CA Identity Manager

Implementation Guide

r12.5 SP6

CA Technologies Product References

This document references the following CA products:

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

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information you need for your Home Office, Small Business, and Enterprise CA Technologies 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 Technologies product documentation, you can send a message to techpubs@ca.com.

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

Contents

Chapter 1: Managing Identities and Access

This section contains the following topics:

User Management and Application Access

The typical Information Technology (IT) department faces a constant demand to maintain user accounts. IT administrators must address urgent needs of users, such as resetting forgotten passwords, creating new accounts, and providing supplies and office equipment.

Simultaneously, IT administrators must provide users with various levels of access to applications. For example, a department manager generates purchase orders and needs an account in a financial application.

To address the escalating demands on IT, CA Identity Manager provides an integrated method of managing users and their access to applications, including:

- Assignment of privileges through roles. Specifically:
  - Roles that enable administrators to create and maintain user accounts
  - Roles that provision additional accounts to existing users (requires provisioning support)
- Delegation of the management of users and application access
- Self-service options so users can manage their own accounts
- Integration of business applications with CA Identity Manager
- Options to customize and extend CA Identity Manager

Role-Based Entitlements

With Identity Manager, you assign privileges to users by assigning roles. A *role* contains tasks that correspond to application functions in Identity Manager or account templates that correspond to additional accounts. When you assign a role to a user, that user can perform the tasks contained in the role or use the accounts associated with the role.

Identity Manager provides these types of roles:

- User management roles are called *admin roles*.
- Account assignment roles are called *provisioning roles*.

Roles simplify privilege management. Instead of associating a user with each task that he performs or each account that he needs, you can assign a role to the user. The user can perform the tasks in the role or use the accounts associated with the role.

You can then edit the role by adding tasks or account templates, which define the accounts. Every user who has the role can now perform the new task or use the new account. If you remove a task or account template from a role, the user can no longer perform that task or use the account.

Tasks enable users to perform Identity Manager functions, such as modifying a profile.

The following illustration shows several tasks which are combined into a single admin role and assigned to multiple users:

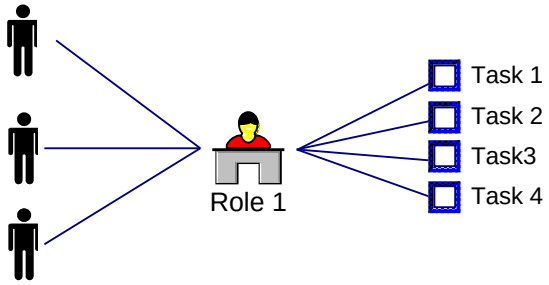

For a provisioning role, you combine several accounts, such as an email account, a database account, and an Active Directory account. You can assign the role to several users, who each need these accounts.

The following illustration shows several accounts which are combined into a single provisioning role and assigned to multiple users.

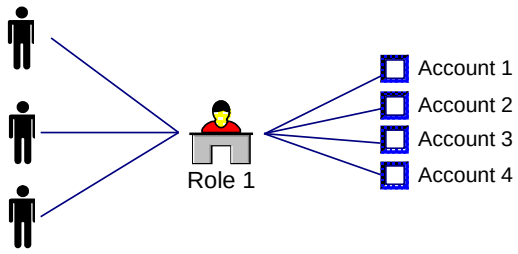

Note: In this figure, each user receives four accounts, when you assign the provisioning role to that user.

Admin Roles for User Account Management

In CA Identity Manager, you manage user store objects (users, groups, and organizations) through admin roles. You also use admin roles to manage the roles and tasks through which you manage user store objects. For example, you use admin roles to modify profile attributes of users, give users options for managing their own accounts, and to approve tasks that use workflow.

Profile Management at the Attribute Level

You can create admin roles for different administrators who need to read or write different profile attributes. For example, a company may have several employees who perform operations on user profiles, each accessing different attributes. The following figure shows three roles and their associated tasks. Each role has different access to profile attributes.

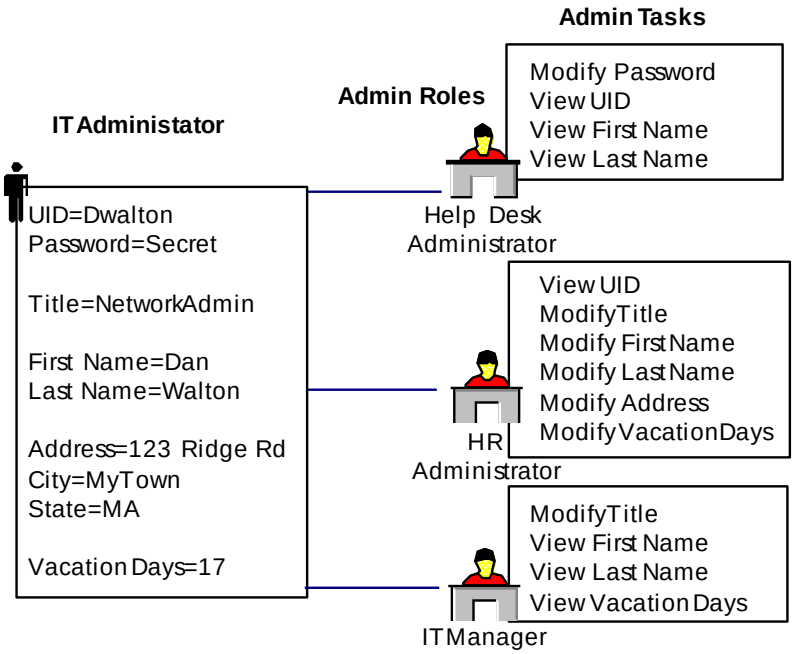

In this example, three roles can manage different attributes for the same user, Dan Walton:

- A Help Desk administrator views user names and addresses and resets user passwords.
- A Human Resources administrator modifies user IDs, user names, addresses, titles, and number of vacation days.
- An IT manager modifies the title of users and views their name and number of vacation days.

Whatever roles you have when you log in to CA Identity Manager, a series of tabs, called categories, appear based on the admin role assigned to your CA Identity Manager account. You click a tab to see the tasks that you can perform in that category as shown in the following figure:

The categories and the tasks in those categories that a user sees are determined by the user's admin roles.

Workflow Approval of Admin Tasks

To help automate business processes, you can design an admin task to generate a workflow process. A *workflow process* automates a well-defined procedure that a company repeats frequently. CA Identity Manager includes the WorkPoint workflow engine.

Workflow processes are triggered by CA Identity Manager events which are part of an admin task. For example, the Create User task includes events called CreateUserEvent and AddToGroupEvent. When an event occurs, the workflow engine can:

- Require approvals--An approver must approve an event, such as modifying a user profile, before CA Identity Manager updates a user store. Approvers are administrators who have the Approver role for a particular task.
- Send notifications--The workflow engine can notify users of an event's status at different stages of a process, such as when a user initiates an event or when an event is approved.
- Generate work lists--Work lists specify the tasks that a particular user must perform. The workflow engine updates administrators' work lists automatically.

For common events, you can use the workflow processes supplied with CA Identity Manager. Alternatively, you can create custom workflow processes.

Provisioning Roles for Additional Accounts

In CA Identity Manager, you provide additional accounts to users by using provisioning roles. Provisioning roles contain account templates, which define accounts that exist in managed endpoints, such as an email server. Once you have users in CA Identity Manager, you can assign provisioning roles to some of those users. The user receives the accounts defined by the templates in the role.

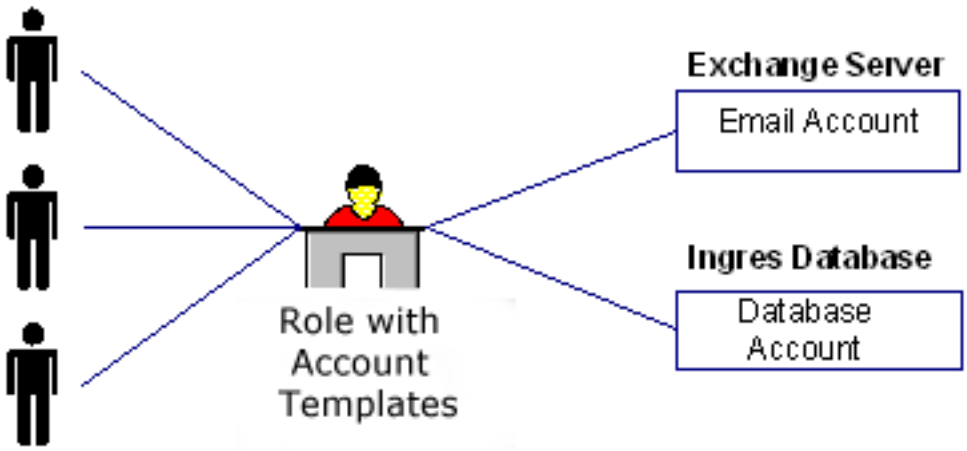

The account templates define the characteristics of the account. For example, an account template for an Exchange account might define the size of the mailbox. The account templates also define how user attributes are mapped to accounts.

To be able to use provisioning roles, you must install the Provisioning Server with the Identity Manager server. Then, you create account templates in the User Console.

Password Management

Identity Manager includes several features for managing user passwords:

- Password Policies—These policies manage user passwords by enforcing rules and restrictions governing password expiration, composition, and usage.

  Note: For advanced password policies, configure integration with SiteMinder. For more information, see the *Installation Guide*.

- Password Managers—Administrators who have the Password Manager role can reset a password when a user calls the Help desk.
- Self-Service Password Management—Identity Manager includes several self-service tasks that allow users to manage their own passwords. These tasks include:
  - Self Registration—Users specify a password when they register at a corporate web site.
  - Change My Password—Users can modify their passwords without help from IT or Help Desk personnel
  - Forgotten Password—Users can reset or retrieve a forgotten password after Identity Manager verifies their identity.
  - Forgotten User ID—Users can retrieve a forgotten user ID after Identity Manager verifies their identity.
- Password Synchronization (for use with provisioning only)—Password changes are synchronized in Identity Manager and in accounts on target systems called endpoints. New passwords are verified against Identity Manager password policies.

Self Service Options for Users

To further reduce the IT workload, CA Identity Manager includes features for registering new users and supplying a forgotten password. These features require no administrator involvement. The user gains access to CA Identity Manager through a *public console*, which requires no login account. Through this console, a user can self-register at a site or request a reminder about a forgotten password.

To save the time of IT administrators, CA Identity Manager users can manage their own accounts. Because users have a self-management role, they can:

- Maintain personal information
- Change their own password
- Join self-subscribing groups

Identity Manager Customization and Extensibility

You customize these CA Identity Manager features:

- The Identity Manager directory, which describes a user store structure to CA Identity Manager.
- The appearance and functionality of the user interface.
- User entry screens, which determine the fields and layout of each task screen.
- Validation of user data entry, through regular expression, JavaScript, or Java implementations.
- Workflow, which defines automated workflow processes. Create or modify processes by linking approvers and actions in the WorkPoint Process Designer.
- Email messages, which inform users of a task's status.
- Task submission, which can be sent by a third-party application to the Identity Manager Task Execution Web Service (TEWS). TEWS processes the remote task request. Remote task requests comply with WSDL standards.

You can extend CA Identity Manager's functionality using the following APIs:

- Logical Attribute API—Enables you to display an attribute differently than how it is stored physically in a user directory.
- Business Logic Task Handler API—Allows you to perform custom business logic during data validation or transformation operations.
- Workflow API—Provides information to a custom script in a workflow process. The script evaluates the information and determines the path of the workflow process accordingly.
- Participant Resolver API—Enables you to specify the list of participants who are authorized to approve a workflow activity.
- Event Listener API—Enables you to create a custom event listener that listens for a specific Identity Manager event or group of events. When the event occurs, the event listener can perform custom business logic.

- Notification Rule API—Lets you determine the users who should receive an email notification.
- Email Template API—Includes event-specific information in an email notification.

Note: For more information on the CA Identity Manager APIs, see the *Programming Guide for Java*.

When Identity Manager includes provisioning, you can also extend provisioning functionality as follows:

- Custom connectors—Enable communication between a Provisioning Server and an endpoint system. The code that makes up a connector can include a GUI plug-in, server plug-in, and agent plug-in.

  A dynamic connector can be generated by Connector Xpress, and a custom static connector can be developed in Java or C++.

  Note: For more information, see the *Connector Xpress Guide*.
- Program exits—Let you reference custom code from the Provisioning Server process flow.

Note: For more information about extending provisioning functionality, see the *Programming Guide for Provisioning,* which is available in the Legacy Components media.

CA RCM Integration

CA RCM is an identity lifecycle management product that enables you to quickly and accurately develop, maintain, and analyze role models. It also provides centralized identity compliance policy controls and automates processes associated with meeting compliance and security demands. Using CA RCM, you can do the following:

- Validate that users have appropriate entitlements
- Ensure that privileges are granted in accordance with security policies
- Monitor the effectiveness of identity management controls
- Understand what roles exist in your organization, and then establish a role model that fits your organization
- Analyze and maintain that role model as business evolves

CA Identity Manager integrates with CA RCM in two ways:

- CA Identity Manager Connector

  The Connector for CA Identity Manager automatically synchronizes the role-based privilege data between CA Identity Manager and CA RCM. By using the connector, you can import data from CA Identity Manager to CA RCM or export data from CA RCM to CA Identity Manager.

  Note: For more information about the CA Identity Manager Connector, see the *CA RCM Connector for CA Identity Manager*.

- Smart Provisioning

  When CA Identity Manager integrates with CA RCM, you can configure additional functionality that allows you to use role and compliance information, which is available in a role model, to support day-to-day identity management operations. Changes made in CA Identity Manager dynamically update the role model in CA RCM.

Note: For more information about the integration with CA RCM, see the *Administration Guide*.

CA Enterprise Log Manager Integration

CA Enterprise Log Manager uses the CA Common Event Grammar (CEG) to map events that originate in various systems in a standard format, and stores all events, even those which are not yet mapped, for review and analysis. Furthermore, CA Enterprise Log Manager provides users with a high-volume solution for managing and reporting on collected data, using configurable database queries and/or reports to search for various types of information and events.

CA Enterprise Log Manager provides better wider and deeper insight into un-managed systems and systems outside of CA Identity Manager's purview and control and also lets you investigate deeper into identities.

Integrating with CA Identity Manager lets you view CA Enterprise Log Manager identity centric reports and/or dynamic queries into CA Enterprise log Manager user Console using the Identity Manager User Console. From the User Console you can configure how existing CA Identity Manager/Enterprise Log Manager reports and/or queries are viewed and modified while you investigate deeper into a specific identity.

CA Enterprise Log Manager Reports

The following CA Enterprise Log Manager Reports are provided with CA Enterprise Log Manager role definitions by default:

Task	Invokes Report
System All Events by User	CA Identity Manager - System All Events filtered by user ID
Account Management by Host	Account Management by Host
Account Creations by Account	Account Creations by Account
Account Deletions by Account	Account Deletions by Account
Account Lockouts by Account	Account Lockouts by Account
Certification Process Activity by Host	CA Identity Manager - Process Activity by Host
Password Policy Modify Activity	CA Identity Manager - Policy Modify Activity

Chapter 2: Addressing Business Needs

This section contains the following topics:

Processing Business Changes

You can automate the processing of certain identity management tasks by using identity policies. An identity policy is a set of business changes that occurs when a user meets a certain condition or rule. You can use identity policy sets to:

- Automate certain identity management tasks, such as assigning roles and group membership, allocating resources, or modifying user profile attributes.
- Enforce segregation of duties (see page 25). For example, you can create an identity policy set that prohibits members of the Check Signer role from having the Check Approver role, and restricts anyone in the company from writing a check over $10,000.
- Enforce compliance. For example, you can audit users who have a certain title and make more than $100,000.

  Identity policies that enforce compliance are called *compliance policies*.

The business changes associated with an identity policy include:

- Assigning or revoking roles, including provisioning roles (when CA Identity Manager includes provisioning)
- Assigning or revoking group membership
- Updating attributes in a user profile

For example, a company may create an identity policy which states that all Vice Presidents belong to the Country Club Member group and have the role Salary Approver. When a user's title changes to Vice President and that user is synchronized with the identity policy, CA Identity Manager adds the user to the appropriate group and role. When a Vice President is promoted to CEO, she no longer meets the condition in the Vice President identity policy so the changes applied by that policy are revoked, and new changes based on the CEO policy are applied.

