

CA SiteMinder®

Implementation Guide

12.52 SP1



This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time. This Documentation is proprietary information of CA and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA.

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

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

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

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

The manufacturer of this Documentation is CA.

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

Copyright © 2014 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Technologies Product References

This document references the following CA Technologies products:

- CA RiskMinder^{<tm>} (formerly CA Arcot RiskFort)
- CA AuthMinder[®] (formerly CA Arcot WebFort)
- CA SiteMinder[®] Federation
- CA Directory
- CA DataMinder[®] (formerly CA DLP) Content Classification Service
- CA SiteMinder[®]

Contact CA Technologies

Contact CA Support

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

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

Providing Feedback About Product Documentation

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

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

Documentation Changes

- No updates have been made to the 12.52 SP1 documentation, as a result of issues that are found in previous releases.
- No updates have been made to the 12.52 documentation, as a result of issues that are found in previous releases.
- The following documentation updates were made to the latest version of the 12.51 documentation:
 - [Periodic Maintenance Tasks](#) (see page 188)—Updated guidance for running XPS Sweeper Utility (169270, 168658, 21175885:01).

Contents

Chapter 1: CA SiteMinder® Components and Stores 11

Components and Stores	11
Policy Server	12
CA SiteMinder® Agents	12
CA SiteMinder® Web Services Security Agents	13
CA Business Intelligence	13
Data Stores	14
CA SiteMinder® Administrative UI	18

Chapter 2: Architectural Considerations 19

Your Enterprise Environment	19
Operating Systems	19
Web Server Vendors	20
Application Server Vendors	21
Enterprise Resource Planning Systems	21
Directory Servers and Databases	21
Architectural Use Cases	22
Simple Deployment	22
Simple Deployment with Optional Components	24
Simple Deployment with Optional Agents	25
Multiple Components for Operational Continuity	26
Clustered Components for Scale	30
Redundancy and High Availability	32
Enhanced Session Assurance Architecture and Performance Considerations	46
Basic Architecture	47
Possible Architecture 1—Use Existing Components	48
Possible Architecture 2—Use Existing Policy Server	49
Possible Architecture 3—Full Separation of the Session Assurance Components	50

Chapter 3: Plan a CA SiteMinder® Implementation 53

Implementation Planning Overview	53
Policy Management Models	53
Policy Management Using Application Objects	54
Policy Management Using Policy Domains and Domain Objects	54
Identify the Applications to Secure	55
Group Resources into Domains or EPM Applications	56

Group Resources into Realms or EPM Components	57
Identify User Stores.....	59
Identify Authentication Methods.....	60
Identify Password Management Options.....	61
Password Policy Considerations.....	62
Identify Who Will Manage Your Web Agents.....	63
Central and Local Configurations Together.....	66
Identify Data Centers	67
Identify Resources to be Secured with Multiple Cookie Domains	68
Load-balancing for SSO between Cookie Provider Domains and Other Cookie Domains	69
Determine if Partnerships Require CA SiteMinder® Federation	70
Determine if Advanced Encryption Standards are Required.....	71
Determine if Virtualization is to be Used	72
Determine how to Manage Policy Servers.....	73
Local Policy Server Management	73
Central Policy Server Management.....	74
Determine how to Manage Web Agents.....	75

Chapter 4: Plan a CA SiteMinder® Web Services Security Implementation **77**

Policy Management Models.....	77
Policy Management Using Application Objects	78
Policy Management Using Policy Domains and Policies	78
Identify the Web Services to Secure	78
Identify User Stores.....	79
Identify Authentication Methods.....	80
Identify Who Will Manage Your SiteMinder WSS Agents	81
Identify Data Centers	83
Determine if Advanced Encryption Standards are Required.....	83
Determine if Virtualization is to be Used	85
Determine how to Manage Policy Servers.....	85
Determine how to Manage SiteMinder WSS Agents	86

Chapter 5: CA SiteMinder® Capacity Planning **87**

Capacity Planning Introduced	87
Use Case: Capacity Planning.....	88
How to Estimate a Sustained Authentication Rate	89
Estimate Daily Authentications	89
Estimate a Sustained Authentication Rate.....	91
Estimate a Peak Authentication Rate	93
How to Estimate a Sustained Authorization Rate	94
Estimate Daily Authorizations	95

Estimate a Sustained Authorization Rate	97
Estimate a Peak Authorization Rate	99

Chapter 6: CA SiteMinder® Web Services Security Capacity Planning 103

CA SiteMinder® Web Services Security Capacity Planning Introduced	103
Use Case: Capacity Planning	104
How to Estimate a Sustained Request Rate	104
Estimate Daily Requests	105
Estimate a Sustained Request Rate	106
Estimate a Peak Request Rate	108
Other Factors to Consider When Capacity Planning	110

Chapter 7: Configuration Considerations 111

Security Zones	111
Multiple Data Centers	113
Best Practices	113
Architectural Considerations	114
Multiple Data Center Use Cases	115
Authentication and a Centralized Login Server	121
Centralize Login Pages	122
Best Practices	123
Login Page Use Cases	124

Chapter 8: Performance Tuning 129

Performance Tuning Introduced	129
Performance Tuning Roadmap	130
Web Tier Performance	131
Server Performance	132
CA SiteMinder® Agent Performance	133
Reduce Traffic between Your Agents and the Policy Server	137
Improve Agent Performance through Load Balancing	144
Web Servers, Web Agents, and Web Server Processes	146
Application Tier Performance	149
CA SiteMinder® Policy Design and Performance	149
CA SiteMinder® Policy Objects and Performance Roadmap	149
Authentication Guidelines	154
Authorization Guidelines	158
Auditing and Performance	164
Load Balancing the Application Tier	164
Data Tier Performance	165

Data Tier Guidelines	165
User Store Capacity Planning	168
User Store Capacity Planning	181
Periodic Maintenance Tasks.....	188

Chapter 9: Diagnose Implementation Issues **191**

Diagnose Issues Introduced	191
Policy Server/Policy Store Connection Issues	192
Work with Support.....	193
Environment Information	193
Log Files.....	194
Policy Server Crash.....	195
Agent Crash	198
Resource Leaks.....	199
Functional Issues.....	200
Random Issues	201
Locate Knowledge Base Articles.....	202
Measure CA SiteMinder® Performance	202
Network Sniffers	203
CA SiteMinder® OneView Monitor.....	203
CA SiteMinder® Test Tool.....	204
Directory Server Utilities and SQL Analyzers.....	204

