits in the Provisioning Manager	150

Chapter 8: Common Program Exit Reference **153**

Program Exit Architecture	153
Program Exit Hierarchy and Order	154
Common Program Exit Structure	155
Program Exit Input Argument	155
Program Exit Return Value	157
Common Exits DLL Interface	160
Common Exits SOAP Interface	162
eExitType	162
Valid Values for eExitType	163
Containment	168
Custom Function Program Exits	171
Obscured Returned Values	173
Sample Flow/Execution Diagram	174
Code Examples	174

Chapter 9: Program Exits In Connectors **175**

Execution Flow (Logic)	175
Pre-Exits	175
Operation	176
Post-Exit	176
Support for Common Exits	177
Parser Table Enhancement	177
GUI Plug-In Enhancement	177
Agent Plug-In Enhancement	177
Support for Native Exits	177
Parser Table Enhancement	177
GUI Plug-In Enhancement	178
Agent Plug-in Enhancement	179
Exit Types	181
Exit Type Functionality	181
Code Examples for Program Exits in Options	182

Chapter 10: Provisioning Maintenance **183**

Back Up and Restore CA Directory	183
Shut Down the Provisioning Server service	183
View and Maintain Log Files	184
Server Event Logging	184
Diagnostic Logging	185
Log Files for High Availability	190
Provisioning Directory Monitoring	190

Monitor Thresholds	190
Observe Router Traffic	191
Enable SSL Encryption	191
Index	193

Chapter 1: Provisioning Manager

This section contains the following topics:

[User Interface for Provisioning](#) (see page 11)

[Provisioning Server](#) (see page 11)

[Administrator Authentication](#) (see page 12)

User Interface for Provisioning

The Provisioning Manager is the user interface for advanced provisioning operations. This interface was formerly called eTrust Admin Manager. To use Provisioning Manager, choose Start, CA, Identity Manager, Provisioning Manager.

Note: The Provisioning Manager method will not be available in future releases, so we recommend using the Identity Manager User Console instead. This guide describes the provisioning features of CA Identity Manager that are not integrated into the Identity Manager User Console. For the features that are integrated into the User Console, such as account management tasks, see the *Administration Guide*.

Provisioning Server

The Provisioning Server is the server that manages additional accounts that are assigned to an Identity Manager user. When you assign a provisioning role to an Identity Manager user, the Provisioning Server creates accounts on endpoints that meet the requirements of the role. For example, if you assign a provisioning role that includes an Exchange account template, the Provisioning Server assigns an Exchange account to the user.

Administrator Authentication

Objects in the provisioning domain are protected at several different levels, but overall access to the domain is protected by authentication security, which requires all administrators to identify themselves. The global user name and password that the administrator enters are checked against information stored in the provisioning directory.

Note: You can configure the Provisioning Server to request authentication with a native system. For more information, see the *Administration Guide*.

Administrator Login

The first time you log on to the Provisioning Manager, you use the etaadmin global user, whose password was set up when the Provisioning Server was installed.

The Provisioning Server also provides other built-in global users that provide authentication information for use only by Provisioning Server components. You will not use these global users to log on; instead, they provide additional authentication information for the domain.

The etaadmin user is similar to a built-in superuser account. You set a password for this user when the Provisioning Server is installed. It is *imperative* to remember that password because you need it to log on as etaadmin the first time you use Provisioning Manager.

When you log on using the etaadmin object, you have access to all the objects in the domain. You should immediately create a global user object for yourself and assign the DomainAdministrator profile to it. When this user has been created, you should log in as that user and not perform any more actions as the non-specific etaadmin administrator. Avoiding the use of etaadmin improves the traceability of actions as seen in various logs. For more security, you can delete or suspend the etaadmin user after you create your own account.

Administrator Authorization

Authorization determines what administrators can do on the Provisioning Server. It defines the privileges that an administrator has in a domain. You can authorize an administrator by assigning an *admin profile* to the administrator's global user object or by assigning an admin profile to one of the administrator global user groups.

Admin Profiles

Admin profiles permit administrators certain types of access and privileges to manage objects in a domain. Admin profiles contain all the privileges that administrators need to perform different tasks. While administrative privileges can be assigned directly to global users, using admin profiles provides several advantages:

- Several administrators can be defined to a profile, and therefore, each receives the same administrative privileges.
- The operations that an administrator is allowed to perform will typically necessitate being granted a long list of administrative privileges. Placing them into an admin profile is less error prone as it lets you define the profile once and then apply those privileges to multiple administrators.
- Admin profiles can be accessed from other domains, making it easy for other administrators to create new profiles from existing profiles in other domains.

Note: When assigning individual administrative privileges, you must give the administrator Read access to the object and its container. This access is necessary to list and search for objects.

Default Admin Profiles

The Provisioning Server provides default admin profiles that control the privileges of an administrator. These profiles give administrators access to the objects in the domain of the profile. Like the default administrator objects, such as etaadmin, the following profiles are created automatically when you install the Provisioning Server:

- **DomainAdministrator and DomainAdministrator-NoWeb**-Gives administrators full access to every object in the domain. Administrators who have this profile in the root domain have full access to all Provisioning Server objects and security information.
- **PasswordAdministrator**-Lets administrators change passwords and activate or suspend global users.
- **UserAdministrator**-Lets administrators manage users in the domain. Administrators with this profile cannot modify provisioning roles or account templates.
- **ReadAdministrator**-Lets administrators read every object in the domain.
- **SelfAdministrator**-Defines the actions that can be performed by self-administrators. By default, this profile authorizes self-administrators to read their own global user object, list their accounts, and modify specific attributes of their own global user or accounts. You can customize this profile to meet your self-administrator authorization requirements.

Chapter 2: Advanced Configuration Options

This section contains the following topics:

[Advanced Configuration Options Overview](#) (see page 15)

[Global Properties](#) (see page 16)

[Domain Configuration](#) (see page 17)

Advanced Configuration Options Overview

The advanced configuration options for the Provisioning Server fall into two categories:

Global Properties

Configuration settings that are saved in the provisioning directory and control the behavior of the Provisioning Server. Click the Global Properties button on the Provisioning Manager System Task frame.

Domain Configuration

Configuration settings that are saved in the provisioning directory and control the behavior of the Provisioning Server. Click the Domain Configuration button on the Provisioning Manager System Task frame.

There are additional places to configure various components of the Provisioning Server, including the following:

PSHOME\Data\im-ps.conf

Parameters control general behaviors of the Provisioning Server service not controllable through Domain Configuration parameters. See the *Installation Guide* for more details and comments in the file itself.

PSHOME\Data\im_ccs.conf

Parameters control general behaviors of the Identity Manager C++ Connector Server service. See the *Installation Guide* for more details and comments in the file itself.

PSAHOME\data\eta_pwdsync.conf

Parameters control an optional Password Synchronization Agent you may choose to install on a Windows system. For more information, see the *Administration Guide*.

Note: *PSAHOME* is the directory where the Password Synchronization Agent is installed.

%DXHOME%\config\knowledge*.dxg and *.dxc

Used to configure CA Directory failover. For more information, see the *Installation Guide*.

Global Properties

Global Properties are stored in the Provisioning Directory. These properties control the entire enterprise and as such are stored outside of any of your domain-specific data. To view these properties you must have a privilege that grants Read access to the SystemSettings object in the ETA domain. All predefined admin profiles in any domain grant this read access.

To change these properties you must have a privilege that grants Modify access to the SystemSettings object in the domain. The DomainAdministrator and DomainAdministrator-NoWeb admin profiles in any domain grant this modify access.

Updates to global properties in most cases take effect immediately, no restarting of services or programs is necessary. Two specific exceptions are:

- Properties that control Manager behaviors, such as UID controls, and Full Name controls, do not affect property sheets that are already displayed. You may have to close a property sheet and reopen it to have the change take effect for your Manager.
- Logging settings are broadcast to affected Provisioning and Identity Manager Provisioning Connector server services. However, this broadcast currently only goes to one Provisioning Server service per domain. If you have installed a failover or load-balancing configuration with multiple Provisioning Server services for a single domain, you will need to restart all Provisioning and Connector Server services for that domain to ensure that the new logging settings are recognized by all affected components.

Domain Configuration

Domain configuration parameters are also stored in the provisioning directory. You manage these from the System Task of Provisioning Manager using the Domain Configuration button. Parameters are organized into a tree hierarchy using folders so that related parameters are easier to manage. These parameters control the Provisioning Server for a single domain.

If you configured multiple alternative servers, each with its own Provisioning Server for the same domain, all servers for the domain share the same configuration parameter settings. There are a few parameters that you might want to set to different values on different servers, even in the same domain. Per-server values are referred to as specializations. Use the Add Specialization or Remove Specialization menu items to work with server-specific values. These server-specific specializations are displayed in the tree hierarchy under the domain parameter. If there is no specialization for a particular server, the domain parameter value applies to that server. In most cases, the Provisioning Server lets you create a specialization for a parameter even if that could result in inconsistent behaviors from the alternative servers for a domain. This lets you have a dedicated server used for a specialized purpose where you actually want that different behavior. However, for a small set of parameters, specializations are not allowed. A typical reason would be because client code needs to know the value of the parameter even when it does not know which server handles its request. In those cases, the Add Specialization menu item is disabled.

To view these parameters, you must have a privilege that grants Read access to Configuration Parameter objects in the domain where they reside. All the predefined admin profiles grant this Read access to the domain configuration parameters in their own domain, including any subordinate domain.

To change these parameters, you must have a privilege that grants Modify access to Configuration Parameter objects in the domain where they reside. The DomainAdministrator and DomainAdministrator-NoWeb admin profiles grant Modify access to the domain configuration parameters in their own domain and any subordinate domain. You can create a custom admin profile that grants Read or Modify access to specific configuration parameters if you need scoping control.

Domain Configuration parameter updates take effect immediately on the provisioning server where the update was processed. However, if you configured multiple alternative provisioning servers for the domain, the other servers will not take the changed parameters into account immediately. The updated parameters are stored in the provisioning directory immediately, but each affected Provisioning Server refreshes its knowledge of the parameter values periodically. By default the update frequency is every 10 minutes; however you may change this value with the parameter Configuration Setup/Parameter Update Time described later. Thus you will need to wait up to 10 minutes for the refresh to take place. The refresh is recorded in the Provisioning Server Trace log with messages that include the text "ETA::Configuration update completed".

Note: For more information, see the [Transaction Log](#) (see page 68) section.

You may choose to restart the affected Provisioning Server services to ensure the parameters are updated. When the service starts, the service writes information to the Provisioning Server trace log about configuration parameters. This log can be valuable in understanding what parameters were in effect at any particular point. The following information is written to the trace log at startup:

- If the Transaction Log/Level domain configuration parameter is set to a value of 0 or greater and Transaction Log/Enabled is Yes, non-default configuration parameters values are written.
- If the Transaction Log/Level domain configuration parameter is set to a value of 1 or greater and Transaction Log/Enabled is Yes, all configuration parameters values are written.

Note: A few parameters do not take effect even after the periodic configuration parameter update. They only take effect on restart of the Provisioning Server service. Such parameters display the following warning on their properties: Changing this parameter will require restarting all affected servers.

Provisioning Directory Parameters

The provisioning directory configuration folder contains parameters you can use if you have a non-default installation of your provisioning directory.

Provisioning Directory/Entry Count Attribute

Values: dxEntryCount (default) or <unset>

Description: Set this attribute to dxEntryCount if CA Directory is being used for the provisioning directory; but clear it otherwise. This attribute is used in queries sent to the provisioning directory to check whether size limits will be reached; but only if the client is not requesting partial results.

Authentication Parameters

The Authentication configuration folder contains parameters you can use to customize user authentication behaviors of the Provisioning Server.

Authentication/Disable Maintenance User

Values: No (default) or Yes

Description: Set this parameter to yes to disable the ability to authenticate to the Provisioning Server using the built-in user with the Distinguished Name cn=etaserver,dc=eta. This user, whose password is controlled by the pwdmgr utility, is used internally during installation.

After installation, this user is only needed for maintenance functions such as resetting an administrator's password. We recommend that you disable this user after installation.

Authorization Parameters

The Authorization configuration folder contains parameters you can use to customize authorization behaviors of the Provisioning Server.

Authorization/Check Owner Access on Indirect Privileges

Values: Yes (default) or No

Description: Controls what access checks are performed when assigning a global user group or admin profile to a global user, global user group or admin profile.

Regardless of the setting of this parameter, the Provisioning Server checks for Modify access to a specific attribute of the object being assigned and Modify access to a specific attribute of the object to which it is assigned.

If this parameter is Yes (the default), the Provisioning Server will also check for Owner access to each of the objects to which the assigned admin profile or global user group grants access. This prevents one from being able to assign privileges through a global user group or admin profile that one could not have assigned directly to the target global user, global user group or admin profile.

If you do not need added protection, this parameter can be set to No to disable additional Owner access checks. Doing so lets you have one set of administrators who define admin profiles and another set of administrators who assign those admin profiles to users.

Cache Parameters

The Cache configuration folder contains parameters that allow you to tune the Provisioning Server's use of its internal caches. Caches are used in the Provisioning Server to save information read from the provisioning directory so that it does not need to be read repeatedly in the same operation or across multiple operations.

Important! Changes to cache parameters do not take effect until the Provisioning Server service is restarted.

Each cache is controlled by the following parameters:

Maximum Age

The maximum time in seconds that an item remains in the cache without being reread from the provisioning directory.

Maximum Size

The maximum number of unused items to retain in the cache. While a cache item is being used by an operation, it is considered in-use, and there is no limit on the number of in-use cache items. However, when all operations finish with the cache item, it is marked unused and retained only when the number of used and unused items in the cache is no more than the configured maximum size.

Cache items are also removed from a cache when explicitly canceled. This occurs when a change is made to the provisioning directory data from which the cache item originates. This cache invalidation only occurs on the Provisioning Server that processed that provisioning directory update. If you have multiple provisioning domains or alternative servers serving a single domain, other servers may have cache items still derived from the prior data. That is why there is a cache maximum age parameter.

Cache items also are canceled when access is to be denied. The privilege caches (Admin Profile, Global User and Global User Group) contain privilege information used to perform authorization checks. If you have recently assigned a privilege to someone, you do not want to have to wait up to 10 minutes (the default cache maximum age for these caches) for that privilege addition to be recognized. Therefore if an authorization check using cached privileges is about to report DENIED, the cache items are canceled and re-initialized from the provisioning directory. If the result is still DENIED, that authorization failure is reported to the administrator.

Important! When you remove a privilege from a global user, admin profile, or global user group, expect that this change will take place at most 10 minutes (the default) from the time of the change. In most cases this is sufficient. However, if the reason for removing the access is to remove an imminent security threat, to ensure immediate enforcement of that privilege change requires you to restart all affected Provisioning Server services.

Admin Profile Privilege Cache

Each admin profile privilege cache item stores information obtained from an admin profile, including administrative privileges, and the names of included admin profiles.

Parameter: Cache/Admin Profile Privilege Cache Maximum Age

Default Value: 600 seconds (equals 10 minutes)

Parameter: Cache/Admin Profile Privilege Cache Maximum Size

Default Value: 10

Domain Cache

Each domain cache item stores information obtained from a Domain (DSA) Registration object. These objects record information that is necessary for one Provisioning Server to talk to another Provisioning Server.

Parameter: Cache/Domain Cache Maximum Age

Default Value: 3600 seconds (equals 1 hour)

Parameter: Cache/Domain Cache Maximum Size

Default Value: 20

Global User Group Privilege Cache

Each global user group privilege cache item stores information obtained from a global user group, including administrative privileges, the names of included admin profiles, and the names of included global user groups.

Parameter: Cache/Global User Group Privilege Cache Maximum Age

Default Value: 600 seconds (equals 10 minutes)

Parameter: Cache/Global User Group Privilege Cache Maximum Size

Default Value: 10

Global User Privilege Cache

Each global user privilege cache item stores information obtained from a global user and its assigned admin profiles and global user groups. This information includes administrative privileges, password, suspension status and the names of included admin profiles and including global user groups.

Unlike the admin profile privilege cache and global user group privilege cache, the global user privilege cache items also store indirect information obtained from referenced items, so it contains a full list of the accesses privileges that a global user has.

Each time a global user privilege cache item is initialized, the global user's full list of effective privileges and assigned admin profiles and global user groups is written to the server trace log. This information, written only if Transaction Log/Level is set to 4 or greater, includes information obtained directly from the global user, information obtained indirectly from assigned global user groups and admin profiles, and information obtained implicitly based on assigned web or workflow privileges. Look for the text "EFFECTIVE PRIVILEGE LIST INITIALIZED" in the server trace log. This will be followed by the distinguished name of a global user and a list of privileges, a list of admin profiles and a list of global user groups.

Parameter: Cache/Global User Privilege Cache Maximum Age

Default Value: 600 seconds (equals 10 minutes)

Parameter: Cache/Global User Privilege Cache Maximum Size

Default Value: 20

Notification Config Cache

The Notification Configuration Cache stores configuration information that drives the Identity Manager Server Notification feature. This configuration information, stored in the provisioning directory and updatable by the service utility `etoloadnotificationconf`, defines the mapping between provisioning actions and notification records that are sent to the IMS.

Parameter: Cache/Notification Config Cache Maximum Age

Default Value: 600 seconds (equals 10 minutes)

Parameter: Cache/ Notification Config Cache Maximum Size

Default Value: 10

Operation Cache

Each operation cache item stores information about an ongoing or recently completed operation. When you do an explore operation, for instance, the displayed Operation detail count value is obtained from an operation cache item.

Parameter: Cache/Operation Cache Maximum Age

Default Value: 600 seconds (equals 10 minutes)

Parameter: Cache/Operation Cache Maximum Size

Default Value: 10

Password Profile Cache

Each password profile cache item stores information from a password profile. Currently there is only one password profile per domain.

Parameter: Cache/Password Profile Cache Maximum Age

Default Value: 600 seconds (equals 10 minutes)

Parameter: Cache/Password Profile Cache Maximum Size

Default Value: 10

Connector Server Cache

Each C++ Connector Server cache item stores a pool of connections between the Provisioning Server and the C++ Connector Server. The C++ Connector Server, also known as the Connector Server service, is the component that loads each endpoint type option's agent module.

Some endpoint types (for example Active Directory) provide a feature to use the administrator's own credentials rather than a configured set of proxy credentials for authenticating to the managed directory. Each C++ Connector Server cache item represents a pool of LDAP connections using a single set of administrator credentials.

Parameter: Cache/Connector Server Cache Maximum Age

Default Value: 3600 seconds (equals 1 hour)

Parameter: Cache/Connector Server Cache Maximum Size

Default Value: 20

Compatibility Parameters

The Compatibility configuration folder contains parameters you can use to provide temporary backwards compatibility with prior releases of eTrust Admin.

Enable Operation Details

Values: Yes (default) or No

Description: Tracks changes to accounts affected by provisioning operations in etutil or the Provisioning Manager, such as Explore and Correlate. If you make no provisioning changes in the Provisioning Manager or etutil, you can set this parameter to no. Disabling it could keep the management of operation detail records from affecting CA Directory performance.

With or without this parameter, inbound synchronization still updates the Identity Manager server. Each provisioning server sends account details to the Identity Manager server or cluster of servers. Those details include attribute-level information not found in the operation details.

Relax Self Q&A Reads

Values: No (default) or Yes

Description: By default (No), global user self authentication Q&A attributes are returned only when explicitly asked for by users. This allows the Provisioning Server to log when these questions and answers are viewed. Some older clients depend on the prior behavior where these attributes could be retrieved along with all other Global User attributes. Set this parameter to Yes to re-instate the prior behavior to allow these applications to work with the current Provisioning Server.

Configuration Setup Parameters

The Configuration Setup configuration folder contains parameters you can use to configure the processing of these domain configuration parameters.

Default Value: 600 seconds (equals 10 minutes)

Description: Configuration parameters are read periodically from the provisioning directory. This parameter defines how often, in seconds, this refresh of parameters occurs. Hence, this parameter defines the maximum amount of time one would need to wait after making a change to parameter before being assured the change has taken effect.

The minimum value for this parameter is 30 seconds.

Connections Parameters

The Connections configuration folder contains parameters you can use to tune the connection management mechanisms within the Provisioning Server.

The Provisioning Server maintains pools of LDAP connections that it uses for communicating with the provisioning directory, with connector servers and with other LDAP servers. A dedicated thread within the provisioning server (the connection monitor thread) wakes up periodically to adjust the pools by closing excess idle connections and attempting to create connections to LDAP servers previously believed to be unavailable.

The configuration parameters in this folder are consulted by the connection monitor thread as it performs its functions.

Connections/CS Pool Maximum Size

Default Value: 200

Description: The maximum size of each of the Provisioning Server's CS Connection Pools. A CS Connection Pool is a reusable set of LDAP connections that are used to communicate with a specific connector server.

Connections/CS Pool Minimum Size

Default Value: 2

Description: The minimum size of each of the Provisioning Server's CS Connection Pools. The connection monitor thread, when it closes expired idle connections, will retain at least this many connections in each CS Connection Pool.

Connections/DB Pool Maximum Size

Default Value: 40

Description: The maximum size of the Provisioning Server's DB Connection Pool. The DB Connection Pool is a reusable set of LDAP connections that are used to communicate with the Provisioning Directory.

Connections/DB Pool Minimum Size

Default Value: 5

Description: The minimum size of the Provisioning Server's DB Connection Pool. The connection monitor thread, when it closes expired idle connections, will retain at least this many connections in DB Connection Pool.

Connections/Expiration Time

Default Value: 1800 seconds (30 minutes)

Description: The time, in seconds, after which an idle connection in the provisioning server's LDAP connection pools will be considered expired. An expired connection is a candidate for being closed by the connection monitor thread.

Connections/Other Pool Maximum Size

Default Value: 20

Description: The maximum size of each of the Provisioning Server's Ad Hoc Connection Pools. Each Ad Hoc Connection Pool is a reusable set of LDAP connections that are used to communicate with a specific LDAP server other than the provisioning directory or regularly used connector servers. For example, changes to endpoint or endpoint type attributes may need to be sent to another provisioning server's connector server, and the connection pool to communicate with that connector server is governed by this parameter.

Connections/Other Pool Minimum Size

Default Value: 0

Description: The minimum size of each of the provisioning server's Ad Hoc Connection Pools. This value is typically zero as there is rarely a need to retain idle connections to these LDAP servers past their normal expiration time.

Connections/Refresh Time

Default Value: 300 seconds (5 minutes)

Description: The time, in seconds, that the provisioning server's connection monitor thread waits between iterations. Each time this thread awakens, it identifies expired connections in its LDAP connection pools and closes them. It also attempts to establish LDAP connections to servers that were believed to be unavailable (but only for pools with a minimum size greater than zero).

Endpoint Parameters

The endpoint configuration folder contains parameters you can use to enable or disable features on a endpoint type-by-endpoint type or endpoint-by-endpoint basis. Each parameter can be set to an ordered list of values, each of which can be one of the following:

Parameter	Description
ALL	Enabled for all endpoints of all endpoint types.
-ALL	Disabled for all endpoints of all endpoint types.
<i>EndpointType</i>	Enabled for all endpoints of the specified endpoint type.
<i>-EndpointType</i>	Disabled for all endpoint of the specified endpoint type.
<i>EndpointType:Endpoint</i>	Enabled for the specified endpoint.
<i>-EndpointType:Endpoint</i>	Disabled for the specified endpoint.

If more than one value for the same parameter specifies the same directory, the last value that specifies the endpoint determines whether the feature is enabled or disabled for that endpoint. This lets you provide a more general rule first (enabled for all directories or all endpoint types) and follow that up with a more specific rule (disabled for endpoint ABC of endpoint type ActiveDirectory).

Endpoint/Check Account Passwords

Default Value: -ALL (disabled for all endpoints of all endpoint types)

Description: When this parameter is enabled for a specific endpoint, the Provisioning Server checks any password in a password change of an existing account on that directory, including attempts to set an empty password.

During account creation, the Provisioning Server performs password quality checking when a password is provided. If no password is provided, no checking is performed unless the Check Empty Account Passwords parameter is also enabled for the directory.

Account password quality checking uses the Password Profile that exists in the domain of the global user that owns the account. If the account is not associated with any global user, then the Password Profile that exists in the domain of the account is used. If the password profile located based on the global user or the account's domain is disabled, account password quality checking is also disabled for that account.

Account password quality checking does not include the checks on self-changes that depend on history of recent password-change activity. Password reuse frequency (history) and minimum time between changes (interval checking) are only applicable to global user password changes where the Provisioning Server retains an accurate history of recent changes. Account passwords and password history are not stored in the Provisioning Server. They are stored only in the managed directory and the Provisioning Server makes no assumption that all password changes are visible to the Provisioning Server.

A synchronized account password is an account password meeting the following criteria:

- Account is correlated to a non-restricted global user
- Account resides on a directory for which Disable password propagation to accounts has not been enabled
- Account has not been deleted (it is not in Delete Pending state)

An attempt to change a synchronized account password to the value of the current global user password will be accepted regardless of the setting of this configuration parameter. Also, the settings of the following configuration parameters can control the effect of password quality checking on synchronized account passwords:

Passwords\Enforce Synchronized Account Passwords

Passwords\Use External Password Policies

Endpoint/Check Empty Account Passwords

Default Value: -ALL (disabled for all endpoints of all endpoint types)

Description: When this parameter is enabled for a specific managed endpoint, the Provisioning Server checks any empty password in an Add request for account on that endpoint. This parameter is ignored if Check Account Passwords is not also enabled for this endpoint.

This parameter is separate from Check Account Passwords because it is acceptable in some endpoint types to create an account with no password.

Endpoint/Use Account Template Status

Default Value: -ALL (disabled for all endpoint of all endpoint types)

Description: Parameter controls whether to ignore a global user status of suspended and allow the creation of accounts in the active state. By default accounts created from account templates for a suspended global user are suspended regardless of the suspended status indicated in the account template.

Endpoint/Validate Endpoint Credentials

Default Value: -ALL (disabled for all endpoint of all endpoint types)

Description: Parameter controls whether the Provisioning Server sends changes to passwords, updated on endpoint property sheets, to the applicable connector for immediate validation or if only the provisioning directory is updated. This functionality is not applicable to CA-provided connectors as they have all been updated to always have the behavior that is controlled by this parameter.

If you have a custom connector written using the Software Development Toolkit from a previous release of eTrust Admin and that connector stores proxy credentials in its endpoint properties, verify the behavior of your endpoint type by enabling this parameter for your endpoints and attempting to change those proxy credentials.

Explore and Correlate Parameters

The Explore and Correlate configuration folder contains parameters you can use to configure the explore and correlate functions used while acquiring managed directories.

Explore and Correlate/Correlation Attribute

Value Syntax:

GUAttrName[=Namespace:AccountAttrName[:Offset,Length]]

Default Values:

GlobalUserName

FullName

Description: Controls the correlation matching algorithm used by the correlate phase of explore/correlate.

The correlation algorithm uses this parameter when determining how accounts are associated with global users.

Each value defines a global user attribute that will be compared against an account attribute. The list is ordered, and only values applicable for a endpoint type are used when correlating accounts from an endpoint of that endpoint type. If there are two defined mappings for the same global user attribute that are applicable to the endpoint type where correlate is being run, then the first parameter value is used.

You can provide this mapping in one of the following ways:

GUAttrName

In this form, you name only the global user attribute and not the corresponding account attribute. This value assumes for the omitted account attribute name the account attribute predefined by the endpoint type to correspond to this global user attribute. For information about the predefined mappings, see the endpoint type's *Connector Guide*.

A parameter value in this form applies to all endpoint types for which an account mapping is defined. All endpoint types define mappings for GlobalUserName (typically the account name). Most endpoint types define mappings for Full Name.

GUAttrName=Namespace:AccountAttrName

In this form, you name the global user attribute and a specific account attribute of a specific endpoint type. A parameter value in this form applies only to the indicated endpoint type. Use this form rather than the first form to match global users on an attribute such as Full Name in one endpoint type but not in all endpoint types.