The change actions that occur based on an identity policy contain events which can be placed under workflow control and audited. In the previous example, the Salary Approver role grants significant privileges to its members. To protect the Salary Approver role, the company can create a workflow process that requires a set of approvals before the role is assigned, and they can configure CA Identity Manager to audit the role assignment.

To simplify identity policy management, identity policies are grouped in an identity policy set. For example, the Vice President and CEO policies may be part of the Executive Privileges identity policy set.

Complying with Business Policies

Compliance is a corporate governance that includes a wide range of procedures that ensure a company and its employees comply with business policies. These compliance procedures often involve documenting, automating, and auditing the allocation of entitlements to applications and systems.

CA Identity Manager includes the following features, which support compliance management:

- **Smart Provisioning**

  Smart Provisioning is a collection of functionality that simplifies provisioning role assignment when CA Identity Manager integrates with CA RCM. This functionality includes:

  - Suggested Provisioning Roles

    CA Identity Manager can provide administrators with a list of provisioning roles that may be appropriate to assign to a user. The list of provisioning roles is determined by CA RCM, based on criteria entered by the administrator.

    Suggested provisioning roles help ensure that users have the correct privileges, while maintaining a company's role model.

  - Compliance and Pattern Messages

    Identity Manager administrators can validate proposed changes against a role model in CA RCM before committing changes. Validating changes before they are committed helps companies maintain the role model that they have defined for their operations.

    Users can validate proposed changes to provisioning roles (assigning or removing them), and changes to user attributes.

CA Identity Manager performs two types of policy validations:

- Compliance

  Proposed changes are validated against the CA RCM role model to see if they violate explicit, predefined business policy rules in CA RCM.

- Pattern

  Proposed changes are compared to the CA RCM role model to see if they cause the subject of the change to become "out of pattern." CA Identity Manager also makes sure that the changes do not significantly alter an established pattern in the role model.

  You can configure CA Identity Manager to perform these validations automatically when users perform certain tasks, or allow users to initiate the validation manually.

You can implement Smart Provisioning in an Identity Manager Environment once there is an established role model, based on Identity Manager data, in CA RCM.

Note: For more information, see the *Administration Guide*.

- **Identity policies**

  You can create a compliance policy, a type of identity policy (see page 21), which prohibits users from having certain privileges if they have other privileges. For example, you can prohibit users who can approve checks from issuing checks.

  Compliance policies enforce a segregation of duties in your environment.

- **Compliance reports**

  CA Identity Manager includes sample reports that display the compliance status for users in your environment. Using these reports, you can see which users are not compliant with your business policies.

Compliance Reports

CA Identity Manager includes the sample reports in the following table that you can use to monitor compliance with corporate business policies.

Report	Description
Role Members	Displays the roles in the report database and lists the members of those roles

Report	Description
Roles	Displays the following information for each role in the report database: ■ Tasks associated with the role ■ Member policies and role members ■ Administrator policies and role administrators ■ Owner policies and role owners
Tasks Roles	Displays the tasks in the reporting database and the roles with which they are associated
User Roles	Displays the users in the reporting database and lists each user's roles
Non-Standard Accounts Trend	Displays non-standard accounts trends for orphan accounts, system accounts, and exception accounts
Non-Standard Accounts	Displays all orphan, system, and exception accounts
Orphan Accounts	Displays all endpoint accounts with no global user in the Provisioning Server
Policies	Displays all identity policies
User Profile	Displays the following information for users: ■ Name ■ User ID ■ Groups where the user is a member or administrator ■ Roles where the user is a member, administrator, or owner
Endpoint Accounts	Displays accounts per endpoint (you can choose which endpoint to view)
Role Administrators	Displays roles and their administrators
Role Owners	Displays roles and their owners

Report	Description
Snapshots	Displays all exported snapshots
User Account	Displays a list of users and their accounts
User Entitlements	Displays user's roles, groups and accounts
User Policy Sync Status	Displays the user's status per policy (which policies should be allocated, deallocated or reallocated)

Note: For more information about reports, see the *Administration Guide*.

Enforcing Segregation of Duties Requirements

Segregation of Duties (SOD) requirements prevent users from receiving privileges that may result in a conflict of interest or fraud. CA Identity Manager provides the following functionality to support SOD:

- **Preventative identity policies**

  These policies, which execute before a task is submitted, allow an administrator to check for policy violations before assigning privileges or changing profile attributes. If a violation exists, the administrator can clear the violation before submitting the task.

  For example, a company can create a preventative identity policy that prohibits users who have the User Manager role from also having the User Approver role. If an administrator uses the Modify User task to give a User Manager the User Approver role, CA Identity Manager displays a message about the violation. The administrator can change the role assignments to clear the violation before submitting the task.

- **Policy Validation through Smart Provisioning**

  CA Identity Manager administrators can validate proposed changes to provisioning roles and user attributes against Business Policy Rules (BPRs) in CA RCM before committing changes. BPRs represent various constraints on privileges. For example, a BPR may prevent users who have a purchasing department role, which allows members to order stock from subcontractors, from also having the subcontractor payment role. A system administrator, business manager, auditor, or role engineer creates BPRs in CA RCM.

  Note: For more information about BPRs, see the *CA RCM Sage DNA User Guide*.

Note: For more information about preventative identity policies and Smart Provisioning, see the *CA Identity Manager Administration Guide*.

Transforming Data in the User Store

In some cases, you may want CA Identity Manager to transform data before it is stored in the user store. For example, you may want to store information in a different format than it is entered, or you may want changes applied when certain types of information are present.

CA Identity Manager includes the following features for transforming data:

- Identity Policies
- Logical Attribute Handlers

Note: You can also use identity policies and logical attribute handlers to implement custom business logic.

Logical Attribute Handlers

Logical attribute handlers are custom Java code that transform user attribute values used on Identity Manager task screens. Using logical attribute handlers, you can control how a physical attribute is displayed on a task screen. You can also use logical attribute handlers to transform a display value, such as cost, on the task screen to one or more physical attributes, such as unit price and quantity, that are stored in the user store.

Note: For more information about logical attribute handlers, see the *Programming Guide for Java*.

Applying Custom Business Logic

You can customize CA Identity Manager to implement the business logic that your company requires. CA Identity Manager includes the following options for implementing custom business logic:

- Identity Policies—You can use identity policies to define a set of of business changes that occur when a user meets a certain condition or rule. For example, identity policies can automate certain identity management tasks, such as assigning roles, or enforce business rules, such as preventing users from signing and approving checks over $20,000.

  Note: For more information about identity policies, see the *Administration Guide*.

- Logical Attribute Handlers—You can associate these handlers with Identity Manager task screens to control the display and modification of attribute values.

  For more information, see the *Programming Guide for Java*.

- Business Logic Task Handlers—Enable you to perform custom business logic, such as the following, during data validation operations for an Identity Manager task:

  - Enforcing custom business rules (for example, an administrator cannot be allowed to manage more than five groups).
  - Validating customer-specific task screen fields (for example, the value of an Employee ID field must exist in the master Human Resources database).

  Business logic task handlers can be implemented in Java or JavaScript.

  Note: For more information, see the *Programming Guide for Java*.

- Workflow—Allows you to create custom process definitions, which are associated with a CA Identity Manager event.

Note: Before deciding whether to implement business logic in a business logic task handler or a workflow process, see the following sections:

- Business Logic Task Handler Considerations (see page 28)
- Workflow Process Considerations (see page 28)

Business Logic Task Handler Considerations

Business Logic Task Handlers perform business logic validation during the synchronous processing phase of the task, which occurs prior to event generation. This allows you to:

- Perform task-level validation. For example, you can add or remove members of a group based on their office location, which is specified in the user profile screen.
- Prevent a task from being submitted if the validation fails.
- Automatically transform all of the information on a task screen so that it conforms to your business policies prior to task submission.

Note: You should not implement activities that take a long time to complete in a Business Logic Task Handler. Long running activities delay the submission of the task and are not well-suited for the synchronous phase where user interaction occurs. Instead, use a workflow process, which executes during the asynchronous phase of the task.

Workflow Process Considerations

Workflow processes are called during the asynchronous phase of the task and are associated with the execution of individual events. This allows you to:

- Execute approval activities based on the individual event data
- Execute long running custom business logic activities

While the Workflow API allows you to obtain task-level data from a Workflow Activity, typically you are operating in the context of that specific event under workflow.

Approving Business Changes

Workflow describes a process that consists of one or more steps that must be performed in order to accomplish some business objective, such as executing a hiring procedure, or obtaining a user's credit score from an external system. Typically, one of the steps in a workflow process involves approving or rejecting the business change.

In CA Identity Manager, a workflow process is associated with an event, an action that occurs during task processing. When an event enters the Pending state in its lifecycle, CA Identity Manager invokes any associated workflow process and pauses the event execution until the process completes. CA Identity Manager then performs or rejects the event based on the results of the workflow process.

This sequence is shown in the following diagram:

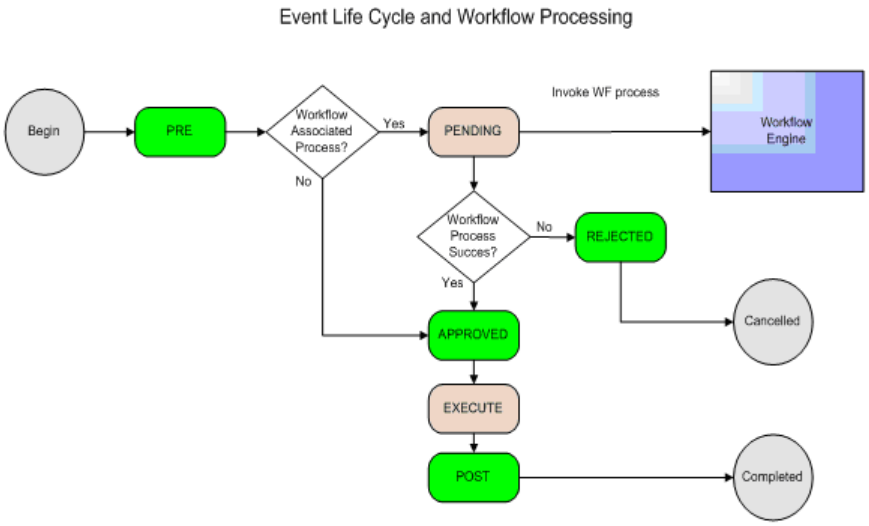

CA Identity Manager includes the InSession WorkPoint workflow engine for creating and managing workflow processes.

Note: For more information, see the *Administration Guide*.

Chapter 3: Identity Manager Architecture

This section contains the following topics:

Identity Manager Components

An Identity Manager implementation may include some or all of the following components:

- Servers
- User Stores
- Databases
- Connectors

Servers

A CA Identity Manager implementation includes one or more types of servers, depending on the functionality you need.

Identity Manager Server (required)

Executes tasks within Identity Manager. The J2EE Identity Manager application includes the Management Console and the User Console.

Identity Manager Provisioning Server

Manages accounts on endpoint systems.

This server is required if the CA Identity Manager installation will support account provisioning.

Note: You must have the Provisioning Directory installed remotely (or locally for a demonstration environment only) on a CA Directory Server before installing the Provisioning Server.

SiteMinder Policy Server

Provides advanced authentication for CA Identity Manager, and provides access to SiteMinder features, such as Password Services and Single Sign-On.

This server is optional.

User Store and Provisioning Directory

To provide options for managing users and automatic provisioning of additional accounts for those users, CA Identity Manager coordinates two user stores:

- The *Identity Manager user store*, the user store maintained by CA Identity Manager. Typically, this is an existing store that contains the user identities that a company needs to manage.

  The user store can be an LDAP directory or a relational database.

  In the Management Console, you create an Identity Manager Directory object to connect to the user store and to describe the user store objects that CA Identity Manager will maintain.

- The *Provisioning Directory*, the user store maintained by the Provisioning Server.

  It is an instance of CA Directory and includes global users, which associate users in the Provisioning Directory with accounts on endpoints such as Microsoft Exchange, Active Directory, and SAP.

  Only some CA Identity Manager users have a corresponding global user. When a CA Identity Manager user receives a provisioning role, the Provisioning Server creates a global user.

Separate User Store and Provisioning Directories

The following figure shows a separate user store and Provisioning Directory. In this figure:

- A CA Identity Manager administrator uses an admin task that edits a user in the user store, which affects the Provisioning Directory.

  This change may also update an endpoint (such as an email server) which has a connector to the Provisioning Server.

- A change made in the Provisioning Server (or an endpoint with a connector to the Provisioning Server) updates the Identity Manager user store and Provisioning Directory.

For example, an endpoint, such as a Human Resources application, might update the email addresses of users.

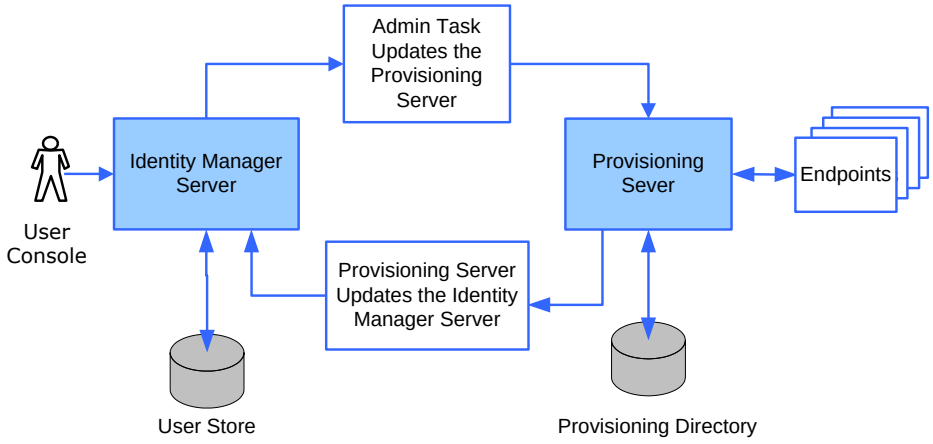

Combined User Store and Provisioning Directory

If you select an instance of CA Directory for both user stores, the directory functions as both a user store and Provisioning Directory. Users you create in Identity Manager are stored in that directory, but you can still modify that directory as a Provisioning Directory. For example, you can select the user attribute to use in the Provisioning Server for a specific user attribute used by Identity Manager.

The following figure shows the use of one directory for both the user store and Provisioning Directory.

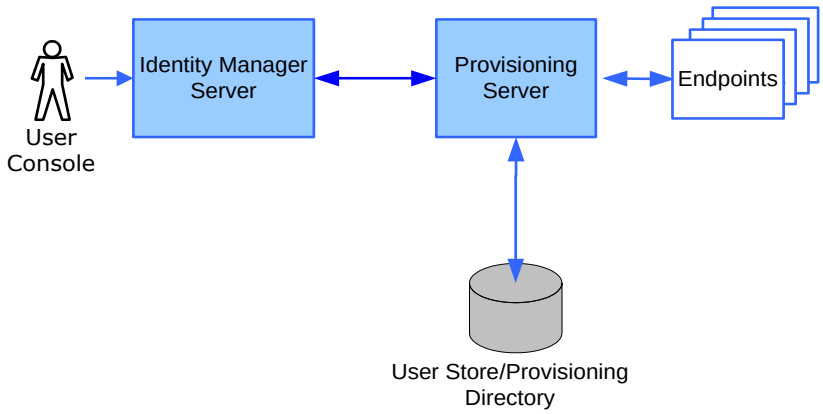

In this situation, changes to the user store/Provisioning Directory can be initiated by CA Identity Manager, the Provisioning Server, or an endpoint.

Note: When CA Identity Manager uses a combined user store and Provisioning Directory, organizations are not supported. If you need to manage organizations, use a separate user store and Provisioning Directory. How to Choose a User Store Solution (see page 49) lists additional differences between combined and separate user store/Provisioning Directory implementations.

Databases

CA Identity Manager uses data sources to connect to databases that store information required to support CA Identity Manager functionality. These databases can reside in a single physical instance of a database, or in separate instances.

Object Database (required)

Contains CA Identity Manager configuration information.

Task Persistence Database (required)

Maintains information about CA Identity Manager activities and their associated events over time. This allows the system to accurately track CA Identity Manager activities, even if you restart the Identity Manager Server.

Workflow Database

Stores workflow process definitions, jobs, scripts, and other data required by the Workflow Engine.