Chapter 10: Product Integrations **205**

CA Arcot WebFort and RiskFort Integration.....	205
Authentication in an On–Premise Arcot Integration	206
Confidence Levels and CA SiteMinder® Authorization	207
Risk Scores and Confidence Levels Compared.....	208
Enable Confidence Level Support for Authorization Decisions.....	210
CA Arcot Integration Use Cases.....	210
User Store Consideration	215
CA Arcot A-OK Integration.....	215
Authentication in a Hosted CA Arcot Integration	216
Confidence Levels and CA SiteMinder® Authorization	217
Risk Scores and Confidence Levels Compared	218
Enable Confidence Level Support	219
CA Arcot A-OK Integration Use Cases	220
User Store Consideration	222
CA DataMinder Content Classification Service Integration.....	223
CA DataMinder Content Classification Service	224
CA DataMinder Content Classification Service Preclassification Agent	225

CA SiteMinder® Policy Server.....	225
CA SiteMinder® Agent for SharePoint.....	225
CA SiteMinder® Session Store.....	226
CA DataMinder Content Classification Service Integration Roadmap	227
CA Identity Manager Roles and Access Control	237

Index	241
--------------	------------

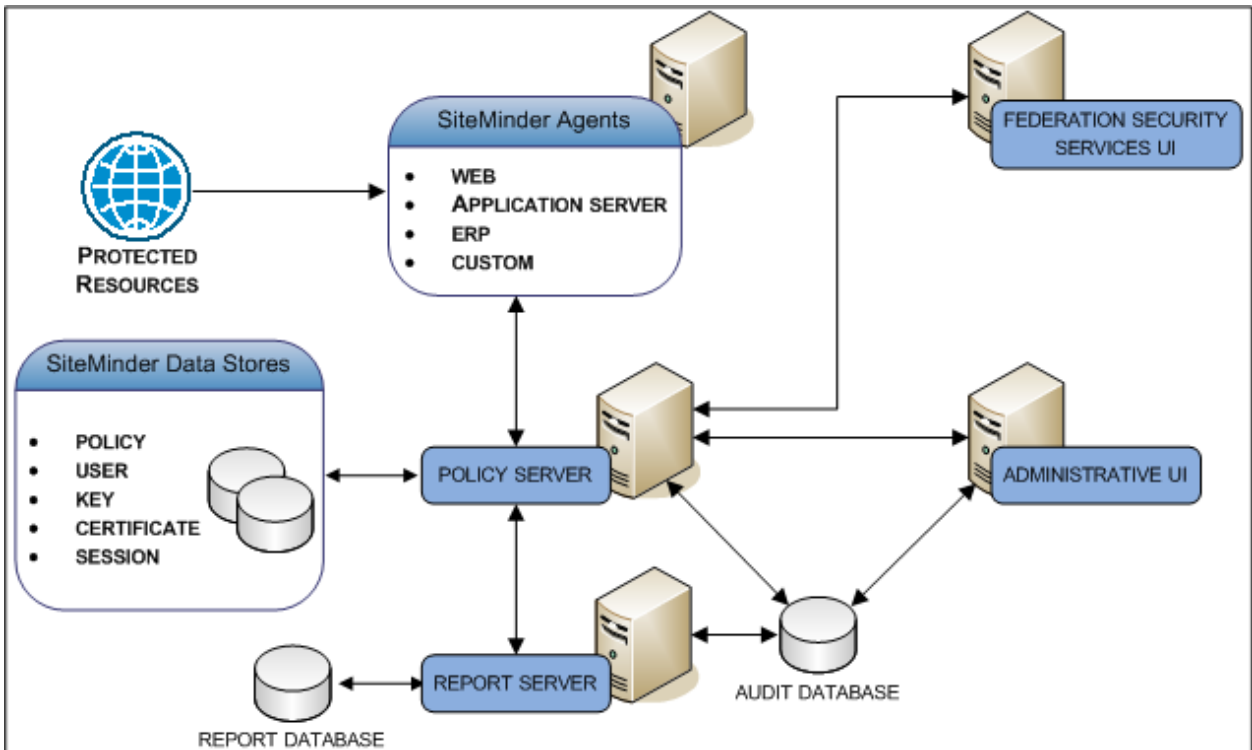
Chapter 1: CA SiteMinder® Components and Stores

Components and Stores

A CA SiteMinder® environment includes multiple components. Some components are required to secure resources, while others are optional, or only required to implement specific features. These components work with the resources, applications, directories, and databases in your organization to provide secure access to resources in your enterprise network.

All CA SiteMinder® components are supported on a number of operating environments. Your CA SiteMinder® implementation is highly dependent on the environment to which you are deploying it. Your implementation does not have to reflect the following diagram. Rather, the purpose of the following diagram is to illustrate the major components in a CA SiteMinder® environment and their general relationships with each other.

Figure 1: Product Components Overview



Use the previous diagram and the following component descriptions as a resource when considering the architectural questions detailed in this guide.

Policy Server

(Required) A CA SiteMinder® Policy Server (Policy Server) acts as the Policy Decision Point (PDP). The purpose of the Policy Server is to evaluate and enforce access control policies, which it communicates to a CA SiteMinder® Agent. A Policy Server provides the following:

- Policy-based user management
- Authentication services
- Authorization services
- Password services
- Session management
- Auditing services

The Policy Server interacts with all other major components to perform these tasks.

CA SiteMinder® Agents

(Required) A CA SiteMinder® Agent can reside on a web server, a J2EE application server, an Enterprise Resource Planning (ERP) system, or custom application. An Agent acts as the Policy Enforcement Point (PEP), intercepting user requests for resources and communicating with a Policy Server to determine if the resource is protected.

If the resource is not protected, the Agent allows access. If the resource is protected, the Agent continues to communicate with the Policy Server to authenticate and authorize users. A successful authorization prompts the Agent to let the resource request proceed to the server. Agents also:

- Provide information to web applications to enable content personalization
- Cache information about authenticated users and protected resources to allow quicker access to resources
- Enable single sign-on (SSO)

CA SiteMinder® Web Services Security Agents

(Required for CA SiteMinder® Web Services Security) CA SiteMinder® Web Services Security (WSS) Agents act as Policy Enforcement Points (PEPs) that work with the following platforms:

- Web server
- J2EE application server
- Custom application

WSS Agents intercept requests for "big" (SOAP-based) web services. The WSS Agents then communicate with a Policy Server to determine whether the resource is protected.

Note: The CA SiteMinder® Agent for JBoss includes CA SiteMinder® *and* WSS agent functionality. .

If the resource is not protected, the agent allows access. If the resource is protected, the agent continues to communicate with the Policy Server to authenticate and authorize users. A successful authorization prompts the agent to let the resource request proceed to the server.

Agents also perform the following other functions:

- Cache information about authenticated users and protected resources to allow quicker access to resources
- Enable single sign-on (SSO)

CA Business Intelligence

(Optional) CA Business Intelligence is a set of reporting and analytic software that various CA products use for the purposes of presenting information and supporting business decisions. CA products use CA Business Intelligence to integrate, analyze, and then present, through various reporting options, information required for effective enterprise IT management.