GUAttrName=Namespace:AccountAttrName:Offset,Length

In this form, you name the global user attribute and a specified substring of an account attribute of a specific endpoint type. *Offset* indicates the start of the substring, the value 1 indicating the start of the attribute value. *Length* indicates the number of characters in the substring value. If the full account value is shorter than $(Length + Offset - 1)$ characters, the substring value used will be shorter than *Length* characters.

A parameter value in this form applies only to the indicated endpoint type. Use this form if you know that an account attribute value (for example, description) has a form where the first 8 characters are known to contain a unique employee identifier that can be matched to a global user attribute value.

For example, assume the configured parameter values are the following:

GlobalUserName
FullName=LDAP Namespace:globalFullName
FullName=ActiveDirectory:DisplayName
CustomField01=ActiveDirectory:Telephone

The following occurs for each previously uncorrelated account found while correlating accounts in an Active Directory container:

1. The Provisioning Server starts with the first parameter value (GlobalUserName) and determines that the Active Directory endpoint type's defined account attribute that maps to GlobalUserName is NT_AccountID (LDAP attribute name eTADsSAMAccountName). It attempts to find the unique global user whose name is equal to the account's NT_AccountID attribute value. If a unique match is found, the Provisioning Server associates the account with the global user. If more than one match is found, the Provisioning Server performs Step 5. If no match is found, the Provisioning Server performs the next step.
2. The Provisioning Server considers the second parameter value (FullName=LDAP Namespace:globalFullName). Since this value is specific to another endpoint type, it is skipped and the Provisioning Server performs the next step.
3. The Provisioning Server considers the third parameter value (FullName=ActiveDirectory:DisplayName). Since this value is specific to Active Directory, it is used. It attempts to find the unique global user whose FullName is equal to the account's DisplayName attribute value. If a unique match is found, the Provisioning Server associates the account with the global user. If more than one match is found, the Provisioning Server performs Step 5. If no match is found, the Provisioning Server performs step 4.

4. The Provisioning Server considers the final parameter value (CustomField01=ActiveDirectory:Telephone). Because this value is specific to Active Directory, it is used. It attempts to find the unique global user whose Custom Field #01 attribute is equal to the account's Telephone attribute value. Note that the name you gave to the custom global user attribute using global properties of the System Task is not displayed here. If a unique match is found, the Provisioning Server associates the account with the global user. If more than one match is found, the Provisioning Server performs Step 5. If no match is found, the Provisioning Server performs the next step.
5. The Provisioning Server associates the account with the [default user] object in the domain specified by the configuration parameter Explore and Correlate/Create Users Domain. If the [default user] object does not already exist, it is created.

Explore and Correlate/Correlation Domain

Values: Root Domain

Description: A value indicating which domain or domains should be searched for global users during the Correlate with existing global users phase of explore/correlate. This parameter is deprecated as the root domain is the only choice now.

Explore and Correlate/Create Users Default Attributes

Value Syntax: *GUAttrName=Value*.

Default Values: None

Description: The parameter provides default values for global user attributes for global users created during the Create Global Users phase of explore/correlate.

For example, use 'SelfAdministration=1' to enable self administration for your new global users. Use this feature to assign constant values to optional global user attributes for global users created during the acquisition of a primary directory.

Explore and Correlate/Create Users Domain

Values: Root Domain

Description: The domain in which global users are to be created during the Create Global Users phase of explore/correlate. A value indicating which domain or domains should be searched for global users during the Correlate with existing global users phase of explore/correlate. This parameter is deprecated as the root domain is the only choice now.

Explore and Correlate/Create Users Verify Not Correlated

Values: Yes (default) or No

Description: Temporarily set this parameter to No to enable the alternate behavior whereby the Create Global Users phase of explore/correlate will skip the check that the account is not already correlated to a global user.

This capability is deprecated since exploring primary endpoints no longer applies.

The Create Global Users function works as follows for each account present in the container being correlated.

- Check to see if the account is already correlated to an existing global user. If so, leave this account still correlated to that global user. On an initial acquisition no accounts will be correlated. However, on a later re-explore/recorrelate this step is important so that the accounts remain correlated to the global user to which they were previously correlated.

Note: In a primary endpoint, you would not expect accounts to be correlated to any global user other than the one named the same as the account. There are various scenarios where this could occur. For instance, you may have renamed the account or the global user at some point after they were correlated to one another. Or you might have some system accounts on your primary endpoint that you do not want to correlate to separate global users - opting instead to correlate them to a single restricted global user.

- Attempt to create a global user named the same as the account. If this global user already exists, go on to the next step. This can happen if the global user was also present in another primary endpoint or you deleted your primary endpoint and re-acquired, correlating to global users created during the prior acquisition.
- Create an inclusion between the account and the global user, correlating the account with the global user.

If this configuration parameter is set to No, the first step is skipped. This can greatly improve the performance since that test is time-consuming and slows down as the Provisioning Directory becomes untuned. If you are acquiring an endpoint with more than thousands of accounts, the Provisioning Directory needs to be tuned, but you do not have the chance to do that tuning in the middle of the long-running Correlate operation.

Important! Set this parameter to No only during the initial acquisition of your primary endpoints. Subsequent use can result in accounts incorrectly correlated to multiple global users. Once the acquisition is completed, set the parameter to Yes.

Explore and Correlate/Map User ID to Lowercase

Values: No (default) or Yes

Description: Map all user IDs to lowercase when creating global users during the Create Users phase of explore/correlate. When you acquire one of your primary directories, you create global users for the accounts on that directory. Two of the global user properties that get set by this creation of global users are the following:

Account Name (LDAP attribute name eTUserId): Set to the corresponding account's account name property.

Global User Name (LDAP attribute name eTGlobalUserName): Set to be the same as the Account Name property, but translated to lowercase.

If you acquire a primary directory with mixed case account names, this will by default result in the created global user's Account Name property also having mixed case. Set the configuration parameter to Yes to force the Account Name property to be the same as the Global User Name property - always in lowercase.

Preserve the original case in the Account Name property by leaving the configuration parameter set to No if you have no endpoint types such as UNIX for which account names are case-sensitive.

An alternative to setting this parameter to Yes is to define your case-sensitive endpoint types' account template rule expressions for account name to use the TOLOWER built-in rule function, `$$$TOLOWER(%AC%)`, instead of the normal account name rule expression, which is `%AC%`.

Note: If your primary directory has only uppercase account names, this configuration parameter has no effect. The global user's Account Name property will already be translated to lowercase.

Explore and Correlate/Explore Compare in Memory

Values: No (default) or Yes

Description: Obtains two lists of objects at a time: one from the endpoint system being explored and one from the provisioning directory. These lists are compared in memory to determine what changes should be applied to the provisioning directory. If this parameter is no only a single list of objects at a time is explored and correlated, which uses far less virtual memory when working with large lists of objects.

Identity Manager Server Parameters

The Identity Manager Server configuration folder contains parameters you can use to control interactions between the Provisioning Server and the Identity Manager Server. Before enabling any of these parameters, you should verify the configuration of the communication between the Provisioning Server and the Identity Manager using the “Identity Manager Setup” button on the System task of Provisioning Manager,

Identity Manager Server/Enable Corporate User Access

Values: No (default) or Yes

Description: Enables/disables retrieval of corporate user attributes from the Identity Manager Server during account template evaluation.

Important! This feature was not available at the publication time for this document. Please check the availability of the feature in the release notes before enabling this parameter.

Identity Manager Server/Enable Notification

Values: No (default) or Yes

Description: Enables/disables the collection of audit data (notifications) by the Provisioning Server for transmission to the Identity Manager Server. When enabled, any changes to data managed by the PS, other than changes directly initiated by the Identity Manager Server, generate notifications which are queued in the Notification DSA and then later sent to the Identity Manager Server. Upon receipt at the Identity Manager Server, certain notifications trigger events, while most are simply added to the full Identity Manager audit data.

Identity Manager Server/Notify Batch Size

Default Value: 100

Description: The number of notifications that are processed in one batch. When sending notifications to the Identity Manager Server, the Provisioning Server will retrieve at most this many records (a batch) from the Notification DSA, process those entries, and then continue with additional batches.

Identity Manager Server/Notify Retry Time

Default Value: 600 seconds (10 minutes)

Description: The time, in seconds, that the notification thread pauses between iterations. The notification thread is a dedicated thread within the Provisioning Server that wakes up periodically and attempts to transmit (or retransmit) any queued notifications.

Identity Manager Server/Notify Timeout

Default Value: 30 seconds

Description: The timeout value, in seconds, for sending notifications or password validations to the Identity Manager Server. A value of zero indicates an unlimited timeout.

Identity Manager Server/Use External Password Policies

Description: When set to Yes, users changing their own global user passwords or one of their synchronized account passwords will have the password validated using externally-defined password rules. Users' synchronized account passwords are the passwords for their accounts on endpoints for which the Disable Password Propagation property is disabled. You should set the parameter Enforce Synchronized Account Passwords to Yes whenever Identity Manager Server/Use External Password Policies is set to Yes. When this parameter is set to Yes, the Provisioning Server password rules that are applicable to users changing their own passwords (Password history checks and Minimum interval between self-changes) are no longer consulted.

Values: No (default) or Yes

Note: Even when integration with Identity Manager password policies is enabled with this configuration parameter, the Provisioning Server uses its per-domain password profiles in various situations. In particular, Administrative password changes, initial global user passwords, changes to unsynchronized account passwords and generating temporary initial passwords all consult the Provisioning Server password profile. In addition, the Locking and Password Expiration features defined in the Provisioning Server password profile are always used. However, the Provisioning Server password profile rules that are applicable to users changing their own passwords (Password history checks and Minimum interval between self-changes) are not consulted when this configuration parameter is Yes.

Operation Details Parameters

The Operation Details configuration folder contains parameters you can use to control the behavior of operation details. Operation Details is the function that tracks the status of child operation spawned from higher-level operations such as Explore, Synchronization or Propagation. When you perform one of these higher-level operations from Identity Manager tasks or from Provisioning Manager, you receive a message in the message summarizing the results of the child operation.

The following is a sample summary message for a User Synchronization request:

(accounts created: 1, updated: 1, recreated: 0, failures: 0)

If you ask to view status details for the task (or double-click on the icon next to the summary message when using Provisioning Manager), this displays a screen with operation details including a series of success, failure, or warning messages corresponding to the statistics present in the summary message.

Operation Details/Maximum Operation Detail

Default Value: 100

Description: The maximum number of operation detail items which can be retrieved in one search of an operation object. When you perform a high-level operation that spawns hundreds or thousands of child operations and you call up the Operation Status window, this parameter controls how the details are returned from the Provisioning Server to the Provisioning Manager or other client application.

Operation Details/Operation Details Expiration Time

Default Value: 96 hours (equals 4 days)

Description: The number of hours to keep operation details in the provisioning directory.

Operation details are maintained in the server in the following parts:

1. An operation object stored in the provisioning directory (one per high-level operation).
2. An XML data file stored in the Operations folder containing the operation details, concatenated one after another.

Both objects are deleted when the operation object is deleted. Some clients delete their operation objects as soon as they retrieve the operation details or when the client terminates. Other clients such as Provisioning Manager leave the operation objects in the directory until they expire and are deleted in four days (by default).

Operation Details/Operations Folder

Default Value: Operations

Description: The name of the folder on the Provisioning Server where the XML data files storing operation details reside. This value can be a simple filename or a relative path name. However, it may not be an absolute path name.

Its value is relative to one of the following file path names:

%ETAVARHOME%
PSHOME
..

Normally, this means that the operations folder is placed alongside the Data and Logs folder with a path name like the following:

```
C:\Program Files\CA\Identity Manager\Provisioning Server\Operations
```

However, to relocate this folder to another drive (so as to be able to run from a read-only drive), you should set the environment variable %ETAVARHOME% to a value such as

D:\ProvisioningData before restarting the Provisioning Server service. Then the operations XML files will be placed instead into the following folder:

```
D:\ProvisioningData\Operations
```

The ETAVARHOME value can also be set as a registry value instead of an environment variable by using the eta-env utility that is installed with the provisioning server:

```
eta-env action=set name="ETAVARHOME" value="D:\ProvisioningData"
```

Important! Changes to this parameter do not take effect until the Provisioning Server service is restarted.

Password Synchronization Parameters

The Password Synchronization configuration folder contains parameters you can use to control the behavior of password synchronization operations. Password synchronization is the feature that involves installing the Password Synchronization Agent on a Windows system or other systems to intercept password changes, send password validation requests, and password notification requests to the Provisioning Server.

Password Synchronization/Agent Response Threshold

Default Value: 600 seconds (equals 10 minutes)

Description: Maximum expected duration (in seconds) of each password change that the Provisioning Server sends to a managed endpoint on which a password synchronization agent is installed. This parameter allows the Provisioning Server to recognize when a Password Synchronization agent is processing a password change sent to it by the Provisioning Server as distinct from a password change originating on that managed endpoint.

When installing a password synchronization agent, you must check that the Password synchronization agent is installed check box on the Endpoint Settings tab. Then when the Provisioning Server sends a password change to the managed endpoint, it records the time when the password was sent. For a number of seconds set by the Agent Response Threshold, any password change notification or password validation request received for this account is assumed to be false. Only password changes originating on the native system initiate password synchronization. Account password changes originating in the Provisioning Server update the account but not the global user or other accounts.

If, during the Agent Response Threshold, a password other than the password just sent to the managed endpoint is provided in a password validation or password change notification, this password is rejected. Two concurrent password changes to the same account are not allowed.

Password Synchronization/Update Only Global User

Values: No (default) or Yes

Description: This parameter controls what action is carried out when the Provisioning Server receives a password change notification. By default, the new endpoint account password received in a password change notification is used first to update the global user's password and then to update all of that global user's account passwords for accounts other than the one from which the notification arrived.

Set this parameter to Yes to change this behavior so that only the global user password is updated. No account passwords will then be updated.

There are various situations in which the global user and affected accounts are not updated, including the following:

- Global user Enable Password Synchronization Agent property is not enabled. Global Users and account passwords are not updated.
- Password change notification occurred during the Agent Response Threshold period and is treated as a false password change notification. Global Users and account passwords are not updated.
- A endpoint containing one of the global user's accounts is marked on its Endpoint Properties for Disable propagation to accounts. The accounts on this endpoint are not updated.
- Global user Restricted property is enabled. Restricted global users such as [default user] are protected from accidentally propagating changes to their associated accounts.

Password Parameters

The Password configuration folder contains parameters you can use to control the behavior of certain password operations.

Passwords/Enforce Synchronized Account Passwords

Values: Yes (default) or No

Description: When Yes, users cannot change any of their synchronized account passwords to a value other than the current value of their global user password. We recommend that you set this parameter to Yes whenever the Identity Manager Server\Use External Password Policies parameter is set to Yes.

Users' synchronized account passwords are the passwords for their accounts on endpoints for which the Disable Password Propagation property is disabled and which have not been marked as Delete Pending.

Passwords/Pre-expire Passwords

Values: No (default) or Yes

Description: Controls having new global users created with their passwords already expired, forcing users to change their passwords during the initial login. If you set this parameter to Yes, global users created from non-interactive interfaces have their password initially set as expired. This is represented in the global user properties as a value of 1 for the property PwdPreExpired. This option appears on the global user's property sheet as the Force one-time expiration (mark password as temporary) check box.

The setting of the password as initially expired occurs when global users are created through the following interfaces:

- Correlate. When acquiring a primary endpoint, global users are created for each account. These users will generally not have a password unless you set a constant value using the Create Users Default Attributes parameter described previously. Enabling the Pre-expire Passwords parameter will cause the global users to be created with passwords that are initially expired. If you set a value for PwdPreExpired using the Create Users Default Attributes parameter, that value takes precedence over one specified by enabling the Pre-expire Passwords parameter.
- Identity Manager Server, ETAUTIL or other on-demand clients. If you create a single global user using the batch utility (ETAUTIL) or some other on demand LDAP client, these users will start out with expired passwords if you enable the Pre-expire Passwords parameter and do not otherwise specify a value of the PwdPreExpired property.

If you create a global user using an interactive client such as Provisioning Manager, whether the global user's password is initially expired or not is determined from the value of the PwdPreExpired property provided when the global user is created. In Provisioning Manager, you control this value by selecting or clearing the check box labeled Force immediate expiration (one-time). Provisioning Manager automatically selects the Force one-time expiration (mark password as temporary) field if the Pre-expire passwords parameter is enabled. To disable this default behavior, clear the field before creating the global user.

Passwords/Store User Passwords

Values: Yes (default) or No

Description: Controls whether the EncryptedPassword global user attribute is stored and whether %P% rule variables are supported.

By default the Provisioning Server encrypts the global user password and stores it in the provisioning directory as a global user attribute named EncryptedPassword. When you later attempt to create an account for that global user using an account template with the %P% expression for the password rule, then the Provisioning Server decrypts the stored EncryptedPassword value and provides it to the endpoint type option as the initial Password attribute for the account being created.

However, if you will not be creating any accounts using account templates with %P% rule expressions, then you can improve security by not storing these passwords.

Note: By not storing the EncryptedPassword attribute, you are only giving up %P% rule evaluation. You can authenticate users by using the global user password. When the Store User Passwords parameter is set to No, the Provisioning Server stores a one-way hash of the password for use in authenticating user passwords during login.

Processes Parameters

The Processes configuration folder contains parameters you can use to control the process behaviors on Windows provisioning servers.

Processes/Catch Program Exit Exceptions

Default Values: Yes (default) and No

Description: This parameter controls the behavior of the Provisioning Server when invoked program exits throw runtime exceptions. By default (yes), the exception is caught and the current operation fails with an uncaught exception error message. However, if you are developing new program exits you may choose to set this parameter to no and allow the uncaught exception to result in server termination, which provides more information about the exception. This parameter only affects common program exits of the DLL type.

Processes/Child Operation Thread Pool Size

Default Value: 200

Description: This parameter defines the maximum number of threads in the server-wide child operation thread pool. When the server decides to split up a single operation into multiple sub-operations, those sub-operations are carried out by the threads in the child operation thread pool. The larger you make the value for this parameter, the more work the Provisioning Server attempts in parallel.

Currently the Provisioning Server only uses the child operation thread pool to carry out the multi-account search and multi-account update functions submitted by the Web interface. For synchronization and propagation operation, regardless of which client submits these requests, and even though they also spawn child operations, the child operations are carried out in series in the main operation's thread.

This parameter does not affect the primary server thread pool used for processing separate requests received from client applications. This thread pool size is controlled by the SLAPD parameter called `threads` in the `eta_slapd.conf` file. See on the Provisioning Manager help for editing parameters in this configuration file.

Important! Changes to this parameter do not take effect until the Provisioning Server service is restarted.

Processes/Parallel Propagation

Default Values: Yes (default) and No

Description: This parameter controls whether account passwords updated as part of global user to account password propagation are carried out in parallel or sequentially. By default, account passwords are updated in parallel, and the degree or parallelization is controlled by the Processes/Child Operation Thread Pool Size parameter.

This parameter has no effect on requests from clients that carry out the account updates explicitly. It only affects those clients that direct the Provisioning Server to update a global user password and propagate that change immediately to all synchronized account passwords.

Processor Parameters

The Processes configuration folder contains parameters you can use to control values related to the use of CPU processors. These parameters control operating system values that are process-wide and as such affect the entire Provisioning Server service. This is specifically relevant if you install any additional backends into the slapd process that runs your Provisioning Server.

Important! Changes to these parameters do not take effect until the Provisioning Server service is restarted.

Processor/Process Affinity Mask

Default Value: 0 (no restrictions)

Description: Specifies the process affinity mask for the Provisioning Server service process. The process affinity mask is a bit vector in which each bit represents the processor of a multiprocessor server on which the threads of the process are allowed to run.

The value 0 (default) signifies no restrictions.

The values 1, 2, 4, 8, 16, 32, 64, or 128 restrict the threads to running on the 1st, 2nd, 3rd, 4th, 5th, 6th, 7th, or 8th processor, respectively.

For example, the values can be combined so the value 5 (1+4) can be used to allow running on processors 1 and 3.

Processor/Process Priority

Default Value: 0 (use system default priority)

Description: Specifies the scheduling priority of the Provisioning Server service process.

The value 0 (default) uses the system default.

The only other recommended value is 16384, indicating below-normal priority. This value should be used when the Provisioning Server runs on the same server as its provisioning directory. This effectively raises the priority of the provisioning directory and consequently increases over-all server performance.

Search Parameters

The Search configuration folder contains parameters you can use to control search behaviors.

Search/Allow Partial Results

Values: Yes (default) or No

Description: Allow search requests to return less than the full number of matched entries when the search size limit is reached. If you set this parameter to No, or if the client fails to request partial results, partial results are not returned and clients receive a size limit exceeded error when the size limit is reached.

For example, if partial results are permitted, the search limit is 200, and the search found 5000 entries, 200 of them are returned. If partial results are not permitted, this search would return *no* results.

Note: The Provisioning Server does not define *which* 200 of the 5000 entries are returned. The Provisioning Server does not necessarily return the *first* 200 entries, either alphabetically or using any other ordering method.

Two settings are required to activate Partial Results:

- The domain configuration parameter Allow Partial Results described here.
- A selection in the Provisioning Manager GUI Search Preference control.

The effective partial result setting is the combination of these two settings. A partial result is returned only if Allow Partial Results is set to Yes *and* the GUI Search Preference control has Show Partial Result Lists selected.

The Provisioning Server is most efficient when partial results are not returned. When partial results are not needed, the Provisioning Server can report quickly when a search would exceed the size limit. However, if required to return the number of entries indicated by size limit, this will add processing load to the provisioning directory, Provisioning Server and client application. This load will take away from the processing load available to other users' queries or modifications.

Search/Max Scope Filter Objects

Default Value: 10

Description: The maximum number of objects that will be placed into a search filter during scoped searches. When a search of a container is initiated by an administrator who has access to only some of the objects in the container, the Provisioning Server augments the client-supplied filter with a scope filter that restricts the objects returned to those for which the administrator has access. However, if the administrator has access to more than Max Scope Filter Objects objects, this is deemed too many to be placed into the filter that will be sent to the provisioning directory and/or managed endpoint and the Provisioning Server uses its backup algorithm. In this case, the Provisioning Server will ask for all objects the client requested and then discard those that the scope filter would have excluded.

Search/Search Size Limit

Default Value: 0 (unlimited)

Description: The maximum number of entries returned by the Provisioning Server in a search request. The effective size limit for a search is the smallest of this value, the SLAPD size limit parameter, and the client-provided size limit operation parameter.

Limiting the number of entries returned in searches is important for good interactive performance of the Provisioning Server. If the effective size limit is too large, poorly formed search requests may return very long lists of global users, accounts or other objects that are not easy for administrators to work with. Conversely, if the effective size limit is too small, it may limit the administrator's ability to browse provisioning roles or other objects whose number is more moderate than that of global users or accounts.

There are three ways to control the effective size limit of a search request:

- The SLAPD *sizelimit* parameter in the following files:

PSHOMEIDatavim_ps.conf

PSHOMEIDatavim_ccs.conf

This value is set to 0 (unlimited) by default and controls all LDAP servers running in the Provisioning and Provisioning Connector server services, respectively. There should be no need to change this parameter. Doing so would limit operations like exploration that perform searches of accounts in managed directories and relies on being able to receive all accounts present in a single container in a single search request.

- The Domain Configuration parameter Search Size Limit

This value is also set to 0 (unlimited) by default. It controls the maximum number of entries returned to provisioning clients in a single search request. If you set this to a non-zero value (500), this will prevent any client from being able to receive more than 500 entries from the Provisioning Server in a search request.

Set this parameter only if you can be sure that no clients require receiving more than this number of entries from any search.

You should leave the Search Size Limit configuration parameter as 0 (unlimited) on any interactive or mixed-use Provisioning Server. If you have an environment where certain Provisioning Servers are dedicated to interactive use and other Provisioning Servers are available for batch activity, you may want to set the Search Size Limit between 500 and 1000 on the interactive servers only. Use the Add Specialization menu item to set a server-specific Domain Configuration parameter.

- The Provisioning Manager's size limit preference setting

To change this preference, select File, Preferences, click the Search tab, and change the value in the Limit on Returned Items for a Search field.

The Provisioning Manager preset return limit is 500. Each user can increase or decrease this preference, but increasing the value above the server's Search Size Limit value has no effect because the server's effective size limit is the smallest of the three size limit controls.

If a search operation encounters a search-limit failure, assuming that you enabled the retrieval of partial results, in some cases the number of items displayed may be different than the actual search limit because a single search operation (from the perspective of the user) might require several searches (from the perspective of the Provisioning Server).

Depending on how a client combines the results of the multiple search operations (for example, through a union or intersection) when displaying the results, the net display may contain more or fewer items than the search-limit. In all cases, an error message is displayed informing you that the results were truncated due to the search limits.

Servers Parameters

The Servers configuration folder contains read-only parameters that identify which Provisioning Servers are installed for your domain. After a default installation, a single server is listed.

If you install alternative servers for failover or load-balancing, multiple servers appear in this list. For each server listed, read-only parameters identify the Build, Patch and Version numbers for the Provisioning Server software. An additional parameter identifies whether the FIPS 140-2 encryption feature is enabled for that Provisioning Server.

Note: The server names listed here are the same server names used when creating specialization parameters. These are also the names you should use for the server parameter in the csfconfig command-line utility when creating specializations for connector server configuration objects.

Statistics Parameters

The Statistics configuration folder contains parameters you can use to control how statistics are maintained by the Provisioning Server. Most objects stored in the provisioning directory have statistic attributes to record when and by whom the object was created; and when and by whom the object was last updated. These statistics are displayed on the Statistics tab of the respective objects' properties.

Statistics/Enabled

Values: Yes (default) or No

Description: When disabled, statistic attributes on objects in the provisioning directory are not updated by the Provisioning Server as those objects are created or updated. This can improve performance during large scale changes or in installations where maintaining creation and update statistics is not necessary.

Certain statistics on global users such as password update date and time, and suspension update date and time are required for correct operation of server functions. These statistics are updated even when the Statistics/Enabled configuration parameter is set to No.

Statistics/Node Stats from Connection

Values: No (default) or Yes

Description: Use the client node name taken from the LDAP connection object when recording node statistics. The default (No) behavior is to take the node name provided by the client application.

The Node statistic displayed on the Statistics tab of the Global User property page and other objects' property pages is not always updated when other statistics such as date, time, userid, and username are. This behavior is a result of the way the Node name is determined. By default, the node name must be provided by client applications in their requests. If the clients fail to do so, or the client submits a high-level operation, such as Synchronize, that spawns child operations to carry out the individual object updates, the server has no Node value to use to update that statistic. To rectify this problem, use the Node Stats From Connection configuration parameter. Change its setting from the default No to Yes to select the alternative Node statistic algorithm.

The alternate algorithm uses the Node information obtained from the LDAP connection, which identifies the host that was the immediate client sending the request. This is often the same as the originating client, but can be another system.

For example, requests originating from the Identity Manager Server would be recorded as originating from the computer where the Identity Manager Server is running, not the computer where the administrator is running a web browser. Also, if the clients connect to a dXRouter process for failover or load-balancing between replicated Provisioning Servers and have the dXRouter send the request to the Provisioning Server, you should not enable the Node Stats from Connection parameter. It will result in all Node statistics indicating the router system.

Synchronization Parameters

The Synchronization configuration folder contains parameters you can use to choose from alternative variants of the Provisioning Server's synchronization functions.

Synchronization/Automatic Correlation

Values: No (default) or Yes

Description: Enable the alternative User Synchronization behavior whereby an attempt to update an existing, uncorrelated account triggers an automatic correlation of the account to the global user prior to the update of the account. If the parameter is No (default), the attempt to update the account will fail with a message indicating the account has not yet been correlated to this global user.

Synchronization/Remove Account Template Values from Accounts

Values: Yes (default) or No

Description: When Yes, the Weak Synchronization algorithm will consider that capability account values (for example, account group membership) prescribed by a account template should be removed when that account template is removed from an account. Set this parameter to No to restore the prior Weak Synchronization behavior where account attribute values are never removed when synchronizing an account with its weak-synchronization account templates. This parameter only affects multivalued attributes. String, integer or Boolean single-valued attributes are only increased in capability by weak synchronization.