Audit Database

Provides a historical record of operations that occur in an Identity Manager environment.

Note: You can configure the amount and type of information that CA Identity Manager stores in the audit database. See the *Configuration Guide* for more information.

Reporting Database

Stores snapshot data, which reflects the current state of objects in CA Identity Manager at the time the snapshot is taken. You can generate reports from this information to view the relationship between objects, such as users and roles.

When you use the Installer, CA Identity Manager configures a connection to a single database, called the Identity Manager Database, which contains the tables for each database type.

Note: You can create a data store for task persistence, workflow, auditing, or reporting in a separate database and configure CA Identity Manager to connect to it. For more information, see the *Installation Guide*.

Connector Components

A connector is the software interface to an endpoint. The Provisioning Server uses the connector to communicate with the endpoint. It translates Provisioning Server actions into changes on the endpoint, such as “Create a new email account on a Microsoft Exchange endpoint.”

Examples of endpoints are UNIX workstation, Windows PC, or an application such as Microsoft Exchange (for email).

A connector server works with multiple endpoints. For example, if you have many UNIX workstation endpoints, you might have one Connector Server that handles all connectors that manage UNIX accounts. Another connector server might handle all connectors that request Windows accounts.

The distributed connector server works with multiple Connector Servers. It provides load balancing when one connector server is busy and high availability when a connector server is down.

Connector Servers

A connector server is a Provisioning Server component that manages connectors. It can be installed on the Provisioning Server system or on a remote system.

There are two types of connector servers:

- The Java Connector Server (Java CS) manages connectors written in Java
- The C++ Connector Server (CCS) manages connectors written in C++

C++ Connector Server

The *C++ Connector Server* is a connector server that manages C++ connectors. It can be installed on the Provisioning Server or on a remote system. The C++ Connector Server provides an object-oriented application framework that simplifies development of connectors, which are responsible for communication between the C++ Connector Server and the endpoint.

Java Connector Server

The *Java Connector Server (Java CS)* is a server component which handles hosting, routing to, and management of Java connectors. The Java CS provides a Java alternative to the C++ Connector Server. It is architecturally and functionally similar to the C++ Connector Server, except that it has a Java API instead of a C++ API, which allows your connectors to be implemented in Java. In addition, the Java CS is data-driven rather than code-driven, which allows more functionality to be addressed by the container (or Java CS) instead of by connectors themselves.

The Provisioning Server handles provisioning of users, and then delegates to connectors (using the C++ Connector Server or Java Connector Server) to manage endpoint accounts, and groups.

Connectors and Agents

Identity Manager Connectors run as part of the wider Provisioning Server architecture and communicate with the systems managed in your environment. A connector acts as a gateway to a native endpoint type system technology. For example, machines running Active Directory Services (ADS) can be managed only if the ADS connector is installed on a Connector Server with which the Provisioning Server can communicate. Connectors manage the objects that reside on the systems. Managed objects include accounts, groups, and optionally, endpoint type-specific objects.

Connectors are installed on the Connector Server and some components are installed on the Provisioning Server (for example, Server plug-in) or Provisioning Manager (user interface plug-ins).

Some connectors require an agent on the systems they manage in order to complete the communication cycle, in which case, they can be installed using the Provisioning Installer. Agents can be separated into the following categories:

Remote Agents

Installed on the managed endpoint systems

Environment Agents

Installed on systems, such as CA ACF2, CA Top Secret, and RACF

Certain components work on UNIX and Windows, including the following C++ Connector Server-based options:

- UNIX (ETC, NIS)
- Access Control (ACC)

  Note: The UNIX ACC connector can manage only UNIX ACC endpoints. The Windows ACC connector is required to manage the Windows ACC endpoints but can also manage UNIX ACC endpoints.
- OpenVMS
- CA-ACF2
- RACF
- CA-Top Secret

The other C++ Connector Server-based connectors can be accessed from the Solaris Provisioning Server by relying on the Connector Server Framework (CSF). The CSF allows a Provisioning Server on Solaris to communicate with connectors running on Windows.

Note: The CSF must run on Windows to use these connectors.

Connector Xpress

Connector Xpress is an Identity Manager utility for managing dynamic connectors, mapping dynamic connectors to endpoints, and establishing routing rules for endpoints. You can use it to configure dynamic connectors to allow provisioning and management of SQL databases and LDAP directories.

Connector Xpress lets you create and deploy custom connectors without the technical expertise generally required when creating connectors managed by the Provisioning Manager.

You can also set up, edit, and remove a connector server configuration (both Java and C++) using Connector Xpress.

The primary input into Connector Xpress is the native schema of an endpoint system. For example, you can use Connector Xpress to connect to a RDBMS and retrieve the SQL schema of the database. You can then use Connector Xpress to construct mappings from those parts of the native schema that are relevant to identity management and provisioning. A mapping describes how the provisioning layer represents an element of the native schema.

Connector Xpress generates metadata that describes, to a dynamic connector, the runtime mappings to a target system.

The output of Connector Xpress is a metadata document produced when you complete your mappings. The metadata is an XML file that describes the structure of your connector to the Java CS.

It describes the Provisioning Server classes and attributes and how they are mapped to the native schema.

The metadata is used to create dynamic endpoint types on one or more Provisioning Servers.

Note: For more information on using Connector Xpress, see the *Connector Xpress Guide*, in the *CA Identity Manager bookshelf.*

Additional Components

Identity Manager includes some additional components, which support Identity Manager functionality. Some of these components are installed with Identity Manager and some must be installed separately.

WorkPoint Workflow

WorkPoint workflow engine and WorkPoint Designer are installed automatically when you install Identity Manager.

These components enable you to place an Identity Manager task under workflow control, and to modify existing workflow process definitions or create new definitions.

Note: For more information about workflow, see the *Administration Guide*.

Provisioning Manager

The Identity Manager Provisioning Manager manages the Provisioning Server through a graphical interface. This is used for administrative tasks such as managing Provisioning Server options. In some cases, you may also use the Provisioning Manager to manage certain endpoint attributes, which you cannot manage in the Identity Manager User Console.

The Provisioning Manager is installed as part of the Identity Manager Administrative Tools.

Note: This application runs on Windows systems only.

For more information about the Provisioning Manager, see *Provisioning Reference Guide*.

IAM Report Server

Identity Manager provides reports that you can use to monitor the status of an Identity Manager environment. To use the reports provided with Identity Manager, you install the IAM Report Server, which is included with Identity Manager.

IAM Report Server is powered by Business Objects Enterprise XI. If you have an existing Business Objects server, you can use that instead of the IAM Report Server to generate Identity Manager reports.

Note: For installation instructions, see the *Installation Guide*.

Sample CA Identity Manager Installations

With CA Identity Manager, you can control user identities and their access to applications and accounts on endpoint systems. Based on the functionality you need, you select which CA Identity Manager components to install.

If you install a Provisioning Server, users can gain access to accounts, such as an email or database account, on endpoint systems. You assign provisioning roles to the CA <idgmr> users. These roles define the accounts on the endpoint systems. To communicate with the endpoint systems, you install connector servers for endpoint-specific connectors, such as an SAP connector.

In all CA Identity Manager installations, the Identity Manager Server is installed on an application server. After you install the application server, you use the CA Identity Manager Installer to install the software you need.

The following sections illustrate some examples of CA Identity Manager implementations at a high level.

Installation with Provisioning Components

CA Identity Manager provisioning allows you to create an Environment that connects to a Provisioning Server for provisioning accounts to various endpoint systems. You can assign provisioning roles to users you create through CA Identity Manager. Provisioning roles are roles with account templates that define accounts that users can receive on endpoint systems. Accounts provide users with access to additional resources, such as an email account.

When you assign a provisioning role to a user, that user receives the accounts defined by the account templates in the role. The account templates also define how user attributes are mapped to accounts. The accounts exist in managed endpoints defined by the account templates.

The following figure is an example of an CA Identity Manager installation with provisioning:

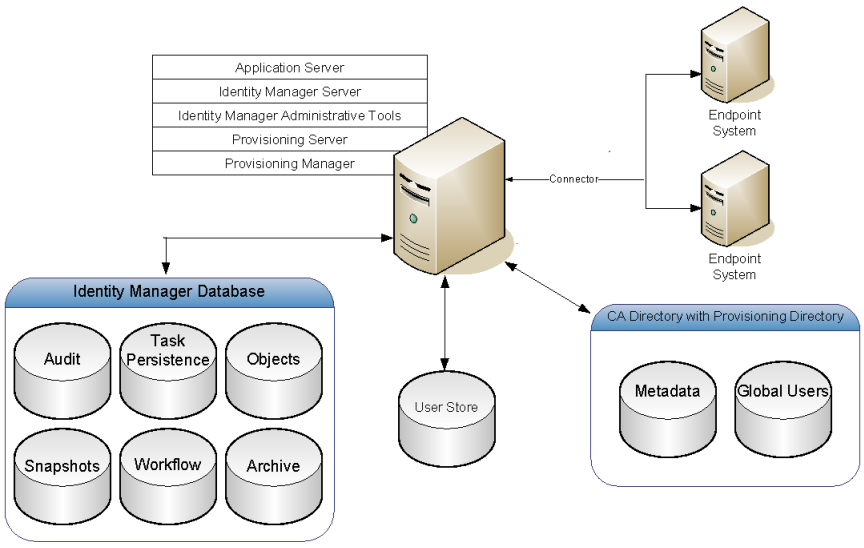

Installation with SiteMinder Policy Server

A SiteMinder Policy Server provides advanced authentication and protection for your Identity Manager environment. The following figure is an example of a CA Identity Manager installation with a SiteMinder Policy Server:

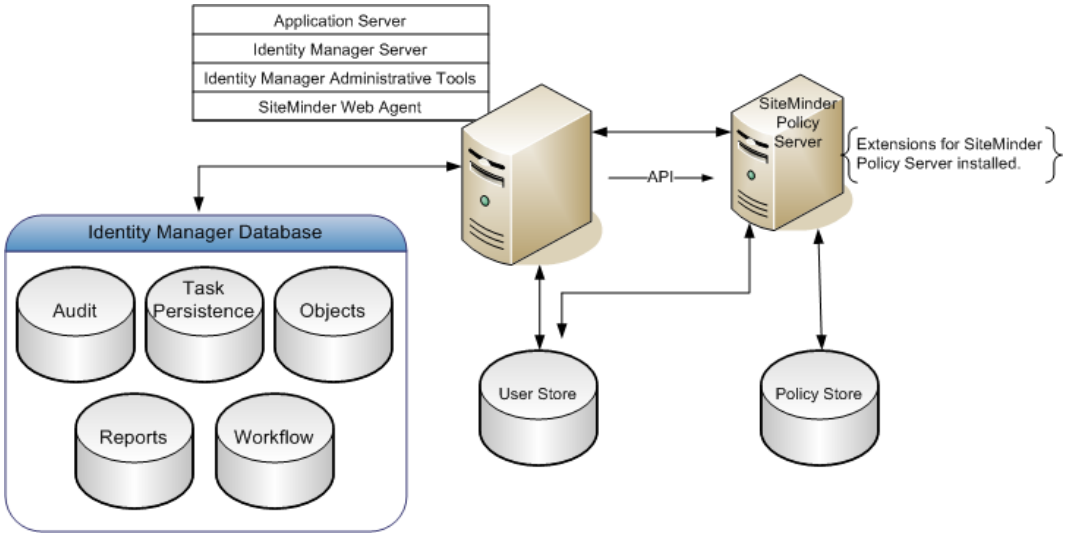

An Identity Manager implementation that includes SiteMinder includes all of the components in the basic installation or the installation with provisioning, plus these additional components:

SiteMinder Web Agent

Works with the SiteMinder Policy Server to protect the User Console. The Web Agent is installed on the system with the Identity Manager Server.

SiteMinder Policy Server

Provides advanced authentication and authorization for Identity Manager, as well as other functionality such as Password Services and Single-Sign On.

Extensions for SiteMinder Policy Server

Enables a SiteMinder Policy Server to support Identity Manager. Install the extensions on each SiteMinder Policy Server system in your Identity Manager implementation.

SiteMinder Policy Store

Stores information that SiteMinder needs to manage access to Web resources.

When Identity Manager integrates with SiteMinder, the policy store also includes information about Identity Manager directories and environments so that SiteMinder can provide advanced authentication.

Chapter 4: Planning Your Implementation

To plan a CA Identity Manager implementation, you decide how CA Identity Manager will manage users and what functionality you need to accomplish your business objectives. Some questions to consider are:

- How do I manage users?
- Do I need account provisioning?
- What are my custom business requirements and should I implement them using workflow?

Based on the decisions you make, you can determine the best way to implement CA Identity Manager for your environment.

This section contains the following topics:

Decide What to Manage

Deciding what you want to manage will help you determine which components you want to install. Using Identity Manager, you can manage the following:

- User identities
- Access to accounts on endpoint systems

User Identities

User identities represent the people that a company needs to manage, such as employees, contractors, suppliers, and others.

To manage user identities, you need to install only the Identity Manager Server and the Administrative Tools.

How to Configure User Management Support

In CA Identity Manager, you manage users with admin roles, which determine the CA Identity Manager tasks that administrators can perform.

Note: Before implementing user management in CA Identity Manager, you should determine which functionality you need and develop a plan (see page 58) for implementing that functionality in stages.

To configure user management support, you complete the following high-level steps:

1. Install the Identity Manager Server and Administrative tools.

   If you need to provision accounts to managed users, you will also need to install support for provisioning (see page 45).

   Note: See the *Installation Guide* for instructions.

2. Create the following in the Identity Manager Management Console:

   - **Identity Manager directory**

     Describes a user store to Identity Manager. It includes the following:

     - A pointer to a user store, which stores managed objects such as users, groups, and organizations.
     - Metadata that describes how managed objects are stored in the directory and represented in CA Identity Manager.

   - **Identity Manager environment**

     Provides a management namespace that lets Identity Manager administrators manage objects such as users, groups, and organizations, with a set of associated roles and tasks. The Identity Manager environment controls the management and graphical presentation of a directory.

   For more information about Identity Manager directories and environments, see the *Configuration Guide*.

3. Modify the default admin roles and tasks to suit your business requirements.

   Typical role modifications include adding or removing default tasks from existing admin roles, or creating new admin roles, which are based on the default roles.

   Typical task modifications include customizing the default user profile tabs to include only the information that you want to manage. (The default profile tabs include all attributes that are defined for users.)

   For information about modifying the default admin roles and tasks, see the *User Console Design Guide*.

4. Assign the admin roles to users who will perform user management tasks.

Provisioning Accounts from Other Applications

The decision to implement provisioning depends on the type of information that you need to manage. If you are managing a central user directory and you do not want to manage user accounts in other systems, you do not need provisioning. If you want to manage user accounts over a variety of systems, then you should implement provisioning support.

Provisioning capabilities are provided through the Provisioning Server, which is integrated with CA Identity Manager. The Provisioning Server provides the following functionality for account provisioning:

- Endpoint Management
- Account Synchronization
- Account Templates
- Explore and Correlate Functionality

Note: Provisioning information is stored in a Provisioning Directory. If CA Identity Manager maintains users in another type of directory, your deployment will include an Identity Manager user store and a provisioning directory.

Endpoint Management

To provision accounts, you define and manage endpoints in the Identity Manager User Console. An *endpoint* is a system for which users need access. Examples of endpoints include Oracle databases, UNIX NIS servers, Windows servers, and Microsoft Exchange servers. *Account templates* (see page 46) create accounts and determine the user capabilities in managed endpoints.

Note: You can also use the Provisioning Manager to define and manage endpoints. Although we recommend using the User Console for most endpoint management tasks, there are some tasks that require the Provisioning Manager, such as managing certain endpoint attributes and managing endpoint objects other than accounts. For more information about the Provisioning Manager, see the *Provisioning Reference*.

Account Synchronization

You can synchronize user accounts across multiple managed endpoints. When account synchronization is enabled, a change made to a user profile in the Provisioning Server is propagated to all of the endpoints where that user has an account.

Note: You specify account synchronization settings on the Profile tab for a CA Identity Manager task. For more information about configuring account synchronization, see the *Administration Guide*.

Account Templates

Account templates define how a user is represented in a managed endpoint. For example, a template for an Exchange account could define the format of a user's email address, such as <first initial><last name>@mycompany.com.