Included in CA Business Intelligence is SAP BusinessObjects Enterprise, a complete suite of information management, reporting, and query and analysis tools. CA Business Intelligence installs SAP BusinessObjects Enterprise as a stand-alone component. In this guide, this stand-alone component is referred to as the Report Server. Installing the Report Server is a separate step within the overall CA SiteMinder® installation process. Installing the Report Server separately from CA SiteMinder®-specific components lets other CA products share the same Business Intelligence Services.

The Report Server compiles reports to help you analyze your CA SiteMinder® environment. The purpose of this component is to create the following types of reports:

- Audit
- Policy analysis

The Report Server communicates with the following components to compile reports:

- The Central Management Server (CMS) database (report database)
- An Administrative UI
- A Policy Server
- A CA SiteMinder® audit database

Data Stores

A CA SiteMinder® implementation contains multiple data stores. Some stores are required, while others are optional, or only required to implement specific features.

The following descriptions detail:

- If the store is required or optional
- The purpose of the store

Policy Store

(Required) The CA SiteMinder® policy store (policy store) is an entitlement store that resides in an LDAP directory server or ODBC database. The purpose of this component is to store all policy-related objects, including the:

- Resources CA SiteMinder® is protecting
- Methods used to protect those resources
- Users or groups that can or cannot access those resources
- Actions that must take place when users are granted or denied access to protected resources

The Policy Server uses this information, collectively known as an Enterprise Policy Management (EPM) application or CA SiteMinder® policy, to determine if a resource is protected and if an authenticated user is authorized to access the requested resources.

User Store

(Required) A CA SiteMinder® user store connection (user store connection) is a connection to an existing user directory or database in your enterprise network. You are not required to use a proprietary CA SiteMinder® user store. The purpose of the user store connection is to make user data available to the Policy Server, which includes the following:

- Organizational information
- User and group attributes
- User credentials, such as passwords
- User attributes, such as first and last name

The Policy Server uses these connections to:

- Verify user credentials when an Agent submits a request for a protected resource
- Retrieve user attributes for the CA SiteMinder® features that require specific user data

Note: For more information about configuring a user store connection, see the documentation roadmap.

External Administrative User Store

(Optional) By default, the Administrative UI uses the policy store as its source for CA SiteMinder® administrator credentials. This default configuration lets you manage the environment immediately after configuring a policy store and installing the Administrative UI. When you configure a policy store, the default CA SiteMinder® super user account (siteminder) is created. This account has maximum system privileges, and is used to access the Administrative UI for the first-time and to create additional CA SiteMinder® administrators.

You can configure the Administrative UI to use an external administrator user store, for example, a corporate directory. An external administrative user store is a connection to an LDAP directory server or ODBC database in your enterprise network. Consider the following:

- An Administrative UI can only connect to a single external administrative user store.
- An Administrative UI can be configured to managed multiple Policy Servers. If an Administrative UI is to manage multiple Policy Servers, a connection to an external administrator user store is required.
- If you configure more than one Administrative UI for high-availability, the same external administrative user store makes all administrators available to each Administrative UI.

Note: For more information about CA SiteMinder® administrators and configuring an external administrative user store, see the documentation roadmap.

Key Store

(Required) The purpose of this component is to store the encryption keys that the Policy Servers and the agents use to encrypt sensitive data, which include:

- The keys that the agents use to encrypt CA SiteMinder® cookies.
- The keys that the Policy Servers use to encrypt sensitive policy store information, such as administrator passwords.
- The keys the Policy Servers use to encrypt CA SiteMinder® session tickets that contain credentials and other information that is related to user sessions.

You can collocate the key store with the policy store or you can store encryption keys in a separate directory or database. The need to deploy a separate key store depends on:

- How you implement Policy Servers and policy stores.
- Single sign-on requirements.

Note: If you use the Policy Server Configuration wizard to configure a policy store, the key store is automatically collocated with the policy store.

Certificate Data Store

(Optional) The CA SiteMinder® certificate data store (CDS) makes the following components and functions available to a CA SiteMinder® environment:

- Certificate authority (CA) certificates
- Public and private keys
- Certificate revocation lists (CRL)
- OCSP revocation checks

Note: CA SiteMinder® federation features use the certificate data store. The user certificates that the X.509 certificate authentication scheme uses for authentication are not stored in the certificate data store. These user certificates are stored in an LDAP/AD user directory or ODBC store.

By default, the certificate data store is automatically configured and colocated with the policy store. As a result:

- A separate external store is not required.
- All Policy Servers sharing a common view into the same policy store have access to the same keys, certificates, and certificate revocation lists.
- All CA SiteMinder® administrators that manage the same policy store can manage the certificate data store centrally using the Administrative UI.

CA SiteMinder® Audit Database

(Optional) By default, the Policy Server writes audit events to a text file, which is known as the Policy Server log. The purpose of audit logs is to track information about all user activity, including:

- All successful authentications
- All failed authentications
- All successful authorization attempts
- All failed authorization attempts
- All administrative login attempts
- All administrative actions, such as changes to administrator passwords, the creation of policy store objects, and changes to policy store objects

However, you can configure a stand-alone CA SiteMinder® audit database (audit database). When deciding where to store audit events, consider that:

- The Report Server requires a connection to an audit database to create audit-based reports. The Report Server cannot create audit-based reports from a Policy Server log written to a text file.
- Storing audit logs to a database is more secure than logging the information to a text file.
- If supported, a policy store can also function as an audit database.

Note: For more information about configuring an audit database, see the documentation roadmap.

Session Store

(Optional) When CA SiteMinder® authenticates a user, the Policy Server issues a session ticket. A session ticket contains basic information about the user and authentication information for the user. By default, CA SiteMinder® implements session management through non-persistent sessions. If non-persistent sessions are enabled, an Agent writes the session ticket to a cookie on the browser of the users. However, some CA SiteMinder® features require persistent sessions.

If persistent sessions are enabled, an Agent must write the session ticket to a stand-alone database.

You deploy a CA SiteMinder® session store (session store) for the following primary reasons:

- If a CA SiteMinder® log off URI is implemented, a session store prevents a CA SiteMinder® session from being used again after a user logs off.
- To provide support for features that require persistent user sessions.

Agents use this information to identify users and provide session information to the Policy Server.

Note: For more information about configuring a session store, see the documentation roadmap.

CA SiteMinder® Administrative UI

(Required) The CA SiteMinder® Administrative UI (Administrative UI) is a web-based administration console that is installed independent of the Policy Server. The Administrative UI is intended for managing all tasks that are related to access control, reporting, and policy analysis.