Only certain multivalued attributes designated as SyncRemoveValues attributes are affected by this feature. Consult the eTaCapability.txt file for a list of which multivalued capability attributes may have values removed by the SyncRemoveValues feature described here.

To generate the eTaCapability.txt file, use the following command:

```
PSHOME\bin\dumpptt -c > eTaCapability.txt
```

Synchronization/Use Existing Accounts

Values: No (default) or Yes

Description: Enable the alternative User Synchronization behavior whereby a global user's set of assigned account templates (through assigned provisioning roles) will only attempt to prescribe one account correlated to the global user on any particular managed endpoint. This behavior can be useful if some accounts already correlated to the global user are named differently or are in different containers than what is prescribed by the account templates included in the global user's provisioning roles and only one account is needed or allowed. If the parameter is enabled and multiple account templates for one endpoint prescribe different names and/or different containers for the account only one account will be created.

If a global user already has multiple accounts on a single endpoint, the User Synchronization function (when Use Existing Accounts is set to Yes) attempts to figure out which account is required by which account template. This is done through a heuristic that attempts to handle situations where a user's provisioning roles do in fact prescribe multiple accounts on one endpoint.

For example, if global users have two accounts (A1 and A2) on endpoint E and their provisioning roles indicate that they should have one account on endpoint E through account template AT1 and one account on endpoint E through account template AT2, User Synchronization pairs each account template (AT1 and AT2) with one of the existing accounts. The pairing is done with the following heuristic:

- Match account template with an account with exactly the DN specified by the account template.
- Match account template with an account already belonging to the specified account template. If more than one account matches, pick the first one.

- Match account template with an account whose endpoint type-specific account name attribute matches the global user's name. In some endpoint types, for instance Active endpoint, the account name is represented by an attribute of the account whereas the name as seen when you list the account is a display name (a full name). This rule accounts for such endpoint types.
- Account with name value matching the name value specified in the account template. That is, it matches an account with the right name but the wrong container. If there is more than one matching account, pick the first one.
- Pick the first account.

Note: When Use Existing Accounts parameter is set to No, only the first of these rules (exact matching based on account DN) is applied.

Continuing with the example, if the previous rules resulted in pairing both account template AT1 and account template AT2 with account A1, then User Synchronization would correct the accounts for this user by doing the following:

Deleting account A2 (assuming the administrator selected the Delete extra accounts or extra account templateaccount template assignments option of User Synchronization); and

Assigning either account template AT1 or AT2 to account A1 that was not already assigned.

These rules ensure that User Synchronization (with Use Existing Accounts enabled) never attempts to create additional accounts on an endpoint where a user already has an account. If your business requires you to create multiple accounts for your users on a single endpoint from provisioning roles, do not enable this configuration parameter. For more information about synchronization, see the Administration Guide.

Transaction Log Parameters

The Transaction Log configuration folder contains parameters you can use to control transaction logging, also known as Provisioning Server trace logging. This is the log you use to monitor activity performed by the Provisioning Server while it processes requests received from its client applications.

Transaction Log/Enable

Values: Yes (default) or No

Description: Set this parameter to No to completely disable the logging of information to the server trace log. Typically, you control the amount of information you want logged using the Level parameter. However, even at level 0 some important items are logged to the server trace log. To disable these items from being logged, set the Enable parameter to No.

Transaction Log/Enable/Configuration

Values: Yes (default) or No

Description: This parameter enables or disables logging of diagnostic output from the Provisioning Server Configuration subsystem. The configuration subsystem checks every 10 minutes (by default) to see whether any of the configuration parameters have been changed. Each time this periodic refresh occurs, a line such as the following is written to the server trace log: ETA::Configuration update completed. No changes found.

Alternative messages are written if actual changes were found. To suppress all of these messages from the server trace log, set this configuration parameter to No.

Transaction Log/Enable/Connector Server Framework

Values: Yes (default) or No

Description: Enables/disables logging of the diagnostic output from the Provisioning Connector Server Framework (CSF).

Transaction Log/Enable/LDAP

Values: Yes (default) or No

Description: This parameter enables or disables logging of diagnostic output from the Provisioning Server LDAP subsystem. The LDAP subsystem manages the communications between each Provisioning Server and other LDAP servers, including the provisioning directory and Connector Servers

Transaction Log/File Name

Default Value: etatrans

Description: This parameter defines the transaction log's base file name. The suffix *YYYYMMDD-HHMM.log* will be appended to this base file name to build the log file name. You can change this parameter to use a different base file name in the *PSHOME*\Logs folder or to relocate the log file to another folder on your server.

The value of this parameter can be any of the following:

- Simple base name (for example, etatrans), the log is created in the Logs folder in the folder where you installed the Provisioning Server. By default, this makes the log file named `C:\Program Files\CA\Identity Manager\Provisioning Server\Logs\etatransYYYYMMDD-HHMM.log`.
- Relative path (for example, `..\Logs\etatrans`), the log path name will be relative to the current directory of the Provisioning Server service (`PSHOME\bin`). For the example given, this will result in the same pathname as before (`C:\Program Files\CA\Identity Manager\Provisioning Server\Logs\etatransYYYYMMDD-HHMM.log`).
- Absolute path (for example, `D:\ProvisioningData\Logs\etatrans`), you can specify an alternative drive for your log file. For this example, the resulting log file would be `D:\ProvisioningData\Logs\etatransYYYYMMDD-HHMM.log`.

The Provisioning Server switches to a new log file every day, every time the Provisioning Server restarts, and any time the log file size exceeds 100 Megabytes.

Transaction Log/Level

Values: 0 through 7 (default)

Description: This parameter lets you set the level of logging for the Server Trace log. Valid values are:

Value	Description
0	No trans logging
1	Log external/child errors
2	Log external operations
3	Log child operations
4	Log informative messages
5	Log DSA (Directory Service Agent) errors

Value	Description
6	Log DSA operations
7	Log search operations

Note: Alert log entries are logged at all logging levels (1 - 7).

After installation, the log level is set to the maximum value (7). This ensures that any problems during or immediately after installation are logged. After installation, you may select alternative logging levels to meet your logging requirements. Many customers run with level 7 for maximum information in the event that problems are reported by users. Other customers select a more modest level such as level 3 that reports failures without much of the internal tracing information associated with the processing of requests. Another useful level is level 6 that removes the many search operations that could dominate the log while maintaining all other information.

Log user-friendly Attribute and Object Class Names

The Identity Manager Provisioning Server currently logs attribute values in its server trace log (etatransYYYYMMDD-hhmm.log) as it logs the attributes in Add and Modify operations, listing the LDAP attribute names and the LDAP objectClass values. For DYN connectors, the LDAP attributes and object classes are generic names (such as eTDYN-str-multi-01, eTDYNObject001) which are not that meaningful. For release 12.5, these log entries are expanded to list the LDAP attribute names and object class values, and the the user-friendly names taken from the metadata.

Chapter 3: SPML Service

This section contains the following topics:

[SPML Overview](#) (see page 73)

[Install SPML](#) (see page 79)

[SPML Support for FIPS 140-2](#) (see page 80)

[Uninstall the SPML Service](#) (see page 81)

[SPML Service Configuration](#) (see page 81)

SPML Overview

The Provisioning Server helps you manage, provision, and de-provision entities. A good provisioning system is vital for security and efficiency. Many companies have multiple provisioning systems. It can often be difficult to configure different provisioning systems to communicate with each other.

OASIS (Organization for the Advancement of Structured Information Standards) has developed a markup language specifically designed to facilitate communications between and within user provisioning software. This is named SPML (Service Provisioning Markup Language).

SPML is an open standard that provides an XML-based protocol for provisioning requests. It facilitates provisioning requests between clients and servers that can be both intranet and extranet.

Benefits of Using SPML

The benefits of the SPML include the following:

- SPML is an open standard and can therefore communicate with other provisioning systems that can process SPML Requests. This lets businesses continue to use and integrate existing systems.
- Data can be shared across different provisioning systems to leverage the best features of each system.
- SPML is especially designed to handle provisioning-related data.
- SPML can easily handle data driven assignments of role-based access control.
- SPML is a best-of-breed technology for user provisioning.
- SPML facilitates business-to-business communications, where appropriate.

- SPML XML requests and responses are more human-readable than LDAP. Requests which is the native language of the Provisioning Server.
- SPML is a web-based technology.

When You Would Use the SPML Service

SPML simplifies provisioning requests and facilitates communication between provisioning systems. You do not need to deploy the SPML Service for a basic Provisioning Server installation, but it is an efficient and elegant provisioning solution.

You would deploy the SPML Service and the associated SPML Configuration Application as part of an installation or upgrade of the Provisioning Server. You can then deploy the SPML clients and tools that come with the Provisioning Server (CMDRA, SPML Manager, and WS-Mapper). You can also write or integrate a third-party clients and requesting authorities if they support the SPML version 1.0 standard.


SPML Architecture

The SPML Service is the server-side component that processes SPML requests. The SPML Service is a Provisioning Server component that uses and processes SPML requests. The SPML Service uses SPML version 1.0.

This section describes the components that make up the SPML architecture of the Provisioning Server.

SPML Architectural Diagram

The following diagram shows the SPML components of the Provisioning Server and how they relate to each other.



SPML Service

The SPML Service is the server-side component that processes SPML requests. The SPML Service is a Provisioning Server component that uses and processes SPML requests. The SPML Service uses SPML version 1.0.

SPML Configuration Application

The SPML Configuration Application is a web-based interface that lets you configure one or more Provisioning Servers as unique instances of SPML services. You should use the SPML Service Configuration tool to configure the SPML Service.

The SPML Configuration Application is automatically installed when you install the SPML service.

To access the SPML Configuration Application, Start Menu, Programs, CA, Identity Manager, IM SPML Service Configuration.

Command Line Requesting Authority (CMDRA)

The Command Line Requesting Authority (CMDRA) is a sample SPML Requesting Authority that can submit well-formed SPML Request XML files to the SPML Web Service. The CMDRA lets advanced users submit SPML requests using the command line or from scripts. It is ideal for sorting and managing large amounts of data using SPML templating, as well as automating requests and large batch jobs.

To download the CMDRA

1. Click Start Menu, Programs, CA, Identity Manager, IM SPML Requesting Authority
2. Click the cmdra.zip link.
3. Unzip the CMDRA to your hard disk to use it.

SPML Manager

The SPML Manager is a graphical user interface that lets administrators create and execute SPML provisioning requests. The SPML Manager can also help advanced users integrate with other provisioning systems. You can design provisioning requests in the SPML Manager then view the SPML requests in its native XML format.


To download the SPML Manager

1. Click Start Menu, Programs, CA, Identity Manager, IM SPML Requesting Authority
2. Click the SPML Manager web link.
3. Unzip the SPML Manager to your hard disk to use it.

How the SPML Service Works

This section explains how an SPML request works from the Requesting Authority through to the provisioning server.

1. Using the SPML Manager, an administrator creates a new user.
2. The SPML Manager creates an SPML-specific XML form and sends it to the SPML service.
3. The SPML Service translates the request to LDAP and passes the LDAP request to the Provisioning Server.
4. The Provisioning Server processes the LDAP request.
5. The Provisioning Server sends confirmation to the SPML Service.
6. The SPML Service sends confirmation to the SPML Manager.



SPML Integration

This section gives you an outline of the following:

- SPML Templating to integrate large amounts of data from external provisioning systems
- The WS Mapper tool to convert protocols
- Requesting Authorities to connect to the Provisioning Server through the SPML Service

SPML Templates

SPML technology lets you create templates that can be applied to data that is being imported into the Provisioning Server through the SPML Service. You can also use SPML Templating to modify the data as it goes into the Provisioning Server.

The SPML template takes records from a CSV or XML files and applies the template to each record. Because the template is applied to each record, you can impose some rule-based conditions to each record that affects the data output. The SPML output is then fed to the SPML Service and then sent to the Provisioning Server.

The templates are written using Velocity. Velocity is an open source Java-based engine that is specifically designed for processing templates. You can create your own templates or use the sample template that comes with the SPML Manager or the CMDRA.

You can write templates using the SPML Manager, or you can code them yourself. You can import mass data into the Provisioning Server using the templating functionality with the SPML Manager, SPML Feed, or the CMDRA.

The SPML service comes with example templates and example data files to help you understand and create your own templates.

For more information on SPML Templates, see using the SPML Manager's Templating Functionality.

WS Mapper

WS Mapper (Web Service Mapper) is a lightweight web service that takes a proprietary web service request and transforms the data into another web service request format. The service can also transform the response to the web service request back into the original format.

This service was designed to allow web service requests from third-party applications to be redirected to this service and mapped into SPML Service provisioning requests that will end up as provisioning tasks in the Provisioning Server.

Requesting Authorities

Any Requesting Authority client that uses standard SPML 1.0 can send provisioning requests to the Identity Manager SPML Service. The SPML Service takes the operations specified in the SPML requests and executes provisioning actions accordingly. The CMDRA and the SPML Manager are both requesting authorities.

Note: Command line requesting authorities such as CMDRA and SPML Feed can accept input from property files. Requesting authorities' process property files using `java.util.Properties` class. For that reason, certain character such as backslash (\) should be escaped. For example, username parameters in property files should be specified as `Domain\\Username`. For more details on usage requirements, see `java.util.properties` documentation.

Before a Requesting Authority can send requests to the SPML Service, it must be authenticated using HTTP basic client authentication. The Requesting Authority must provide the login credentials of a valid eTrust Admin user.

Important! When the client prompts for the username and password, the Username must include the user's domain name and a backslash in the format `Domain\Username`.

Install SPML

If you intend to use SPML in a FIPS-compliant environment, [additional instructions](#) (see page 80) apply.

To install the SPML service

1. Locate the Provisioning Component installation media.
2. Run the SPML installer under Clients.

Answer the questions to provide information about your system.

SPML Support for FIPS 140-2

The SPML server is FIPS 140-2 compliant. We recommend deploying the SPML service on:

- Apache Tomcat Server 4.1.36 or a higher version of 4.1
- JDK 1.5.11 or a higher version of JDK 1.5. Note that Tomcat must be enabled to run in SSL mode. For details, see the Apache's administrator guide for Tomcat 4, (<http://jakarta.apache.org/tomcat/>) section "SSL Configuration HOW-TO."

If you use CA Tomcat instead of Apache Tomcat, CA Identity Manager requires these workarounds for SPML:

- If you are using JDK 1.4.xx with CA Tomcat, FIPS 140-2 must be disabled. JDK 1.4.xx is incompatible with CA Tomcat because the RSA Jsafe CryptoJ 4.0 library needed for FIPS 140-2 support cannot be placed as the first security provider in JDK1.4.

To disable FIPS 140-2 support, pass the JVM flag
"-Dcom.ca.commons.security.fips=false" during Tomcat start up.

- If you are running Tomcat from the command line, you can include the JVM flag catalina.bat. More details exist in the batch file itself.
- If you are running Tomcat as windows service, pass the flag as follows:
 - a. Using the registry editor, navigate to
"HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CA Tomcat 4.1.29 eTrustIAMWebServer\Parameters"
 - b. Add a String Value called "JVM Option Number n" where 'n' is the number following on from the previous JVM parameter. For the value, specify:

Dcom.ca.commons.security.fips=false
 - c. Increase by one the value of Edit DWORD Value "JVM Option Count" to account for the newly added parameter.
- If you are using JDK 1.5 with CA Tomcat, an incompatibility problem exists. To work around this problem:
 - a. Manually remove the two Xerces libraries (xercesImpl.jar and xmlParserAPIs.jar) from %TOMCATHOME%\common\endorsed.
 - b. Restart Tomcat.

Uninstall the SPML Service

To un-install SPML, use the Windows Add or Remove Programs option and remove CA Identity Manager.

Important! This action removes all CA Identity Manager products.

SPML Service Configuration

During the installation of the Identity Manager SPML Service, you specified a single Provisioning Server. We can now run a Requesting Authority to connect to the SPML Service and start sending provisioning requests targeting that Provisioning Server. But you can also configure the SPML Service to manage multiple Provisioning Servers.

This section explains how to perform functions using the application. It covers topics such as the following:

- Log on to the SPML Configuration Application
- Add a new SPML service
- Modify, rename and delete an existing service

Log On to the SPML Configuration Application

These instructions assume that you have installed the SPML Service on the local computer.

To access and use the SPML Configuration Application, follow these steps:

1. Click Start Menu, Programs, CA, Identity Manager, IM SPML Service Configuration.

The SPML Configuration Login page appears.

2. Log in to SPML Configuration by entering your Provisioning Server login credentials. A user name with Delegated Administration (DAWI) privileges must be used:

Username

Provisioning Server user with administrator privileges.

Password

Password for this username.

Service

Admin Service to authenticate against.

Domain Name

Enter the domain name to which the username belongs.

Note: The domain name and password are case-sensitive.

3. Click Enter.

You are now logged in to the SPML Service Configuration application.

Add a New SPML Service

To add a new SPML service using the SPML configuration application, follow these steps:

Note: These instructions assume that you have installed the SPML Service on the local computer.

1. Log on to the SPML Configuration application.
2. Enter the following fields in the Admin Service Details form, on the right side of the screen:

Service Name

Specifies a reference name to Provisioning Server service. This name must not appear in the list of available Provisioning Servers on the left side or this will modify the existing service rather than create a new one.

Admin Hostname

Specifies the name of the computer running the Provisioning Server.

Clear Port Number

Specifies the LDAP port number used for communication with the Provisioning Server. By default, this port number is 20389.

SSL/TLS Port Number

Specifies the LDAP TLS port number if you are securing communication with TLS encryption. By default, this port number is 20390.

User SSL/TLS

Select this option. For security reason, the TLS encryption must be used.

3. Click the Add/Modify Service button to save the new service.

Note: If you enter *adminserver* in the Service Name field and *adminserver.yourcompany.com* in the Admin Hostname field then the Requesting Authority client would need to use the URL *http://spmlserver.yourcompany.com:8443/iamspml/spml/adminserver* when connecting to the SPML Server in order to send provisioning requests to this Provisioning Server.

You should replace *spmlserver.yourcompany.com* and *adminserver.yourcompany.com* with the names of the computers on which the SPML Service and the Provisioning Server are running.

Modify an Existing Service

To modify an existing Provisioning Server Service, perform the following steps:

1. Log on to the SPML Service Configuration application.
2. Select the service from the list of Available Admin Services
3. In the Admin Service Details form, make any modifications to the following fields:
 - Admin Hostname
 - Clear Port Number
 - SSL/TLS Port Number
 - User SSL/TLS
4. Click the Add/Modify Service button.

Note: If you modify the Service Name field you will create a new instance of an SPML service rather than modifying an existing one.

Rename an Existing Service

To rename an existing Provisioning Server Service, perform the following steps:

1. Log on to the SPML Service Configuration application.
2. Select the service from the list of Available Admin Services.
3. Click the Remove the Selected Service button to remove the service from the list.
4. On the form on the right, enter the new name in the Service Name field.
5. Click the Add/Modify Service button.

The new service is added to the list of Available Admin Services with the new name.

Delete an Existing Service

To delete an existing Provisioning Server Service, perform the following steps:

1. Log on to the SPML Service Configuration tool.
2. Select the service from the list of Available Admin Services.
3. Click the Remove the selected Service button.

The service is now deleted.

Configure SSL Support for Tomcat Servers

The Secure Socket Layer (SSL) is a technology that helps ensure the authentication, integrity, and confidentiality of SPML messages. For information on setting up the SSL, see the Configuration HOWTO at <http://jakarta.apache.org/tomcat/>.

Note: The following procedure is provided for reference only. You may want to configure your SSL certificate differently or change your keystore password to one of your own choosing for better security. Also, if you have installed JDK version 1.5, you should refer to <http://jakarta.apache.org/tomcat/> for details.

To install and configure SSL support for Tomcat using a self-signed certificate, perform the following steps:

1. Verify that JDK version 1.4.2_04 is installed by selecting the Add/Remove Programs list in your Control Panel for the program Java 2 SDK, SE v 1.4.2_04.
2. Create a new keystore containing one self-signed certificate by entering the appropriate command from the command prompt.

On Windows systems, you should enter:

```
%JAVA_HOME%\bin\keytool -genkey -alias tomcat -keyalg RSA -keystore %path%\keystore_filename
```

On UNIX systems, you should enter the following:

```
%JAVA_HOME%/bin/keytool -genkey -alias tomcat -keyalg RSA -keystore %path%\keystore_filename
```

The keystore creation process begins.

3. Enter the keystore password when prompted.

Note: The default password used by Tomcat is changeit (all lowercase). If preferred, you can specify a custom password, but you must then specify the custom password in the server.xml configuration file also (see Step 8).

The keystore creation process continues.

4. Enter general information for the certificate when prompted. The general information includes company, contact name, and so on. This information displays to users who attempt to access a secure page in your application, so make sure that the information provided here is appropriate.

The keystore creation process continues.

5. Enter the key password when prompted. This password is created specifically for this certificate (not for any other certificates stored in the same keystore file). You must use the same password for this and the keystore password.

A keystore file with a certificate that your server can use is created.

6. Browse to the <Tomcat_installation_directory>\conf\ directory and open the server.xml file in a text editor.
7. Ensure that the SSL Coyote HTTP/1.1 Connector entry is not commented out in the file. The connector information looks similar to the following:

```
<!-- Define an SSL HTTP/1.1 Connector on port 8443 -->
<!--
<Connector className="org.apache.catalina.connector.http.HttpConnector"
  port="8443" minProcessors="5" maxProcessors="75"
  enableLookups="true" acceptCount="10" debug="0" scheme="https"
  secure="true">
  <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
    clientAuth="false" protocol="TLS"/>
</Connector>
-->
```

If the Connector element is commented out, you must remove the comment tags, defined as less than sign, exclamation point, hyphen, hyphen (<!--) and hyphen, hyphen, greater than sign (-->) around it.

8. Configure the SSL Coyote HTTP/1.1 Connector entry to include the keystoreFile and keystorePass attributes for the Factory element.

keystoreFile

Specifies the location where the keystore file is located.

keystorePass

Specifies the keystore (and certificate) password.

The connector information should look similar to the following:

```
<Connector className="org.apache.catalina.connector.http.HttpConnector"
  port="8443" minProcessors="5" maxProcessors="75"
  enableLookups="true" acceptCount="10" debug="0" scheme="https"
  secure="true">
  <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
    keystoreFile="your_keystore_full_path"
    keystorePass="your_keystore_password"
    clientAuth="false" protocol="TLS"/>
</Connector>
```

9. Save the file and close it.

SSL support and self-signed certificates are configured for Tomcat.

10. Restart the Tomcat server.

Configure SPML Client Computer to Support SSL Security

The SPML Web Service requires that the Secure Socket Layer (SSL) be enabled. The SPML clients, the CMDRA, SPML Manager, and SPML Feed must trust the SSL server certificate to communicate with the server.

Note: Third party requesting authorities will need to support SSL to communicate with the SPML Web Service.

To configure the SPML client computer to use SSL security, perform the following steps:

1. Install the SSL certificate to the user's trusted keystore on the computer where the Requesting Authority runs. (By default, the SSL certificate will be added to the .spmlkeystore file in the user's home directory, as determined by the %HOMEPATH% system property.)
 - a. In a web browser, open the following URL:

`https://spmlserver.yourcompany.com:8443`
 - b. Double click on the SSL certificate icon at the bottom right corner of web browser to view the certificate.
 - c. On the Certificate Viewer window, select the Details tab and click Copy to File.
 - d. Save the server certificate.
 - e. Run the following command:

```
<drive>:\<JRE-File-Path>\bin\keytool -import -file <Certificate-File-Path> -keystore  
"%HOMEDRIVE%%HOMEPATH%\spmlkeystore" -storepass changeit -noprompt
```

This command creates a new keystore called .spmlkeystore, located in user's home directory (as determined by "%HOMEDRIVE%%HOMEPATH%"). The batch files that launch the RA clients (SPMLManager, Command Line RA, and SPML Feed) read this file to allow SSL communication.

Note: By default the batch files use the truststore path and password as defined by the keytool command described in step 1e. To use different path and password, variables set in the batch files for each client have to be modified accordingly. For example:

```
set TRUSTSTORE=%HOMEDRIVE%%HOMEPATH%\spmlkeystore  
  
set TRUSTSTORE_PASSWORD=changeit
```

2. Test the SPML Service with the Command Line RA:

- a. Open the login.properties file, in the Command Line RA directory, to make sure that HTTPS version of the Server URL is used and user logon details are correct.
- b. Open the command line prompt.
- c. From the Command Line RA directory, type:

```
RA.batsampleXML\schemaRequest.xml
```

CMDRA Commands

CMDRA Command Options

This table shows you the command options that you can use in the CMDRA.

Command	Full Command	Explanation
-c	--check	Check the request is a valid SPML request. This will not send the request to the SPML Server. If you are using the SPML Templating feature, the records will be expanded and the resulting SPML request will be checked.
-e	--explodedOutputFile	Specify a file to contain the exploded request XML output, overwriting any file by that name.
-f	--propertyFile	Specify a file that contains default command settings. You would create a property file that contains frequently used values to avoid having to specify them manually every time. For example: mappingFile=C:\SPMLdata\Mapping1.csv templateFile=C:\SPMLdata\ImportUsers.xml dataFile=C:\SPMLdata\Users.csv serverURL=https://spmlserver.yourcompany.com:8443/iamspml/spml/adminserver You would not typically include the data file in the property file because it is highly variable.

Command	Full Command	Explanation
		The CMDRA looks for the property file <i>login.properties</i>. If your property file is named <i>login.properties</i> you do not need to specify -f. When you specify information in the property file, you should use the full command but remove the two dashes (--).
-h	--help	Display the command line help page that gives you a summary of these commands.
-i	--inputFileNames	Specifies a file to read data/request file names from instead of putting file names on the command line
-m	--mappingFile	Use this to match your Velocity template variables with your data file variables. Typically this is generated using the Save Mapping button in the SPML Manager. You can include multiple mappings that are separated by commas if you include multiple -t options. Note: This must always be used in conjunction with the Data File and the Template File.
-o	--outputFile	Specify where you want the response from the SPML Service stored. By default, this is written to stdout. This output can be redirected to a file, for example: <pre>RA.bat yourparameters > SPMLresponse.text</pre> This can contain information other than the raw XML response.
-p	--password	Specify the Provisioning Server password for the user. Therefore, this must be used in conjunction with the user (-u). You do not need to specify the user or password if you are just checking a request.
-q	--quiet	Specify that you want minimum detail in the output.

Command	Full Command	Explanation
-R	--explodePerRowOutputFileNames	Specify file that will list the names of files that contain the exploded request XML output, overwriting any file of the same name. One file is created per template per datafile record.
-s	--serverUrl	Specify the SPML server URL that you want to send the request to.
-S	--csvRuntimeStatistics	Name of .CSV file to write time to complete each request
-t	--templateFile	Specify the SPML Template file. Typically you would create the template file using the SPML Manager. If the -R option is included, the -t option can include a comma separated list of filenames. Note: This must always be used in conjunction with the Data File and the Mapping File.
-u	--user	Specify the Provisioning Server user. This must be used in conjunction with password (-p). You do not need to specify the user or password if you are just checking the request. You must include the domain of the user, for example: YOURDOMAIN\\user
-V	--verbose	Include detailed information in the output.
-v	--version	The version and build number of the CMDRA. This must be lowercase.

CMDRA Examples

The first example creates several small requests:

```
Ra.bat -t 01_add_user.xml.vpp,02_modify_user.xml.vpp -R req_file_names.txt -m -,funny_mapping.csv  
data10.csv
```

- When used with the -R option, the -t template and -m mapping options accept multiple files separated by commas.
- The -R option creates a request file per template per datafile record and the names of resulting files are collected in the named file. So this command creates file names 01_add_user0000000.xml and 02_modify_user0000000.xml through 01_add_user0000009.xml and 02_modify_user0000009.xml.

These files are written to the same directory as the -R file (.) and appended to ./req_file_names.txt.

The second example creates many requests and reports basic performance metrics:

```
Ra.bat -S stats.csv -i req_file_names.txt
```

- The -i option takes the name of a file containing names of SPML request files. In this case, the filename is output by the first example.