Account templates also determine the privileges a user has within a managed system. For example, in addition to defining the format of an email address, a template for an Exchange account may also limit a user's mailbox size.

You create and manage account templates in the User Console.

Explore and Correlate Functionality

The Explore and Correlate features simplify endpoint management by discovering and synchronizing changes in managed sytems.

The Explore feature finds objects, including accounts, in endpoints, and stores references to them in the Provisioning Directory. You can use the Explore feature to detect any new objects to be managed. For example, if you provision accounts in an LDAP directory and new organizations are added in that directory, you can use the Explore feature to introduce those new organizations for use in account templates.

The Correlate feature associates an account in a managed endpoint with a global user in the Provisioning Directory. When a change is made to the account through the endpoint, the Correlate feature can synchronize those changes with the global user account.

Note: For more information about the Explore and Correlate functionality, see the *Administration Guide*.

How to Configure Support for Provisioning

After deciding to implement provisioning, you complete the following high-level steps.

1. Use the CA Identity Manager Server installer to install the Identity Manager Server, the Provisioning Server, the Provisioning Directory Initialization, and the Administrative Tools.

   Note: For more information about installing CA Identity Manager components, see the *Installation Guide*.

2. Configure the Provisioning Manager to connect to the Identity Manager Server.

3. Configure Provisioning in the Identity Manager Management Console:
   a. Enable Provisioning.
   b. Configure an environment for Provisioning by completing the following:
      - Importing custom role definitions
      - Configuring an inbound administrator
      - Connecting the environment to the Provisioning Server.

   Note: For more information, see the *Configuration Guide*.

4. Create endpoints in the User Console.

   This allows Identity Manager to manage the endpoint. When you create an endpoint, it populates the provisioning directory with accounts and other objects found in the endpoint.

   Note: For more information about endpoint management, see the *Administration Guide*.

5. Explore and correlate the endpoint.

   When you explore an endpoint, CA Identity Manager finds the objects in the endpoint and stores instances of them in the provisioning directory. When you correlate accounts on an endpoint, CA Identity Manager associates them with a global user in the provisioning directory. You may choose whether the correlate function creates any global users that are not present or whether it associates accounts with no matching global user to the [default user] global user.

6. Create and maintain endpoint accounts by using account templates, which contain the attributes that are used to create accounts.

7. Associate the account templates with provisioning roles.

   When you assign provisioning roles to users, CA Identity Manager creates accounts in the associated endpoints for those users.

   Note: For information about account templates and provisioning roles, see the *Administration Guide*.

Determine Audit Requirements

Identity Manager includes auditing capabilities that allow you to monitor activities in an Identity Manager environment.

This information is stored in an audit database. The amount and type of information that is stored in the audit database is configurable.

You view audit data in the User Console through a task called View Submitted Tasks. This task allows administrators to search for and view tasks that occur in the system. Administrators can view task information at a high level or view task and event details.

Identity Manager Auditing Considerations

Audit data provides a historical record of operations that occur in an Identity Manager environment. To audit data in CA Identity Manager, you need the following:

- An auditing database
- An audit settings file

Audit Database

When you use the CA Identity Manager Installer, Identity Manager configures a connection to a single database, called the Identity Manager Database, and creates a data source to connect to the database tables for auditing.

Note: The Identity Manager Database also includes data that is used by other CA Identity Manager functionality, including task persistence, workflow, and reporting. For scalability purposes, you can create a new, separate instance of a database for auditing.

Note: For more information about the auditing database, see the *Installation Guide*.

Audit Settings

You configure audit settings in an audit settings file. An audit settings file determines the amount and type of information that CA Identity Manager audits. You can configure an audit settings file to do the following:

- Enable auditing for an Identity Manager environment.
- Enable auditing for some or all of the CA Identity Manager events generated by admin tasks.
- Record event information at specific states, such as when an event completes or is cancelled.
- Log information about attributes involved in an event. For example, you can log attributes that change during a ModifyUserEvent event.
- Set the audit level for attribute logging.

Note: For more information about configuring auditing, see the *Configuration Guide*.

CA Audit Considerations

CA Audit is an audit management system that enables you to collect and store security related data for auditing, reporting, compliance verification and event monitoring.

To integrate with CA Audit, you install the iRecorder component when you install the Identity Manager Server. The iRecorder retrieves events from CA Identity Manager. Based on policies in the CA Audit Policy Manager, the iRecorder ignores the event or routes it through to CA Audit.

Decide User Store Requirements

An Identity Manager implementation must include a user store that contains the user identities that Identity Manager maintains. Typically, this is an existing user store that an enterprise uses to store information about its users, such as employees and customers.

If your implementation includes provisioning, Identity Manager also requires a provisioning directory that includes global users, which are associated with accounts on endpoints such as Microsoft Exchange, Active Directory, and Ingres.

How to Choose a User Store Solution

To manage an existing user store, use CA Identity Manager to manage it, and then create a separate Provisioning Directory. However, in implementations where you can choose a user store option, consider the key differences in the following table.

Capability	**Separate User Store and Provisioning Directory**	**Single User Store and Provisioning Directory**
User Object Support	Yes	Yes (custom schema)
Group Object Support	Yes	Yes (custom schema)
Group Membership Support	Yes	Yes (proprietary model)
Organization Object Support	Yes	No
Direct Provisioning Support	No	Yes
Additional user store required for provisioning	Yes	No

Capability	Separate User Store and Provisioning Directory	Single User Store and Provisioning Directory
Recommended number of users	Millions of users	Tens of thousands of users

Managing Multiple User Stores

An enterprise may maintain multiple user stores. In each user store, the user identity allows access to different corporate resources. You can use one of the following methods to manage multiple user stores:

- Use CA Identity Manager to directly manage the Provisioning Directory and use the Provisioning Server to indirectly manage the users and accounts in the different user stores.

  This approach allows you to:

  - Centrally manage users who can be assigned various enterprise resources from one location
  - Implement common security and business rules across enterprise resources. This may include the following:
    - Role-based access control
    - Delegated administration
    - Tasks and screens that are customized based on the type of corporate identities they manage
    - Identity policies for rule-based identity management
    - Customization and extensibility

  Note: For information on these features, see the *Administration Guide*.

- Create a separate Identity Manager environment to manage each user store

  With this method, information is not shared between environments.

Select Components to Install

The following table lists the components to install to support the functionality that you want to implement.

Note: For instructions on installing these components, see the *Installation Guide*.

If you want to...	Install these components
Manage user identities in an existing corporate user store	■ Identity Manager Server
Provision accounts in endpoint systems	■ Provisioning Server ■ Provisioning Directory ■ Provisioning Manager ■ Connectors ■ **Note**: For instructions on installing connectors, see the *Connector Guide* for the type of connectors that you want to install.
Implement one or more of the following features: ■ Advanced authentication ■ Advanced password policies ■ Console skins for different sets of users ■ Configure locale preferences	■ SiteMinder Policy Server ■ Policy store ■ SiteMinder Web Agent ■ Identity Manager Extensions to the Policy Server **Note**: For instructions on installing the SiteMinder Policy Server and policy store, see the *CA SiteMinder Web Access Manager Policy Server Installation Guide.* For instructions on installing the Web Agent, see the *CA SiteMinder Web Access Manager Web Agent Installation Guide.*
Generate reports on Identity Manager activity	IAM Report Server

Decide Hardware Requirements

The hardware that you need for an Identity Manager installation depends on the functionality that you want to implement and the size of your deployment.

The following sections describe typical Identity Manager implementations and their required hardware.

Deployment Types

When planning the hardware needed for a CA Identity Manager deployment, consider the features that you want to implement and the initial size of the deployment. Use one of the following categories to estimate the size of the deployment.

Note: The deployment type you select determines the size of the DxGrid file that is used by the Provisioning Directory. You specify the deployment type when you install the Identity Manager Server.

Demonstration

A single server deployment for use in demonstrations or basic testing in a Development environment. A demonstration deployment supports up to 10,000 provisioned accounts.

Note: This implementation type does not support production implementations.

Basic

A high availability implementation that is suitable for most small to medium size implementations. A basic implementation supports up to 400,000 provisioned accounts.

This type of implementation requires two servers for running the Identity Manager application and its components and two servers for running the Identity Manager database and the user store.

Intermediate

A high availability implementation that is suitable for medium size implementations. An intermediate deployment supports up to 600,000 provisioned accounts.

Large Enterprise

A high availability implementation that includes additional server clusters to address additional users and an increased number of transactions. A large deployment supports more than 600,000 provisioned accounts.

Note: For more information about high availability implementations, see the *Installation Guide*.

Additional Requirements for Provisioning

In addition to the components required for a basic CA Identity Manager implementation, the following additional components are required when CA Identity Manager includes provisioning:

- Provisioning Server

  Can be installed on the same machine as the Identity Manager server.

- Provisioning Directory Initialization

  Important! The Provisioning Directory Initialization must be installed on CA Directory.

- Provisioning Manager

  Can be installed on any Windows machine that can access the Provisioning Server.

Note: In a development environment, these components can be installed on one machine that also includes the basic installation components.

Additional Requirements for SiteMinder Integration

When CA Identity Manager integrates with SiteMinder, the implementation must include the components in the basic CA Identity Manager installation, plus the following additional components:

- **Policy Server**

  Provides policy management, authentication, authorization, and accounting services.

  The Policy Server can be installed on the same machine as the Identity Manager Server, if the Policy Server is dedicated to CA Identity Manager. If the Policy Server is protecting other applications, we recommend installing it on a separate machine to ensure best performance.

- **Policy Store**

  Contains all of the Policy Server data. You can configure a policy store in a supported LDAP or relational database. In high availability implementations, we recommend installing the policy store on a separate server.

- **Extensions to the Policy Server**

  Enables a SiteMinder Policy Server to support CA Identity Manager. Install the extensions on each SiteMinder Policy Server system in your CA Identity Manager implementation.

- **SiteMinder Web Agent**

  Works with the SiteMinder Policy Server to protect the User Console. Installed on the system with the Identity Manager Server.

Choose a Method to Import Users

If you need to import users into an existing user store, the method you select should be based on your business requirements.

The following sections describe options for importing users.

How to Import Users into a New User Store

After you decide how to store user data, you may need to import users from one store to another. Depending on your implementation, you can use different methods to import users.

Note: After importing users into a new user store, you can use identity policies (see page 56) to apply changes to imported users.

Import Users Through Identity Manager

CA Identity Manager provides the following methods for adding users to a user store that it manages directly.

Method	Features	Limitations
Bulk Loader	Allows you to use the Bulk Loader task in the User Console to upload feeder files that are used to manipulate large numbers of managed objects simultaneously. The advantage of the Bulk Loader method is that you can automate the process of manipulating a large number of managed objects using an information (feeder) file. The Bulk Loader task can also be mapped to a workflow process.	If you are using the Bulk Loader, you may see out-of-memory exceptions depending on the number of users you are importing. To address this issue, increase the JVM memory settings.
Remote Task Invocation via Task Execution Web Service (TEWS)	Allows execution of any CA Identity Manager task that is enabled for Web Services, including the Create User task. If the task is configured for User Synchronization, CA Identity Manager will execute any applicable identity policies.	Performance characteristics of web service model may not be well-suited for high-throughput requirements of bulk import operations
IM API	■ Provides User-based APIs that can be invoked directly for creating users via a Java client ■ Provides the highest throughput capabilities.	■ Bypasses audit and security mechanisms provided by the Task Server. ■ Does not support execution of Identity Policies.

Note: For more information about the Bulk Loader, see the *Administration Guide*. For more information about TEWS and the IM API, see the *Programming Guide for Java*.

Execute Identity Policies on Imported Users

An *identity policy* is a set of business changes that occur when a user meets a certain condition or rule. These changes can include assigning or revoking roles (including provisioning roles for users in the provisioning directory), assigning or revoking group membership, and updating attributes in a user profile.

You can use identity policies to apply changes to user accounts after they have been imported to a new user store.

This section describes methods for executing identity policies for imported users in one or two steps.

One-Step Approach

You can use the following import methods to execute identity policies on users that you import into a new user store in a single step:

- Bulk Loader in the User Console
- Create User Task Execution via TEWS
- Inbound Synchronization

Two-Step Approach

Using a two step approach, you first import users and then execute identity policies against the those users. You can use this method when CA Identity Manager manages users in the Provisioning Server. This method may provide more flexibility, depending on your import requirements.

1. Use one of the import tools for adding users into the Provisioning Directory.
2. Invoke the Identity Manager Synchronize User Task through TEWS on each of the imported users.

Import Users Through the Provisioning Server

The Provisioning Server includes bulk import options for adding and managing users in the Provisioning Directory. The following tables describe the methods for importing users into the Provisioning Directory.

Method	Features	Limitations
Batch utility (etautil)	A command line interface utility that allows you to manage objects in the Provisioning Directory	■ Currently supported for Windows systems only

Method	Features	Limitations
Explore and Correlate	■ Discovers new objects that the Provisioning Server can manage in a known endpoint (including users) ■ Provides correlate capabilities for object instances that exist in the endpoint and the Provisioning Server. Additional information exists in Explore and Correlate Functionality.	■ By default, the Explore and Correlate functionality is available for the currently supported connectors. Can be extended with custom connectors ■ The Correlate option may affect scalability when working with large user populations. If you select this import option, be sure to evaluate the performance and scalability implications.

After importing users, you can use the method in the following table to make changes to user accounts in the Provisioning Server:

Method	Features	Limitations
Service Provisioning Markup Language (SPML)	You can generate SPML requests that are submitted directly to the Provisioning Server without having to convert user data to a CSV or XML input file.	■ Static solution that requires an application that can generate and submit SPML-formatted XML payloads as SOAP requests ■ Requires export of HRMS/ERP data into a file; additional data manipulation required

Synchronize Global Users with the Identity Manager User Store

After you import users into the Provisioning Server, you can use the following methods to add those users to the Identity Manager user store:

- **Inbound Synchronization**

  Inbound Synchronization keeps Identity Manager users up to date with changes that occur in the Provisioning Directory. Changes in the Provisioning Directory include those made using Provisioning Manager or systems with connectors to the Provisioning Server.

  Note the following when using inbound synchronization to import users:

  - In the Identity Manager Management Console, you can customize how the attributes from the inbound request are mapped to attributes in the Identity Manager task.

    Note: For more information, see the *Administration Guide*.

  - Consider which Provisioning Server changes require synchronization with the corporate user store. Synchronizing a large number of changes may impact performance and scalability.

- **Provisioning Roles and Account Templates**

  The Provisioning Server can manage accounts in the Identity Manager user store using provisioning roles and account templates. This requires that a managed endpoint, which points to the Identity Manager user store has been acquired and the appropriate account templates and roles exist. In this case, global users created through one of the options described in Import Users Through the Provisioning Server (see page 56) can be assigned a provisioning role that creates the user account in the Identity Manager user store.

Develop a Deployment Plan

When planning a large implementation, you should deploy CA Identity Manager functionality in stages. The following deployment order allows you to gain significant value from CA Identity Manager quickly, evaluate the changing needs of your implementation over time, and carefully construct your environment for best performance and scalability:

- Self-service and password management
- Identity policies
- Workflow approvals

- Delegated administration for user, group, and organization objects
- Delegated administration for role administration

After each deployment stage, be sure to evaluate performance and make adjustments before proceeding to the next stage. Optimizing Identity Manager (see page 69) provides information on performance, tuning, and scalability.

Deploy Self-Service and Password Management

Deploy self-service tasks and password management before deploying other Identity Manager features for the following reasons:

- Self-service tasks and password management are easy to deploy and provide significant value quickly.
- These features are independent of the delegated administration model and can be reconfigured as needed to address changing business needs.
- These features typically generate the highest volume of tasks that Identity Manager processes on a regular basis. Because of this, they provide a way to test the scalabilty of your implementation before you deploy additional features.

To deploy self-service tasks, you complete the following steps:

1. Configure the self-registration task.

   This is a public task, which is enabled by default during installation. To configure this task, you add or remove fields from the default self-registration task, as needed.

2. Deploy the Self Manager role.

   The member rule for this role should be configured to apply to all users, or should include a member rule that automatically assigns the role to new users. For example, you can create a member rule that assigns the Self Manager role to all full-time employees. When a user self-registers, Identity Manager can set the employee type to full-time (by using a logical attribute handler, or business task handler). The user meets the criteria in the member rule and receives the Self Manager role automatically.