Chapter 2: Architectural Considerations

This section contains the following topics:

[Your Enterprise Environment](#) (see page 19)

[Architectural Use Cases](#) (see page 22)

[Enhanced Session Assurance Architecture and Performance Considerations](#) (see page 46)

Your Enterprise Environment

CA SiteMinder® implementations are highly dependent on the environments in which you deploy them. We recommend that you develop a plan that breaks your implementation into steps that make sense for your enterprise. As you plan the deployment, there are many questions for consideration.

The answers to these questions are critical to planning your CA SiteMinder® implementation.

Operating Systems

CA SiteMinder® components are supported across multiple platforms, the details of which are located in the CA SiteMinder® Platform Support Matrix. Which of the following operating systems has your enterprise deployed?

- Microsoft® Windows®
- Oracle® Solaris™
- Red Hat® Enterprise Linux®
- Novell® SUSE® Linux
- Hewlett–Packard Company UNIX (HP–UX)
- IBM® AIX®
- IBM z/OS®

Note: For the specific versions of supported operating systems, see the CA SiteMinder® Platform Support Matrix.

Use the following table to help determine if the operating systems your enterprise has deployed are currently supported for the required CA SiteMinder® components.

Component	Required?	Operating Systems
Policy Server	Yes	
Agent	Yes	
Administrative UI	Yes	
Report Server	No	

Note: Additional non–platform requirements exist for each of these components, for example, minimum memory requirements. For more information about non–platform requirements for the Policy Server, the Administrative UI, and the Report Server, see the *Policy Server Installation Guide*. For more information about non–platform requirements for an Agent, see the specific CA SiteMinder® Agent documentation.

Web Server Vendors

You can install and configure CA SiteMinder® Agents to protect resources on web servers. Which of the following supported server vendors has your enterprise deployed?

- Apache™ HTTP Server
- Apache Tomcat
- Hewlett–Packard Company (HP) Apache
- IBM HTTP Server
- IBM Lotus® Domino
- Microsoft IIS
- Oracle® HTTP Server
- Red Hat Apache
- Sun Java™ System

Note: For the specific versions of the supported web servers, refer to the CA SiteMinder® Platform Support Matrix.

If you use other web servers not listed here, consider using CA SiteMinder® SPS to protect the resources on these web servers.

Application Server Vendors

You can install and configure CA SiteMinder® agents to protect resources on J2EE application servers. Which of the following supported server vendors has your enterprise deployed?

- Oracle WebLogic®
- IBM WebSphere®
- RedHat JBoss®

Note: For the specific versions of supported application servers, refer to the CA SiteMinder® Platform Support Matrix.

Enterprise Resource Planning Systems

You can install and configure Agents to protect resources on ERP systems. Which of the following supported ERP vendors has your enterprise deployed?

- Oracle PeopleSoft®
- Oracle Siebel®
- SAP®

Note: For the specific versions of supported ERP systems, see the CA SiteMinder® Platform Support Matrix.

Directory Servers and Databases

CA SiteMinder® data stores are supported across multiple directory servers and databases. Which supported vendors has your enterprise deployed?

Note: For more information, see the Platform Support Matrix.

Use the following table to help determine if the directory server and database types your enterprise has deployed are currently supported for the components your implementation requires.

Component	Required	LDAP?	Database?
Policy store	Yes		
User store connection	Yes		
Administrative user store	No		
Audit database	No	N/A	

Component	Required	LDAP?	Database?
Key store	No		
Session store	No	N/A	

Architectural Use Cases

The purpose of the following use cases is to get you thinking about your CA SiteMinder® architecture in terms of high availability and performance. The use cases begin with a simple deployment and progress into more complex scenarios. Each case is based on the idea of a logical "block" of CA SiteMinder® components and illustrates how an environment can contain multiple blocks to address the following architectural considerations:

- Redundancy
- Failover
- Capacity and scale
- Multiple cookie domains

Extrapolate the necessary infrastructure from these cases to:

- Determine how to implement redundancy and high availability between CA SiteMinder® components
- Determine how to implement multiple data centers
- Support the usage metrics you gather from capacity planning
- Support your implementation considerations
- Begin the iterative process of performance tuning the environment

More information:

[Capacity Planning Introduced](#) (see page 87)

[Performance Tuning Introduced](#) (see page 129)

[Redundancy and High Availability](#) (see page 32)

Simple Deployment


The simplest CA SiteMinder® deployment requires one "block" of components. A block of components is a logical combination of dependent components that include:

- A Web Agent
- A Policy Server

- A user store
- A policy store
- An Administrative UI

You protect web-based resources by deploying at least one block.

The following diagram illustrates a simple deployment:



Each component has a specific role with resource protection.

Note: For more information about the primary purpose of each component, see CA SiteMinder® Components.

Simple Deployment with Optional Components

You can extend the functionality of a simple deployment through the use of optional CA SiteMinder® components. The decision to implement optional components is determined by the CA SiteMinder® features your enterprise requires. For example:

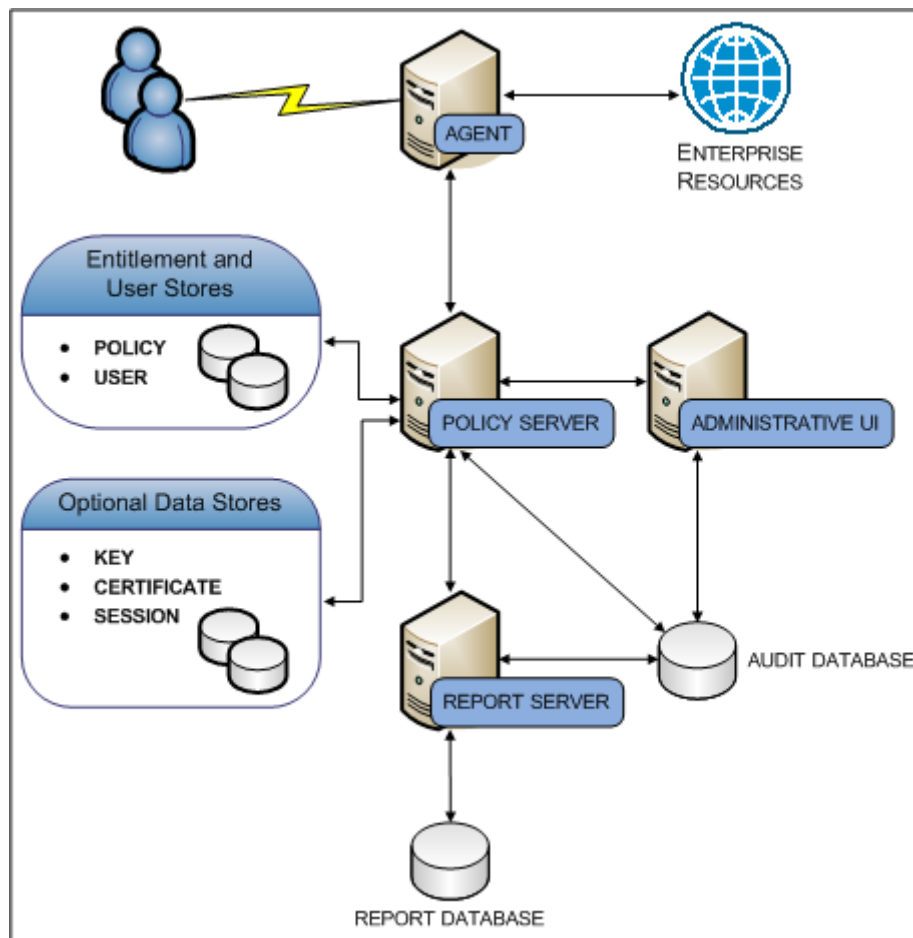
- If you are planning to implement Federation–based functionality, your environment requires a certificate data store and a session store.
- If you are planning to create audit-based reports, your environment will require a Report Server and an audit database.