The -S option can be included when you submit SPML requests without a template (whether using file names on the command line, or the -i option, or both). The option records the execution time in milliseconds for each request and presents summary data after execution.

SPML Feed

SPML Feed is a command line application that you start by using the feeder.cmd (on Windows) or feeder.sh (on Unix). All command options can be set by using a properties file; some options can be specified on the command line. Options specified on the command line override those in the property file. A property file is used when you include the -f or -propertyFile command line option.

You can run SPML Feed in standard or daemon mode. In standard mode, the application runs once over the specified input files and then quits. In daemon mode, the application does not terminate. Instead, it periodically wakes up and checks if any input files has been modified. If a file has been modified, it is processed as it would be in standard mode.

The template supplied to the SPML Feed by using the -template option is applied to every record in the input CSV or XML files. The template produces a single SPML add, modify, or delete request.



The template is in the cmdra.zip file. To access it, go to Start, Program CA, Identity Manager, IM SPML Requesting Authority.

SPML Feed Command Options

Options can have three names: the short command line name (-s), the long command line name (--serverURL) and the property file name (serverURL).

-b, --batchSize, batchSize

Integer value that defaults to 1. If this number is greater than 1, requests will be submitted as batch requests, with this value determining how many requests are placed in each batch.

By batching multiple requests together, performance can be improved because less time is spent performing TCP socket and SOAP setup, at the expense of additional memory usage. The memory used depends on the size of your requests, but setting batchSize to several hundred should pose no problem.

Note: batchSize should be always be 1 if your template is for a search or batch request, as these cannot be placed inside a batch request.

-d, --daemon, daemon

Causes the application runs in daemon mode.

-x, --explodeOnly, explodeOnly

The template is exploded using the data files as normal. However, instead of sending the resulting request to the SPML server, it is written to stdout (or a file if output or daemonResponseDir is set).

-h, --help

Display help for the command line switches to stdout.

-i, --inputFileNames, inputFileNames

Specifies a text file. Each line in the text file is used as the name of an input file. This is an alternative method of listing input files on the command line itself. The set of input files is the union of those listed in this file and those listed on the command line.

-l, --logging, logging

Specifies a properties file to configure the Java logging system instead of using the default logging system. See the documentation for the `java.util.logging.LogManager` class for details.

-m, --mappingFile, mappingFile

Specifies a CSV file in a special format that maps parts of the input files to Velocity variable names. The variable names defined in this file can be referenced in the Velocity template file. If there is a mapping to a variable named timestamp, this has a special meaning and is used to determine which records have been changed since the last run.

-o, --outputFile, outputFile

Specifies a file to record the output, overwriting any existing file. If you omit this option, and the daemonResponseDir property is not set, output is written to stdout.

-p, --password, password

Specifies the password to use to authenticate with the SPML server.

-f, --propertyFile

Specifies a property file with the options to use. You can specify any command line option (except -h and -f) in the properties file by setting a property that matches the option's long name without the - prefix. For example, -mapping becomes simply mapping in the properties file. Some options for daemon mode can only be set via a properties file. Any option specified on the command line replaces a setting from the properties.

-q, --quiet, quiet

Causes no output to appear unless a catastrophic failure occurs. In that case, an error message is output to stderr before the program exits.

-s, --serverUrl, serverUrl

The URL of the SPML server to send the request to.

-t, --templateFile, templateFile

Specifies a Velocity template file that can be merged with data from XML or CSV input files to produce an SPML request. The application runs in exploder mode; the input files are either XML or CSV files instead of SPML requests. Each record in the input files is applied to the template (after mapping to variable names via the mapping file) to produce an SPML request that is sent to the server.

Note: the template should not contain `<?xml?>` processing headers.

-u, --user, user

Specifies a user name to authenticate with the SPML Server. For the SPML Server, this option's value should be of the form domain\username.

-v, --verbose, verbose

Outputs additional information to stderr about the application's actions.

-V, --version, version

When the application starts, its name and version number is written to stderr.

Property File Only

timestampFile

Specifies a file used to record when input files have been processed. This is mainly useful in daemon mode to keep track of the latest run times when the daemon is shutdown temporarily or restarted.

daemonResponseDir

Specifies a directory to write SPML responses to when running in daemon mode. The SPML response from each run over an input file by the daemon is written to a new file in this directory. The output files have the same name as the input data file, with a digit appended. For example, if the file test.csv is processed for the first time, the response is written to test.csv.1, the second run to test.csv.2, and so on.

If you omit this option, all responses are written to the file specified with -o, or stdout if neither -o nor -q are present.

daemonSleep

Specifies the length of time in milliseconds to sleep between polling for data file changes. If this parameter is not specified, the length of sleep time is 30 seconds.

Flow of the SPML Feed Command

The flow of the SPML Feed command is:

1. This command is invoked with these parameters:
 - Mapping file
 - Velocity template that produces a valid SPML request.
 - Data files to watch
 - Output file or output directory
 - Length of time to sleep between polling for file data file changes.
2. When a data file changes, the running daemon begins processing the file on the next poll.
3. The data file is locked to prevent writes.

4. The daemon reads the data file one record at a time.
 - If the mapping file has an entry mapping to the special value timestamp, that field will be retrieved from the record. If the record's timestamp is earlier than the last time the file was processed, the record has not changed and is skipped.
Note: The timestamp should be specified in GMT time zone and YYYY-MM-DD HH:MM:SS format, such as 2006-02-14 21:02:03. If the timestamp is not in this format, or is absent, the request is skipped.
 - If the record is not skipped, it is loaded into a Velocity context and merged with the template to produce an SPML request.
 - If the mapping file does not have an entry mapping to the special value timestamp, the daemon submits all requests from the data file without exception.
5. The SPML requests are submitted to the SPML Service.
6. The response from the SPML Service is classified and appended to the appropriate output files.

The SPML Feed can use the Velocity templates created for the SPML Manager and CMDRA applications if the batchSize argument is set to 1. Ideally, you should modify templates to match SPML Feed requirements, which are less strict.

Using the SPML Manager's Templating Functionality

The SPML Manager is a graphical user interface that lets administrators create and execute SPML provisioning requests. You can design provisioning requests using the SPML Manager then view the SPML Request in XML format.

Note: The SPML Manager is an unsupported technology preview.

Download the SPML Manager

1. Download the SPML Manager from the following location:
<https://spmlserver.yourcompany.com:8443/iamspml/download/techpreview/SPMLManager.zip>
2. Unzip SPMLManager.zip to your hard disk.
3. To launch the SPML Manager, navigate to the SPMLManager folder and double-click SPMLManager.bat.


Create an SPML Template Request

To create an SPML Template request, perform the following steps:

1. Create an XML file containing some sample data.

You can find a sample in the SPMLManager\sample Templates directory.

2. Using the SPML Manager, connect to an SPML service.



3. Use one of the tabs (Add, Modify, Delete, or Extended) to construct an SPML request that conforms to the schema of the SPML service.

The name given to the Exploder Velocity context is REC_, so whenever you want to refer to a variable in the data file, use the syntax `${REC_.variable}`.

To hard-code a constant in the generated requests, type the data you want into the attribute fields.

4. Click the Raw XML tab to see the SPML Request in XML format.
5. Click the Add to Batch Request button when you are happy with your simple variable replacement.

Note: The Add to Batch Request button is only available when viewing the tab in Request mode and is therefore not available in Raw <XML> mode.

6. Use the Batch tab to see the addition to batch.

A batch request can contain as many individual requests as you like.

7. Save the SPML Batch Request Template to a file.
8. Click the Template tab and load your simple SPML Template from your saved Batch Request File.
9. Click the Raw <XML> tab to edit your SPML Template:
10. In the XML code, insert a velocity directive for each loop at the point in your batch request where you want the request to cycle through loading each row in the CSV into the context and put a #end statement where you want the cycling to end. The format of the syntax for each loop is `#foreach ( $REC_ in $RECS_`.

11. Click Save Template to save your modified template file.
12. Load the XML data file that will fill in the data in your request.
13. Click the Save button if you want to Save the resulting batch request to a file for inspection or click Submit if you want to submit the resulting Batch Request to the SPML Service.

For each XML record in the data file, a corresponding SPML Request will be generated to initiate a provisioning operation inside the Provisioning Server.

The example files used in this tutorial are in the SPMLManager/sampleTemplates/simple directory.

Using Velocity Templates

The SPML Manager, SPML Feed, and CMDRA use the same templating system, which parses references to variables and performs data transformations.

List Templating Variables

The templating system deduces which variables are ArrayLists by parsing references to them in the Velocity template you provide. The templating system looks for an attribute references or method invocations against the variable which show that it should be bound to a java.util.ArrayList (which can have 0, 1 or more values) rather a single value. The SPML Manager and CMDRA can use this capability.

All the read-only methods on the class java.util.ArrayList are looked for:

- isEmpty
- get
- size
- contains
- indexOf
- lastIndexOf
- subList
- iterator
- listIterator

In mapping file entries, such list variables have the suffix [] after their names. For instance the variable `comments` is a single valued variable, but `comments[]` would be a list.

For example references, see the `sampleTemplates\simple\template.xml.vpp` template, including the `$comments` variable. Also, included are example mapping files: `map_csv_datafile.csv` against the CSV datafile and `map_xml_datafile.csv` against the XML datafile.

Data Transformations

SPML Templating offers several routines for manipulating data from a unique or proprietary format into the format required by the SPML Service which is generally the same as the standard XML Schema Data Types (see <http://www.w3.org/tr/xmlschema-2> for more information). These tools are provided by the Velocity Tools project which is in the SPML Requesting Authority classpath and therefore available for reference in your SPML Template.

SPML Manager, SPML Feed, and Requesting Authority.

Useful tools for manipulating and transforming data inside of your SPML Template include the following:

- **Date Tool:** A tool for manipulating and formatting dates.
- **Math Tool:** A tool for performing floating point math.
- **Number Tool:** A tool for formatting numbers.
- **Iterator Tool:** A tool to use with `#foreach` loops. It wraps a list to let the designer specify a condition to terminate the loop, and reuse the same list in different loops.
- **Render Tool:** A tool to evaluate and render arbitrary strings of VTL (Velocity Template Language).

The templating system loads data from XML and CSV files in string format. You can use these tools to convert data in your template to the type required for the operation you want to perform with the data.

Example Data Transformation

This example takes a numeric value which has been converted to a string inside the templating system and then converts it into an integer value to perform preprocessing on the data before sending it to the SPML Service.

In this example, the imported CSV file contains only three pieces of data.

- username,expirydate,priority
- user3,20101001,999
- user4,20101005,333

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<batchRequest onError="urn:oasis:names:tc:SPML:1:0#resume"
processing="urn:oasis:names:tc:SPML:1:0#sequential"
execution="urn:oasis:names:tc:SPML:1:0#synchronous" xmlns="urn:oasis:names:tc:SPML:1:0">
  #foreach ( $REC_ in $RECS_ )
    #set ( $userhandle =
"User=${REC_.username},Domain=YOUR_USER_DOMAIN,Server=Server")
    ""
    #set ( $datetimeobject = $date.toDate('yyyyMMdd',${REC_.expirydate}))
    #set ( $formatdate = $date.format('yyyy-MM-dd', $datetimeobject))
    #set ( $formattime = $date.format('H:m:s', $datetimeobject))"
    <addRequest requestID="batchAdd${REC_.INDEX}">
      <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
        <id>$userhandle</id>
      </identifier>
      <attributes>
        <attr name="accountid">
          <ns1:value
xmlns:ns1="urn:oasis:names:tc:DSML:2:0:core">${REC_.username}</ns1:value>
        </attr>
        <attr name="roleHandles">
          <ns2:value
xmlns:ns2="urn:oasis:names:tc:DSML:2:0:core">Role=ntRole,Domain=YOUR_USER_DOMAIN,Server=S
erver</ns2:value>
        </attr>
        <attr name="enableDate">
          <ns3:value
xmlns:ns3="urn:oasis:names:tc:DSML:2:0:core">"${formatdate}T${formattime}"</ns3:value>
        </attr>
      </attributes>
    </addRequest>
    ""#if ($math.toInteger(${REC_.priority}) > 500)""
```

```
<extendedRequest xmlns="urn:oasis:names:tc:SPML:1.0">
  <operationalAttributes/>
  <providerIdentifier providerIDType="urn:oasis:names:tc:SPML:1.0#URN">
    <providerID/>
  </providerIdentifier>
  <operationIdentifier operationIDType="urn:oasis:names:tc:SPML:1.0#GenericString">
    <operationID>User-SyncWithRolesAddAccounts</operationID>
  </operationIdentifier>
  <identifier type="urn:oasis:names:tc:SPML:1.0#DN">
    <id>$Userhandle</id>
  </identifier>
  <attributes/>
</extendedRequest>
#end
#end
</batchRequest>
```

This request formats a proprietary data format into a format that the Provisioning Server can understand. The results of this request are:

- user3 is created with an XSD Enable Date of 2010-10-01T0:0:0 that has been converted from the yyyyymmdd format of 20101001
- user3 has been considered a high priority case and has synced with the NT role immediately upon creation to create accounts. This is because user3 had a high priority of 999 which has been evaluated to see if it was greater than 500.
- user4 is created with an XSD Enable Date of 2010-10-05T0:0:0 that has been converted from the yyyyymmdd format of 20101005
- user4 has not been considered a high priority case for account creation because their priority code was 333 which is less than 500. Their account will not be created until a later stage, perhaps when a nightly sync operation occurs.

```

<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<batchRequest onError="urn:oasis:names:tc:SPML:1.0#resume"
processing="urn:oasis:names:tc:SPML:1.0#sequential"
execution="urn:oasis:names:tc:SPML:1.0#synchronous" xmlns="urn:oasis:names:tc:SPML:1.0">
  <addRequest requestID="batchAdd0">
    <identifier type="urn:oasis:names:tc:SPML:1.0#DN">
      <id>User=user3,Domain=YOUR_USER_DOMAIN,Server=Server</id>
    </identifier>
    <attributes>
      <attr name="accountID">
        <ns1:value xmlns:ns1="urn:oasis:names:tc:DSML:2.0:core">user3</ns1:value>
      </attr>
      <attr name="roleHandles">
        <ns2:value
xmlns:ns2="urn:oasis:names:tc:DSML:2.0:core">Role=ntRole,Domain=YOUR_USER_DOMAIN,Server=S
erver</ns2:value>
      </attr>
      <attr name="enableDate">
        <ns3:value xmlns:ns3="urn:oasis:names:tc:DSML:2.0:core">"2010-10-01T0:0:0"</ns3:value>
      </attr>
    </attributes>
  </addRequest>
  <extendedRequest>
    <operationalAttributes/>
    <providerIdentifier providerIDType="urn:oasis:names:tc:SPML:1.0#URN">
      <providerID/>
    </providerIdentifier>
    <operationIdentifier operationIDType="urn:oasis:names:tc:SPML:1.0#GenericString">
      <operationID>"User-SyncWithRolesAddAccounts"</operationID>
    </operationIdentifier>
    <identifier type="urn:oasis:names:tc:SPML:1.0#DN">
      <id>"User=user3,Domain=YOUR_USER_DOMAIN,Server=Server"</id>
    </identifier>
    <attributes/>
  </extendedRequest>
  <addRequest requestID="batchAdd1">
    <identifier type="urn:oasis:names:tc:SPML:1.0#DN">
      <id>User=user4,Domain=YOUR_USER_DOMAIN,Server=Server</id>
    </identifier>

```

```
<attributes>
  <attr name="accountId">
    <ns4:value xmlns:ns4="urn:oasis:names:tc:DSML:2:0:core">user4</ns4:value>
  </attr>
  <attr name="roleHandles">
    <ns5:value
xmlns:ns5="urn:oasis:names:tc:DSML:2:0:core">Role=ntRole,Domain=YOUR_USER_DOMAIN,Server=S
erver</ns5:value>
  </attr>
  <attr name="enableDate">
    <ns6:value xmlns:ns6="urn:oasis:names:tc:DSML:2:0:core">"2010-10-05T0:0:0"</ns6:value>
  </attr>
</attributes>
</addRequest>
</batchRequest>
```

Retrying SPML Requests

You can configure certain requests to be retried on failure. Request retrying is attempted if all of the following are true:

- The request is flagged for asynchronous execution.
- The object on which the request is acting at the time of failure resides on a remote endpoint system, such as an account/container/native group.
- The request is causing a change and not a query.
- The failure occurs after a request has reached the Provisioning Server. It is the client's responsibility to retry if either:
 - The communication channels between the client and the web server on which the SPML server is executing is broken.
 - The communication channel between the SPML server and Provisioning Server is broken.
- The failure is a soft failure between the Provisioning Server and the targeted endpoint system.

Note: Batch requests will not support retrying on their constituent sub-requests.

On successful completion or hard failure, the standard SPML success/failure conditions are returned in the status response's result.


The SPML server is involved in the configuration of the retry persistence mechanism in JIAM but only JIAM actually adds and deletes records from it.

Retry Architecture

You can configure the SPML service to retry asynchronous requests that act on objects residing on Endpoint Systems which cannot be contacted due to a transient failure. The flow of retries follows the standard asynchronous request processing in SPML.

- The client submits an asynchronous request to the SPML server with an additional operational attribute indicating that retries should be attempted on failure.
- If the request is processed without error, the request is marked as successfully processed.
- If the request's processing terminates with a hard failure, or its target object is not considered retrievable by JIAM, an immediate failure results.
- If the request's processing terminates with a soft failure and the client has flagged the request as retrievable, JIAM stores the request in the retry database and attempts to process it again after a configured interval. Prior to storage the current operation is simplified by removing any sub-operations which succeeded, so that only the sub-operations which failed due to soft failures will be retried.
- If a request has been retried a configured maximum number of times, then it is considered to have suffered a hard failure and JIAM discontinues the retries.

The following diagram gives an overview of the components of the retries:



Note that the SPML server is involved in the configuration of the retry persistence mechanism in JIAM but only JIAM actually adds and deletes records from it.

Retry Configuration Files

Interface

Two operational attributes control the SPML retry capability:

- `caIamRetry` can be passed in with an asynchronous request, requesting that retry functionality is activated where supported by JIAM. `caIamRetry` is operational attribute and its value can be true or false.
- `caIamRetryDetails` is returned when an SPML status request targeting a retried request is received. It provides status information to the user.

For an example of a request flagged for retry, see the appendix "Sample SPML Requests."

Hivemodule-plugin-deploy.xml

The file `hivemodule-plugin-deploy.xml` in the SPML deployment directory includes configuration items:

- A new database called `db_spml_retry` configured by the `retryBasicDataSource` service-point, and advertised to JIAM via the `retryPersister` service-point. This database (as well as the pre-existing `db_qrtz` and `db_delegate`) should not grow and shrink as requests are processed and time out.

If a problem occurs in this area, the databases can be truncated by shutting down the SPML server and simply deleting their directories (the SPML server will automatically create new blank directories on start-up).

- The `retryMaximumCount` and `retryDelayMinutes` settings which control the JIAM retry behavior using the `JIAMService` class, which is also provided with the `retryPersister`.

`retryMaximumCount` tells the SPML server the maximum number of times it should retry an operation before giving up and returning an error. When a retry is needed, the SPML server waits before doing the retry, to give the network or server failure a chance to correct itself. The time it waits before retrying is controlled by the `retryDelayMinutes` setting.

- The `scheduledHoldingIntervalMinutes` setting dictates the time that status requests targeting an asynchronous request can be submitted after its processing is completed, and consequently applies to retried requests too.

Any manual changes made to a `hivemodule-plugin-deploy.xml` file from a previous release will need to be reapplied to the file for this release.

To cancel a retried request, submit a SPML cancel request referring to its identifier. If a retry attempt has begun, it will continue, but no further retry attempts will be made.

Note: Retry functionality is not supported on SPML modify requests that change an object's Distinguished Name (DN).

Access Credentials

Access credentials for the databases used by the SPML server are stored in the `spml_quartz.properties` and `hivemodule-plugin-deploy.xml` configuration files. Each database supports two different styles of URLs in their datasources:

- `jdbc:hsqldb:file:<dir>\<db>\<db>` : databases can only be accessed by the SPML server process itself (the default).
- `jdbc:hsqldb:hsqldb://localhost/<db>` : allows SQL access to the database by other processes.

Configure Retry for a Request

The process of configuring an SPML request to be retried is the following:

1. In the SPML Manager or Command-line RA, you submit a request for asynchronous execution including the `caIamRetry` operational attribute. This attribute can be referenced in Java as `com.ca.commons.spml.IAMSpmlUtil.CA_IAM_ATTR_RETRY`.
2. The SPML server schedules the request for immediate execution. If the processing is completely successful or any sub-operation fails with a hard failure, a success or failure SPML response is stored.

Otherwise if one or more sub-operations fail with a soft failure, the SPML server will return `true` from the `actionFailed()` method of the `com.ca.iam.spml.ProcessingDetails` instance assigned as the `IAMCommitObserver` for the current session, which informs JIAM that there is an interest in retrying.

3. You can track processing of an asynchronous request by submitting an SPML status request quoting the request ID. The status eventually changes to complete, signifying either complete success or a non-retriable failure. The non-retriable state could be due to the retry limit being exhausted as set by `retryMaximumCount` in `hivemodule-plugin-deploy.xml`.
 - a. Results for asynchronous requests are cached for a configurable period after their processing completes (refer `scheduledHoldingIntervalMinutes` in `hivemodule-plugin-deploy.xml`)
 - b. When the status of a request which is being, or has been, retried is queried, the operational attribute `caIamRetryDetails` appears in its operational attributes and provides a rough summary of progress suitable for a human reader.
4. JIAM then analyses the operation that failed and determines if it does indeed support retrying for it. If retrying is supported, JIAM will remove any successful sub-operations and use the `HSQLDBPersister` configured by the SPML server to save the retry operation to the database named `db_spml_retry`. The `IAMSession.commit()` call then completes and a pending SPML response is returned to the RA that submitted the request.
5. The JIAM retry subsystem periodically retries the request based on the `retryDelayMinutes` setting in `hivemodule-plugin-deploy.xml`. assuming `retryMaximumCount` limit is not exhausted.
 - a. If a soft failure is encountered, the retry process repeats at step 4.
 - b. If a hard failure is encountered or the limit is exhausted, the retry sequence terminates and the SPML server is informed via the `failed()` method of the registered `QueueObserver`. Future status requests targeting the retried request will report failure.
 - c. If no failures are encountered, JIAM informs the SPML server that the request was processed successfully via the `completed()` method of the registered `QueueObserver`. Future status requests targeting the retried request will report success.

Chapter 4: Sample SPML Requests

This appendix describes the sample SPML requests that a Requesting Authority can use to send provisioning requests to the SPML Service.

For a detailed description of the format of these requests, see the SPML v1.0 specification at:

<http://www.oasis-open.org/committees/download.php/3032/cs-pstc-spml-core-1.0.pdf>

This section contains the following topics:

[Request Execution Types](#) (see page 107)

[Request Types](#) (see page 108)

[Global Settings](#) (see page 120)

[Account Containers](#) (see page 121)

[Complex Attributes](#) (see page 123)

[Request Retries](#) (see page 124)

[Propagate Global User Changes](#) (see page 125)

[Escaping Special Characters in Object Identifiers](#) (see page 126)

[Escaping Special Characters in Search Filters](#) (see page 127)

Request Execution Types

The Identity Manager SPML Service supports both synchronous and asynchronous request models.

By default if no execution attribute is set with a request to the SPML Service the request will be treated as having a synchronous execution mode.

If *synchronous* is specified, any request that you send must be completed and a response sent back before the next request can be sent. This can sometimes cause delays.

If *asynchronous* is specified as the execution mode, the SPML Service will schedule the request to be executed asynchronously and return immediately with a unique request ID. The Requesting Authority can later look up the corresponding result of the request by specifying the request ID in a Status Request. Any pending asynchronous request can be canceled by specifying the request ID from the asynchronous request in a Cancel Request.

Request Types

Add Request

The add request is used by a Requesting Authority to create new entity instances such as User, Role, Group, Profile, Policy, EndPoint, or Account objects.

Fields in an Add Request

Objects may contain mandatory fields that must be populated in order for the object to be created.

An add request contains the following fields:

identifier

Specifies the ID of the new object to be created

attributes

(Optional) Specifies initial values for some of the attributes as appropriate

Example of an Add Request

The following request creates a new User object with the unique identifier User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<addRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm="urn:oasis:names:tc:DSML:2:0:core">
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <attributes>
    <attr name="accountid">
      <dsm:value>_spml_user</dsm:value>
    </attr>
    <attr name="suspended">
      <dsm:value>true</dsm:value>
    </attr>
  </attributes>
</addRequest>
```

Batch Request

The batch request collates multiple SPML operations into a single request.

Example of a Batch Request

The following batch request executes these SPML operations in this order:

- Adds the user *_spml_user*
- Modifies *_spml_user* to update the comments attribute
- Deletes the *_spml_user* object
- Performs a CheckSync operation on the user *administrator*

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<batchRequest execution="urn:oasis:names:tc:SPML:1.0:synchronous"
onError="urn:oasis:names:tc:SPML:1.0:resume" processing="urn:oasis:names:tc:SPML:1.0:sequential"
xmlns="urn:oasis:names:tc:SPML:1.0" xmlns:dsm="urn:oasis:names:tc:DSML:2.0:core">
  <addRequest requestID="batchAdd">
    <identifier type="urn:oasis:names:tc:SPML:1.0:DN">
      <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
    </identifier>
    <attributes>
      <attr name="accountId">
        <dsm:value>_spml_user</dsm:value>
      </attr>
    </attributes>
  </addRequest>

  <modifyRequest requestID="batchModify">
    <identifier type="urn:oasis:names:tc:SPML:1.0:DN">
      <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
    </identifier>
    <modifications>
      <modification name="comments" operation="replace">
        <dsm:value>new comments</dsm:value>
      </modification>
    </modifications>
  </modifyRequest>

  <deleteRequest requestID="batchDelete">
    <identifier type="urn:oasis:names:tc:SPML:1.0:DN">
      <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
    </identifier>
  </deleteRequest>
```

```
<extendedRequest requestID="batchExtended">
  <providerIdentifier providerIDType="urn:oasis:names:tc:SPML:1:0#URN">
    <providerID>urn:ca.com:etrust:iam</providerID>
  </providerIdentifier>
  <operationIdentifier operationIDType="urn:oasis:names:tc:SPML:1:0#GenericString">
    <operationID>User-CheckSync</operationID>
  </operationIdentifier>
  <attributes>
    <attr name="IAMUser">
      <dsml:value>User=adminstrator,Domain=EXAMPLE_DOMAIN,Server=Server</dsml:value>
    </attr>
  </attributes>
</extendedRequest>
</batchRequest>
```

Cancel Request

The cancel request allows a client to request the cancellation of an asynchronous request from the SPML Service.

Example of a Cancel Request

For example, a previously-sent asynchronous with the request ID A4DF567HGD can be cancelled with the following request:

```
<cancelRequest requestID="A4DF567HGD" xmlns="urn:oasis:names:tc:SPML:1:0"
xmlns:dsml="urn:oasis:names:tc:DSML:2:0:core"/>
```

Delete Request

Use the delete request to delete any object that is available through the Provisioning Server, except a domain or namespace object. Domain and namespace objects cannot be created or deleted through SPML requests.

Fields in a Delete Request

A delete request contains only one field:

identifier

Specifies the ID of the object to be deleted

Example of a Delete Request

The following request deletes the *_spml_user* object from the Provisioning Server:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<deleteRequest xmlns="urn:oasis:names:tc:SPML:1:0">
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
</deleteRequest>
```

Extended Request

Use the extended request to perform actions that are unique to the Provisioning Server, such as explore, correlate, sync, and checkSync.

Fields in an Extended Request

An extended request contains the following fields:

operationIdentifier

Specifies the type of the extended operation performed.

identifier

Identifies the object that the extended operation is to be applied to.

attributes

Passes parameters specific to the extended operation. The parameters required by an extended operation are in the core Schema Response.