   Note: When you configure member rules for the Self Manager role, do not allow administrators to add or remove role members. Since the role is assigned automatically, there is no need for an administrator to explicitly assign the role.

To deploy password management capabilities, you complete the following steps:

1. Configure the public password management tasks, such as the Forgotten Password task.
2. Create password policies that determine how passwords are created and when they expire.
3. Deploy the Password Manager role, which enables role members to reset user passwords.

Note: For information on roles, tasks, and password management, see the *Administration Guide*.

Deploy Identity Policies

An identity policy is a set of business changes that occurs when a user meets a certain condition or rule. You can use identity policies to provide business-driven entitlements before a complete delegation model is deployed. For example, you can create an identity policy that assigns the Sales Manager provisioning role, which grants access to sales applications, to all users whose title is Sales Manager. When a sales representative is promoted to Sales Manager, he automatically receives access to all of the systems he needs to do his job without waiting for administrator involvement.

To deploy identity policies, you complete the following steps:

1. Configure identity policies that are triggered by changes to user profile attributes.
2. Configure the User Manager role to allow a small number of administrators to use user tasks, such as Create User and Modify User, to change the attributes that trigger the identity policies.

   Be sure to configure the scope rules in the User Manager member policies to determine the set of users that role members can manage.

Note the following when deploying identity policies:

- Consider initially creating identity policies that grant entitlements that do *not* require workflow approvals. This allows you to deploy identity policies without having to define workflow processes, approval forms, and approver models.
- Before creating identity policies, you should be familiar with other methods of implementing business rules in CA Identity Manager, such as data validation rules, logical attributes, business logic task handlers, and workflow processes, to determine which method provides the best solution.

  Note: For more information about these methods, see the *Administration Guide* and the *Programming Guide for Java*.
- Identity policies are an efficient way to assign entitlements in CA Identity Manager, however, they may significantly impact performance (see page 84).
- For the initial deployment of user tasks, consider removing or hiding relationship tabs, such as Roles tabs, that manage the same entitlements as identity policies. This prevents the risk of unauthorized entitlements and prevents the potential performance impact of improperly constructed roles.

Note: For more information about identity policies, see the *Administration Guide.*

Deploy Workflow Approvals

Workflow approvals can add an additional level of security and automation to your Identity Manager implementation.

Deploying workflow approvals requires the following tasks:

1. Decide which events or tasks require approvals.
2. Define the set of approvers, called participants, for each workflow process.

   Note: All participants are determined dynamically by participant resolvers. To maintain good performance, limit the number of participants to thirty users.
3. Configure approval forms.
4. Define custom workflow processes, if needed.

Environment and Task Level Workflow Approvals

Identity Manager supports two types of approvals: environment-level approvals and task-level approvals. Environment-level approvals are defined for all instances of an event, regardless of the tasks they are associated with. Task-level approvals are defined for a specific event associated with a specific task. Task-level approvals take precedence over environment-level approvals.

Most approvals are defined at the environment level to ensure that the same workflow activities occur for an event, regardless of the task that it is associated with. However, in the following situations, consider implementing task-level workflow:

- You have specialized tasks that execute specific business changes that generate events, which do not require approvals.
- You have changes actions, triggered by identity policies, that generate events that do not require workflow approval.
- You need the flexibility to associate specific workflow processes with task-specific changes.

Environment-level approvals may require significant processing and system resources as the volume of transactions increases. This may eventually introduce performance and scalability issues. Using task-level approvals, where appropriate, may reduce or eliminate these issues.

Deploy Delegated Administration for Users, Groups and Organizations

Delegated administration is the management of users and their entitlements by having different Identity Manager users perform the functions of modifying, assigning, and using a role.

Note: Delegation models must be carefully constructed to ensure good performance and scalability in your Identity Manager implementation.

Delegation is enforced by scope rules, which are defined in member and admin policies for admin roles. A scope rule determines the objects on which a role member can use the role. For example, a scope rule may enable a User Manager to manage users in his department, but not in other departments.

Generally, scope rules should reflect the logical structure of the user store. For example, in a hierarchical LDAP user store, scope may be defined by organizations. In a relational database, scope can be defined using attributes such as department ID.

Note the following when deploying delegated administration for users, groups, and organizations:

- Limit access to relationship tabs, such as the Admin Roles and Provisioning Roles tabs, in user-related tasks. These relationship tabs are included in default user tasks, such as Create User and Modify User. Consider removing them from the default tasks and using them only in specialized tasks which are associated with a small number of admin roles.
- Identity Manager evaluates each scope rule dynamically; scope information is not cached. Consider creating scope rules that contain simple directory queries to ensure good performance.
- Evaluate the performance of scope rules by determining how long it takes Identity Manager to return the objects an administrator can manage.

Deploy Delegated Administration for Roles

Delegated administration of roles grants the most significant privileges in Identity Manager and can have the greatest affect (see page 70) on performance. For these reasons, you should consider deploying delegated administration for roles after you have deployed all other functionality.

When deploying delegated administration for roles, note the following:

- Limit the number of admin roles, admin role members, and admin role administrators to protect the environment and ensure good performance.
- Once you deploy delegated administration for roles, conduct performance and scalability tests. Optimize the environment as needed.

Chapter 5: Integrating with SiteMinder

This section contains the following topics:

SiteMinder Integration

When CA Identity Manager integrates with SiteMinder, SiteMinder can add the following functionality to an Identity Manager environment:

Advanced Authentication

CA Identity Manager includes native authentication for Identity Manager Environments by default. CA Identity Manager administrators enter a valid username and password to log into an Identity Manager Environment. CA Identity Manager authenticates the name and password against the user store that CA Identity Manager manages.

When CA Identity Manager integrates with SiteMinder, CA Identity Manager uses SiteMinder basic authentication to protect the Environment. When you create an Identity Manager Environment, a policy domain and an authentication scheme are created in SiteMinder to protect that Environment.

When CA Identity Manager integrates with SiteMinder, you can also use SiteMinder authentication to protect the Management Console.

Access Roles and Tasks

Access roles enable Identity Manager administrators to assign privileges in applications that are protected by SiteMinder. Access roles include access tasks, which represent a single action that a user can perform in a business application, such as generating a purchase order in a finance application.

Directory Mapping

An administrator may need to manage users whose profiles exist in a different user store from the one that is used for authenticating the administrator. In other words, when logging in to the Identity Manager Environment, the administrator must be authenticated using one directory and authorized to manage users in a second directory.

When CA Identity Manager integrates with SiteMinder, you can configure an Identity Manager Environment to use different directories for authentication and authorization.

Advanced Password Policies

CA Identity Manager enables you to create basic password policies that manage user passwords by enforcing rules and restrictions governing password expiration, composition, and usage.

If you configure CA Identity Manager to integrate with SiteMinder, you can create advanced password policies that enable you to define the additional rules and restrictions.

Note: For more information, see the *Administration Guide*.

Skins for Different Sets of Users

A skin changes the look of the User Console. When CA Identity Manager integrates with SiteMinder, you can enable different sets of users to see different skins. To accomplish this, you use a SiteMinder response to associate a skin with a set of users. The response is paired with a rule in a policy, which is associated with a set of users. When the rule fires, it triggers the response to pass information about the skin to CA Identity Manager, to build the User Console.

Note: For more information, see the *User Console Design Guide*.

Locale Preferences for a Localized Environment

When CA Identity Manager integrates with SiteMinder, you can define a user's locale preference using an imlanguage HTTP header. In the SiteMinder Policy Server, you set this header within a SiteMinder response and specify a user attribute as the header's value. This imlanguage header acts as the highest priority locale preference for a user.

Note: For more information, see the *User Console Design Guide*.

Note: For information on the components in a CA Identity Manager implementation that includes SiteMinder, see Installation with SiteMinder Policy Server (see page 41).

SiteMinder Authentication

CA Identity Manager includes the following consoles, which should be protected:

User Console

Enables CA Identity Manager administrators to perform tasks in an Identity Manager environment.

Management Console

Enables CA Identity Manager administrators to create and configure an Identity Manager directory, Provisioning Directory, and an Identity Manager environment.

CA Identity Manager includes native authentication, which protects the User Console by default. The Management Console is not protected by default.

To protect the Management Console or to configure other types of authentication for the User Console, such as certificate or key authentication, CA Identity Manager must integrate with SiteMinder.

Note: For more information, see the *Configuration Guide*.

Password Policies with SiteMinder

CA Identity Manager enables you to create basic password policies that manage user passwords by enforcing rules and restrictions governing password expiration, composition, and usage.

If CA Identity Manager integrates with SiteMinder, you can create advanced password policies that enable you to define these additional rules and restrictions:

- Directory filters
- Password expiration:
  - Track failed or successful logins
  - Authenticate on login
  - Password expiration if not changed
  - Password inactivity
  - Incorrect password
  - Multiple regular expressions
- Password restrictions:
  - Minimum days before reuse
  - Minimum number of passwords before reuse
  - Percent different from last password
  - Ignore sequence when checking for difference
  - Profile attribute matching
  - Dictionary matching

Chapter 6: Optimizing Identity Manager

This section contains the following topics:

Identity Manager Performance

Identity Manager performance depends on the individual performance of different features and components.

You can optimize the following functionality in an Identity Manager environment:

- Roles
- Tasks
- Group membership and management
- Identity policies

For additional performance gains, you can also tune the following components:

- User store
- Provisioning components
- Runtime components, including the databases, such as the task persistence database, and application server settings

To ensure best performance, configure the Identity Manager functionality using the guidelines in the following sections. Then, measure performance and tune components, as needed. Because the components work together, it may take several iterations before you find the best tuning settings for your environment.

Role Optimizations

Identity Manager includes three types of roles:

- Admin roles

  Determine the privileges a user has in the User Console.

  When a user logs into an Identity Manager environment, the user's account has one or more admin roles. Each admin role contains tasks, such as Create User, that a user can complete in that Identity Manager environment. The admin roles that a user has determine the presentation of the User Console, therefore users see only the tasks that are associated with their roles.

- Provisioning roles

  Give users accounts in managed endpoints, such as an email system.

- Access roles

  Offer an additional way to provide entitlements in Identity Manager.

Roles include policies that determine the following:

- Who can use the role (for admin and access roles only) and where they can use it
- Who can manage role members and administrators
- Who can modify the role definition

Evaluating roles and their associated privileges can have a significant impact on Identity Manager performance.

How Role Evaluation Affects Performance at Login

When an Identity Manager user attempts to log into the User Console, the following actions occur:

1. Identity Manager prompts the user to supply credentials, such as a user name and password.
2. The user's credentials are authenticated using one of the following methods:
   - Identity Manager native authentication
   - SiteMinder authentication, if the Identity Manager implementation includes SiteMinder

3. Identity Manager evaluates every member policy for every admin role in the environment to determine which admin roles apply to the user.

   Note: This evaluation occurs only once for a given user. After the initial evaluation, Identity Manager caches the results. Identity Manager uses the cached information until a change occurs to the user or to the set of member policies, which causes Identity Manager to refresh the information in the cache.

4. The Identity Manager User Console displays the categories that the user can view based on his roles.

This process occurs for every user that logs into the User Console. If an Identity Manager environment contains a large number of roles, or inefficient member policies, role membership evaluation can significantly impact performance. In this case, the initial screen that users see when they log into the User Console may display slowly.

Note: Identity Manager does not need to evaluate member policies when a user accesses a public task to self-register or to request a forgotten password. In these cases, Identity Manager does not need a list of the user's roles because it does not display the complete User Console.

Role Objects and Performance

To support each role, Identity Manager creates a number of objects in the Identity Manager object store (see page 34), depending on the role configuration.

Identity Manager creates one base object for each role. In addition to the base object, Identity Manager creates two objects for each member, admin, and owner policy rule (see page 73), and two objects for each scope rule. The rule objects include:

- Rule definition object

  Contains metadata about the rule, such as rule type

- Rule data object

  Contains the expression to be evaluated

The following table illustrates the objects created for a single admin role.

Object Type	Base Object	Member Policy Objects	Admin Policy Objects	Owner Policy Objects
Admin Role	1	Member rules: 2 (1 rule definition object and 1 rule data object) Scope rules: 2 (1 rule definition object and 1 rule data object) Total: 4 objects	Admin rule: 2 (1 rule definition object and 1 rule data object) Scope rules: 2 (1 rule definition object and 1 rule data object) Total: 4 objects	Owner rule: 2 (1 rule definition object and 1 rule data object) Total: 2 objects

Note: This table assumes that there is only one member, admin, and owner policy.

For any admin role, Identity Manager creates at least 11 objects. If the member policy included 3 scope rules, the number would increase to 15 objects.

Large numbers of role objects may impact the performance of the object store searches and policy evaluations.

Object Store Performance

Identity Manager stores information that it needs to manage users and entitlements in an object store. Having a large number of role objects in the object store may cause the following issues:

- Searches for managed objects on Identity Manager tasks screens may take longer.

  To reduce the impact on searches, index attributes used in searches (see page 89).

- Role management tasks may execute slowly.

  Some examples of role management tasks that are affected by a large object store include the following:

  - A Create Admin Role task is slow because Identity Manager must confirm that the role name is unique in the object store.
  - The Delete Admin Role task must remove all objects created to support the role and the object cache must be updated.

- Identity Manager takes a long time to evaluate role policies.

Identity Manager caches information in the object store to improve performance.

Optimize Role Policy Evaluation

For each admin role, you can create three types of policies:

- Member policies

  Define a member rule, which determines the users who receive the role, and scope rules, which determine the objects that role members can manage

- Admin policies

  Define admin rules, scope rules, and administrator privileges for a role

- Owner policies

  Define who can modify a role

To optimize performance when Identity Manager evaluates role policies, consider the following:

- Limit the number of admin roles in an Identity Manager environment.
- Follow the guidelines for creating policy rules (see page 73).
- Tune the user store. the user store.
- Tune the policy store, if Identity Manager includes SiteMinder.

Guidelines for Policy Rule Creation

One of the key factors in determining the overall performance of role policy evaluations is the amount of time it takes to evaluate any single policy rule. To improve policy rule evaluation time, note the following when you create a policy:

- When possible, limit the number of policy objects that Identity Manager creates and the number of user store searches that it performs by creating policy rules with complex expressions.

  A single rule with a complex expression is more efficient than multiple rules with simple expressions.

- When possible, select the most efficient and scalable type of policy rule.
- Enable the in-memory evaluation option for policy rules.

  The in-memory evaluation option significantly reduces policy evaluation time by retrieving information about a user to be evaluated from the user store and storing a representation of that user in memory. Identity Manager uses the in-memory representation to compare attribute values against policy rules.

  Note: For more information about the in-memory evaluation option, see the *Configuration Guide*.

- Tune the user store.
- Tune the policy store, if your Identity Manager implementation includes SiteMinder.

Limit Policy Objects and User Store Searches

Each rule in a role policy requires a set of objects in the object store. When Identity Manager evaluates a rule, it loads these objects and performs any required user store searches.

The following example shows a member policy that includes three member rules. Each rule includes four scope rules.

In this example, Identity Manager creates the objects and performs the user store searches described in the following table when evaluating and applying the member policy.

Rule	Policy Objects	Potential User Store Searches
■ Member rule: where (Department = "Administration") ■ User scope: City = "Boston" ■ Group scope: Group Name = "Product Team" ■ Provisioning role scope: Name = "Employee" ■ Access Task Scope:Name = "Development"	10 (one for each rule definition object, one for each rule data object)	5 (one for each rule definition object)
■ Member rule: where (Department = "Engineering") ■ User scope: City = "Boston" ■ Group scope: Group Name = "Product Team" ■ Provisioning role scope: Name = "Employee" ■ Access Task Scope:Name = "Development"	10	5
■ Member rule: where (Department = "Human Resources") ■ User scope: City = "Boston" ■ Group scope: Group Name = "Product Team" ■ Provisioning role scope: Name = "Employee" ■ Access Task Scope:Name = "Development"	10	5

In this example, Identity Manager creates 30 objects and executes 15 directory searches to determine membership and scope.

To limit the number of policy objects and user store searches that Identity Manager performs, combine rules into complex expressions. The following example specifies the same entitlements in the first example as one member rule.