The following diagram illustrates the optional components and their required dependencies:

- A Report Server
- A report database
- An audit database
- A key store

- A session store
- A certificate data store

Figure 2: Simple deployment with optional components



Each component has a specific role in resource protection.

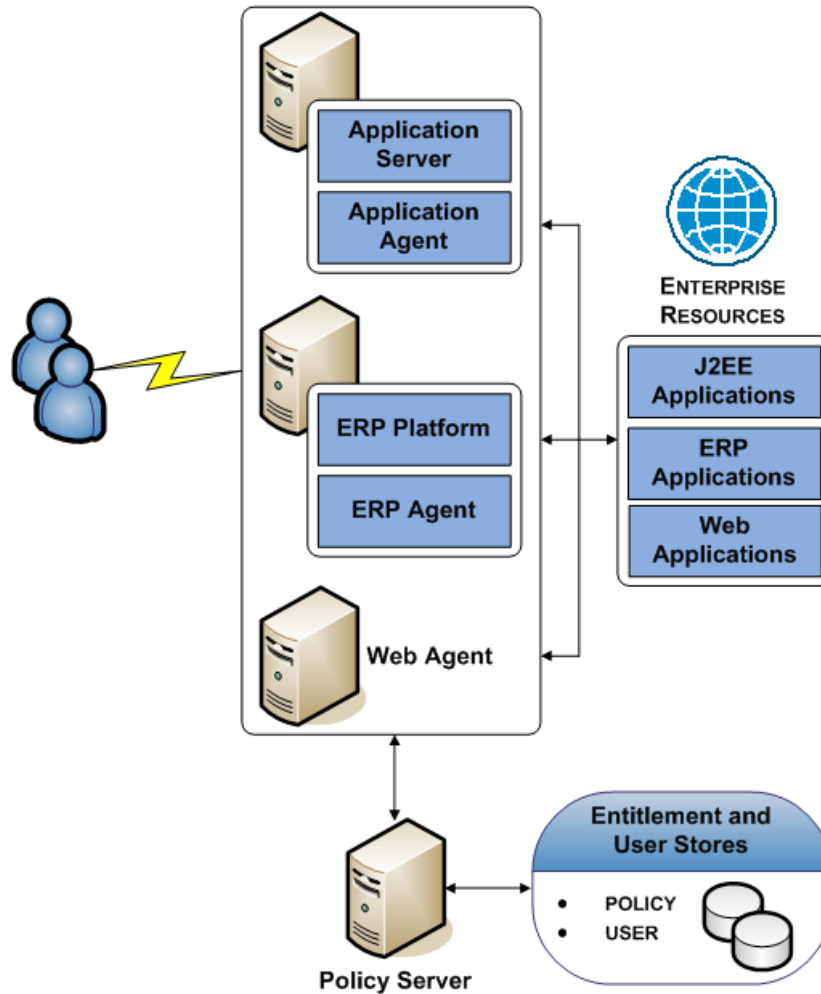
Note: For more information about the primary purpose of each component, see CA SiteMinder® Components.

Simple Deployment with Optional Agents

You can extend the functionality of a simple deployment your environment to protect resources that do not reside on a Web Server. For example, if your environment hosts resources on an:

- Application server, you can implement Application Server Agents to protect them.
- ERP system, you can implement ERP Agents to protect them.

The following diagram illustrates optional Agents:



Each component has a specific role with resource protection.

Note: For more information about primary purpose of each component, see CA SiteMinder® Components.

Multiple Components for Operational Continuity

The following use cases show how you can implement multiple blocks of components to build redundancy and failover into the environment using the following methods:


- SiteMinder round robin load balancing
- A hardware load balancer

Multiple Components for Operational Continuity Using SiteMinder Load Balancing

You can implement multiple blocks of components to build redundancy and failover into the environment using CA SiteMinder® round robin load balancing. This use case builds on a simple deployment to explain how you can begin thinking about operational continuity. The following diagram illustrates:

- Multiple Agent instances intercepting user requests. As illustrated, each Agent is configured to initialize and communicate with a primary Policy Server and failover to the second Policy Server.
- A Policy Server cluster evaluating and enforcing access control policies. Load is dynamically distributed between each Policy Server in the cluster.
- Multiple user store connections. Each Policy Server is configured to communicate with a primary user store. The primary user store connection is configured with a secondary user store connection. The Policy Servers load balance requests for user information across both connections. If the primary connection becomes unavailable, Policy Servers failover to the secondary connection.

- A single policy store instance. Each Policy Server connects to the same policy store for a common view of policy information. The primary policy store connection is configured with a secondary connection to which the Policy Servers can failover.



Each component has a specific role with resource protection.

Note: For more information about the primary purpose of each component, see CA SiteMinder® Components. For more information about CA SiteMinder® redundancy and high availability, see Redundancy and High Availability.