Example of an Extended Request

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<extendedRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm1="urn:oasis:names:tc:DSML:2:0:core">
  <providerIdentifier providerIDType="urn:oasis:names:tc:SPML:1:0#URN">
    <providerID></providerID>
  </providerIdentifier>
  <operationIdentifier operationIDType="urn:oasis:names:tc:SPML:1:0#GenericString">
    <operationID>User-CheckSync</operationID>
  </operationIdentifier>
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>User=Administrator,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <attributes></attributes>
</extendedRequest>
```

Extended Request Types

Request Type	Function
Account-CheckSync	Check an account against its assigned policies for out-of-sync attributes.
Account-ForcedDelete	Delete this account and clear all references to it.
Account-Relocate	Move the selected account to the correct container as specified by a given policy.
Account-SyncWithPolicies	Synchronize an account with its assigned policies.
Container-CheckAccountSync	Check whether the accounts in a container need to be synchronized with associated policies.
Container-Correlate	Perform a correlation on an endpoint or on a container of a hierarchical endpoint. Scope The search scope. ONELEVEL_SCOPE is represented by the integer 1 and SUBTREE_SCOPE is represented by the integer 2. ONELEVEL_SCOPE Correlates one level the managed accounts on an endpoint system with the Global Users. SUBTREE_SCOPE Correlates all managed accounts on an endpoint system with Global Users. CreateUserAsNeeded Use "true" to create Global Users as needed.
Container-Explore	Perform an explore operation on an endpoint or a container of a hierarchical endpoint. Scope The search scope. ONELEVEL_SCOPE is represented by the integer 1 and SUBTREE_SCOPE is represented by the integer 2. ONELEVEL_SCOPE Searches one level for managed objects on the given Container. SUBTREE_SCOPE Searches for all managed objects on the given Container.
Container-SyncAccountsWithPolicies	Synchronize the accounts in a container with their assigned policies.
Container-UpdateUserFields	Update the global users' attributes with their

Request Type	Function
elds	correlated accounts' attributes, according to the attribute mappings defined in the defaultUserUpdateMap field of the container object. Scope The search scope. ONELEVEL_SCOPE is represented by the integer 1 and SUBTREE_SCOPE is represented by the integer 2. ONELEVEL_SCOPE Performs the updateUserFields operation for the accounts directly below the given container. SUBTREE_SCOPE Performs the updateUserFields operation for all accounts of the given container.
EndPoint-IncludeContainer	Bring a top-level container into the database (but not its contents). This is required for the exploration operation to work on some hierarchical endpoints such as ADS endpoints where the normal ONELEVEL exploration does not add the top-level container to the provisioning directory. This is also useful to manage only a portion of the hierarchical endpoint with the Provisioning Server while the remaining portion is completely hidden to Provisioning Server users. ContainerName The name of the container to include. ContainerType Option-specific types of containers. See the JIAM OptionDescriptor Javadoc for available container types for each option
Group-ListMembers	Search for Global Users that are members of this Group. Scope The search scope. ONELEVEL_SCOPE is represented by the integer 1 and SUBTREE_SCOPE is represented by the integer 2. ONELEVEL_SCOPE Searches one level for group members. SUBTREE_SCOPE Searches recursively through all nested child groups for members. UserNameMatchString The string to match the user name against. Use null to return all the users of this group. CountLim

Request Type	Function
	The maximum number of users to return. If 0, return all entries that satisfy the above matching expression.
Policy-CheckAccountSync	Check the accounts against a given policy.
Policy-ForcedDelete	Delete this policy and clear all references to it.
Policy-RelocateAccounts	Move the accounts associated with a policy to the correct containers.
Policy-SyncAccounts	Synchronize the accounts associated with a policy.
Role-CheckAccountSync	Check whether the accounts of multiple users need to be synchronized against the policies assigned to this role.
Role-CheckUserSync	Check whether the users of a role need to be synchronized with their associated policies.
Role-DeleteWithPolicies	Delete the role and all associated policies. ForcedDelete Set to <i>true</i> to delete the role and clear all references to it as well as associated Policies.
Role-ForcedDelete	Delete this role and clear all references to it.
Role-SyncAccountsWithPolicies	Synchronize the accounts of multiple users against the policies assigned to a role.
Role-SyncUsers	Synchronize the users of a role with all assigned policies.
User-CheckAccountSync	Check whether the accounts of a user need to be synchronized with associated policies.
User-CheckSync	Check whether the user requires synchronization with associated Roles.
User-DeleteWithAccounts	Delete the user and all associated accounts. ForcedDelete Set to <i>true</i> to delete the user and clear all references to it as well as associated accounts.
User-ForcedDelete	Delete this user and clear all references to it.
User-GeneratePassword	Generate a random password that conforms to the password quality rules for a global user.
User-RequestPasswordReset	Register a user's password reset request.

Request Type	Function
User-SyncAccountsWith Policies	Synchronize the accounts of a user with their assigned policies.
User-SyncWithRolesAdd Accounts	Synchronize a user with roles and create accounts.
User-SyncWithRolesDeleteAccounts	Synchronize a user with roles and delete accounts.

Modify Request

Use the modify request to update all objects provisioned by the SPML server including the namespace and domain objects.

Fields in a Modify Request

A modify request contains the following fields:

identifier

Identifies the object to be modified

modifications

Lists the attributes to be modified. Attribute values can be added, deleted, or replaced as specified in the *operation* flag.

Example of a Modify Request

The following request sets the *comments* attribute of User `_spml_user` to the string *new comment*.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<modifyRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm1="urn:oasis:names:tc:DSML:2:0:core">
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <modifications>
    <modification name="comments" operation="replace">
      <dsm1:value>new comment</dsm1:value>
    </modification>
  </modifications>
</modifyRequest>
```

Propagate Global User Changes

Modifications made to global user attributes can be propagated to the global user's accounts, by adding syncAccounts attribute to the modify request and setting it to true. The same rule applies if you modify global user complex attributes such as address.

Note: By default the SPML manager does not display the syncAccounts attribute in the Modify tab for any object except from the global user. To propagate changes made to global users complex attributes, add a field by clicking New Modification. Then, specify the name of the field to be syncAccounts and set it to true.

Example of a Modify/Propagate Request

```
<?xml version="1.0" encoding="UTF-8"?>
<modifyRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm="urn:oasis:names:tc:DSML:2:0:core">
  <operationalAttributes></operationalAttributes>
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <modifications>
    <modification name="comments" operation="replace">
      <dsm:value>new comment</dsm:value>
    </modification>
    <modification name="syncAccounts" operation="replace">
      <dsm:value>true</dsm:value>
    </modification>
  </modifications>
</modifyRequest>
```

Example of a Modify/Propagate Complex Attribute Request

```
<?xml version="1.0" encoding="UTF-8"?>
<modifyRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm="urn:oasis:names:tc:DSML:2:0:core">
  <operationalAttributes></operationalAttributes>
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>address@User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <modifications>
    <modification name="city" operation="replace">
      <dsm:value>new city</dsm:value>
    </modification>
    <modification name="syncAccounts" operation="replace">
      <dsm:value>true</dsm:value>
    </modification>
  </modifications>
</modifyRequest>
```

Schema Request

Use the schema request to exchange provisioning schema between the Requesting Authority and SPML Service.

The schema request is used by the Requesting Authority to determine the specific data structures and extended operations that the SPML Service provides access to.

The core eTrust SPML Service schema is identified by the provider identifier *urn:ca.com:etrust:iam* and schema identifier *core*.

Example of a Schema Request

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<schemaRequest xmlns="urn:oasis:names:tc:SPML:1.0">
  <providerIdentifier providerIDType="urn:oasis:names:tc:SPML:1.0#OID">
    <providerID>urn:ca.com:etrust:iam</providerID>
  </providerIdentifier>
  <schemasIdentifier schemasIDType="urn:oasis:names:tc:SPML:1.0#GenericString">
    <schemasID>core</schemasID>
  </schemasIdentifier>
</schemaRequest>
```

Search Request

Use the search request to read the attributes of objects provisioned by the Provisioning Server.

Each object is uniquely identified by an ID, which is similar to an LDAP distinguished name. For example, the following identifier is the unique identifier representing the EXAMPLE_DOMAIN domain:

```
"Domain=EXAMPLE_DOMAIN,Server=Server"
```

The following identifies the user *Administrator* in this domain:

```
"User=Administrator,Domain=EXAMPLE_DOMAIN,Server=Server".
```

A search request allows you to look up objects in a container. In particular, a domain can contain the following objects:

- User
- Group
- Dynamic Group
- Profile
- Password Profile
- Role
- Child Domain
- Namespace

A namespace contains policies and end points, and an end point contains accounts and account containers.

Search Filters

Use an LDAP filter in the search request to identify the objects that you wish to return.

For example, "(name=*)" would list all the objects in this domain.

You can combine several expressions to form a sophisticated filter such as "(&(name=admin*)(role=*newrole*))" in accordance with the LDAP filter search described in RFC 2254.

Fields in a Search Request

A search request contains the following fields:

searchBase

Specifies the starting point for the search operation using the ID string of the container object

filter

Specifies the search criteria

attributes

Lists the attributes to be returned in the search response

Example of a Search Request

The following search request queries the EXAMPLE_DOMAIN to list all objects and return the name and description attributes for all of these objects:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<searchRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsml="urn:oasis:names:tc:DSML:2:0:core">
  <searchBase type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </searchBase>
  <filter>
    <dsml:equalityMatch name="name">
      <dsml:value>*</dsml:value>
    </dsml:equalityMatch>
  </filter>
  <attributes>
    <dsml:attribute name="name"></dsml:attribute>
    <dsml:attribute name="description"></dsml:attribute>
  </attributes>
</searchRequest>
```

Status Request

The Requesting Authority uses a Status Request to query the processing status of an asynchronous operation.

Example of a Status Request

The example below is for requesting the status of a previously-sent asynchronous with the request ID "A4DF567HGD".

```
<statusRequest requestID="A4DF567HGD"
xmlns="urn:oasis:names:tc:SPML:1:0"
xmlns:dsml="urn:oasis:names:tc:DSML:2:0:core"/>
```

Global Settings

Global settings are settings that affect the Provisioning Server Domain. Global settings are changed at the corporate level for entire company. Some of the properties that can be defined in global settings include

- Enabling and disabling self-authentication preferences
- Setting the number of questions
- Setting the number of optional fields

In Provisioning Manager, these settings can be seen from the System (task frame), Global Properties.

Example: Search for Attributes Defined in Global Settings

Here is an example of how to search for attributes defined in global settings.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<searchRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm1="urn:oasis:names:tc:DSML:2:0:core">
  <searchBase type="urn:oasis:names:tc:SPML:1:0:DN">
    <id></id>
  </searchBase>
  <filter>
    <dsm1:equalityMatch name="name">
      <dsm1:value>*</dsm1:value>
    </dsm1:equalityMatch>
  </filter>
  <attributes>
    <dsm1:attribute name="selfAuthEnabled"></dsm1:attribute>
    <dsm1:attribute name="numberSelfAuthQuestions"></dsm1:attribute>
    <dsm1:attribute name="numberOptionalSelfAuthProperties"></dsm1:attribute>
  </attributes>
</searchRequest>
```

Note: You should not specify a value for the searchBase field. If you are doing this search from the SPML Manager you will need to leave the searchBase field empty.

Example: Modify Attributes in Global Settings

Here is an example of how to modify attributes defined in global settings. Note that the ID string "Server=Server" identifies the global settings object.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<modifyRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsmi="urn:oasis:names:tc:DSML:2:0:core">
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>Server=Server</id>
  </identifier>
  <modifications>
    <modification name="autoGenerateUIDs" operation="replace">
      <dsmi:value>true</dsmi:value>
    </modification>
    <modification name="autoGenerateUIDsForNewUsers" operation="replace">
      <dsmi:value>true</dsmi:value>
    </modification>
    <modification name="numberOptionalSelfAuthProperties" operation="replace">
      <dsmi:value>5</dsmi:value>
    </modification>
    <modification name="numberSelfAuthQuestions" operation="replace">
      <dsmi:value>4</dsmi:value>
    </modification>
    <modification name="selfAuthEnabled" operation="replace">
      <dsmi:value>1</dsmi:value>
    </modification>
  </modifications>
</modifyRequest>
```

Account Containers

Hierarchical namespace accounts, such as ADS, LDAP, eWac, NDS, and PLS (CA SSO WAC Namespace), are stored in containers.

An account object in a hierarchical namespace is identified by the following ID string:

```
Account=xyzaccount,Container=ChildContainer,Container=ParentContainer,EndPoint=EXAMPLE_ENDPOINT,
Namespace=EXAMPLE_NAMESPACE,Domain=EXAMPLE_DOMAIN,Server=Server
```

If there is more than one container type in the endpoint, specify "Container.type=" instead of just "Container=" in the ID string. Container types are set to be the same as the LDAP objectClass of the container entry.

Example: Create an Account Container

The following example creates an account container on an ADS endpoint system of type ADSOrgUnit:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<addRequest xmlns="urn:oasis:names:tc:SPML:1.0">
  <identifier type="urn:oasis:names:tc:SPML:1.0#DN">

    <id>Container.ADSOrgUnit=ADSSubContainer,Container.ADSOrgUnit=ADSContainer,EndPoint=EXAMPLE_A
    DS_ENDPOINT,Namespace=ActiveDirectory,Domain=EXAMPLE_DOMAIN,Server=Server</id>
    </identifier>
    <attributes/>
  </addRequest>
```

Example: Create an Account within a Sub-Container

The following example creates an account on an ADS endpoint system within an existing sub-container:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<addRequest xmlns="urn:oasis:names:tc:SPML:1.0" xmlns:dsm="urn:oasis:names:tc:DSML:2.0:core">
  <operationalAttributes></operationalAttributes>
  <identifier type="urn:oasis:names:tc:SPML:1.0#DN">

    <id>Account=EXAMPLE_ACCOUNT,Container.ADSOrgUnit=ADSSubContainer,Container.ADSOrgUnit=ADSC
    ontainer,EndPoint=EXAMPLE_ADS_ENDPOINT,Namespace=ActiveDirectory,Domain=EXAMPLE_DOMAIN,S
    erver=Server </id>
    </identifier>
    <attributes>
      <attr name="objectClass">
        <dsm:value>user</dsm:value>
      </attr>
      <attr name="password">
        <dsm:value>test123</dsm:value>
      </attr>
      <attr name="NT_AccountID">
        <dsm:value>egaccount</dsm:value>
      </attr>
    </attributes>
  </addRequest>
```

Complex Attributes

Objects, such as User, have attributes. Most attributes are of simple types like string, integer, or Boolean. The "address" or "createStatistics" attributes, however, are of complex types as they contain nested elements. For example "street," "city," "country," "state" and "postcode" are nested fields of an "address" attribute. When you add a complex attribute, the Identifier of the complex attribute has the following special format:

attributeName@ID_Of_The_Actual_Object

Some complex attributes are multi-valued, such as the list of self authentication questions and answers for a Global User.

For multi-valued complex attributes, the Identifier format is as follows, with the #index to indicate the index of the attribute value. The index always start from 0:

attributeName#index@ID_Of_The_Actual_Object

Simple attributes can be populated when you add an object:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<addRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm="urn:oasis:names:tc:DSML:2:0:core">
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>User=new_global_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <attributes>
    <attr name="accountid">
      <dsm:value>new_global_user</dsm:value>
    </attr>
    <attr name="firstName">
      <dsm:value>new_global_user</dsm:value>
    </attr>
  </attributes>
</addRequest>
```

But complex attributes must be populated afterwards in a separate AddRequest.

Example: Add a Single-Valued Complex Attribute

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<addRequest xmlns="urn:oasis:names:tc:SPML:1.0" xmlns:dsm="urn:oasis:names:tc:DSML:2.0:core">
  <identifier type="urn:oasis:names:tc:SPML:1.0#DN">
    <id>address@User=new_global_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <attributes>
    <attr name="street">
      <dsm:value>123 Church St </dsm:value>
    </attr>
    <attr name="postcode">
      <dsm:value>3121</dsm:value>
    </attr>
  </attributes>
</addRequest>
```

Example: Add a Multivalued Complex Attribute

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<addRequest xmlns="urn:oasis:names:tc:SPML:1.0" xmlns:dsm="urn:oasis:names:tc:DSML:2.0:core">
  <operationalAttributes></operationalAttributes>
  <identifier type="urn:oasis:names:tc:SPML:1.0#DN">
    <id>selfAuthQA#0@User=new_global_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <attributes>
    <attr name="answer">
      <dsm:value>Sample Answer</dsm:value>
    </attr>
    <attr name="question">
      <dsm:value>Sample Question</dsm:value>
    </attr>
  </attributes>
</addRequest>
```

When you search for an object, asking about a complex attribute, the attribute value returned is a special attribute Identifier that refers to the real attribute value stored in a separate search result entry.

Request Retries

SPML requests, such as add, modify, delete and rename, can be flagged for retry. The request should be asynchronous and should be given a unique requestID. In addition, operational attribute caIamRetry should be set to true.

For more information about operation retries, see the chapter "SPML Service."

Example: N16 Account-Add Request Flagged for Retry

```
<?xml version="1.0" encoding="UTF-8"?>
<addRequest execution="urn:oasis:names:tc:SPML:1:0#asynchronous" requestID="AddN16Account"
xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm="urn:oasis:names:tc:DSML:2:0:core">
  <operationalAttributes>
    <attr name="calamRetry">
      <dsm:value>true</dsm:value>
    </attr>
  </operationalAttributes>
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>Account=new_account,EndPoint=LocalHost,Namespace=Windows
      NT,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <attributes>
    <attr name="password">
      <dsm:value>myPassword</dsm:value>
    </attr>
  </attributes>
</addRequest>
```

Propagate Global User Changes

Modifications made to global user attributes can be propagated to the global user's accounts, by setting the "syncAccounts" attribute to "true" in the modification request.

Example: Modify a Global User and Propagate Changes to Associated Accounts

```
<?xml version="1.0" encoding="UTF-8"?>
<modifyRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm="urn:oasis:names:tc:DSML:2:0:core">
  <operationalAttributes></operationalAttributes>
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <modifications>
    <modification name="comments" operation="replace">
      <dsm:value>new comment</dsm:value>
    </modification>
    <modification name="syncAccounts" operation="replace">
      <dsm:value>true</dsm:value>
    </modification>
  </modifications>
</modifyRequest>
```

But because the SPML Manager doesn't display "syncAccounts" attribute in the "Modify" tab for any object apart from the Global User. So, to propagate changes made to Global Users' complex attributes like "address", the user will have to manually add the operational attribute "syncAccounts" to the modification request. This is done by pressing the "Show Hidden Attributes" button then select "New Operational Attribute". Once the new field is added, specify the name of the field to be "syncAccounts" and set it to true.

Example: Modify Complex Attribute and Propagate Changes to Accounts

```
<?xml version="1.0" encoding="UTF-8"?>
<modifyRequest xmlns="urn:oasis:names:tc:SPML:1:0" xmlns:dsm="urn:oasis:names:tc:DSML:2:0:core">
  <operationalAttributes>
    <attr name="syncAccounts">
      <dsm:value>true</dsm:value>
    </attr>
  </operationalAttributes>
  <identifier type="urn:oasis:names:tc:SPML:1:0#DN">
    <id>address@User=_spml_user,Domain=EXAMPLE_DOMAIN,Server=Server</id>
  </identifier>
  <modifications>
    <modification name="city" operation="replace">
      <dsm:value>new city</dsm:value>
    </modification>
  </modifications>
</modifyRequest>
```

Escaping Special Characters in Object Identifiers

There are two special characters in the Provisioning Server object ID that need to be escaped when using SPML. If there is a comma character in the object name then you will have to use a backward slash to escape the comma. If there is a backward slash in the name then you have to escape it by another backward slash.

For example the identifier

"User=\\new\\,user,Domain=EXAMPLE_DOMAIN,Server=Server" identifies the user by the name "\\new,user" inside the domain EXAMPLE_DOMAIN.

Escaping Special Characters in Search Filters

If you need to search for a pattern that includes a special character `*`, `)`, `(`, `\` or `NULL`, it must be escaped using the format `'\code'` (the code is actually the 2 hexadecimal characters representing the ASCII character) as follows:

- `\2a` replaces or escapes `*`
- `\28` replaces or escapes `(`
- `\29` replaces or escapes `)`
- `\5c` replaces or escapes `\`
- `\00` replaces or escapes `NULL`

Escaped Search Examples

```
(name=*\2a*)           # searches for * anywhere in the name  
  
(file=d:\5cmyfile.html) # searches for d:\myfile  
  
(description=*\28*\29) # searches for both ( and ) anywhere and in that order  
  
(bin=\5b\04)           # searches for binary values (or unicode characters)  
5b04
```


Chapter 5: etautil Batch Utility

You use the etautil batch utility to perform the same tasks as you do with the Provisioning Manager, but from a command line interface. This utility is especially useful for performing repetitive and time-consuming tasks. This chapter explains etautil and provides examples of its use.

Note: etautil sometimes uses the original terminology associated with eTrust Admin, such as namespace and policy, instead of endpoint type and account template. This occurs when you use actual LDAP schema items (object class names, attribute names, attribute values) which retain the original terminology for backwards compatibility. However, etautil also allows the use of user-friendly attribute names as specified in parser tables. These names use the new Identity Manager terminology.

This section contains the following topics:

[Tasks You Can Perform](#) (see page 129)

[etautil Syntax](#) (see page 130)

[Use DeletePending](#) (see page 137)

[Common Error Messages](#) (see page 138)

[Obtain Operation Details](#) (see page 139)

[DOS Output from etautil](#) (see page 140)

Tasks You Can Perform

You can use etautil to maintain property sheets and inclusion pages for Provisioning Manager objects. The following are tasks you can perform with the Batch Utility:

- Create a batch file to explore and correlate endpoint accounts
- Synchronize several accounts with the account template to which they are assigned
- Search and replace attribute values for a large set of objects

For more information about the rules for control statements to use with etautil, see the Provisioning Manager help. For details on using etautil with a specific connector, see the *Connectors Guide*.

etutil Syntax

This is the generalized syntax of etutil. For an explanation of the syntax and use of etutil, see the Provisioning Manager help.

```
etutil [-n] [-d domain] [-u user [-p password]] [-y password-file] [options] control_statements
```

Note: Using an input file provides better performance. A single bind executes all commands in the file.

-n

Verifies the syntax of the command you entered, without executing the command.

-d domain

Specifies the name of the provisioning domain.

-u user

Specifies the global user name for authentication.

-p password

Specifies the password of the named global user for authentication.

Cannot be specified with -y password_file option

-y password-file

Specifies a file name that contains a global user password. Cannot be specified with -p password option. Please see the "Important" section below for more information.

options

Includes any of the following:

-f filename

Reads the control statements in the indicated file and executes them. Use semicolons (;) to delimit multiple control statements.

-i

Invokes the etautil interactive mode, which lets you enter control statements at the prompt. (Use <Ctrl+D> or <Enter> to terminate the interactive mode).

-o

Displays operation details to stdout. See the section Obtain Operation Details.

-h

Displays etautil help.

control statements

For more information about control statements, see etautil Control Statements.

Important: Enter all DNs in the same case as stored in the provisioning directory. DNs are strings that etautil often requires in your commands. In most cases, an incorrect-case DN supplied to the Provisioning Server is processed as is. Authorization errors are common as most permission checking is done by a case-sensitive comparison of the DN of an object being operated upon with the DN specified in a privilege. Copying DN strings from logs or the JXplorer utility ensures the DN is in the correct case.

On UNIX, we strongly recommended you include the -y password-file option to specify an authentication password. For example, if "\$HOME/.pwdfile" contains myglobaluser's password, then you can use etautil command as follows:

```
$ etautil -u myglobaluser -y "$HOME/.pwdfile" <other-options>
```

The command disregards any newline character if one exists at the end of the password file, but it uses the rest of the content as the authentication password.

etutil Control Statements

Control statements tell etutil the procedures to carry out; this is the request that is sent to the Provisioning Server. Use semicolons to delimit multiple control statements in a single etutil command.

Each statement must begin with a verb followed by a base distinguished name (base dn), an object's class name, and the object's operands.

verb basedn classname operands

Note: For more information about control statements see the Provisioning Manager help. For endpoint type-specific details, see the *Connectors Guide*.

The following are examples of the etutil control statements:

ADD

The following example creates role-based accounts for a user:

```
add 'eTGlobalUserContainerName=Global Users,eTNamespaceName=CommonObjects'  
    eTGlobalUser globalusername=denro01 in  
    'eTRoleContainerName=Roles,eTNamespaceName=CommonObjects' eTRole  
    RoleName=TeamManager
```

The following example register a UNIX endpoint:

```
add 'eTNamespaceName=UNIX - etc' eTETCDirectory name=hpdevsrv eTETCHost=hpdevsrv  
    eTETCUnicenterSec=0 eTETCUnicenterUser=0
```

The following example creates a global user named HAAS14 and assigns the values of myvalue1 and myvalue2 to the custom fields with the IDs of 01 and 02.