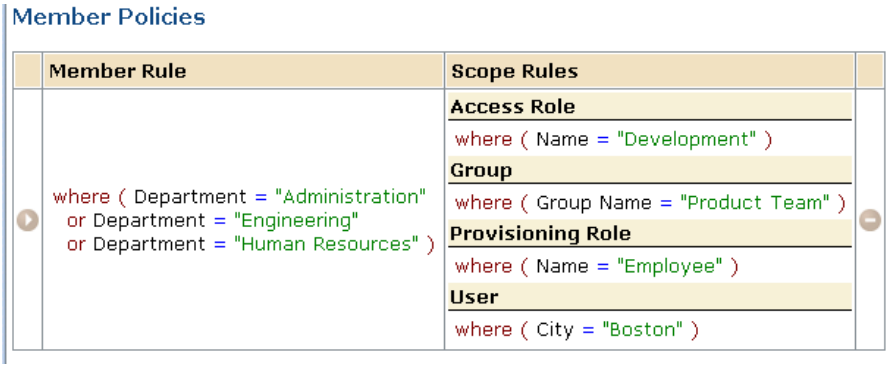

In this example, Identity Manager creates only ten policy objects and performs only five user store searches.

Rule	Policy Objects	Potential User Store Searches
■ Member rule: where (Department = "Administration") OR where (Department = "Engineering") OR where (Department = "Human Resources") ■ User scope: City = "Boston" ■ Group scope: Group Name = "Product Team" ■ Provisioning role scope: Name = "Employee" ■ Access Task Scope:Name = "Development"	10	5

Select Scalable Policy Rule Types

In addition to the number of policy rules, the type of policy rule may also impact performance. Typically, policy rules are constructed based on how the user store is structured and how entitlements are determined. For example, you may create policy rules based on group membership, organization, or user attributes. However, when there are multiple ways to construct policy rules, consider the performance guidelines in the following table before deciding which type of rule to construct.

Note: The policy rule types in the following table are listed in order of performance, beginning with the most efficient rule type.

Policy Rule Type	Performance Notes
Organization	■ Best overall performance ■ Does not require a search in LDAP directories. Identity Manager uses the DN of the user being evaluated and the DN of the organization in the policy rule
Role	■ Identity Manager stores role object information and previous evaluations in the object store cache ■ In most cases, performance will be as good as organization policy rules
User attribute	■ Provides the best user store search performance, and is the least affected by large user populations ■ Allows you to enable in-memory evaluation for significant performance gains
Group Membership	■ Performance depends on group size and user store type

Task Optimizations

In Identity Manager, the tasks that a user sees in the User Console depend on that user's specific privileges. To display and execute tasks, Identity Manager must perform multiple security evaluations, which may have a significant impact on performance when applied over all of the users in an Identity Manager environment.

Identity Manager performs security evaluations when the following actions occur:

- A user logs into the User Console

  In this case, Identity Manager must evaluate a user's roles to determine which tasks that user can access in the User Console.

- A user invokes a task

  When a task is invoked, Identity Manager must determine which objects that user can manage with that task.

- A user accesses a relationship tab

  A relationship tab is any tab where a user can view or manage a one-to-many relationship between the task's subject and a set of entitlements. An example of a relationship tab is the Admin Roles tab, which displays the roles that a user has.

- A user adds objects on a relationship tab

  For example, Identity Manager performs additional security checks when a user adds additional roles to another user on the Admin Roles tab.

Task performance is affected by the following:

- Task scope, which determines where an administrator can use a task
- Relationship tabs, which display an object's relationship to other objects

Task Scope Evaluation and Performance

When an administrator uses an admin task that involves searching for a managed object, such as a user, group, organization, task, or role, Identity Manager evaluates and applies task scope rules. These rules can significantly impact the amount of time Identity Manager takes to display the list of objects to select for the task.

Note: Unlike member, admin, and owner policy evaluations, information about scope rule evaluations is not stored in a cache.

Task scope is determined by the following:

- The type of object that the task manages.
- Scope rules that apply to the admin role that includes the task. Scope rules are defined in member, owner, and admin policies.
- Any user-defined search criteria.

For example, consider a Modify User task, which is included in the User Manager role. The User Manager role has a member policy with a scope rule that allows User Managers to manage users in the Employees organization. An administrator opens the Modify User task and enters the search criteria: Last Name starts with A. In this case, the scope for the Modify User task is all users in the Employees organization whose last name starts with A.

How Identity Manager Renders Relationship Tabs

A relationship tab allows users to view and manage the relationship that a task's subject has with a set of entitlements. For example, the Provisioning Roles tab shows the provisioning roles that a user has.

To determine the objects that appear on a relationship tab, CA Identity Manager performs numerous security evaluations, which can significantly impact performance.

The following example shows the steps that CA Identity Manager takes to render the Provisioning Roles tab:

1. An administrator clicks the Provisioning Roles tab in the Modify User task.
2. CA Identity Manager retrieves the provisioning roles where the selected user is a member.
3. If the tab is configured to allow management of role administrators, CA Identity Manager makes a second call to retrieve the list of provisioning roles where the selected user is an administrator.
4. CA Identity Manager evaluates each provisioning role that the user has to see if the administrator who initiated the task can manage membership for that role.

   If the administrator can manage role members, CA Identity Manager displays an active check box in the Membership column for that role in the list of roles on the tab.
5. CA Identity Manager evaluates each provisioning role that the user has to see if the administrator who initiated the task can manage administrative rights for that role.

   If the administrator can manage administrative rights, CA Identity Manager displays an active check box in the Administrator column for that role in the list of roles on the tab.

   CA Identity Manager must complete steps 2-5 to display the provisioning roles the user currently has. If the administrator needs to assign a new provisioning role, the following additional steps are required.
6. The administrator clicks the Add button to locate new provisioning roles to assign.
7. CA Identity Manager displays a search screen that the administrator can use to search for the role to add.
8. The administrator enters a search filter to find the role to add.
9. CA Identity Manager returns the list of provisioning roles that meet following criteria:
   - The roles match the search filter entered by the administrator.
   - The administrator can manage membership for the roles.
   - The user is in the administrative scope of the administrator for the roles.
   - The user does not already have the provisioning roles.
10. CA Identity Manager repeats step 9 to determine the roles where the administrator can manage administrative privileges.

Relationship Tabs and Performance

Because of the number of security evaluations Identity Manager performs, rendering a relationship tab can significantly impact performance. The factors that determine performance vary depending on the type of tab.

For role relationship tabs, the following factors can impact performance:

- Number of roles where the task's subject is a member
- Number of roles where the task's subject is an administrator
- Number of total objects in the system that Identity Manager requires to calculate the subject's roles
- Number of member/admin policies per role
- Complexity of the member/admin policy scope rules
- The ability to maintain cached authorizations for the task invokers to limit the effect of the security enforcements

To determine group membership and administrative privileges on group relationship tabs, Identity Manager must search all of the groups in the user store. Performance of these searches depends on the following factors:

- Number of group objects in the user store
- Number of members in any group
- Performance of the database or directory where the user store exists

Task Processing and Performance

Admin tasks include events, actions that CA Identity Manager performs to complete the task. A task may include multiple events. For example, the Create User task may include events that create the user's profile, add the user to a group, and assign roles.

When CA Identity Manager processes a task, it processes each event associated with the task. During event processing, CA Identity Manager saves each event four times. This allows CA Identity Manager to preserve in-process actions in the event of an unexpected system shutdown.

When CA Identity Manager processes multiple events at the same time, the events are added to a queue. When the first event completes the first stage of its lifecycle, it is saved, and then moved to the back of the queue to wait for the second stage processing to begin. CA Identity Manager then completes the first processing stage for the next event in the queue, and that event moves to the end of the queue. The process continues until all of the events in the queue have completed the first processing stage. Then, the first event in the queue begins the second processing phase. This continues until all of the events in the queue complete all four processing stages.

Under normal load conditions, this behavior does not impact performance. However, if the system is processing a large number of tasks and events, such as during a bulk load of a large user population, each event and task must wait longer in the queue and, therefore, has a longer completion time.

To prevent performance issues under load conditions, consider the following actions:

- Use the Task Priority setting on a task's Profile tab.

  The Task Priority setting allows you to set the priority of a task to High, Medium, or Low.

  Tasks that need to be processed immediately should be set to High. Tasks involved in a bulk load should be set to Low.

  If a task priority is set, the events associated with the task are processed with other tasks that have the same priority. For example, if the Modify User task is set to High priority, and an administrator modifies a user profile, CA Identity Manager processes that task before tasks with Medium or Low priority. If there are other High priority tasks, CA Identity Manager completes the first processing stage for the first High priority event and then moves that event to the end of the list of other High priority events.

- Install a separate, dedicated Identity Manager Server to handle bulk load operations

Guidelines for Optimizing Tasks

The default tasks, which Identity Manager deploys when you create an Identity Manager environment, are configured to support a wide range of administration use cases. Most Identity Manager implementations do not require all of the functionality provided in the default tasks. After creating an Identity Manager environment, modify these tasks to suit specific administration needs.

The following steps provide guidelines for modifying tasks:

- **Create specialized user management tasks**

  The default Create User, Modify User, and View User tasks provide full administrative capabilities. In most implementations, only a small number of administrators need all of the available capabilities.

  Create new tasks that include only the required capabilities. For example, if most user management tasks involve only profile and group management, create a new Modify User task that includes only the Profile and Group tabs. Remove the Admin Roles, Access Roles, and Provisioning Roles tabs, which are available in the default Modify User task.

  Unused tabs can cause significant overhead if they are left in frequently used tasks. This is especially true when using a Task Execution Web Service (TEWS) client, where these tabs may be inadvertently activated through the tab java class, which is provided with Identity Manager.

  The specialized tasks that you create should match the delegated administration model (see page 63) that you defined for your environment.

- **Disable Manage Administrators in relationship tabs**

  By default, all relationship tabs provide the ability to manage administrative rights for the object that the tab manages, such as roles and groups. Most implementations do not need to provide this functionality to administrators.

  To eliminate the additional overhead that occurs when Identity Manager evaluates administrative rights, clear the Manage Administrators option on the following tabs, if this functionality is not required:

  - Admin Roles
  - Provisioning Roles
  - Access Roles
  - Groups

  To enable users to manage administrative rights on specific tabs, create copies of the default tabs, enable the Manage Administrators option, and disable the Manage Members option. Add the new tabs to specialized tasks, which are only used by the administrators who need them.

- **Enable scoped searches in role relationship tabs**

  You can configure each role tab to include searches that allow administrators to specify criteria for new roles to assign to a user. Role searches limit the number of member and admin policy rules that Identity Manager must evaluate to determine which roles an administrator can assign to a user.

- **Set task synchronization options**

  For each Identity Manager task, you can specify a user synchronization option, which synchronizes users with identity policies, and a provisioning account synchronization option, which synchronizes users with provisioned accounts. The options enable you to synchronize users when a task completes, or when an event completes.

  To eliminate evaluation and processing time, set the synchronization to occur when a task completes, instead of when events complete.

Guidelines for Group Member\Administrator Optimizations

To improve performance of searches for group members and administrators, consider the following:

- Define well-known attributes in the directory configuration file (directory.xml), which describes the user store structure and contents to Identity Manager.

  A well-known attribute is an attribute that has a special meaning in Identity Manager.

  To improve group member\administrator searches, define the following well-known attributes for the user object:

  %MEMBER_OF%

  Identifies an attribute on the user object that stores a list of groups where the user is a member.

  When defined, this attribute can prevent Identity Manager from searching all of the members in all of the groups in the user store. Group searches can significantly affect performance in very large groups.

  %ADMINISTRATOR_OF%

  Identifies an attribute on the user object that stores a list of groups where the user is an administrator.

  Like the %MEMBER_OF% attribute, this well-known attribute can eliminate lengthy group searches.

- Specify the Group Type in the directory configuration file

  Identity Manager supports three types of groups: standard groups, nested groups, and dynamic groups.

When you define the group object in the directory configuration file, you can specify the type of groups that the user store supports. If your implementation does not support nested or dynamic groups, set the Group Type attribute as follows:

GroupType = NONE

The setting NONE specifies support for standard groups.

The default Group Type setting is ALL, which may impact performance.

Note: For more information about well-known attributes and group types in the directory configuration file, see the *Configuration Guide*.

- Set the Provisioning Directory cache indices to improve GlobalGroup performance

  For Identity Manager implementations that include a combined user store and Provisioning Directory, GlobalGroup membership can be optimized for policy rule evaluation for roles and identity policies.

  To enable this optimization, you index the following attributes, which the Provisioning Server uses to resolve group membership, in the Provisioning Directory cache:

  eTID

  The unique object ID attribute. For group membership lookups, the value is a specific user or group involved in the lookup.

  eTPID

  The parent ID of the object used when searching for membership relationships.

  eTPID

  The child ID of the object used when searching for membership relationships.

  Additionally, add the following hash entries:

  eTSuperiorClass

  The type of the parent object in a membership lookup

  eTSubordinateClass

  The type of the child object in a membership lookup

  Note: For more information about the Provisioning Directory cache, see the *Installation Guide*.

Identity Policy Optimizations

An *identity policy* is a set of business changes that occurs when a user meets a certain condition or rule. These changes can include assigning or revoking roles, assigning or revoking group membership, and updating attributes in a user profile.

Identity Manager evaluates identity policies when user synchronization occurs.

Identity policy performance is affected by the following:

- How the identity policies are configured
- How often user synchronization occurs

How Users and Identity Policies Are Synchronized

When using identity policies, it is important to understand how CA Identity Manager evaluates and applies the policies to users. Without a thorough understanding of the user synchronization process, you may configure identity policy sets that yield unexpected results.

The following procedure describes how CA Identity Manager evaluates and applies identity policies:

1. The user synchronization process begins:
   - **Automatically**—You can configure CA Identity Managertasks to automatically trigger user synchronization
   - **Manually**—Use the Synchronize User task in the User Console to synchronize a user.
2. CA Identity Manager determines the set of identity policies that apply to a user.
3. CA Identity Manager compares the set of identity policies that apply to a user with the list of policies that have already been applied to that user.

   Note: The list of policies that have been applied to a user is stored in the %IDENTITY_POLICY% well-known attribute in the user profile. For information on configuring this attribute, see the *Configuration Guide*.
   - If an identity policy is on the list of applicable policies, *and* the policy has *not* been applied to the user previously, then CA Identity Manager adds the policy to an allocation list.
   - If an identity policy is on the list of applicable policies, the policy has been previously applied to the user, and the Apply Once setting for the policy is disabled, CA Identity Manager adds the policy to a reallocation list.
   - An identity policy is not on the list of applicable policies, and the policy has been applied to the user, the user no longer matches the policy condition. CA Identity Manager adds these policies to a deallocation list.
4. After CA Identity Manager evaluates all of the policies for a user, it applies policies in the following order:
   a. Identity policies from the deallocation list
   b. Identity policies from the allocation list
   c. Identity policies from the reallocation list

5. After the identity policies have been applied, CA Identity Manager reevaluates the policies to see if any additional changes are needed based on changes that occurred in the first synchronization process (steps 2-4).

   This is to ensure that changes made by applying identity policies did not trigger other identity policies.

6. CA Identity Manager continues to reevaluate and apply identity policies until the user is synchronized with all applicable policies, or until CA Identity Manager reaches the maximum recursion level, which is defined in the Management Console.

   For example, an identity policy may change a user's department when the user is assigned a role. The new department triggers another identity policy. However, if the recursion level is set to 1, the subsequent change is not made until the user is synchronized again.

   For more information about setting the recursion level, see the Management Console Online Help.

Design Efficient Identity Policies

Use the following guidelines when you create identity policies:

- **Limit the number of policy objects**

  Identity Manager creates objects in the object store that support identity policies. To reduce the number of objects in the object store, create identity policies with complex expressions.
  A similar approach is recommended for role policies (see page 74).

- **Limit identity policy set iterations**

  You can configure the recursion level for an identity policy, which determines the number of times that CA Identity Manager evaluates and applies identity policies when a user is synchronized. For example, an identity policy may change a user's department when the user is assigned a role. The new department triggers another identity policy. However, if the recursion level is set to 1, the subsequent change is not made until the user is synchronized again.

  Setting the recursion level limits the number of times that CA Identity Manager must evaluate identity policies.

- **Limit dependencies between identity policy rules**

  You can create an identity policy where the change action (Action on Apply Policy or Action on Remove Policy) of one policy is used in the identity policy condition of another policy as shown in the following table.