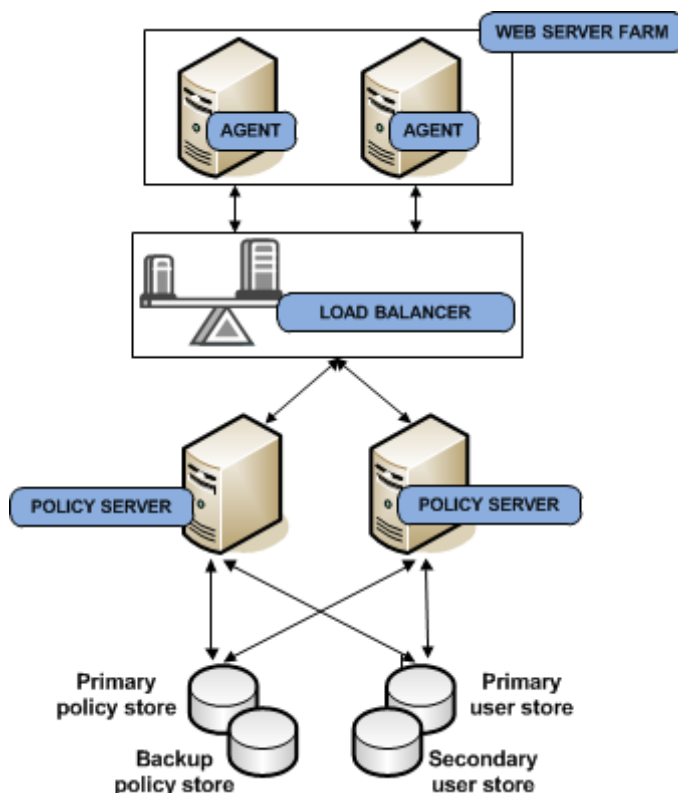
More information:

[Redundancy and High Availability](#) (see page 32)

Multiple Components for Operational Continuity Using Hardware Load Balancing

You can implement multiple blocks of components to build redundancy and failover into the environment using hardware load balancing. This use case builds on a simple deployment to explain how you can begin thinking about operational continuity. The following diagram illustrates:

- Multiple Agent instances intercepting user requests. As illustrated, each Agent is configured to initialize and communicate with a primary Policy Server and failover to the second Policy Server.
- A hardware load balancer configured to expose multiple Policy Servers through a virtual IP address (VIP). The hardware load balancer dynamically distributes load between all Policy Servers associated with that VIP.
- Multiple Policy Servers evaluating and enforcing access control policies.
- Multiple user store connections. Each Policy Server is configured to communicate with a primary user store. The primary user store connection is configured with a secondary user store connection. The Policy Servers load balance requests for user information across both connections. If the primary connection becomes unavailable, Policy Servers failover to the secondary connection.
- A single policy store instance. Each Policy Server connects to the same policy store for a common view of policy information. The primary policy store connection is configured with a secondary connection to which the Policy Servers can failover.



Each component has a specific role with resource protection.

Note: For more information about the primary purpose of each component, see CA SiteMinder® Components. For more information about CA SiteMinder® redundancy and high availability, see Redundancy and High Availability.

More information:

[Redundancy and High Availability](#) (see page 32)

Clustered Components for Scale

You can implement additional clusters to help performance levels remain high as you scale to extend throughput. This use case builds on the multiple components for operational continuity use case to explain how you can begin thinking about your architecture in terms of scale.

The initial deployment section of the diagram illustrates:

- A load balancer distributing user requests across multiple Agent clusters.
- Multiple Agent instances intercepting user requests for specific applications. Agents are configured to initialize and communicate with a primary Policy Server in the cluster. If enough Policy Servers in the cluster become unavailable, the Agents failover to another Policy Server cluster.

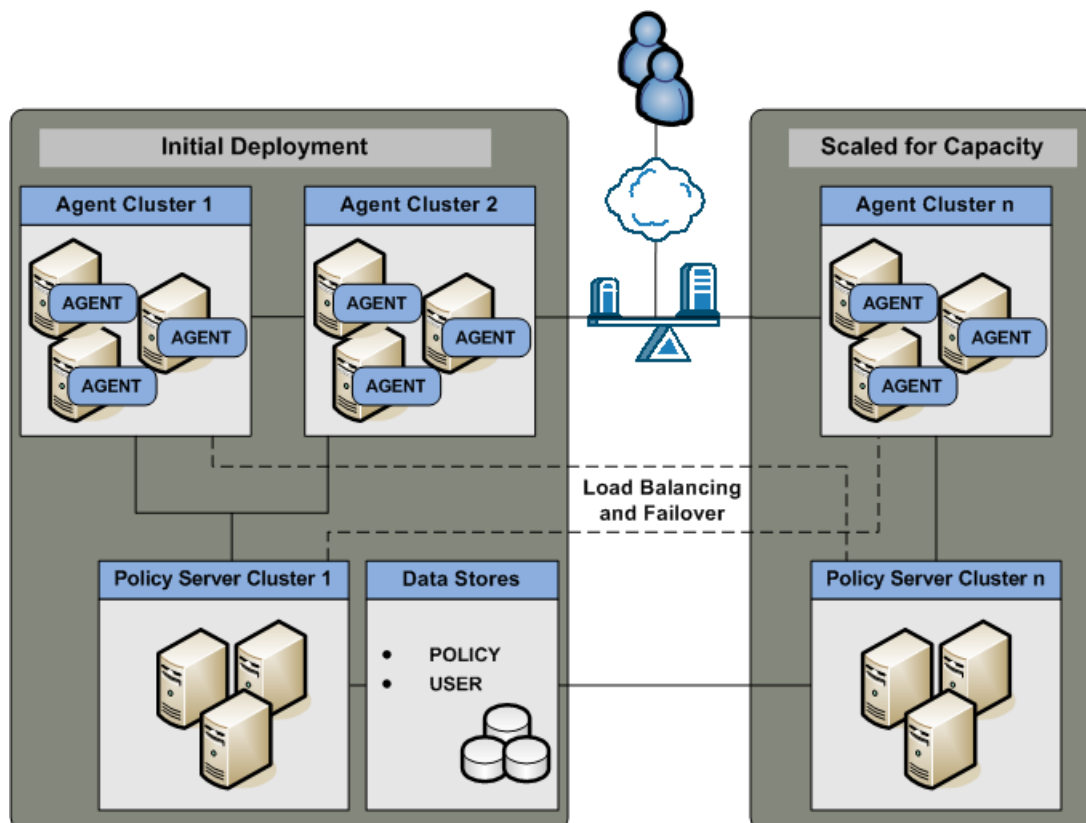
Note: For more information about Agent and Policy Server redundancy and high availability, see Redundancy and High Availability.

- A Policy Server cluster evaluating and enforcing access control policies. Load is dynamically distributed between each Policy Server in the cluster.
- Multiple user store connections. Each Policy Server is configured to connect to a primary user store. The primary user store connection is configured with a secondary user store connection. The Policy Servers load balance requests for user information across both connections. If the primary connection becomes unavailable, Policy Servers failover to the secondary connection.

Note: For more information about Policy Server and user store redundancy and high availability, see Redundancy and High Availability.

- A single policy store instance. Each Policy Server in the cluster connects to the same policy store for a common view of policy information. The primary policy store connection is configured with a secondary connection to which the Policy Servers can failover.

Note: For more information about Policy Server and policy store redundancy, see Redundancy and High Availability.



Each component has a specific role with resource protection.

Note: For more information about the primary purpose of each component, see CA SiteMinder® Components.