```
etutil -u etaadmin -p super**s add 'eTGlobalUserContainerName=Global  
    Users,eTNamespaceName=CommonObjects'  
    GlobalUserName=user14 eTCustomField01=myvalue1 eTCustomField02=myvalue2  
    eTPassword=super**s eTUserId=user14
```

Note: You cannot use the ADD statement to add mainframe endpoints to the Provisioning Server.

COPY/COPYALL

Copy creates a new global user with the same properties as an existing global user, including the same roles.

Copyall performs the same function as Copy but also copies the existing user's relationships (inclusions) to the new global user.

Syntax:

```
copy|copyall 'eTGlobalUserContainerName=Global Users,eTNamespaceName=CommonObjects' eTGlobalUser
globalusername=existing_user[.domain] to
globalusername=new_user eTFullName='new fullname'
[property1=value property2=value ... propertyn=value]
```

Example:

```
copyall 'eTGlobalUserContainerName=Global Users,eTNamespaceName=CommonObjects'
eTGlobalUser globalusername=user01 to globalusername=user12 FullName='John Doe'
Password=password EmailAddress=JohnDoe@mycompany.com
```

DELETE

Deletes a global user and its relationships from an endpoint. To delete an object and its inclusion objects, the syntax is:

```
delete basedn classname namingattribute=value
```

To delete an inclusion object, the syntax is:

```
delete childbasedn childclass childnamingattribute=value in parentbasedn parentclass
parentnamingattribute=value [relationship=rel]
```

Note: The deletion of a global user and its accounts can be done using the Update control statement described later in this chapter.

EXPLORE

Finds objects in a registered endpoint and stores them in the provisioning directory. Optionally, correlates or creates a global user in the Provisioning Server for the person associated with each account in the endpoint.

Syntax:

To explore an entire endpoint, the syntax is:

```
explore dirbasedn dirclassname dirnamingattribute=value list [explore options]
```

To explore only a specific container, the syntax is:

```
explore base_dn_container_class_name name=container_name [scope=value] list explore_options
```

The *explore_options* include the following:

- **ExploreUpdateEtrust**-Retrieves all managed objects.
- **ExploreCorrelateUsers**-Correlates accounts with global users using existing ones.
- **ExploreCreateUsers**-Creates global users as needed during the correlation. **ExploreUpdateUsers**-Sets/refreshes the global user attributes using account attribute values.

Note: Combining explore, correlate, and update actions into a single request is not supported.

Examples:

To explore and correlate an entire UNIX endpoint using existing global users:

```
explore 'eTNamespaceName=UNIX - etc' eTETCDirectory
name= hpserve01 list eTExploreUpdateEtrust

explore 'eTNamespaceName=UNIX - etc' eTETCDirectory
name= hpserve01 list eTExploreCorrelateUsers
```

To explore a specific NDS container:

```
explore 'eTNDSOrganizationName=Org1,eTNDSTreeName=SampleTree,eTNamespaceName=NDS Servers'
eTNDSEntityName=OrgUnit1 scope=1
list ExploreUpdateEtrust
```

MASSCHANGE

Sets the same attribute values on a set of objects or searches and replaces attribute values on a set of objects.

Syntax:

```
masschange basedn class criteria [scope=value] to property0=value [property1=value... propertyn=value]
```

where:

criteria-Is the filter for selected target objects.

Scope-Specifies the scope of the search operation (1 for 1-level, 2 for sub-tree level; the default is 1).

propertyn=value-Specifies the attribute to be updated and its new value.

Example:

This example replaces the string (310) with (424) in the eTTelephone value and sets the eTStreetAddress to 15 Software Street for the global users who have eTCity equal to Santa Monica and a name beginning with u:

```
masschange 'eTGlobalUserContainerName=Global Users, eTNamespaceName=CommonObjects'
eTGlobalUser City='Santa Monica' GlobalUserName=u* to Telephone=#sp(310)p(424) StreetAddress='15
Software Street'
```

REPORT

Use REPORT to check account or user synchronization. For more information, see Report Accounts that Do Not Comply with Account Templates.

Syntax:

```
report basedn class namingattr=value list reporting_attribute
```

reporting_attribute-Must be eTSyncAccounts, eTSyncUsers, or eTSyncDelete.

Example:

This example reports all existing accounts that do not comply with the account templates to which they are assigned, for the global user ayrton02:

```
report 'eTGlobalUserContainerName=Global Users,  
eTNamespaceName=CommonObjects' eTGlobalUser globalusername= user02 list eTSyncAccounts
```

UPDATE

Use the Update control statement to do the following:

- Synchronize accounts with account templates.
- Suspend and resume a global user. You can specify that all accounts associated with the global user are also suspended or resumed.
- Change the attributes of an account template and apply those changes to the associated accounts. To propagate those changes to each account assigned to the account template, specify the phrase eTSyncAccounts=1.
- Delete a global user, its relationships, and its accounts.
- Update the attributes values of an existing object.

Syntax:

```
update basedn class namingattribute=value to entries
```

Examples:

To synchronize an account synchronization for a role:

```
update 'eTRoleContainerName=Roles,  
eTNamespaceName=CommonObjects'  
eTRole RoleName=F1Drivers to eTSyncAccounts=1
```

To delete a global user and its accounts:

```
update 'eTGlobalUserContainerName=Global Users,eTNamespaceName=CommonObjects'  
eTGlobalUser globalusername=user02 to DeleteUserAndAccounts=1
```

To remove a value of a multivalued attribute such as eTRoleDN:

```
update 'eTGlobalUserContainerName=GlobalUsers,eTNamespaceName=CommonObjects' GlobalUser  
GlobalUserName=y272705 to -eTRoleDN=  
'eTRoleName=LNDSuspended,eTRoleContainerName=Roles,eTNamespaceName=  
CommonObjects,yourdomainsuffix'
```

This command example uses a plus (+) or minus (-) sign operator in the update section to add or remove values of a multivalued attribute. In this example, there is a minus sign (-) operator before the eTRoleDN attribute to delete an association between Global User and Role.

Multivalued Attributes

Each provisioning role, account template, and global user is an object. Each object has attributes, some of which are multivalued. For example, a Global User may belong to multiple Roles. You may need to update or delete the values for these attributes using plus sign (+) or minus sign (-) operators with the UPDATE command using the following syntax:

```
+attribute_name=attribute_value for adding a value  
-attribute_name=attribute_value for removing a value  
attribute_name=attribute_value for replacing existing value(s) by a new one  
attribute_name="" for clearing existing value(s)
```

Multivalued attributes include the following:

- Account Objects
 - eTPolicyDN (the list of account templates assigned to an account)
 - endpoint type-specific group membership attributes
- Account Template Objects
 - endpoint type-specific group membership attributes
- Global Users Objects
 - eTUserAdminProfile (the list of assigned admin profiles)
 - eTCustomField01 through eTCustomField99 (all global user custom attributes are multivalued)
 - eTAccessControlList (the list of privileges the global user has)

Note: For more information about multivalued attributes, see the *Connectors Guide*.

Use DeletePending

To designate an account as DeletePending, you set two endpoint attributes:

eTAccountDeletable

controls what action the Provisioning Server performs when accounts on an endpoint are deleted. The values are:

0--Enable DeletePending to suspend an account and mark it for later deletion.

1--Disable DeletePending and physically delete the account from the managed endpoint. This is the default value.

2--Enable an alternate delete behavior to remove an account from the Provisioning Server but leave the account unchanged on the managed endpoint.

eTAccountForcedDeletable

controls whether or not an account marked for DeletePending can be deleted through the Forced Delete operation. Use these values:

0--Disable ForcedDelete on DeletePending accounts. This is the default.

1--Enable ForcedDelete on DeletePending accounts.

To track accounts that have been suspended or are in a DeletePending state, use these attributes:

- eTSuspendedDate is the date the account was suspended using the Provisioning Server.
- eTSuspendedTime is the time the account was suspended using the Provisioning Server.
- eTSuspendedReason is either DeletePending or AdminSuspended.

Note: These attributes are only set when an account is suspended using the Provisioning Server. If an account is acted upon by the native endpoint type tools, these attribute values will be stale. If you are taking action based on these attributes, use eTSuspended to confirm whether an account is actually suspended.

Common Error Messages

The following are some common error messages associated with etautil:

Unknown error nnn opening Common Object Repository

This message appears when the authentication to the Provisioning Server fails. If the nnn value in the message is:

102, the user/password is wrong (-u/-p).

96, the domain is wrong (-d).

End of file reached while expecting an operator

Etautil cannot parse the control statements correctly. This message appears when the grammatical syntax is not respected.

Object 'XXXX' operation failed: DB operation failed: Target DN not found.

The target object cannot be reached. This message appears when the base dn is not correct (wrong value for the components of the dn).

Object 'XXXX' operation failed: No server plug-in found for operation

The Provisioning Server is not able to find the connector server corresponding to XXXX. This message appears when XXXX is not correct.

Class 'classname' is not a valid class name

The LDAP name or user-friendly name class name is not defined in the corresponding endpoint type.

Could not find keyword xxxxx for class classname

The LDAP name or user-friendly name xxxxx does NOT correspond to an attribute defined in the class classname. The problem is on 'xxxxx', not on classname.

Obtain Operation Details

You can use the following methods to obtain operation details:

- Specify the -o argument with etautil to display operational details to your standard output device (stdout).
- Specify the OpDetail attribute to control operation detail on a command-by-command basis. By setting this attribute to 0 or 1 you can determine the commands for which you receive operation details. For SELECT and EXPLORE commands, you must set OpDetail in the filter, for example:

```
explore 'eTNamespaceName=Windows NT' eTN16Directory name='My NT Directory' OpDetail=1 List
eTExploreUpdateEtrust
```

For other commands such as UPDATE, you must set OpDetail in the attribute list, for example:

```
update 'eTGlobalUserContainerName=Global Users,eTNamespaceName=CommonObjects' eTGlobalUser
GlobalUserName=gluser01 to
```

```
LastName='gluser01 lastname' OpDetail=1 eTSyncAccounts=1
```

Note: You can combine both methods to obtain operation details for only some of the commands defined in an input file.

This example makes use of a batch input file to run multiple commands that explore a Windows NT endpoint and update a global user name:

```
etautil -o -u etaadmin -p password -f myinputfile
```

where myinputfile contains the following syntax:

```
explore 'eTNamespaceName=Windows NT' eTN16Directory name=My NT Directory OpDetail=0 List
eTExploreUpdateEtrust;
update 'eTGlobalUserContainerName=Global Users,eTNamespaceName=CommonObjects'
eTGlobalUser GlobalUserName=gluser01 to LastName='gluser01 lastname' eTSyncAccounts=1
```

Note: This example makes use of the -o flag to display operation details to stdout. To control the amount of information displayed, use the OpDetail attribute. In this example, by setting OpDetail=0 in the Explore command, only the details of the Update command are displayed.

DOS Output from etautil

When commands are issued directly from the DOS command prompt, non-ASCII 7-bit (ENU) characters are not converted correctly in etautil. This problem occurs because the character set used by DOS (EOM) and Windows (ANSI) are different. The following is a workaround:

- For single-byte non-ASCII (ENU) characters, redirect the output of the etautil command to a text file.

```
etautil [-d DomainName] -u UserName -p Password control statement > Output.txt
```

- For multi-byte non-ASCII (ENU) characters, use an input file that contains etautil control statements you want to execute. Also, redirect the output to a text file.

```
etautil [-d DomainName] -u UserName -p Password -f Input.txt > Output.txt
```

Chapter 6: Provisioning Servers on UNIX

The Provisioning Server, the C++ Connector (SuperAgent) Server and various utilities that work with these servers can run on either Windows or UNIX platforms. For the most part the servers and utilities behave the same and therefore are documented with a single description. This appendix describes the major differences based on the operating system.

This section contains the following topics:

[No UNIX GUI Clients or Utilities](#) (see page 141)

[Command Line Examples](#) (see page 142)

[Libraries and Executables](#) (see page 142)

[Registry Access](#) (see page 143)

[Parser Tables](#) (see page 144)

[UNIX Services for Provisioning](#) (see page 144)

[Working with Hung or Crashed Servers](#) (see page 144)

[Scheduling Periodic Actions](#) (see page 145)

[Passwords on Command Lines](#) (see page 145)

[Server Event Logging Destinations](#) (see page 145)

[Program Exit Definitions](#) (see page 146)

[C++ Connector Server on Solaris](#) (see page 146)

No UNIX GUI Clients or Utilities

Other than installation and web clients, no graphical user interface (GUI) clients or utilities exist on UNIX. For example, Provisioning Manager (etadmin.exe) with its endpoint type-specific GUI plug-ins run only on a Windows system and access the UNIX Provisioning Server remotely. Also, the pwdmgr utility has a different format on UNIX, without a GUI, and must be run from Windows.

Command Line Examples

The Identity Manager documentation includes examples of invoking commands from a command prompt. On Windows, this prompt is a Command (DOS) Window; on UNIX, the prompt use one of various shells. Except where noted, you can assume the examples are for Windows. You can understand the environment variables and path separators that would be necessary to use a given commands on UNIX by replacing Windows pathnames such as

`%VARIABLE%\data\im_ps.conf`

with

`$VARIABLE/data/im_ps.conf`

Also, on UNIX nearly all directories (folders) and file names are in lower-case. Since case is significant in file names on UNIX but insignificant on Windows, some examples in the documentation that refer `PSHOME\Data` function correctly on Windows, even though new installations name that folder `data` instead of `Data`. If you are unsure about the case used for a directory on UNIX, use the `ls` command to locate the exact directory name.

When UNIX examples are given, they apply to any command shell, but were specifically tested to work with the Bourne shell (`/bin/sh`).

Also, note that quoting rules are different in Windows and UNIX command interpreters. Consult the respective interpreter documentation for how to quote or escape data that requires quoting or escaping.

Libraries and Executables

Libraries and executables differ on UNIX and Windows as follows:

Windows

Libraries are named `LibraryName.DLL` (mixed-case and `dll` suffix) and typically installed into a folder such as `PSHOME\bin`.

Executable programs are called `ProgramName.EXE` (mixed-case, `exe` suffix) and installed into `PSHOME\bin`.

Scripts are named `Script.bat` (mixed-case, `bat` suffix).

Message files are named `FileName.DLL` (mixed-case and `dll` suffix) and installed in `PSHOME\bin`.

UNIX

Libraries are named `liblibraryname.so` (lower-case, `lib` prefix and `so` suffix) and installed into a directory such as `$PSHOME/lib`.

Executable programs are called `programname` (lower-case, no suffix) and installed into `$PSHOME/bin`.

Scripts are named `script` or `script.sh` (lower-case, optional `sh` suffix).

Message catalogs are named `filename.res` (lower-case and `res` suffix) and installed in `PSHOME\data`.

Registry Access

On Windows, configuration information is stored in the Windows registry and edited with a native Windows utility such as regedt32 or regedit. On UNIX, the registry is emulated as files in the file system (/opt/CA/SharedComponents/EnterpriseCommonServices/registry). Protect these files as you would the contents of other configuration files. Installation will protect the Identity Manager keys by default. Only imps group users can read them and only the imps user can write them.

To dump out the entire registry, you can use the command `eCSoption /r`. To view, modify or delete specific registry settings that are specific to Identity Manager, use the Identity Manager utility `eta-env`.

For example to view a registry setting, you could use these commands:

```
eta-env action=get name="etrust_bindtodb_need_tls" type=int
eta-env action=get name="logging/calldap_client_logfile" eta-env action=get
name="/enterprise_common_services/installpath"
```

and to set a registry value

```
eta-env action=set name="etrust_bindtodb_need_tls" value=1 type=int
eta-env action=set name="logging/calldap_client_logfile" value=my_file_name
```

Names that begin with / (slash), are relative to:

```
[HKEY_LOCAL_SYSTEM]\SOFTWARE\ComputerAssociates
```

Simple names, without a / (slash), are relative to:

```
[HKEY_LOCAL_SYSTEM]\SOFTWARE\ComputerAssociates\Identity Manager\Provisioning Server
```

Consequently, these two command invocations

```
eta-env action=get name="/Identity Manager\Provisioning Manager/etrust_bindtodb_need_tls" type=int
eta-env action=get name="etrust_bindtodb_need_tls" type=int
```

refer the same configuration parameter.

Note: The registry path of "Identity Manager\Provisioning Server" is set in the \$ETAHOME/data/reg_path.conf.file. The preceding eta-env commands are valid on UNIX and Windows; however, Windows has multiple eta-env.exe commands installed. If you run the eta-env.exe command from the provisioning server installation, it consults the reg_path.conf file from that installation and the registry keys and values are as shown in this section with UNIX. However, if you run the eta-env.exe command that is installed with the provisioning manager installation, it consults the reg_path.conf file from that installation and the registry keys and values being accessed are those under "Identity Manager\Provisioning Manager" instead.

Parser Tables

Parser table files are compiled files with suffix ptt that are installed in PSHOME\data on Windows (\$PSHOME/data on UNIX). They are read by the Provisioning Server and various utilities, such as dumpptt, etautil, showpttdit, and schemagen. The format is a platform-neutral format so that it can be freely copied between Windows and UNIX systems.

UNIX Services for Provisioning

The Provisioning Server (im_ps.exe) and C++ Connector Server (im_ccs.exe) are typically run as services on Windows. Thus you would typically start and stop them on Windows by going to the Services application. Alternatively you could start and stop them from the command line with commands such as `net start im_ps` and `net stop im_ccs`.

On UNIX, the Provisioning Server executable is called slapd and both servers normally start automatically through control files installed in /etc/rc*.d. To view, start, and stop the services manually, you can use commands such as `imps status`, `imps start im_ps`, and `imps stop im_ccs`. The command “imps” is also available as the command “eta” for backwards compatibility with prior eTrust Admin installations.

Working with Hung or Crashed Servers

On Windows, a crashed server may cause information to be written to the system's drwtsn32.log file, a file that CA Customer Support may ask you to send to help analyze the problem.

On UNIX, a crashed server creates a core file in \$PSHOME/bin unless you have configured your server not to generate core files. If a core file is generated, please do not send it to CA unless instructed to do so. Instead, run the command `pstack core > pstack.txt` to capture the stack traces of all threads running within the crashed application. This output is valuable in diagnosing the failure.

On Windows, a hung server (one where one or more requests did not run to completion) can generally only be debugged using the provisioning server trace log (*PSHOME\logs\etatranyyyyymmdd-hhmm.log*) in conjunction with the analyzelog utility. You will generally be asked to capture the provisioning server trace log (at logging level 7 if at all possible) and CA will use “analyzelog” to locate operations that have not yet completed. The C++ Connector Server trace log (*satransyyyymmdd-hhmm.log*) and sometimes other logs are also useful to collect.

On UNIX, capturing the provisioning server trace log and running `analyzelog` is still useful. But another option that often provides additional information is once again the `pstack` command. Locate the process ID (pid) of the hung service by reading the contents of the file `$PSHOME/data/pid/servicename.pid`, and then issue the command `pstack pid > pstack.txt` to capture the stack traces of all active threads within the running process. Please include this output file along with the provisioning server trace logs.

Scheduling Periodic Actions

The UNIX cron command is useful for scheduling periodic tasks such as script invocations. This includes invocations of `etutil` commands (for checking or performing synchronization of accounts or users or performing refresh explore or correlate operations) and invocations of other utilities, such as `etadailybatch` and `etacreateouglocalgroups`. However, using `etutil` for invoking period explore or correlate operations is no longer recommended. Instead you can configure explore/correlate tasks directly within Identity Manager.

Passwords on Command Lines

In UNIX, command-line arguments are public to anyone who can use the `ps` command on the UNIX system. Therefore, you should never supply a password or other sensitive information as a command-line argument. Each Identity Manager command accepts input from a file so you can avoid entering data on the command line. Often, the command-line parameter is still allowed for backwards compatibility with Windows.

Server Event Logging Destinations

Some of the server event logging destinations behave differently on UNIX from how they behave on Windows. In particular, the System log destination logs to `syslogd` on UNIX and to the Windows Event Viewer on Windows. Also, the eTrust Log destination should be avoided on UNIX. It logs to a local file that cannot be viewed locally since there is no UNIX version of the eCS Log Viewer utility. It is not recommended that you run the eCS Log Daemon on UNIX to export the log contents to a remote Windows system because you cannot control who can view the log remotely. These same logging destinations also apply to directory-level logging.

Program Exit Definitions

When defining a common program exit, you enter fields that are interpreted by the Provisioning Server, which invokes the program exit routine. In entering these fields, consider the operating system (Windows or UNIX) of that domain's Provisioning Servers so that these fields work on that operating system. If the domain includes Windows and UNIX Provisioning Servers, be sure that these fields work on both operating systems.

For SOAP program exits, the WSDL can be specified by a URI or its fully qualified name as seen from the Provisioning Server or by a pathname relative to PSHOME\bin (\$PSHOME/bin on UNIX), which is the current working directory of the Provisioning Server.

For DLL program exits, the library name can be a fully qualified name or it can be a common name with or without the lib prefix or the .dll or .so suffix. When only one Provisioning Server exists for a domain, no restrictions exist for how you specify the library name. But when a domain has multiple servers, the library name must be valid for all the servers and since UNIX and Windows have different path syntaxes, files systems, prefixes, and suffixes, the library name should be defined as a common name without any prefix or suffix.

Thus the preferred way to define a program exit object for the CommonExit sample exit is to enter the string CommonExit for library name (in this exact case). The UNIX server will search LD_LIBRARY_PATH for a library named libCommonExit.so. On Windows, this will locate CommonExit.dll by searching the PATH environment variable.

C++ Connector Server on Solaris

The C++ Connector Server installed on Solaris can manage only Solaris UNIX ETC and ACC endpoints. For all other Connectors, install the C++ Connector Server on a Windows system and register it with the Provisioning Server installed on Solaris. During installation, specify that this Connector Server is your default C++ Connector Server.

Chapter 7: Program Exits

This section contains the following topics:

[Program Exits Overview](#) (see page 147)

[Ordering of Program Exit Invocations](#) (see page 148)

[Basic Structure of Program Exits](#) (see page 150)

[Define Common Exits in the Provisioning Manager](#) (see page 150)

Program Exits Overview

Program exits let you write software that executes during certain Provisioning Server actions. Program exits let you reference custom code from in the Provisioning Server process flow, extending the framework of the Provisioning Server to allow additional functionality that changes or augments standard behavior. Numerous exit points are available where custom code can be referenced, depending on the type of object. For example, you may want to install some files on a system every time a UNIX account is created. You could write a program exit that performs the file creations, and specify that it be run whenever a UNIX account is created.

There are two types of program exits:

- *Common Exits* are executed from the Provisioning Server core infrastructure.
- *Native Exits* are executed from the managed endpoint types.

The type of program exit is determined by where it is handled, not where it is referenced.

Note: For information about native exits, see the endpoint type-specific *Connector Guide*.

Program exits are implemented as separate objects, allowing you to define the necessary exits and associate them at the points where they need to be referenced. The following objects reference program exits:

- Common Configuration Objects
- Provisioning Roles
- Account Templates
- Endpoints

Each of these objects can reference multiple program exits, including multiple exits of the same type. For example, a directory can reference two exits that handle routines to be executed before creating an account.

Ordering of Program Exit Invocations

A single request processed by the Provisioning Server may make multiple program exit invocations. The order in which these program exits are invoked depends on:

- The type of program exit
- The location of the program exit reference
- The priority number assigned to the program exit reference

Each program exit type identifies a place in the Provisioning Server's control flow where that particular type of program exits gets a chance to affect the Provisioning Server's behavior. Therefore, to understand the order in which different program exit types are invoked requires understanding how requests are processed.

For instance, a single request to the Provisioning Server might change a global user password and then propagate that password to one or more of that user's accounts. The processing of this request is done as a high-level global user operation that spawns separate account operations.

This results in invoking program exits in the following order:

1. PRE_CHANGE_GLOBAL_USER_PWD
2. PRE_CHANGE_ACCOUNT_PASSWORD
3. POST_CHANGE_ACCOUNT_PASSWORD
4. PRE_CHANGE_ACCOUNT_PASSWORD
5. POST_CHANGE_ACCOUNT_PASSWORD
6. POST_CHANGE_GLOBAL_USER_PWD

In some cases, multiple program exit types apply to the same object class (PRE_CHANGE_GLOBAL_USER_PWD and PRE_MODIFY_GLOBAL_USER) and could potentially be applicable to the same request (a single global user modification that changes both the password and full name, say). In such a case, all exits of one of these exit types will be called before all exits of another of these exit types. But the order is unspecified and you shouldn't assume that the ordering will remain unchanged in future versions of the Provisioning Server.

For a single exit type, sometimes you have a choice as to the class of object on which you define the program exit reference. In particular, references to some exit types that affect global users can be defined on a provisioning role (affecting only users in that role) or on the common configuration object (affecting all users in the domain). If you define exit references on both kinds of objects, then the Provisioning Server invokes the ones defined on the provisioning roles before invoking the ones that are defined on the common configuration object.

Similarly, references to some exit types that affect accounts can be defined on an account template (affecting only accounts assigned to that account template) or on the endpoint (affecting all accounts in the endpoint). If you define exit references on both kinds of objects, then the Provisioning Server invokes the ones defined on the account templates before invoking the ones that are defined on the endpoint.

Finally, exit references of the same type and defined on the same class of object are invoked in priority order using the priority number you assigned when you created the program exit reference, such as priority 1 first, priority 2 second, and so on. Two program exit references of the same priority are invoked in unspecified order.

Basic Structure of Program Exits

Program exits are referred to as pre-operation or post-operation (that is, some operation that Provisioning Manager is used to performing). Program exits have a single common interface for their calling structure. This interface consists of a single argument and a single return value. The single argument is an XML buffer representation, encoded in Unicode Transformation Format 8 (UTF-8), of the object being acted upon combined with any custom information from the definition of the exit. The return value is status information about the result of the program exit execution as well as any documented custom information that is required for a particular exit.

Define Common Exits in the Provisioning Manager

You can define a common exit that will be executed in the Provisioning Server core infrastructure by using the Program Exit property sheet.

Note: To define a Native Exit to be executed in a managed endpoint type, see the specific *Connector Guide*.

To define a common program exit

1. Click Endpoints.
2. Select Common Program Exit from the drop-down list in the Object Type field.
3. Click New. The Common Program Exit dialog appears.
4. Fill in the name and description of the exit to be invoked on the Program Exit tab. If the Disabled box is selected, the program exit will not be invoked, even if it is referenced in another object.
5. Specify whether the Provisioning Server uses the Simple Object Access Protocol (SOAP) or a Dynamic Link Library (DLL) file to invoke the exit on the Common Parameters tab. You can specify that the information be sent securely by selecting SSL Enabled.
6. Enter the path that points to the DLL file or the address of the SOAP service in the Location field. In the Method field, provide the name of the exported function in the DLL file or the name of the function defined by the SOAP service.

For DLL program exits, you can enter for location either a full path, such as:

`c:\yourfolder\yourlibrary.dll`

or you can enter just the common name of the library

`yourlibrary`

If you enter just the common name, you can have more than one Provisioning Server for the domain, where the library does not have to appear at exactly the same path. This is important if the domain has a mix of Solaris and Windows servers, because these operating systems have different pathname syntax.

If you provide just a common name (yourlibrary), the Provisioning Server will locate the library in the following way:

- For Windows, locate yourlibrary.dll file on the Provisioning Server service's execution path as defined by the PATH environment variable. We recommend that you place the library in the PSHOME\bin folder which is known to already be on PATH.
- For Solaris, locate the libyourlibrary.so file on the Provisioning Server service's library path as defined by the LD_LIBRARY_PATH environment variable. We recommend that you place the library into the \$PSHOME/lib directory which is known to already be on LD_LIBRARY_PATH.

7. Select an Authentication Type to provide information for authentication data to be passed to the invoked program exit on the Authentication tab:

- Select None to pass no authentication data to the exit.
- Select Current User to pass the authentication data of the global user who is logged on at the time the exit is invoked.
- Select Proxy User to select a specific global user that will be used for authorization of the operations.

Note: When you select Proxy User, the information you provide must be that of a valid global user.

- Select Other to enable the Authentication Details group field, which lets you select an arbitrary name and password. The exit code uses this information for authentication.

8. Click OK to complete the definition of the common exit.

Chapter 8: Common Program Exit Reference

Common program exits let you write software to run during certain Provisioning Server actions, thereby extending the framework of the Provisioning Server with added functionality. Common program exits are called by the Provisioning Server during processing of user-provisioning operations.

Native program exits are optional exits that may be present in some connectors where such facilities are available and their use is warranted (for example the OS400 connector). Native exits are called from their respective connector plug-in running under the C++ Connector Server.

This section contains the following topics:

[Program Exit Architecture](#) (see page 153)
[Program Exit Hierarchy and Order](#) (see page 154)
[Common Program Exit Structure](#) (see page 155)
[eTExitType](#) (see page 162)
[Custom Function Program Exits](#) (see page 171)
[Sample Flow/Execution Diagram](#) (see page 174)
[Code Examples](#) (see page 174)

Program Exit Architecture

Program exits let you reference custom code from the Provisioning Server process flow. Many entry points are available where custom code can be referenced. In addition, you can invoke program exits as custom functions during policy rule evaluation so you can write custom logic to compute account attribute values. For example, to install some files on a system every time a UNIX account is created, you can write a program exit that creates the file, and indicate that the program exit be run whenever a UNIX account is created.

Program exits belong to the following types:

- *Common exits* are executed in the Provisioning Server core infrastructure.
- *Native exits* are executed in the managed endpoints. For more information about native exits, see the connector guide for the specific endpoint.

Where the program exit is handled determines which type of exit it is, not where the exit is referenced. Program exits are implemented as separate objects in the endpoint and are referenced in these objects, allowing you to define only the exits that are necessary and associate them where they need to be referenced.

The following objects reference program exits:

- Common configuration objects
- Provisioning roles
- Account templates
- Endpoints

Each object can reference multiple program exits, including multiple exits of the same type. For example, an endpoint can reference two PRE_CREATE_ACCOUNT exits.

Program Exit Hierarchy and Order

Program exits are serialized in the Provisioning Server process flow and are both hierarchical and ordered, as described below:

- In terms of hierarchy, the exits are called as referenced from the following objects in the following order:
 - Common configuration objects
 - Provisioning roles
 - Account templates
 - Endpoints
- In terms of order, in a given hierarchy exits are called in a specified order.

An operation on a global user checks for exits to be invoked in the common object and all roles to which the global user belongs. Exits referenced by the common object are invoked before exits referenced by the provisioning role. Similarly, an operation on an account checks the account templates and the endpoints to which the account belongs for exits to be invoked. Exits referenced by the account templates are invoked before exits referenced by the endpoint.

Common Program Exit Structure

Common program exits are referred to in terms of “pre” or “post” in relation to an operation that the Provisioning Server commonly performs. Program exits have a single common interface for their calling structure. This interface consists of a single argument and a single return value. The input argument is an XML buffer representation, encoded in UTF-8, of the object being acted upon combined with any custom information from the definition of the exit. The return value is status information on the result of the program exit execution as well as any documented custom information that is required for a particular exit.

There are two types of common exits:

- DLL deployed
- SOAP executable

Program Exit Input Argument

Program exits have a single interface consisting of a single input argument, which is an XML buffer. All program exits are passed to the XML buffer with the following format:

```
<eTExitInvoke eTExitType={one of the exit types}>
  <{the objectclass of the object being processed}>
  <dn>{the full DN of the object}</dn>
  <name>{the name, that is, RDN value, of the object}</name>
  <{attribute type}>{attribute value}</{attribute type}>
  ...
</{the objectclass of the object being processed}>
<Authentication>
  <Type> </Type>
  <User> </User>
  <Password> </Password>
</Authentication>
</eTExitInvoke>
```

For example:

```
<eTExitInvoke eTExitType=PRE_ADD_ACCOUNT>
  <eTSDKAccount>
    <dn>eTSDKAccountName=test1, eTSDKAccountContainerName=SDK Accounts,
      eTSDKDirectoryName=Team1, dc=Dev</dn>
    <name>test1</name>
    <eTSDKCity>Renton</eTSDKCity>
  </eTSDKAccount>
  <Authentication>
    <Type>GLOBAL_USER</Type>
    <User>{the DN of the global user}</User>
    <Password>{the password of the global user}</Password>
  </Authentication>
</eTExitInvoke>
```

For modify operations, the modify mode is specific in each tag. The possible modify modes are ADD, DELETE, and REPLACE. For example:

```
<eTExitInvoke eTExitType=PRE_MODIFY_ACCOUNT>
  <eTSDKAccount>
    <dn>eTSDKAccountName=test1, eTSDKAccountContainerName=SDK Accounts,
      eTSDKDirectoryName=Team1, dc=Dev</dn>
    <name>test1</name>
    <eTSDKCity modify-mode="replace">Kirkland</eTSDKCity>
    <eTSDKDescription modify-mode="delete">
      Old description</eTSDKDescription>
    </eTSDKAccount>
  </eTExitInvoke>
```

The program exit parses this input argument to get the data it needs to perform its specific task.

If a program exit is defined to handle only a specific type of exit, it should check the eTExitType to make sure that it can handle that specific type. For example, if a program exit is designed to handle exit type PRE_ADD_ACCOUNT, it should check eTExitType and perform only its task, if the exit type is correct. If the exit type is not handled by the program exit, it should do nothing and return a warning.

Input XML Buffer Authentication Type

Each input XML buffer may contain an optional authentication XML block. The format of the authentication XML block is always defined as follows.