Identity Policy Condition	Action on Apply Policy	Action on Remove Policy
where (Job Code = "100")	Make member of (provisioning role "Account Manager")	Remove member of (provisioning role "Account Manager")
Who are members of (provisioning role "Account Manager")	Make member of (group "Account Managers")	Remove member of (group "Account Managers")

  When CA Identity Manager evaluates this type of policy, it must evaluate and apply changes at least twice to ensure that both conditions are met. The recursion level, which is set for an entire Identity Manager environment, must be greater than 1, which then causes additional evaluations for each identity policy set.

Limit the Tasks that Trigger User Synchronization

Identity policies are evaluated and applied during the user synchronization process. You can configure automatic synchronization by specifying one of the following user synchronization options for a task:

On Task Completion

CA Identity Manager starts the user synchronization process after all of the events in a task have completed.

On Every Event

CA Identity Manager starts the user synchronization process when each event in a task completes.

For best performance, limit the number of tasks that trigger automatic user synchronization.

Consider the following when configuring user synchronization:

- **Disable user synchronization for password tasks**

  In most cases, passwords are not used in identity policy conditions.

- **Disable user synchronization for the Synchronize User task**

  Since the Synchronize User task triggers identity policy evaluations, CA Identity Manager performs the evaluations again if the user synchronization option is enabled for this task.

- **Create specialized tasks**

  When possible, create tasks that execute modifications that trigger identity policy conditions and enable user synchronizations for those tasks only.

Optimize Identity Policy Rule Evaluation

To reduce the evaluation time for identity policy conditions that include user-attributes, you can enable an in-memory evaluation option. When the in-memory evaluation option is enabled, Identity Manager retrieves information about a user to be evaluated from the user store and stores a representation of that user in memory. Identity Manager uses the in-memory representation to compare attribute values against policy conditions. This limits the number of calls Identity Manager makes directly to the user store.

Note: For more information about the in-memory evaluation option, see the *Configuration Guide*.

User Store Tuning

User store tuning involves a number of steps, including the following:

- Optimizing the structure of the user store
- Tuning underlying stores
- Implementing load balancing and replication

These steps depend on the type of user store that you are using. For tuning information in these areas, see the documentation for the database or directory that contains the user store.

In addition to the general tuning considerations, the following tuning considerations are specific to CA Identity Manager:

- **Measure user store search performance**

  For optimum performance, CA Identity Manager policy evaluation searches should complete within 10-20 milliseconds.

  To ensure that CA Identity Manager can consistently complete these searches in the recommended time, consider testing search performance under multiple load conditions.

  You can also use this measurement to determine when a user store reaches its physical limits and additional servers are required for load balancing.

- **Index attributes**

  Index each attribute that is used in a role policy or identity policy. Indexing attributes can provide significant performance improvements.

  Note: For information about indexing attributes, see the documentation for the LDAP directory or relational database that contains the user store.

- **Cache LDAP Binds**

  In CA Identity Manager, all directory LDAP binds are executed by the proxy user defined on the Identity Manager Directory object. For each connection, the same LDAP bind occurs for this same user repeatedly.

  If you are using an LDAP directory as a user store, configure the directory to cache LDAP binds (or sessions), if the directory supports it.

- **Enable user store caches**

  When CA Identity Manager evaluates the policy decisions for a user, that information is stored in an authorization cache. When the cached information expires, CA Identity Manager evaluates all policies for that user again.

  To improve performance of user store searches in subsequent policy rule evaluations, enable the user store to cache searched data, if your user store supports it.

CA Directory includes a cache, called dxCache, which is an in-memory database implementation that can search across cached data.

Note: For more information about CA Directory, see the *CA Directory Administrator Guide*.

Tuning for Provisioning Components

When a CA Identity Manager implementation includes provisioning, use the following optimizations to ensure the best performance:

- Optimize the connection between the Identity Manager Server and the Provisioning Server

  CA Identity Manager communicates with the Provisioning Server using the Java IAM (JIAM) API. To improve communication performance, you configure the following:

  - JIAM session pool for multiple connections to the Provisioning Server

    Note: CA recommends setting the initial sessions value to 8, and the maximum sessions to 128.

  - JIAM cache for objects retrieved from the Provisioning Server

  Note: For information on JIAM configuration settings, see the *Administration Guide*.

- Set account synchronization to occur at the end of a task (see page 82), instead of the end of each event
- Tune the Provisioning Server

  Note: See the *Administration Guide* and *Installation Guide* for more information.

Runtime Components Tuning

Business changes in Identity Manager are accomplished through tasks. A task includes one or more events, which represent activities that Identity Manager performs to complete the task. For example, a Create User task may include the CreateUserEvent and the the AddToGroupEvent.

Identity Manager includes the following components, which process tasks and events at runtime:

- Identity Manager databases, which support Identity Manager functionality
- JMS messages, which are responsible for processing events

Tuning Identity Manager Databases

When executing tasks, Identity Manager uses the following databases:

- Task persistence

  Maintains information about Identity Manager tasks and events over time. This allows Identity Manager to restore the last known state of events and tasks in the case of system failure.

  Note: This database has the most significant impact on Identity Manager performance because the task and its events are saved and retrieved from the database during state transitions.

- Audit

  Provides a historical record of operations that occur in an Identity Manager environment.

- Workflow

  Stores workflow process definitions, jobs, scripts, and other data required by the workflow engine.

- Reporting

  Stores snapshot data, which reflects the current state of objects in Identity Manager at the time the snapshot is taken.

Identity Manager communicates with each database through a JDBC connection pool. You create and configure a JDBC connection pool in the application server that hosts Identity Manager. When you configure the JDBC connection pool, note the following:

- Consider the number of concurrent tasks that will execute at any one time.
- Consider the other runtime components when you configure the JDBC connection pool size. Each runtime component works in conjunction with the other runtime components.

  Note: CA recommends setting the initial connection pool value to 128.

- For the task persistence database, the number of database connections in the pool must allow each executing task to retrieve and update task and event data throughout the lifetime of the task.
- The task persistence database uses prepared statements. Be sure to configure the prepared statement cache for the database that you are using to store task persistence data.

  Note: See the documentation for the database that you are using for task persistence for information on configuring the prepared statement cache.

JMS Settings

A CA Identity Manager task includes events, actions that CA Identity Manager performs to complete a task.

During an event's life cycle, it transitions through the following states:

- BEGIN
- APPROVED
- EXECUTING
- COMPLETED

Identity Manager uses JMS messages to control these state transitions.

How JMS Messages Drive Event Transactions

CA Identity Manager uses JMS messages to drive an event’s state transitions. The following procedure describes the steps involved:

1. A user submits a task.
2. The task generates one or more events.
3. When an event is ready for processing, CA Identity Manager sets the event's state to BEGIN and the event is persisted in the task persistence database.
4. CA Identity Manager creates a JMS message containing the event ID and posts that message to the Event Message Queue.
5. Upon receiving the message, JMS then invokes an instance of the Event Message Driven Bean, which is an implementation of the Event Controller.
6. The Event Controller uses the event ID in the message to retrieve the event from the task persistence database, and executes the actions for the event’s current state.
7. Upon completion of that state, the event is set to the next state, persisted in the task persistence database, and a new JMS message is posted for processing the next state.

   This cycle continues until the event has completed its state machine.

JMS Messages and Performance

For any event, there are four states that require JMS messages for state transition:

- BEGIN
- APPROVED
- EXECUTING
- COMPLETED

To process a single event, the following actions take place:

- Four posts to the Events Message Queue
- Four invocations of the Message Driven Bean
- Eight connections to the task persistence database (one read action and one write action per state)

Theses actions may impact the amount of time it takes Identity Manager to process a task.

To ensure best performance during state transitions, tune the JMS resources in the application server that hosts Identity Manager so that adequate JMS resources are available.

Tuning JMS Settings

The following application server JMS tuning parameters define Queue connections and Message Driven Bean instance pools.

- **WebSphere JMS Tuning**

  WebSphere provides two parameters that you can configure to improve performance of event state transactions:

  - The number of connections in the Queue Connection Factory that are available for posting JMS messages
  - The number of instances of Message Driven Bean objects for processing the posted messages

  Use the WebSphere Administration Console to set the following properties:

  - Set the maximum session pool connections for the EventMsgQueue Queue Connection Factory to 128
  - Set the maximum sessions for the SubscriberMessageEJBPort Listener Port in the Message Listener Service to 128

- **WebLogic Tuning**

  In WebLogic application servers, Queue Connection Factories obtain connection-handling threads from the server's JMS Thread Pool or the default execute pool, depending on the JMS Thread Pool size. If the JMS Thread Pool size is 0, then WebLogic uses the threads in the execute pool.

  CA recommends setting the number of JMS Thread Pool threads equal to the maximum Bean Pool size for the Identity Manager Event Message Driven Bean, which is set to 128 by default.

You use the WebLogic Server Console to set the JMS Thread Pool size in the JMS Services properties for the domain and server where CA Identity Manager is installed.

The Identity Manager Event Message Driven Bean pool size is set by modifying the max-beans-in-free-pool setting in the descriptor file in the following location:

WebLogic_home\domain\applications\IdentityMinder.ear\identityminder_ejb_jar\META-INF\weblogic-ejb-jar.xml

```
<weblogic-enterprise-bean>
    <ejb-name>SubscriberMessageEJB</ejb-name>
    <message-driven-descriptor>
      <pool>
        <max-beans-in-free-pool>128</max-beans-in-free-pool>
        <initial-beans-in-free-pool>16</initial-beans-in-free-pool>
      </pool>
      <destination-jndi-name>com.netegrity.ims.msg.queue</destination-jndi-name>
    </message-driven-descriptor>
  </weblogic-enterprise-bean>
```

- **JBoss Tuning**

  In JBoss application servers, Queue Connection Factories obtain connection-handling threads from the server's Standard JMS Pool session factory. By default, the number of maximum threads is set to 15.

  CA recommends setting this value to match the maximum size value of the Standard Message Bean Container.

  The JMS Session Pool section factory is set in the MaximumSize element of the JMSContainerInvoker in the following file:

  jboss_home\server\default\conf\standardjboss.xml

```
<invoker-proxy-binding>
    <name>message-driven-bean</name>
    ….
    <proxy-factory-config>
      <JMSProviderAdapterJNDI>DefaultJMSProvider</JMSProviderAdapterJNDI>
      <ServerSessionPoolFactoryJNDI>StdJMSPool</ServerSessionPoolFactoryJNDI>
      <MaximumSize>128</MaximumSize>
      <MaxMessages>1</MaxMessages>
      ….
    </proxy-factory-config>
  </invoker-proxy-binding>
```

The Identity Manager Event Message Driven Bean pool size is set by modifying the maximum size value in the following descriptor file:

jboss_home\server\default\conf\standardjboss.xml

```
<container-configuration>
      <container-name>Standard Message Driven Bean</container-name>
      <call-logging>false</call-logging>
      <invoker-proxy-binding-name>message-driven-bean</invoker-proxy-binding-name>
     ……
      <container-pool-conf>
        <MaximumSize>128</MaximumSize>
      </container-pool-conf>
  </container-configuration>
```

Chapter 7: Creating a Disaster Recovery Plan

This section contains the following topics:

Loss of Service from a Disaster

In the event of a disaster, users can lose access to services that are critical to their jobs. As a result, these users cannot provide services to other users.

The urgency to restore access to services depends on the actual use of CA Identity Manager. In some organizations, users require uninterrupted access to services provided by CA Identity Manager while other users require system restoration within a day. In either case, we recommend that you make preparations to protect your CA Identity Manager implementation from an event that causes partial or complete loss of your systems.

By configuring a redundant architecture for CA Identity Manager, you can ensure that services are highly available to users. When a primary component fails, the alternate component continues to provide the same service. In addition, you can routinely back up critical systems and software, so you can restore any system or data that is completely lost.

This document provides general planning guidelines for these scenarios. We recommend that you use these guidelines to develop specific disaster recovery procedures that address your organization's requirements.

How to Plan for Disaster Recovery

To develop an effective disaster recovery plan, you engage in the following phases, which are detailed in this chapter.

✓	Phase
	1. Define Disaster Recovery Requirements (see page 99) Based on your organizational needs, identify what types of disaster to anticipate and how quickly you would need to restore services.
	2. Design a Redundant Architecture (see page 99) According to your requirements, design an architecture with redundant components at a remote location.
	3. Develop Backup Plans (see page 101) To protect your installation, develop plans for backing up components.
	4. Develop Restore Procedures (see page 103) Develop procedures for restoring lost components.
	5. Document the Recovery Plan (see page 106) Document your plans for recovering CA Identity Manager from a disaster.
	6. Test the Recovery Plan (see page 106) Based on your disaster recovery procedures, verify that you can reinstate your CA Identity Manager implementation as it existed before the event.
	7. Provide Disaster Recovery Training (see page 108) Complete the effort by making sure that the people responsible to recover systems from a disaster are trained to do so.

Define Disaster Recovery Requirements

The following are some general guidelines to consider for defining requirements for a disaster recovery plan:

1. Assemble a team with the following knowledge:
   - Knowledge of the architecture and systems that support CA Identity Manager
   - Knowledge of how to back up the relational databases and LDAP user stores used by CA Identity Manager
2. Identify potential disaster scenarios to address, including partial or complete loss of systems at one or more sites.
3. List the systems that are critical to be available to support your installation.
4. Define the acceptable maximum downtime for each of these systems.

   For example, systems that support an alternate server may have a lower priority for restoration.

Design a Redundant Architecture

To protect against failure of a critical component, consider the following protective actions using alternate components (servers and directories) and redundant databases at remote locations.

Configure redundancy for CA Identity Manager, using the *Installation Guide.* Include the following components:

- Redundant CA Identity Manager application server nodes as part of a cluster
- A Policy Server cluster provides failover (if you are using CA SiteMinder to protect CA Identity Manager)
- Alternate Provisioning Servers, Provisioning Directories, and connector servers. If a primary component is lost, the system switches over to the alternate component.

Configure redundancy for databases including the following:

- Any of the runtime databases that are part of CA Identity Manager such as the workflow or audit database.

  See the documentation supplied with ORACLE or Microsoft SQL Server.
- The Business Objects database if you are using the Report Server.

  See the Business Objects Enterprise, Release 2 and Release 2 SP 4 documentation on the SAP documentation web site.

Alternate Identity Manager Servers

Providing redundant Application Server nodes for the Identity Manager server provides scalability and performance benefits as well as disaster recovery if individual servers fail. The most common method of providing failover for an application server is to create a cluster. The procedures for creating the cluster are covered in the cluster section of the *Installation Guide*.

Note: For CA Identity Manager r12.0 and higher releases, an application server cluster is the only valid method to implement a multi-node deployment. CA Identity Manager environments require the industry standard J2EE cluster architecture, which uses JMS queues for the backbone. As a result, the only valid method of using multiple nodes in a CA Identity Manager configuration is an application server cluster.

Alternate Provisioning Components

Several provisioning components have the option of an alternate component to provide high availability. The alternate component should be at a remote site for the highest protection.

See the High Availability Provisioning chapter of the *Installation Guide* for specific configuration details of alternate servers and directories.

Multi-Site Provisioning Directories

You can create primary and alternate provisioning directories with the alternate directories at a remote location. CA Directory recommends that you install three Provisioning Directories, one primary and two alternate directories.

Multi-Site Provisioning Servers

To protect against failure of the primary Provisioning Server, you can configure an alternate Provisioning Server. The difference between primary and alternate Provisioning Servers is that the primary server installation populates the Provisioning Directory container entries. Also, uninstalling a primary server removes those entries. Apart from installation and uninstallation, the primary and alternate servers function in the same manner.

Multi-Site Connector Servers

For either the Java or C++ Connector Server, you can configure multiple connector servers to serve the same endpoint or endpoint type.

For each connector server you configure, you should configure an alternate connector server at a remote location to handle the same endpoints. If the connector server fails, the alternate server immediately manages the communication with the endpoints.

Redundant Databases

The supported database software, Microsoft SQL Server and Oracle, provide the capability to provide redundant databases. If the main database fails, the redundant database is immediately available. The redundant database should be at a remote site in case the entire site is affected.

Develop Backup Plans

To protect against the loss of any or all systems, use off-site storage for all data that you back up and a backup schedule that meets your maximum downtime requirements. The backup and restore procedures use different applications, so they should be coordinated for recovery of the CA Identity Manager system as a whole.