```
<Authentication>
  <Type> </Type>
  <User> </User>
  <Password> </Password>
</Authentication>
```

The data in the authentication XML block depends on the type of authentication defined for the program exit. The following are the possible authentication types:

NONE

No credentials are passed to the method being invoked. Thus, the input XML buffer does not contain an authentication block.

GLOBAL_USER

The credentials of the currently logged on global user are passed to the program exit being invoked. The <User> tag contains the DN of the global user. The <Password> tag contains the password for that global user.

Note: The password is not encrypted.

PROXY

The credentials of a specific global user are passed to the program exit being invoked. The <User> tag contains the DN of the specific global user. The <Password> tag contains the password for that global user.

Note: The password is not encrypted.

OTHER

Indicates that the <User> and <Password> tags are program-exit specific. The <User> and <Password> tags can be any free form text. It is up to the program exits to define what these fields mean.

Program Exit Return Value

Program exits have a single return value, which is an XML buffer. Program exits must return an XML buffer, which has the following format:

```
<eTExitReturn>
  <eTExitReturnCategory> </eTExitReturnCategory>
  <eTExitReturnNative> </eTExitReturnNative>
  <eTExitLogMsg> </eTExitLogMsg>
  <eTExitContinue> </eTExitContinue>
  <eTExitCustom> </eTExitCustom>
  <eTPersistentFailure> </eTPersistentFailure>
</eTExitReturn>
```

eTExitReturnCategory XML

Requirement

This value is not required.

Purpose

Groups various native return codes into one of three categories for the purpose of simplifying process flow.

Valid Values

SUCCESS

WARNING

FAILURE

Default Values

If no value is specified, SUCCESS is assumed.

eTExitReturnNative XML

Requirement

This value is not required.

Purpose

Specifies the return value from the native program exit call.

Valid Values

This value is a string representation of what occurred.

Default Values

None.

eTExitLogMsg XML**Requirement**

This value is not required. It is, however, highly recommended to enter a value for failure or warning responses:

- Without eTExitLogMsg value, the server will send the eTExitReturnNative code for logging.
- Without eTExitLogMsg and eTExitReturnNative values, the server will make up a generic message indicating no message present and that there was an error/ warning.

Purpose

Specifies a string value that the native program exit wants the server to log.

Valid Values

This value will be a UTF-8 string.

Default Values

None.

eTExitContinue XML**Requirement**

This value is not required.

Purpose

Specifies whether to continue the process flow after the return from the program exit. This value overrides the default behavior. See Default Values.

Valid Values

TRUE - Continue Execution.

FALSE - Stop Execution.

Default Values

The default values are based on the eTExitReturnCategory attribute.

TRUE - If eTExitReturnCategory is SUCCESS or WARNING.

FALSE - If eTExitReturnCategory is FAILURE.

eTExitCustom XML

Requirement

This value is not required.

Purpose

For common program exits, this value is reserved for future use.

For native exits, this value is connector-specific.

Valid Values

Any valid XML document.

Default Values

None.

The program exit parses this input argument to get the data it needs to perform its specific task.

eTPersistentFailure

Requirement

This value is not required.

Purpose

Used only in responses from IMS Notifications, which share with program exits the same XML buffers for encoding requests and responses. A persistent failure is a notification that is rejected based on a problem in the content (likely a programming error) rather than based on some retry-able situation.

Valid Values

TRUE - Indicates a persistent failure.

FALSE - Indicates a transient failure, one that might succeed later if retried.

Default Values

FALSE

Common Exits DLL Interface

DLL deployed program exits must export the function with the following prototype:

```
int function_name(  
    char    Input_XML,  
    char*   Return_XML,  
    int*    Return_Buffer_Length)
```

The following list describes the parameters for the DLL deployed program exit prototype:

function_name

Name of the program exit. One DLL can export multiple program exits, where each program exit is an exported function with the prototype defined above.

InputXML

Character buffer in UTF-8 format. It contains the XML buffer that the Provisioning Server passes to the program exit.

ReturnXML

Character buffer in UTF-8 format. It is an empty buffer that the Provisioning Server passes to the program exit for it to send a return value back to the Provisioning Server. The size of the buffer is passed to the program exit is the Return_Buffer_Length parameter.

Return_Buffer_Length

Both an input and output parameter:

- On input, Return_Buffer_Length indicates the maximum length, in characters, that the Return_XML buffer can contain. The program exits must not exceed this length when building the return XML buffer.
- On output, Return_Buffer_Length contains the actual length of the return XML buffer the program exit built. That is, after the program exit builds the return XML buffer, it sets Return_Buffer_Length to the actual length of the buffer being returned.

Common Exits SOAP Interface

SOAP-deployed program exits must present an external interface like the following prototype:

```
char* function_name ( char* Input_XML )
```

The following list describes the parameters for the SOAP-deployed program exit prototype:

function_name

Name of the program exit. The Web Services Description Language (WSDL) file that describes the program exit contains a definition of the `function_name`. This is also the name of the character buffer in UTF-8 format that is returned. This buffer is allocated by the referenced program, but must be cleared from the calling program.

Input_XML

Character buffer in UTF-8 format. `Input_XML` contains the XML buffer the Provisioning Server passes to the program exit.

The definition of this interface needs to be presented in the following ways:

- A way that the SOAP client can understand.
- A way that the SOAP server can understand.

The SOAP client relies upon WSDL to specify the interface. For a description of WSDL, see <http://www.w3.org/tr/wsdl.html>.

The SOAP server described here is the Apache SOAP server. The Apache SOAP server requires an XML document known as a Deployment Descriptor. The Deployment Descriptor indicates to the SOAP server what the interface to the SOAP program is. For a more complete description of deployment descriptors, see the Deployment Descriptors section in the *User Guide* at <http://ws.apache.org/soap/docs/index.html>.

An example of a SOAP exit can be found in the following folder:

Samples/ProgramExitSOAP

eTExitType

Exit types determine the circumstances under which an exit is called. A value is entered for `eTExitType`, in the input XML buffer passed to the program exit.

The exit types with `ACCOUNT` in their names can be common or native exits, meaning that common code and connector code can be triggered to process them.

All other exit types, however, must be common exits. It should also be noted that not all program exit types are referenced from the various object types.

Notes:

- In all cases, the name of the object being passed is sent. This is formatted in both DN and Common Name format.
- To have complete control over passwords, either at the global user or the account levels, you must provide exits both for create user/account and for change password user/account. In other words, for new global users (accounts), the change password exit is not called. For new global users (accounts), the password is passed in as part of the attribute for the create exit (for example, PRE_CREATE_GLOBAL_USER).

More information:

[Valid Values for eTExitType](#) (see page 163)

Valid Values for eTExitType

The following values are valid for eTExitType:

PRE_ADD_ACCOUNT

The account information that is being passed to the create account request is also passed to this program. Unlike the Modify operation, the password is passed to the Create operation as part of the account information.

POST_ADD_ACCOUNT

The account information that is being passed to the create account request is also passed to this program. Unlike the Modify operation, the password is passed to the Create operation as part of the account information.

PRE_MODIFY_ACCOUNT

The account information that is being passed to the modify account request is also passed to this program. The only exclusion to this is the password attribute.

POST_MODIFY_ACCOUNT

The account information that is being passed to the modify account request is also passed to this program. The only exclusion to this is the password attribute.

PRE_CHANGE_ACCOUNT_PASSWORD

A special case of MODIFY. This exit is triggered when the password attribute for the account changes. If the password attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The account name and the password attributes contain the only information available to this program exit.

POST_CHANGE_ACCOUNT_PASSWORD

A special case of MODIFY. This exit is triggered when the password attribute for the account changes. If the password attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The account name and the password attribute contain the only information available to this program exit.

PRE_ENABLE_ACCOUNT

A special case of MODIFY. This exit is triggered when the enable attribute for the account changes. If the enable attribute is the only change, the other modify code is not triggered. If other attributes change, this is triggered. The account name is the only account attribute available to this exit.

POST_ENABLE_ACCOUNT

A special case of MODIFY. This exit is triggered when the enable attribute for the account changes. If the enable attribute is the only change, the other modify code is not triggered. If other attributes change, this is triggered. The account name is the only account attribute available to this exit.

PRE_DISABLE_ACCOUNT

A special case of MODIFY. This exit is triggered when the disable attribute for the account changes. If the disable attribute is the only change, the other modify code is not triggered. If other attributes change, this is triggered. The account name is the only account attribute available to this exit.

POST_DISABLE_ACCOUNT

A special case of MODIFY. This exit is triggered when the disable attribute for the account changes. If the disable attribute is the only change, the other modify code is not triggered. If other attributes change, this is triggered. The account name is the only account attribute available to this exit.

PRE_DELETE_ACCOUNT

Triggered prior to a DELETE request. The account name is the only account attribute available to this exit.

POST_DELETE_ACCOUNT

Triggered after a DELETE request. The account name is the only account attribute available to this exit.

PRE_ADD_GLOBAL_USER

The global user information that is being passed to the create request is also passed to this program.

POST_ADD_GLOBAL_USER

The global user information that is being passed to the create request is also passed to this program.

PRE_MODIFY_GLOBAL_USER

The global user information that is being passed to the modify request is also passed to this program. The only exclusion to this is the password attribute.

POST_MODIFY_GLOBAL_USER

The global user information that is being passed to the modify request is also passed to this program. The only exclusion to this is the password attribute.

PRE_CHANGE_GLOBAL_USER_PWD

A special case of MODIFY. This exit is triggered when the password attribute for the global user changes. If the password attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The global user name, the password, and optionally the password clue attributes are the only information available to this program exit.

POST_CHANGE_GLOBAL_USER_PWD

A special case of MODIFY. This exit is triggered when the password attribute for the global user changes. If the password attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The global user name, the password, and optionally the password clue attributes are the only information available to this program exit.

PRE_ENABLE_GLOBAL_USER

A special case of MODIFY. This exit is triggered when the enable attribute for the global user changes. If the enable attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The global user name is the only attribute available to this exit.

POST_ENABLE_GLOBAL_USER

A special case of MODIFY. This exit is triggered when the enable attribute for the global user changes. If the enable attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The global user name is the only attribute available to this exit.

PRE_DISABLE_GLOBAL_USER

A special case of MODIFY. This exit is triggered when the disable attribute for the global user changes. If the disable attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The global user name is the only attribute available to this exit.

POST_DISABLE_GLOBAL_USER

A special case of MODIFY. This exit is triggered when the disable attribute for the global user changes. If the disable attribute is the only change, the other modify code is not triggered. If other attributes change, this code is triggered. The global user name is the only attribute available to this exit.

PRE_DELETE_GLOBAL_USER

Triggered prior to a DELETE request. The global user name is the only attribute available to this exit.

POST_DELETE_GLOBAL_USER

Triggered after a DELETE request. The global user name is the only attribute available to this exit.

PRE_ASSOCIATE_ROLE

Refers to the changing of provisioning role membership in the Provisioning Server, regardless of what happens at the account level. The global user name and the provisioning role name is the only information available to this program exit.

POST_ASSOCIATE_ROLE

Refers to the changing of provisioning role membership in the Provisioning Server, regardless of what happens at the account level. The global user name and the provisioning role name is the only information available to this program exit.

PRE_DISASSOCIATE_ROLE

Refers to the changing of provisioning role membership in the Provisioning Server, regardless of what happens at the account level. The global user name and the provisioning role name are the only information available to this program exit.

Note: This value is called only for incremental provisioning role changes. If you use a replace-mode modification of global user's provisioning roles to replace one set of provisioning roles with another, this value calls the associate-role exits only. The exit would not read the database to find out which provisioning roles were previously included to see which were being set that were previously set and which were being removed.

POST_DISASSOCIATE_ROLE

Refers to the changing of provisioning role membership in the Provisioning Server, regardless of what happens at the account level. The global user name and the provisioning role name are the only information that is available to this program exit.

Note: This value is called only for incremental provisioning role changes. If one uses a replace-mode modification of global user's provisioning roles to replace one set of provisioning roles with another, this value calls the associate-role exits only. The exit would not read the database to find out which provisioning roles were previously included to see which were being set that were previously set and which were being removed.

PRE_ADD_GLOBAL_GROUP

The global group information that is being passed to the add request is also passed to this program.

POST_ADD_GLOBAL_GROUP

The global group information that is being passed to the add request is also passed to this program.

PRE_MODIFY_GLOBAL_GROUP

The global group information that is being passed to the modify request is also passed to this program.

POST_MODIFY_GLOBAL_GROUP

The global group information that is being passed to the modify request is also passed to this program.

PRE_DELETE_GLOBAL_GROUP

Triggered prior to a delete request. The global group name is the only attribute available to this exit.

POST_DELETE_GLOBAL_GROUP

Triggered after a delete request. The global group name is the only attribute available to this exit.

PRE_ADD_ROLE

The provisioning role information that is being passed to the add request is also passed to this program.

POST_ADD_ROLE

The provisioning role information that is being passed to the add request is also passed to this program.

PRE_MODIFY_ROLE

The provisioning role information that is being passed to the modify request is also passed to this program.

POST_MODIFY_ROLE

The provisioning role information that is being passed to the modify request is also passed to this program.

PRE_DELETE_ROLE

Triggered prior to a delete request. The provisioning role name is the only attribute available to this exit.

POST_DELETE_ROLE

Triggered after a delete request. The provisioning role name is the only attribute available to this exit.

CUSTOM_FUNCTION

Triggered when a program exit is invoked through a policy rule expression such as %\$funcname(%UN%,%AC%)%.

Containment

Containment refers to the allowed combination of objects and program reference type, and not to X.500 containment.

Common Configuration Object

The Common Configuration Object is used to assign program exits for the global user, global user group, or provisioning role object classes. For example, if certain program exits should be called when global users are processed, the common configuration object should reference those exits. In addition, the common configuration object is needed to provide a way to call exits during the add operation.

- PRE_ADD_GLOBAL_USER
- POST_ADD_GLOBAL_USER
- PRE_MODIFY_GLOBAL_USER
- POST_MODIFY_GLOBAL_USER
- PRE_CHANGE_GLOBAL_USER_PWD
- POST_CHANGE_GLOBAL_USER_PWD
- PRE_ENABLE_GLOBAL_USER
- POST_ENABLE_GLOBAL_USER
- PRE_DISABLE_GLOBAL_USER
- POST_DISABLE_GLOBAL_USER
- PRE_DELETE_GLOBAL_USER
- POST_DELETE_GLOBAL_USER
- PRE_ADD_GLOBAL_GROUP
- POST_ADD_GLOBAL_GROUP
- PRE_MODIFY_GLOBAL_GROUP
- POST_MODIFY_GLOBAL_GROUP
- PRE_DELETE_GLOBAL_GROUP
- POST_DELETE_GLOBAL_GROUP
- PRE_ADD_ROLE
- POST_ADD_ROLE
- PRE_MODIFY_ROLE
- POST_MODIFY_ROLE

- PRE_DELETE_ROLE
- POST_DELETE_ROLE

Provisioning Roles

A role object can reference program exits to assign the exits that are invoked for various operations on global users associated with that provisioning role. If a provisioning role references a program exit, those exits are called in addition to the exits referenced by the Common Configuration Object. The exits defined on the common configuration object are invoked before exits defined on the provisioning role (hierarchy order).

When adding a global user (PRE_ADD_GLOBAL_USER and POST_ADD_GLOBAL_USER exit types), program exits are invoked based on the initial set of provisioning roles being assigned to the user.

When associating or disassociating a provisioning role with a global user (PRE_ASSOCIATE_ROLE, POST_ASSOCIATE_ROLE, PRE_DISASSOCIATE_ROLE and POST_DISASSOCIATE_ROLE exit types), the program exits referenced by the provisioning roles being associated or disassociated are invoked.

When modifying an existing global user in other ways (using the exit types listed below), all provisioning roles to which the global user belongs are consulted to identify program exits to invoke.

Other Exit Types

The common configuration object handles add exits for a global user.

- PRE_ADD_GLOBAL_USER
- POST_ADD_GLOBAL_USER
- PRE_ASSOCIATE_ROLE
- POST_ASSOCIATE_ROLE
- PRE_DISASSOCIATE_ROLE
- POST_DISASSOCIATE_ROLE
- PRE_MODIFY_GLOBAL_USER
- POST_MODIFY_GLOBAL_USER
- PRE_CHANGE_GLOBAL_USER_PWD
- POST_CHANGE_GLOBAL_USER_PWD
- PRE_ENABLE_GLOBAL_USER
- POST_ENABLE_GLOBAL_USER

- PRE_DISABLE_GLOBAL_USER
- POST_DISABLE_GLOBAL_USER
- PRE_DELETE_GLOBAL_USER
- POST_DELETE_GLOBAL_USER

Account Templates

An account template object can reference program exits to affect the accounts associated with that template.

If an account template references program exits, these exits are called in addition to the exits that are referenced by the endpoint to which the account belongs. The exits defined on the account template are invoked before the exits on the endpoint (hierarchy order).

If an account is being created from one or more account templates (PRE_ADD_ACCOUNT and POST_ADD_ACCOUNT exit types), those template exits are called.

When working with an existing account, whether the current set of assigned account templates is being adjusted or not, it is the initial set of assigned templates whose program exits are invoked.

- PRE_ADD_ACCOUNT
- POST_ADD_ACCOUNT
- PRE_MODIFY_ACCOUNT
- POST_MODIFY_ACCOUNT
- PRE_CHANGE_ACCOUNT_PASSWORD
- POST_CHANGE_ACCOUNT_PASSWORD
- PRE_ENABLE_ACCOUNT
- POST_ENABLE_ACCOUNT
- PRE_DISABLE_ACCOUNT
- POST_DISABLE_ACCOUNT
- PRE_DELETE_ACCOUNT
- POST_DELETE_ACCOUNT

Endpoints

Endpoint objects are used to assign program exits to accounts. For example, if certain program exits should be called when accounts are processed, the endpoint objects should reference those exits.

In addition, endpoint objects are needed to provide a way to call exits during an add operation.

- PRE_ADD_ACCOUNT
- POST_ADD_ACCOUNT
- PRE_MODIFY_ACCOUNT
- POST_MODIFY_ACCOUNT
- PRE_CHANGE_ACCOUNT_PASSWORD
- POST_CHANGE_ACCOUNT_PASSWORD
- PRE_ENABLE_ACCOUNT
- POST_ENABLE_ACCOUNT
- PRE_DISABLE_ACCOUNT
- POST_DISABLE_ACCOUNT
- PRE_DELETE_ACCOUNT
- POST_DELETE_ACCOUNT

Custom Function Program Exits

A custom function program exit is invoked from an account template rule expression. Custom function program exits share the following characteristics:

- The exit type is always CUSTOM_FUNCTION. There are no PRE or POST variants.
- The exit must be a common program exit (DLL or SOAP). Native exits cannot be used to compute the custom function.
- The exit must be registered in the same domain as the account being created or updated from the policy. The reference to a custom function program exit (%\$funcname(...)%) contains the name of the exit (funcname), but there is no rule string syntax to let you specify the domain of the program exit so it is always presumed to be in the domain of the account.
- The input to the program exit includes zero or more single- or multi-valued parameters.

For example, if the global user for the account being created or updated has the following attribute settings:

eTCustomField01: { value1a, value1b } (that is, two values assigned)
eTCustomField02: value2

the `%%*$FuncName(%*UCU01%, %UCU02%)%` rule expression is evaluated.

The input XML passes all values of eTCustomField01 and the value of eTCustomField02 as follows:

```
<eTExitInvoke eTExitType=CUSTOM_FUNCTION>
  <eTFunction>
    <eTFuncParam1>value1a</eTFuncParam1>
    <eTFuncParam1>value1b</eTFuncParam1>
    <eTFuncParam2>value2</eTFuncParam2>
  </eTFunction>
  <Authentication>
    <Type>GLOBAL_USER</Type>
    <User>{the DN of the global user}</User>
    <Password>{the password of the global user}</Password>
  </Authentication>
</eTExitInvoke>
```

The output from the program exit can indicate an error (as with any other program exit) so that the creation or update of the account is not attempted, or can contain a single- or multi-valued output parameter. For example, a program exit could return the following XML block to indicate two values (ReturnValue1 and ReturnValue2) to set for the corresponding account attribute:

```
<eTExitReturn>
  <eTExitReturnCategory>SUCCESS</eTExitReturnCategory>
  <eTExitReturnNative>0</eTExitReturnNative>
  <eTExitContinue>TRUE</eTExitContinue>
  <eTExitCustom>
    <eTFuncReturn>ReturnValue1</eTFuncReturn>
    <eTFuncReturn>ReturnValue2</eTFuncReturn>
  </eTExitCustom>
</eTExitReturn>
```

The function rule expression controls the number of values to set as follows:

- If the `$FuncName` in the rule expression is preceded by `*` (asterisk) as in the example above, this will set 0 or more values of the attribute depending on what is included in the output XML document.
- If the function rule expression does not have the `*` preceding `$FuncName`, only the first value returned is relevant. Additional values are ignored.

Obscured Returned Values

The program exit returns information inside the eTFuncReturn XML block, for example:

```
<eTFuncReturn>Returned value from program exit</eTFuncReturn>
```

If logging is enabled, then this XML block can be read.

However, if the program exit returns information like a password, then you may not want the information to be logged. In this case, you can flag the returned value as obscured to prevent it from being logged.

The format of the obscured value is:

```
<eTFuncReturn obscured="yes">MyPassword</eTFuncReturn>
```


This tells the Provisioning Server to replace the value with the string **** NOT SHOWN **** as it does for attribute names that are recognized as storing sensitive attributes. For example:

```
<eTFuncReturn obscured="yes">** NOT SHOWN **</eTFuncReturn>
```

Note: The obscured attribute is case-sensitive. For example, the result will not be replaced correctly if the attribute is set to "YES".

Sample Flow/Execution Diagram

The following illustration provides a sample flowchart of the execution logic of a program exit.



Code Examples

Use the code examples located in the following directory as a guide when coding your Common Program Exits:

Samples\ProgramExits
Samples\ProgramExitSOAP

Chapter 9: Program Exits In Connectors

This chapter covers information about supporting program exits you have created in the eTrust Admin SDK for the connectors you have created.

This section contains the following topics:

[Execution Flow \(Logic\)](#) (see page 175)

[Support for Common Exits](#) (see page 177)

[Support for Native Exits](#) (see page 177)

[Exit Types](#) (see page 181)

[Code Examples for Program Exits in Options](#) (see page 182)

Execution Flow (Logic)

Program exits are referred to as either pre-exits or post-exits, depending on the operation that the Provisioning Server performs.

Pre-Exits

The Provisioning Server framework invokes a pre-exit before executing a particular operation. For common exits, the Provisioning Server framework interprets the return XML buffer to determine whether to continue execution of the particular operation. For native exits, the agent plug-in must interpret the return XML buffer.

If the agent plug-in returns a success status code (LDAP_SUCCESS), the Provisioning Server continues to perform the operation. If the agent plug-in returns an error code, the operation is aborted.

Note: For native program exits, your custom connector agent plug-in must interpret the return XML buffer. Furthermore, your agent plug-in must return either success or failure to the Provisioning Server framework. Success lets the operation continue. Failure aborts the operation.

If one pre-exit fails, the operation will be aborted even if other pre-exits let the operation continue. In addition, as soon as a pre-exit aborts the operation, other pre-exits (with the same or lower priority) will not be called.

Priority of Pre-Exits

Pre-exits are called in order of hierarchy and priority. If two program exits are referenced with the same priority, the order in which they are called is undefined. Priority guarantees only that higher priority exits are called before lower priority exits. The highest priority is "1" and the larger the number, the lower its priority.

More information:

[Program Exit Hierarchy and Order](#) (see page 154)

Operation

Once the Provisioning Server framework has invoked all pre-exits associated with the operation and each pre-exit lets the operation continue, the Provisioning Server framework executes the operation. Execution of the operation may result in another request to your agent plug-in.

Note: The agent plug-in receives multiple requests: one for the pre-exit, one for the operation, and one for the post-exit.

Post-Exit

Once the operation is completed successfully, the Provisioning Server framework invokes the post-exits referenced for that operation. Post-exits should be handled the same way as pre-exits. Your agent plug-in should not differentiate between a pre-exit and a post-exit, and you can use the exact same code to handle pre- and post-exits.

Priority of Post-Exits

Post-exits are called in order of hierarchy and priority. If two program exits are referenced with the same priority, the order in which the exits are called is not guaranteed. Priority only guarantees that higher priority exits are called before lower priority exits. The highest priority is "1" and the larger the number, the lower its priority.

More information:

[Program Exit Hierarchy and Order](#) (see page 154)

Support for Common Exits

Common exits are processed by the Provisioning Server framework. Thus, supporting common exits requires minimal changes. You only need to enhance your GUI plug-in. No agent plug-in change is needed.

Parser Table Enhancement

The parser table files already define new attributes for exit reference. Option parser table files include these parser table files, so you do not need to make any changes to your parser table.

GUI Plug-In Enhancement

Provisioning Manager provides a standard property page for referencing program exits. To support common exits, you only need to add this property page to your account and endpoint property sheets. This enables your account and endpoint objects to reference program exits.

The standard exit reference property page is managed by the `CosExitRefPage` class in the COS module. For example, to add the property page to your account and directory property sheets, add the following line to your property sheet code:

```
propertyPages->AddTail(new CosExitRefPage(this));
```

Agent Plug-In Enhancement

No agent plug-in change is needed to support common exits.

Support for Native Exits

Native exits are processed by the endpoint. Thus, to support native exits, the endpoint must enhance both its GUI plug-in and agent plug-in.

Parser Table Enhancement

The parser table files already define new attributes for exit reference. Parser table files include these parser table files, so you do not need to make any changes to your parser table.

Since you are providing native program exits, you need to define your custom program exit object in the parser table. The exit object must have the following CLASS definition line:

```
CLASS Exit.<exit class name>,eTExit.<exit class name>,etavlc,secobjar
```

For example, the SDK defines the exit object class as follows:

```
CLASS Exit.SDKExit,eTExit.eTSDKExit,etavlc,secobjar
```

For a complete example of how to define an exit object, see the SDK `sdkparse.pty` sample file, which is provided with the SDK.

GUI Plug-In Enhancement

Exit Reference

The Provisioning Server GUI framework provides a standard property page for referencing program exits. To support common exits, you need to add the property page to your account and directory property sheets. This lets your account and endpoint objects reference program exits.

The standard exit reference property page is managed by the `CosExitRefPage` class in the COS module.

For example, to add `CosExitRefPage` to your account and directory property sheets, add the following line to your property sheet code:

```
propertyPages->AddTail(new CosExitRefPage(this));
```

Exit Definition

A property sheet must be provided to define a native program exit. This endpoint-specific exit definition property sheet is used to enter specific information that will be passed to the agent plug-in during exit invocation.

Your endpoint must define an XML format for this data, which is stored as an XML buffer in the `eTExitPayload` attribute in the custom program exit object. When an exit is invoked, the Provisioning Server framework sends this data to the agent plug-in. The agent plug-in parses the `eTExitPayload` attribute to get the data it needs to invoke the program exit.

Agent Plug-in Enhancement

Program Exit Invocation Request

The Provisioning Server framework sends all native program exit invocation requests to the endpoint. Even if the exit is referenced by an account object, the invocation request is sent to the code that manages the endpoint modify operation. Specifically, the directory `DEmodify()` function is called.

On a program exit invocation request, the Provisioning Server framework includes the `eTExitPayload` and `eTExitInvoke` attributes.

`eTExitPayload` contains the data regarding the definition of the exit. The value of `eTExitPayload` is the XML buffer stored in the program exit object that was defined by the program exit definition property sheet that you added to your GUI plug-in.

`eTExitInvoke` is an XML buffer. This data should be passed to the program exit, which needs to process it. The agent plug-in can process this information; however, often it does not need to do so.

The agent plug-in must be enhanced for performing the following tasks to support the program exit:

1. Determine whether an operation is an exit invocation request.
2. Invoke the program exit.
3. Interpret the result from the program exit.

More information:

[Program Exit Input Argument](#) (see page 155)

Determine Exit Invocation Request

Typically, the directory `DEmodify()` function processes requests to change values in the directory object. To support native program exits, you must enhance the directory `DEmodify()` function to also handle native exit invocation.

For a program exit invocation request, the Provisioning Server framework sends the `eTExitInvoke` attribute as part of the modify operation. The presence of the `eTExitInvoke` attribute is an indication that the request is an exit invocation request and not a normal modify request, as shown in the following example:

```
/*
|| The special attribute UTFEXITINVOKE indicates a request to invoke a
|| program exit.
*/
if (pMods->find_mod(UTFEXITINVOKE)) {

    // Invoke the exit.