Include the following components in your backup plans:

Component	Description	Backup Method
The Identity Manager User Store	An LDAP user directory or a relational database that contains the records for Identity Manager users	See the documentation supplied with your database or LDAP software.
The Identity Manager User Store	An LDAP user directory or a relational database that contains the records for Identity Manager users	See the documentation supplied with your database or LDAP software.
The Identity Manager User Store	An LDAP user directory or a relational database that contains the records for Identity Manager users	See the documentation supplied with your database or LDAP software.
The Identity Manager Databases	The databases for Task Persistence, Workflow, Auditing, Object Store, Reporting, and Task Persistence archive Workflow, Task Persistence, and Auditing have the highest frequency of change and backups should be scheduled accordingly.	See the documentation supplied with your database software.

Component	Description	Backup Method
SiteMinder Policy Store	An LDAP user directory or a relational database with objects for the SiteMinder Policy Server, if you are using SiteMinder	See the documentation supplied with your database or LDAP software.
Provisioning Directory	An LDAP user directory that contains the records for provisioning users and provisioning objects	See the CA Directory documentation.
Application Server JMS persistent stores	The stores used to hold CA Identity Manager Task Event processing messages	See the Application Server documentation.
Reporting databases	Snapshot database Business Objects database	See the documentation supplied with your database software.
Custom reports	Custom reports and related XML files	See the Business Objects Enterprise, Release 2 and Release 2 SP 4 documentation on the SAP documentation web site.

Include the following components in your backup plans using a file system backup program:

Component	Description
Web Server Components	Configuration of deployed Web Server components, such as Application Server plug-ins and SiteMinder Web Agents. A Web Server front-end is required if you are using load balancing of if you are using SiteMinder to protect access to the Identity Manager User Console.
XML data files	All Identity Manager Directory and Environment files that are used to create, maintain, and archive Identity Manager Object Store objects.
Identity Manager customization components	Files found under the following deployed IdentityMinder.ear folders: ■ Config ■ User_console.war WEB-INF\web.xml

Component	Description
Scripts and Programs	TEWS scripts, programs, program exits
Connector Xpress components	Custom connectors Connector Xpress project files
Disaster Recovery Documentation	Once you create your own documentation for disaster recovery, back it up regularly in case the instructions change.

Develop Restore Procedures

The restore procedures depend on the backup method. The recovery process for a failed system depends on the circumstances. However, in many cases, reinstalling the software is the restore method. See the High Availability Provisioning chapter of the *Installation Guide* for details.

Restore the Identity Manager User Store

To restore the Identity Manager user store, see the documentation supplied with your database or LDAP software. Verify that the data store from backup is intact including access to all user stores.

Restore the Identity Manager Databases

To restore the Identity Manager databases, see the documentation supplied with your database. Verify that the data store from backup is intact including access to all databases.

Restore the SiteMinder Policy Store

To restore the SiteMinder policy store, see the documentation supplied with your database or LDAP software. Verify that the data store from backup is intact including access to all user stores.

Restore the Identity Manager Server

If you lose a cluster node for an Identity Manager server, perform the following steps:

1. Use the standard documented procedure to add a node.

   See the *Installation Guide* chapter on cluster installation.

2. Update the connection to the Provisioning Server.

   See the section on Provisioning failover in the High Availability chapter of the *Installation Guide* for details.

Restore a Provisioning Server and Directory

You can restore a lost Provisioning Server by installing an alternate server. If all systems have failed, restore the data lost during the disaster.

Use the following steps:

1. Copy any custom schema files to CA Directory config\schema directory.
2. Install the new Provisioning Directory.

   The datastores will be empty.

3. Restore the data from the backup location.
4. Use the Provisioning Server installer, providing the details for the newly restored Provisioning Directory.

   The domain information should be there already.

5. Restore any custom connector and configuration files from backup.

Note: For more detail, see the CA Directory documentation.

Restore Connector Servers

If you lose a connector server, perform the following steps:

1. Use the Connector Server installer to install a new connector server

   Register it with the Provisioning Server during installation.

2. Remove the registration of the lost connector server using csfconfig or Connector Xpress.

Restore a Report Server

If you lose the report server, see the Business Objects documentation for the procedures that apply. On the SAP documentation web site, check for Business Objects Enterprise, Release 2 and Release 2 SP 4 documentation.

Restore Admin Tasks

If an admin task was in process at the time of disaster, it can be recovered under the following conditions.

- Any admin task that was in a Pending state waiting on approvals continues to be available if the stores used to maintain that state information are preserved. The stores include the Task Persistence database, the JMS store that holds the task and event JMS Messages, and the Workflow database.
- Tasks in the In Progress state (any state other than Pending) are subject to additional conditions.

  A task in this state requires the posting of a new JMS Message to the CA Identity Manager Event message queue to continue being processed. Outages that occur before that event being posted to the queue prevent the task from continuing upon recovery.

  In this situation, two options exist to recover the task:

  - If the task is present in the View Submitted Tasks task in the failed state, go to the task details page and use the Resubmit Task option to resubmit the task.
  - Submit a new task with the same changes.

Document the Recovery Plan

Based on the guidelines in this chapter, we recommend that you develop specific disaster recovery documentation that applies to your organization.

Consider the following approach:

1. Identify the names and locations of systems in your architecture and alternate components for each system.

   For each system, list the software installed, such as the specific JDK installed, the fix release of an application server, and the amount of memory installed. This detail is necessary for any system that you decide you need to rebuild completely.

2. Write procedures for recovering each component or for rebuilding a complete system, if necessary.
3. Identify a method of locating or resetting usernames and passwords to systems and CA Identity Manager user interfaces if they are known only to one or two people.
4. Protect your disaster recovery documentation for loss by creating a backup copy that you store at a well-known off-site location.

Test the Recovery Plan

To help ensure a successful recovery from a disaster, you can schedule a simulated disaster, where certain systems become unavailable. Consider the following tests, which are described in the following sections.

1. Test the failover process.
2. Test restoration of systems.

Test the Failover Process

All servers or directories should have an alternate server or directory at a remote site, including these components:

- Identity Manager server
- Provisioning Server
- Provisioning Directories
- C++ and Java Connector Servers
- Report Server
- Policy Server

Manually stop each component and verify that all operations continue to function, using the alternate component. For example, you could perform the following test of the Provisioning Server:

1. On a system with the Primary Provisioning Server, stop the Provisioning Service services from the Windows services dialog.

   The Primary Provisioning Server is stopped.

2. In the User Console, perform the following actions:

   a. Assign a Provisioning Role to a user.

   b. Verify that the endpoint accounts are created for that user.

      The accounts being created depend on the alternate Provisioning Server handling the communication with the Identity Manager server.

This procedure is an example of one test. For each component that you stop, develop similar tests to verify that the alternate component is in use.

Test the Restore Procedures

According to your disaster recovery documentation, perform a test of each critical component to confirm that you can restore the lost system.

Provide Disaster Recovery Training

Once you believe that the recovery procedures are reliable, you help ensure that the people who must implement the recovery are able to do so. Your organization may require other steps, however, the following are some general guidelines:

1. Publicize the location of the recovery documentation.
2. Perform a dry-run of the training.
3. Incorporate feedback from the training to help ensure the final disaster recovery procedures are sufficient.

Note: You may also choose to use the training as an opportunity to assign recovery coordinators, including one person as the recovery coordinator and a second person as an alternate coordinator. These people should be instructed to meet at a documented location to begin the disaster recovery plan.

Chapter 8: Transitioning From eTrust Admin to Identity Manager

This section contains the following topics:

Develop a Transition Plan

When planning to transition from eTrust Admin to CA Identity Manager, consider the following:

1. Become familiar with the features and functionality available in CA Identity Manager.
2. Identify the components and functionality that you use in your current eTrust Admin implementation. Common components and functionality include:
   - Management Interfaces
   - Feed Options and Web Services Interfaces
   - Advanced Workflow
   - Custom Endpoint Connectors
3. Compare the functionality in your existing eTrust Admin implementation with the functionality in CA Identity Manager.
4. Decide where to implement the functions that your business requires.

For assistance in making transition decisions, see the following sections:

- eTrust Admin Management Interfaces (see page 111)
- Custom Endpoint Connectors (see page 113)

Business Changes in Identity Manager

The Identity Manager Server, which users access through the Identity Manager User Console or through public tasks, typically manages a company's business tasks. These tasks may include user management, delegation, self service, approvals, business logic and password management.

The Identity Manager Server offers the following benefits:

- A single location for user management.
- Fine-grained privilege management—Roles, scope rules, and delegated administration allow administrators to share their workload while maintaining control over which objects a delegated administrator can manage.

  Note: For more information, see the *Administration Guide*.
- Customized user interfaces—You can customize the appearance and content of the CA Identity Manager task screens to suit your business needs. For example, you can create two sets of tasks—one for employees and one for customers. Each set of tasks can include different branding and can collect different information.

  Note: For more information, see the *Configuration Guide*.
- Support for Multiple Platforms—CA Identity Manager runs on WebLogic, WebSphere, and JBoss application servers running on Windows or Solaris systems.
- Support for multiple user store types—CA Identity Manager can manage users in multiple types of LDAP user directories and relational databases.

  For a complete list of supported user store types, see the CA Identity Manager support matrix on CA Support.

Terminology Changes

Existing eTrust Admin customers may notice certain terms have changed now that eTrust Admin is part of CA Identity Manager. The following table shows these changes:

eTrust Admin Term	New Term in Identity Manager
eTrust Admin Server	Provisioning Server
eTrust Admin Manager	Provisioning Manager
Directory	Endpoint, Endpoints
Namespace	Endpoint Type
Policy or Provisioning Policy	Account Template
Roles	Provisioning Roles
Distributed SuperAgent Framework	Connector Server Framework

eTrust Admin Term	New Term in Identity Manager
SuperAgent	C++ Connector Server
Option	Connector
Administrative Directory or Administrative Repository	Provisioning Directory
Identity Manager Corporate Directory	Identity Manager User Store
Corporate User	Inbound Administrator

eTrust Admin Management Interfaces

An eTrust Admin implementation may include user interfaces for user self-service, delegated administration, and endpoint and account management.

In a CA Identity Manager implementation, most of these functions are performed in the Identity Manager Server. For migration considerations, see the following sections:

- SAWI/DAWI Considerations (see page 111)
- Password Management Considerations (see page 111)
- IA Manager Considerations (see page 112)
- Provisioning Manager Considerations (see page 112)

SAWI/DAWI Considerations

In some eTrust Admin implementations, self service and delegated administration are provided through the Self Administration Web Interface (SAWI) and the Delegated Administration Web Interface (DAWI) respectively.

Identity Manager self-service tasks and delegated administration replace the functionality in the SAWI/DAWI components. You should not use the SAWI/DAWI components in an Identity Manager implementation.

Password Management Considerations

Like delegated administration and user self-service, passwords are managed in Identity Manager instead of eTrust Admin. For information on Identity Manager password management, see Password Management.

IA Manager Considerations

eTrust Admin included the IA Manager, which replaced SAWI/DAWI in some eTrust Admin implementations. The IA Manager allowed users to do the following:

- Self-register
- Manage their profiles, including changing passwords
- Reset forgotten passwords
- Delegate user administration tasks
- Access Advanced Workflow requests and approval worklists

CA Identity Manager enables users to perform self-service tasks, including password management, and delegated administration through the User Console.

Advanced Workflow is not supported in CA Identity Manager r12. However, CA Identity Manager enables you to perform all of the functions that you can perform using Advanced Workflow.

Provisioning Manager Considerations

In an Identity Manager implementation, most user management tasks are performed in the User Console. For example, you use the User Console to manage global users and groups, and to create and assign provisioning roles. Although you can still perform these functions in the Provisioning Manager (previously called the Admin Manager), the User Console provides a single location for performing business changes.

After migrating user management tasks to the Identity Manager User Console, the following provisioning tasks will still require the Provisioning Manager:

- Some synchronization tasks, such as synchronizing users with roles, and synchronizing accounts with account templates
- Management of the Entrust PKI, CA SSO, CA EEM, Novell Netware, Ingres and NSK Safeguard endpoint types
- Creation of direct associations between accounts and account templates
- Management of objects other than accounts, account Templates and endpoints
- Management of nested Provisioning Roles
- Program Exit configuration
- Provisioning Server administration

Batch Processing in Identity Manager Implementations

If you used the Universal Feed Option or the PeopleSoft Feed Option for batch loading users in eTrust Admin, consider the following methods, which allow you to batch load and modify user information without feed options:

- Bulk Loader—You can use the Bulk Loader tab in the User Console to upload feeder files that are used to manipulate large numbers of managed objects simultaneously. The advantage of the Bulk Loader method is that you can automate the process of manipulating a large number of managed objects using an information (feeder) file. The Bulk Loader task can also be mapped to a workflow process.

Note: CSV is the default file format for the feeder, but you can create a custom feed for other file formats.

- Service Provisioning Markup Language (SPML)—You can generate SPML requests that are submitted directly to the Provisioning Server without having to convert user data to a CSV or XML input file.
- TEWS—CA Identity Manager provides Task Execution Web Services (TEWS) which expose CA Identity Manager tasks as web services. TEWS requests are submitted to CA Identity Manager as if they were initiated in the Identity Manager User Console.

  CA Identity Manager provides documentation and samples that illustrate how to submit batch changes using TEWS. See the *Programming Guide for Java* for details.
- etautil Batch Utility—You can use this command line utility to submit batch changes to the Provisioning Server.

CA Identity Manager can apply identity policies or workflow processes to any batch feed transaction initiated using one of these methods.

Custom Endpoint Connectors

Existing eTrust Admin implementations may manage custom endpoints. To enable CA Identity Manager to read custom namespace policies and accounts that are associated with provisioning roles, there are additional procedures required.

Note: For instructions, see the appendix on custom endpoint options in the *Programming Guide for Provisioning*.

Deprecated Provisioning SDKs and Utilities

The following Provisioning Server SDKs and interfaces are deprecated in CA Identity Manager r12.5 SP1; however, they continue to function as documented.

To use the C++ Connector SDK and the JIAM SDK, download and install the CA Identity Manager Legacy Components package. It includes the *Programming Guide for Provisioning*, which describes these SDKs.

- **C++ Connector SDK**

  This SDK allows you to write custom static C++ Connectors. Existing C++ Connectors will continue to work with CA Identity Manager [assign the value for rn in your book].

  Note: New connectors should be developed using the Java Connector SDK, which is described in the *Programming Guide for Java Connector Server*.

- **Java Identity and Access Management (JIAM) SDK**

  The JIAM SDK provided the following functionality in previous versions of CA Identity Manager:

  - Java interface to the Provisioning Server
  - An abstraction of Provisioning Server functionality to develop custom client applications
  - A single interface to supply multiple clients with access to Identity and Access Management functionality

  This API is being deprecated because it only provides access to a subset of CA Identity Manager functionality.

This functionality is replaced by the following CA Identity Manager 12.5 functionality:

- Admin tasks in the User Console

  You can use admin tasks to manipulate most of the objects that Identity Manager manages.

- Task Execution Web Services (TEWS)

  The CA Identity Manager Task Execution Web Service (TEWS) is a web service interface that allows third-party client applications to submit remote tasks to CA Identity Manager for execution. This interface implements the open standards of WSDL and SOAP to provide remote access to CA Identity Manager.

- Managed Object interfaces

  CA Identity Manager provides interfaces for managed objects, which are accessible through CA Identity Manager APIs.

For more information about admin tasks, see the *Administration Guide*. For more information about TEWS and managed object interfaces, see the *Programming Guide for Java*.

- **etautil**

  You use the etautil batch utility to perform the same tasks as you do with the Provisioning Manager, but from a command line interface. It is described in the *Provisioning Reference Guide*.

  This functionality is replaced by the Task Execution Web Services (TEWS), which is described in the *Programming Guide for Java.*

- **Universal Provisioning Connector (UPC)**

  The UPC provides a mechanism for Identity Manager to invoke user-specified external programs when user provisioning requests are received. It uses program exits to send alerts regarding non-managed systems (non-managed mode) so that administrators can manually carry out the request and update the account request status. It also uses exits in a synchronous mode (managed mode) to provide a direct management interface to remote endpoint types.

Next Steps

After deciding which components and functionality to implement, refer to the following guides for installation and configuration instructions:

- *Installation Guide*
- *Administration Guide*

Index

I

J

L

M

N

O

P

R

S

T

U

W