}
else {

    // Normal directory modify request.

}
```

Invoke the Program Exits

You must define how your custom connector agent plug-in invokes the program exit.

The common exit has the following invocation methods:

- Through the DLL function call
- Through the SOAP method invocation

Your agent plug-in will probably define some other form of program exit invocation. That data is passed to the agent plug-in in the eTExitPayload attribute, which is an XML buffer.

The method of invoking program exits is to execute a command line utility. Thus the only information it needs is the utility name (including the path). You can define the SDK exit object as having a payload that only contains the full path to the utility.

The sample SDK exit payload is the following:

```
<eTSDKExit>
  <Program>program to execute</Program>
</eTSDKExit>
```

Interpret the Result from the Program Exit

Each program exit must return an XML buffer.

The agent plug-in must interpret this return XML buffer and return an appropriate status code to the Provisioning Server framework. For pre-exits, returning a success status to the Provisioning Server framework lets the operation continue. Returning a failure status will abort the operation.

Note: If the operation has multiple pre-exits, one pre-exit might return a success, which would let the operation continue; but another pre-exit could return failure, thus aborting the operation. If one pre-exit aborts the operation, it is aborted, even if other pre-exits let the operation continue. In addition, as soon as a pre-exit aborts the operation, other pre-exits with the same or lower priority will not be called.

More information:

[Program Exit Return Value](#) (see page 157)

Exit Types

The Provisioning Server framework only sends an exit invocation request to the agent plug-in if the exit type is one that can be handled by the agent plug-in. In general, your agent plug-in does not need to handle exit types. However, if your custom connector only permits certain program exit types, it must check the eExitType tag attribute in the eExitInvoke attribute.

Exit Type Functionality

Exit type is a value that determines the circumstances under which an exit is called. One of the types of exits is entered for eExitType (in the input XML buffer passed to the program exit). The first 12 exit types (values) can be common or native exits, that is, common code and namespace (connector) code can be triggered to process them. The remaining exit types, however, can be common exits only. It should be noted that not all program exit types are referenced from various object types.

Notes: In all cases the name of the object being passed is sent. This is formatted in both DN and Common Name format.

To have complete control over passwords (either at the global user or the account levels), you must provide exits both for create user/account and for change password user/account. In other words, for new global users (accounts), the change password exit is not called. For new global users (accounts), the password is passed in as part of the attribute for the create exit (for example, PRE_CREATE_GLOBAL_USER).

More information:

[Valid Values for Exit Types](#) (see page 163)

Code Examples for Program Exits in Options

See the SDK sample connector. The exit handling code is in the SDKDirectory.cpp file.

The method SDKDirectory::DEmodify() determines whether a request is an exit invocation request, and if it is, calls the SDKDirectory::InvokeExit() method.

Chapter 10: Provisioning Maintenance

This section contains the following topics:

[Back Up and Restore CA Directory](#) (see page 183)

[Shut Down the Provisioning Server service](#) (see page 183)

[View and Maintain Log Files](#) (see page 184)

[Provisioning Directory Monitoring](#) (see page 190)

Back Up and Restore CA Directory

To ensure that data is coherent across your entire organization, regular backups should be done. Regular backups of CA Directory prevent data loss and damage caused by network disasters and failures. CA Directory provides an online backup utility (and the `dxdumpdb` and `dxloaddb` utilities for offline backup) to back up and restore CA Directory.

Note: For information about these utilities, see the *CA Directory Administrator Guide*.

Shut Down the Provisioning Server service

If the Provisioning Server service does not shut down, you can manually shut it down as follows:

1. Open a command prompt and enter the following command:

```
net stop im_ps
```

2. If Services indicates that the Provisioning Server service is still in the stopping state, issue the following commands:

```
net start im_ps
```

```
net stop im_ps
```

A similar procedure can be used to manually shut down the Provisioning Connector Server service, whose service name is `im-ccs`.

If the service still does not stop, open the Task Manager, select `im_ps.exe` (or `im_ccs.exe`) on the Processes tab, and click End Process.

View and Maintain Log Files

The provisioning components (Provisioning Server, Connector Servers, Provisioning Manager) can be configured to log information about all transactions that they process. You can use this information to predict and identify the sources of system or security problems. For example, if the warning messages in log files show that some accounts on an endpoint could not be explored, you can use the logged information to investigate those accounts and determine why they were not explored. Use a text editor to view and edit provisioning log files.

Server Event logs track messages generated by the Provisioning Server. You can log messages to several optional destinations, including CA Audit.

The provisioning components provide other types of logging to diagnose specific problems. Other than the provisioning server trace log, these logs are usually not enabled unless you need them to trace a particular event. They include provisioning server logs, slapd logs, and C++ Connector Server logs. You can also diagnose problems that occur when communicating with the provisioning server by enabling Provisioning Manager logging.

Messages from all logs are written to text files in the *PSHOME*\Logs directory and are named accordingly:

- Provisioning Server Event Log — *etayyyyymmdd.log*
- Provisioning Server Trace Log — *etatransyyyymmdd-hhmm.log*
- Provisioning Server IMS Notification Log — *etanotifyyyyymmdd-hhmm.log*
- Provisioning Server SLAPD Log — *im_ps.log*
- Provisioning Manager Log — *etaclientyyyymmdd.log*
- C++ Connector Server Endpoint Log — *sayyyyymmdd.log*
- C++ Connector Server Trace Log — *satransyyyymmdd-hhmm.log*
- C++ Connector Server SLAPD Log — *im_ccs.log*

Server Event Logging

Server Event logs record important events generated from the Provisioning Server. These events consist of all *severity levels* (success, information, warning, fatal, and error). The logs record every client-initiated operation and its success or failure, including generated sub-operations.

In the System Task frame of the Provisioning Manager, under Global Properties, use the Logging tab to configure Server Event logging. Server Event logs typically only need to be configured once.

In some cases, you can turn logging on or off, or you can configure the severity levels of the messages logged. Thus, this Server Event logging can serve to audit the activities that are taking place within the Provisioning Server. However, the preferred auditing of provisioning activity is to enable the IMS Notifications features. The IMS Notifications feature sends detailed audit records to the IMS server for inclusion in the full audit record of Identity Manager activity. The notification records sent to the IMS can also trigger events for additional Identity Manager Server processing.

Endpoint Logging

In the Endpoint Task frame, you can configure endpoint-specific logging. Endpoint logs track messages that a connector generates when it processes requests for objects residing in that endpoint. Each endpoint can be configured separately so you can turn logging on or off for just the endpoints where you need to learn additional information to diagnose problems.

You can also specify the severity (success, information, warning, fatal, and error) of the messages that get logged.

To turn logging on or off and to set the logging destinations and the severity levels of the messages logged for each directory, use the Logging tab of the endpoint's property sheet in the Provisioning Manager. For detailed instructions, see Setting Endpoint Logging in the Provisioning Manager help.

Endpoint logging is sent to a log file for the connector server in which the connector for the endpoint runs. For C++ connectors, the default log file name is *PSHOME*\Logs\saYYYYMMDD.log. The C++ connector server also adds some additional messages to this log. You control the log file name in the *im_ccs.conf* using the *BaseLogFileName* parameter. And you control which severities of these other messages are logged in the same conf file using the *LogSeverities* parameter.

Endpoint logging from connectors which run directly within the provisioning server (for example, the CA ACF2 connector) log to the provisioning server's event log which has the default name of *PSHOME*\Logs\etaYYYYMMDD.log.

Diagnostic Logging

To diagnose specific problems, you can enable the provisioning server trace log, slapd logs, or C++ Connector Server logs. These are typically not enabled unless you need them to trace a specific type of event. Provisioning Manager logging also is used for diagnosing problems in the Provisioning Manager or client utilities.

Provisioning Server Trace Log

Enable this logging component to generate a special transaction log file that records the details of every transaction processed by the Provisioning Server. You can choose from several logging levels to match the level of logging detail you prefer using the domain configuration parameter Transaction Log/Level.

The Provisioning Server trace log writes messages to *PSHOME\Logs\etaTransyyyymmdd-hhmm.log*. To change the base part of the file name (the part before the date) or to relocate this log file to another drive, modify the domain configuration parameter Transaction Log/File name. For more information about the *etaTransyyyymmdd.log* file, see the Provisioning Manager help.

Note: Unlike most logging which is turned off by default, Provisioning Server logging is fully enabled as the component is installed. If you choose not to run with maximum trace logging of the provisioning server, you need to change the domain configuration parameters that control this logging. These parameter are located in the "Transaction Log" parameter folder in the Provisioning Manager on the System task under Domain Configuration.

Provisioning Server IMS Notification Log

The Provisioning Server is typically configured to send notifications (global user and other object change records) to the Identity Manager Server for integration with the IMS event system and audit data base. A notification thread running within the Provisioning Server reads notification records from the local notify DB and transmits them to the IMS. This activity is captured in the IMS Notification log, whose name is *PSHOME\Logs\etanotifyyyyymmdd-hhmm.log*.

You configure the severity of log messages included in this log on the Identity Manager Setup screen in the System Task of Provisioning Manager.

The format of this log is similar to the Provisioning Server and Connector Server trace logs.

SLAPD and C++ Connector Server Logs

On Windows, you can enable SLAPD logging for advanced debugging tasks such as LDAP protocol packet handling and search-filter processing. You can set the log level in the Windows registry by assigning a value to the DebugLevel key. There are two registry keys, each controlling the logging for one of the services:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\slapd\im_ps\CurrentVersion\DebugLevel

The im_ps registry key controls logging for im_ps.exe, run by the Provisioning Server service.

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\slapd\im_ccs\CurrentVersion\DebugLevel

The im_ccs registry key controls logging for im_ccs.exe, run by the Connector Server service.

Important! The preferred method for enabling SLAPD logging is by setting the loglevel parameter in im_ps.conf or im_ccs.conf, for both Windows and Solaris. Each file contains configuration instructions.

The DebugLevel registry key or loglevel configuration file parameter specifies the amount of information the server writes to its log file, which is one of the following, depending upon your type of slapd service:

- PSHOME\Logs\im_ps.log
- PSHOME\Logs\im_ccs.log

Note: A "TLS: can't accept" error message may appear in the im_ps.log file when running in FIPS mode due to a low-level initialization problem that clears up after the first connection from a client. Since clients retry connections, you can ignore this message.

You can select a debug level to match the type of debugging you want to perform. The debug levels are listed in the following table:

Value	Debug Information
1	Trace function calls
2	Debug packet handling
4	Heavy trace debugging
8	Connection management
16	Print out packets sent and received
32	Search filter processing

Value	Debug Information
64	Configuration file processing
128	Access control list processing
256	Stats log connections/operations/results
512	Stats log entries sent
1024	Print communication with shell back-ends
2048	Entry parsing
65535	All tracing

C++ Connector Server Trace Logging

C++ Connector Server Trace Logs record the activity of the C++ Connector Server, which is a module used to help manage many endpoint types. This log performs the following functions:

- Logs trace and debug messages for the C++ Connector Server.
- Monitors all statuses returned by its connectors. For example, if a connector returns fatal LDAP errors, the C++ Connector Server logs these errors with severity LOG_FATAL.

To set the log file name and logging levels in `im_ccs.conf` set the `SATransLog` and `SATransLogLevel` parameters. The supported logging levels are 0 (for off) and 1 (for on). The default is 0. These parameters must exist in the file after the database superagent line.

Provisioning Manager Logging

To diagnose problems communicating with the server, you can set logging to record events that transpire between the Provisioning Manager and the Provisioning Server to which it is connected. Use the Logging tab under File, Preferences to trace all requests sent to any server from the Provisioning Manager.

This logging is actually logging within the C/C++ client library used by the Provisioning Manager and some other clients (batch utility, password manager, `csfconfig`, `bindeta`, `pingeta`). Once logging is enabled and configured using Provisioning Manager, those log settings apply for these other clients as well. Each client logs its command name as it logs messages so you can identify which log messages are specific to which client.

However, for this to work the client being run must reside in same file system folder as the Provisioning Manager's etadmin.exe program. When this is not the case (such as when running on Solaris where there is no Provisioning Manager install, or even on Windows when you run utilities from the Provisioning Server's installation), the client library consults registry settings specific to the Provisioning Server instead of specific to the Provisioning Manager. Set these other registry settings by running these eta-env commands using the eta-env program included in the Provisioning Server installation:

```
eta-env
  action=set
  name=Manager/LogMaster
  type=int
  value=1

eta-env
  action=set
  name=Manager/LogDestinations
  type=int
  value=16

eta-env
  action=set
  name=Manager/LogSevFile
  type=int
  value=31
```

These have the effect of configuring the C/C++ client library for the provisioning server's installation, setting the destination to "text file" and logging all message severities.

Finally, the csfconfig command has a "debug=yes" command-line parameter you can specify to turn this logging on for one command invocation overriding any registry settings configured with Provisioning Manager or eta-env.

Use AnalyzeLog

The command line utility, AnalyzeLog, takes as input a Provisioning Server trace log (etatransyyyyymmdd-hhmm.log) and produces different views of the information depending on what options you set. You can use this information to diagnose functional or performance problems reported by users.

Note: For more details on this utility, see the Provisioning Manager online help.

Log Files for High Availability

To ensure proper operation of your high-availability configuration, you should monitor the following log files:

- Alarm
- Warn
- Stats
- Diag
- Summary
- Trace

All logs can be flushed through the DXserver console. Only the SUMMARY and TRACE logs can be closed from the console.

Provisioning Directory Monitoring

Monitor Thresholds

Monitoring is one of the most significant activities as part of the system availability and security. Monitoring is needed to determine the source of performance problems, fault detection, and to take corrective action.

Important! Regular monitoring of the multiwrite queue is highly recommended. For more information on using multiwrite, see the information on replication in *CA Directory Administrator Guide*.

CA Directory can provide SNMP counters to an SNMP-aware Enterprise Management Application. CA Directory can also provide SNMP traps under the following conditions:

- Authentication Failure
- Alarms
- Directory Updates

You can also use the CA Directory Statistics logs to gauge application load over time. This can be invaluable in providing data to both measure and show the growth of the new service.

Observe Router Traffic

CA Directory includes the DXconsole utility, which you use to monitor the network traffic going through an CA Directory router.

The router knowledge files, such as `imps_router.dxc`, define the console port for a router. You can adjust the console port number to satisfy your enterprise needs. To monitor traffic sent through a router, start up DXconsole and connect it to the localhost with the port number defined in the router knowledge file.

Note: Because the Provisioning Server uses LDAP for communication between various components, you can set the trace level to `ldap`. For more information, see the *CA Directory Administrator Guide*.

Enable SSL Encryption

CA Directory includes a `dxcertgen.exe` utility, which you can use to generate certificates and personality files. Using this utility is the recommended method for enabling encryption for high availability solutions.

For more information, see the *CA Directory Administrator Guide*.

Index

A

- Account Containers • 121
- Account Templates • 170
- Add a New SPML Service • 82
- Add Request • 108
- Admin Profile Privilege Cache • 23
- Admin Profiles • 13
- Administrator Authentication • 12
- Administrator Authorization • 13
- Administrator Login • 12
- Advanced Configuration Options • 15
- Advanced Configuration Options Overview • 15
- Agent Plug-in Enhancement • 179
- Agent Plug-In Enhancement • 177
- Authentication Parameters • 20
- Authentication/Disable Maintenance User • 20
- Authorization Parameters • 21
- Authorization/Check Owner Access on Indirect Privileges • 21

B

- Back Up and Restore CA Directory • 183
- Basic Structure of Program Exits • 150
- Batch Request • 108
- Benefits of Using SPML • 73

C

- C++ Connector Server on Solaris • 146
- C++ Connector Server Trace Logging • 188
- CA Product References • iii
- Cache Parameters • 22
- Cancel Request • 110
- Class 'classname' is not a valid class name • 139
- CMDRA Command Options • 87
- CMDRA Commands • 87
- CMDRA Examples • 89
- Code Examples • 174
- Code Examples for Program Exits in Options • 182
- Command Line Examples • 142
- Command Line Requesting Authority (CMDRA) • 76
- Common Configuration Object • 168
- Common Error Messages • 138
- Common Exits DLL Interface • 160

- Common Exits SOAP Interface • 162
- Common Program Exit Reference • 153
- Common Program Exit Structure • 155
- Compatibility Parameters • 28
- Complex Attributes • 123
- Configuration Setup Parameters • 29
- Configure Retry for a Request • 105
- Configure SPML Client Computer to Support SSL Security • 86
- Configure SSL Support for Tomcat Servers • 84
- Connections Parameters • 29
- Connections/CS Pool Maximum Size • 30
- Connections/CS Pool Minimum Size • 30
- Connections/DB Pool Maximum Size • 30
- Connections/DB Pool Minimum Size • 30
- Connections/Expiration Time • 31
- Connections/Other Pool Maximum Size • 31
- Connections/Other Pool Minimum Size • 31
- Connections/Refresh Time • 32
- Connector Server Cache • 27
- Contact CA • iii
- Containment • 168
- Could not find keyword xxxxx for class classname • 139
- Create an SPML Template Request • 96
- Custom Function Program Exits • 171

D

- Data Transformations • 98
- Default Admin Profiles • 14
- Define Common Exits in the Provisioning Manager • 150
- Delete an Existing Service • 84
- Delete Request • 110
- Determine Exit Invocation Request • 179
- Diagnostic Logging • 185
- Domain Cache • 24
- Domain Configuration • 17
- DOS Output from etautil • 140
- Download the SPML Manager • 95

E

- Enable Operation Details • 28
- Enable SSL Encryption • 191

- End of file reached while expecting an operator • 138
- Endpoint Logging • 185
- Endpoint Parameters • 32
- Endpoint/Check Account Passwords • 33
- Endpoint/Check Empty Account Passwords • 35
- Endpoint/Use Account Template Status • 35
- Endpoint/Validate Endpoint Credentials • 36
- Endpoints • 171
- Escaping Special Characters in Object Identifiers • 126
- Escaping Special Characters in Search Filters • 127
- etutil Batch Utility • 129
- etutil Control Statements • 132
- etutil Syntax • 130
- eTExitType • 162
- Example
 - Add a Multivalued Complex Attribute • 124
 - Add a Single-Valued Complex Attribute • 124
 - Create an Account Container • 122
 - Create an Account within a Sub-Container • 122
 - Modify a Global User and Propagate Changes to Associated Accounts • 125
 - Modify Attributes in Global Settings • 121
 - Modify Complex Attribute and Propagate Changes to Accounts • 126
 - Search for Attributes Defined in Global Settings • 120
- Example Data Transformation • 98
- Example of a Batch Request • 109
- Example of a Cancel Request • 110
- Example of a Delete Request • 111
- Example of a Modify Request • 115
- Example of a Modify/Propagate Request • 116
- Example of a Schema Request • 117
- Example of a Search Request • 118
- Example of a Status Request • 119
- Example of an Add Request • 108
- Example of an Extended Request • 111
- Execution Flow (Logic) • 175
- Exit Definition • 178
- Exit Reference • 178
- Exit Type Functionality • 181
- Exit Types • 181
- Explore and Correlate Parameters • 36
- Explore and Correlate/Correlation Attribute • 37
- Explore and Correlate/Correlation Domain • 40

- Explore and Correlate/Create Users Default Attributes • 40
- Explore and Correlate/Create Users Domain • 41
- Explore and Correlate/Create Users Verify Not Correlated • 41
- Explore and Correlate/Explore Compare in Memory • 43
- Explore and Correlate/Map User ID to Lowercase • 42
- Extended Request • 111
- Extended Request Types • 112

F

- Fields in a Delete Request • 110
- Fields in a Modify Request • 115
- Fields in a Search Request • 118
- Fields in an Add Request • 108
- Fields in an Extended Request • 111
- Flow of the SPML Feed Command • 94

G

- Global Properties • 16
- Global Settings • 119
- Global User Group Privilege Cache • 24
- Global User Privilege Cache • 25
- GUI Plug-In Enhancement • 177, 178

H

- How the SPML Service Works • 77

I

- Identity Manager Server Parameters • 44
- Identity Manager Server/Enable Corporate User Access • 44
- Identity Manager Server/Enable Notification • 44
- Identity Manager Server/Notify Batch Size • 45
- Identity Manager Server/Notify Retry Time • 45
- Identity Manager Server/Notify Timeout • 45
- Identity Manager Server/Use External Password Policies • 46
- Input XML Buffer Authentication Type • 156
- Install SPML • 79
- Interpret the Result from the Program Exit • 180
- Invoke the Program Exits • 180

L

- Libraries and Executables • 142

- List Templating Variables • 97
- Log Files for High Availability • 190
- Log On to the SPML Configuration Application • 81
- Log user-friendly Attribute and Object Class Names • 71

M

- Modify an Existing Service • 83
- Modify Request • 114
- Monitor Thresholds • 190
- Multivalued Attributes • 136

N

- No UNIX GUI Clients or Utilities • 141
- Notification Config Cache • 26

O

- Object 'XXXX' operation failed
 - DB operation failed
- Target DN not found. • 138
 - No server plug-in found for operation • 138
- Obscured Returned Values • 173
- Observe Router Traffic • 191
- Obtain Operation Details • 139
- Operation • 176
- Operation Cache • 26
- Operation Details Parameters • 47
- Operation Details/Maximum Operation Detail • 47
- Operation Details/Operation Details Expiration Time • 48
- Operation Details/Operations Folder • 48
- Ordering of Program Exit Invocations • 148

P

- Parser Table Enhancement • 177
- Parser Tables • 144
- Password Parameters • 51
- Password Profile Cache • 27
- Password Synchronization Parameters • 49
- Password Synchronization/Agent Response Threshold • 50
- Password Synchronization/Update Only Global User • 51
- Passwords on Command Lines • 145
- Passwords/Enforce Synchronized Account Passwords • 52
- Passwords/Pre-expire Passwords • 52

- Passwords/Store User Passwords • 54
- Post-Exit • 176
- Pre-Exits • 175
- Priority of Post-Exits • 176
- Priority of Pre-Exits • 176
- Processes Parameters • 54
- Processes/Catch Program Exit Exceptions • 55
- Processes/Child Operation Thread Pool Size • 55
- Processes/Parallel Propagation • 56
- Processor Parameters • 56
- Processor/Process Affinity Mask • 57
- Processor/Process Priority • 57
- Program Exit Architecture • 153
- Program Exit Definitions • 146
- Program Exit Hierarchy and Order • 154
- Program Exit Input Argument • 155
- Program Exit Invocation Request • 179
- Program Exit Return Value • 157
- Program Exits • 147
- Program Exits In Connectors • 175
- Program Exits Overview • 147
- Propagate Global User Changes • 115, 125
- Provisioning Directory Monitoring • 190
- Provisioning Directory Parameters • 20
- Provisioning Directory/Entry Count Attribute • 20
- Provisioning Maintenance • 183
- Provisioning Manager • 11
- Provisioning Manager Logging • 188
- Provisioning Roles • 169
- Provisioning Server • 11
- Provisioning Server IMS Notification Log • 186
- Provisioning Server Trace Log • 186
- Provisioning Servers on UNIX • 141

R

- Registry Access • 143
- Relax Self Q&A Reads • 28
- Rename an Existing Service • 83
- Request Execution Types • 107
- Request Retries • 124
- Request Types • 108
- Requesting Authorities • 79
- Retry Architecture • 103
- Retry Configuration Files • 104
- Retrying SPML Requests • 102

S

- Sample Flow/Execution Diagram • 174

- Sample SPML Requests • 107
- Scheduling Periodic Actions • 145
- Schema Request • 116
- Search Filters • 118
- Search Parameters • 57
- Search Request • 117
- Search/Allow Partial Results • 58
- Search/Max Scope Filter Objects • 59
- Search/Search Size Limit • 59
- Server Event Logging • 184
- Server Event Logging Destinations • 145
- Servers Parameters • 62
- Shut Down the Provisioning Server service • 183
- SLAPD and C++ Connector Server Logs • 187
- SPML Architectural Diagram • 75
- SPML Architecture • 74
- SPML Configuration Application • 76
- SPML Feed • 90
- SPML Feed Command Options • 92
- SPML Integration • 78
- SPML Manager • 76
- SPML Overview • 73
- SPML Service • 73, 75
- SPML Service Configuration • 81
- SPML Support for FIPS 140-2 • 80
- SPML Templates • 78
- Statistics Parameters • 62
- Statistics/Enabled • 63
- Statistics/Node Stats from Connection • 63
- Status Request • 119
- Support for Common Exits • 177
- Support for Native Exits • 177
- Synchronization Parameters • 64
- Synchronization/Automatic Correlation • 64
- Synchronization/Remove Account Template Values from Accounts • 65
- Synchronization/Use Existing Accounts • 66

T

- Tasks You Can Perform • 129
- Transaction Log Parameters • 68
- Transaction Log/Enable • 68
- Transaction Log/Enable/Configuration • 68
- Transaction Log/Enable/Connector Server Framework • 68
- Transaction Log/Enable/LDAP • 69
- Transaction Log/File Name • 69
- Transaction Log/Level • 70

U

- Uninstall the SPML Service • 81
- UNIX Services for Provisioning • 144
- Unknown error nnn opening Common Object Repository • 138
- Use AnalyzeLog • 189
- Use DeletePending • 137
- User Interface for Provisioning • 11
- Using the SPML Manager's Templating Functionality • 95
- Using Velocity Templates • 97

V

- Valid Values for eTExitType • 163
- View and Maintain Log Files • 184

W

- When You Would Use the SPML Service • 74
- Working with Hung or Crashed Servers • 144
- WS Mapper • 79