The Scaled for Capacity section of the diagram details another component block and illustrates:

- A load balancer distributing requests to the new Agent cluster.
- Multiple Agent instances intercepting user requests. In addition to their connections to Policy Servers in Cluster n, each Agent can also be configured to failover to any Policy Server in the environment. As illustrated by the dotted line, the Agents in Agent Cluster n are configured to failover to the Policy Servers in Policy Server Cluster 1.

- A Policy Server cluster evaluating and enforcing access control policies. As illustrated by the dotted line, each Policy Server cluster is configured with a failover threshold. When the number of available Policy Servers falls below the specified threshold, all requests that the failed Policy Server would otherwise service are forwarded to another cluster.

Note: For more information about failover thresholds for Policy Server cluster failover thresholds, see the *Policy Server Administration Guide*.

More information:

[Redundancy and High Availability](#) (see page 32)

[Multiple Components for Operational Continuity Using SiteMinder Load Balancing](#) (see page 27)

[Multiple Components for Operational Continuity Using Hardware Load Balancing](#) (see page 29)

Redundancy and High Availability

You configure redundancy and high availability between logical blocks of CA SiteMinder® components to maintain system availability and performance.

Agent to Policy Server Communication

When you configure a CA SiteMinder® Agent, a Host configuration file (named `SmHost.conf` by default), is created on the host server. The Agent uses the connection information in this Host configuration file to create an initial connection with a Policy Server.

After the initial connection is established, the Agent obtains subsequent Policy Server connection information from the Host Configuration Object (HCO) on the Policy Server.

You can configure the HCO to include multiple Policy Servers and specify the method the Agent uses to distribute requests among multiple Policy Servers.

A CA SiteMinder® Agent can distribute requests among multiple Policy Servers in the following ways:

- Failover
- Round-robin load balancing
- Round-robin load balancing over one or more clusters of Policy Servers

Alternatively, you can configure the HCO to include a single virtual IP address configured on a hardware load balancer to expose multiple Policy Servers. In this case, the load balancer is responsible for failover and load balancing, rather than the Agent software.

Failover

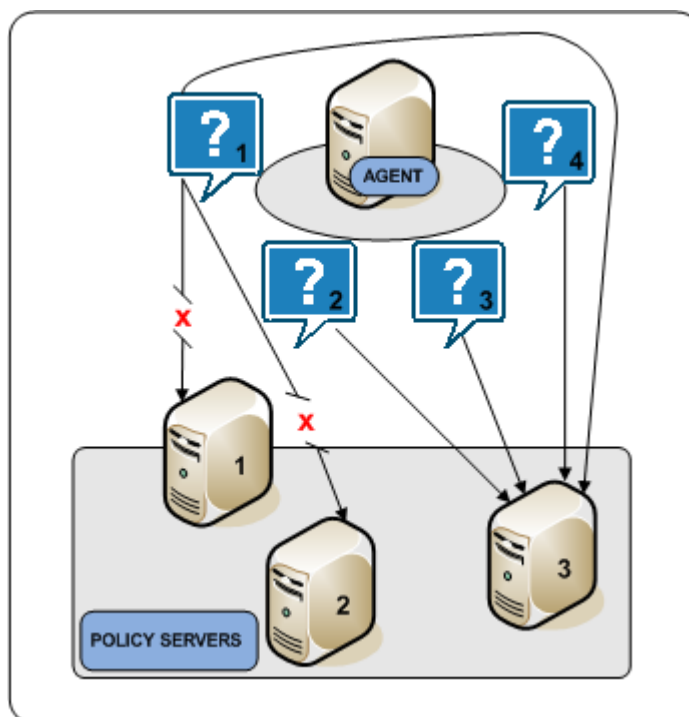
Failover is the default HCO setting. In failover mode, a CA SiteMinder® Agent delivers all requests to the first Policy Server that the HCO lists and proceeds as follows:

1. If the first Policy Server does not respond, the Agent deems it unavailable and redirects the request, and all subsequent requests, to the next Policy Server that the HCO lists.
2. If the first two Policy Servers do not respond, the Agent deems both of them unavailable and redirects the request, and all subsequent requests, to the next Policy Server that the HCO lists.

Note: For more information about configuring an HCO with multiple Policy Servers, see the *Policy Server Configuration Guide*.

If an unresponsive Policy Server recovers, which the Agent determines through periodic polling, the Policy Server is automatically returned to its original place in the HCO list and begins receiving all Agent requests.

The following diagram illustrates the Agent failover process:



Round Robin Load Balancing

Round robin load balancing is an optional HCO setting. Round robin load balancing distributes requests evenly over a set of Policy Servers, which:

- Results in more efficient user authentication and authorization
- Prevents a single Policy Server from becoming overloaded with Agent requests


Note: For more information about configuring an HCO for round robin load balancing, see the *Policy Server Configuration Guide*.

In round robin mode, an Agent distributes requests across all Policy Servers that the HCO lists. An Agent:

1. Sends a request to the first Policy Server that the HCO lists.
2. Sends a request to the second Policy Server that the HCO lists.
3. Sends a request to the third Policy Server that the HCO lists.
4. Continues sending requests in this way, until the Agent has sent requests to all available Policy Servers. After sending requests to all available Policy Servers, the Agent returns to the first Policy Server and the cycle begins again.

If a Policy Server does not respond, the Agent redirects the request to the next Policy Server that the HCO lists. If the unresponsive Policy Server recovers, which the Agent determines through periodic polling, the Policy Server is automatically restored to its original place in the HCO list.

The following diagram illustrates the round robin process:



Policy Server Clusters

Round robin load balancing evenly distributes CA SiteMinder® Agent requests to all Policy Servers that the HCO lists. Although an efficient method to improve system availability and response times, consider that:

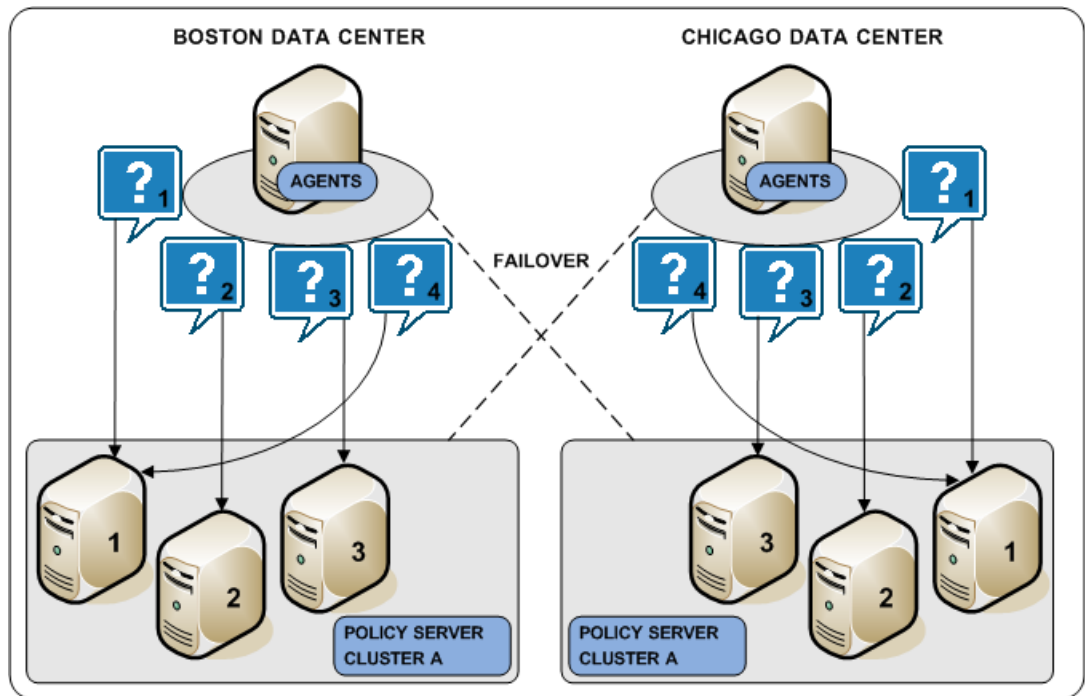
- Round robin load balancing is not the most efficient distribution method in a heterogeneous environment where computing capacity can differ. Each Policy Server receives the same number of requests, regardless of capacity.
- Round robin load balancing to Policy Servers that are located in different geographical locations can degrade performance. Sending Agent requests to Policy Servers outside a certain locale can result in increased network communication overhead and network congestion.

A Policy Server cluster is a group of Policy Servers to which Agents can distribute requests. Policy Server clusters provide the following benefits over round robin load balancing:

- A cluster can be configured to include Policy Servers only in a specific data center. Grouping Agents with distinct Policy Server clusters avoids the network overhead involved with load balancing across geographically separate regions. Network overhead is only incurred if Agents failover to another Policy Server cluster.
- A cluster can failover to another cluster based on a Policy Server failover threshold.
- Agents dynamically distribute requests across all Policy Servers based on response time, instead of distributing requests evenly.

Note: For more information about configuring a Policy Server cluster, see the *Policy Server Administration Guide*.

The following diagram illustrates two Policy Server clusters. Each cluster is geographical separated to avoid the network overhead that can be associated with round robin load balancing.




Hardware Load Balancing

CA SiteMinder® supports the use of hardware load balancers configured to expose multiple Policy Servers through one or more virtual IP addresses (VIPs). The hardware load balancer then dynamically distributes request load between all Policy Servers associated with that VIP. The following hardware load balancing configurations are supported:


- Single VIP with multiple Policy Servers exposed by each VIP
- Multiple VIPs with multiple Policy Servers exposed by each VIP

Single VIP, Multiple Policy Servers Per VIP



In the configuration shown in the previous diagram, the load balancer exposes multiple Policy Servers using a single VIP. This scenario presents a single point of failure if the load balancer handling the VIP fails.

Multiple VIPs, Multiple Policy Servers Per VIP



In the configuration shown in the previous diagram, groups of Policy Servers are exposed as separate VIPs by one or more load balancers. If multiple load balancers are used, this amounts to failover between load balancers, thus eliminating a single point of failure. However, all major hardware load balancer vendors handle failover between multiple similar load balancers internally such that only a single VIP is required. If you are using redundant load balancers from the same vendor, you can therefore configure Agent to Policy Server communication with a single VIP and still have robust load balancing and failover.

Note: If you are using a hardware load balancer to expose Policy Servers as multiple virtual IP addresses (VIPs), we recommend that you configure those VIPs in a failover configuration. Round robin load balancing is redundant as the hardware load balancer performs the same function more efficiently.

Policy Server to User Store Communication

The Policy Server can distribute queries over multiple LDAP or ODBC user stores to enable the following:

- Failover
- Round robin load balancing

Note: For more information about configuring user store connections, see the *Policy Server Configuration Guide*.

More information:

[User Store](#) (see page 15)

Failover


Failover is an optional setting in the CA SiteMinder® user store object. In failover mode, a Policy Server distributes all requests to the primary user store and proceeds as follows:

1. If the primary user store does not respond, the Policy Server deems it unavailable and redirects the request, and all subsequent requests, to the next user store that the CA SiteMinder® user store object lists.
2. If the first and second user stores do not respond, the Policy Server deems them both unavailable and redirects the request, and all subsequent requests to the next user store that the CA SiteMinder® user store object lists.

Note: For more information about configuring user store failover, see the *Policy Server Configuration Guide*.

If an unresponsive user store recovers, the user store is automatically returned to its original place in the failover list and begins receiving all Policy Server requests.

The following diagram illustrates the user store failover process:



Round Robin Load Balancing

Round robin load balancing is an optional CA SiteMinder® user store object setting. Round robin load balancing distributes requests evenly over a set of user stores, which:

- Results in more efficient user store queries
- Prevents a single user store from becoming overloaded with Policy Server requests

Note: Consider the following:


- For more information about configuring load balancing between LDAP user stores, see the *Policy Server Configuration Guide*.
- The Administrative UI does not include settings to configure round robin load balancing between ODBC user stores. However, the Policy Server installation includes:
 - The CA SiteMinder® Oracle Wire Protocol. This protocol supports load balancing over multiple Oracle stores. You can configure Oracle user store load balancing at the data source level.
 - The CA SiteMinder® SQL Server Wire Protocol, which you can use to configure SQL Server or SQL Server Cluster Enterprise. You can configure SQL Server user store load balancing at the database level.

In round robin mode, a Policy Server distributes requests across all user stores that the CA SiteMinder® user store object lists. A Policy Server:

1. Sends a request to the first user store that the user store object lists.
2. Sends a request to the second user store that the user store object lists.
3. Sends a request to the third user store that the user store object lists.
4. Continues sending requests in this way, until the Policy Server has sent requests to all available user stores. After sending requests to all available user stores, the Policy Server returns to the first user store and the cycle begins again.

Note: Configure load balancing with failover to add the benefit of redundancy in the event of a user store failure. For more information about configuring load balancing and failover, see the *Policy Server Configuration Guide*.

The following diagram illustrates the user store round robin process:



Policy Server to Policy Store Communication

All Policy Servers must connect to the same policy store for a common view of policy information. However, we recommend that the deployment includes multiple "hot" policy stores to which Policy Servers can failover.

The following are policy store failover scenarios:

- A master policy store configured with replicated versions
- Multi-mastered policy stores

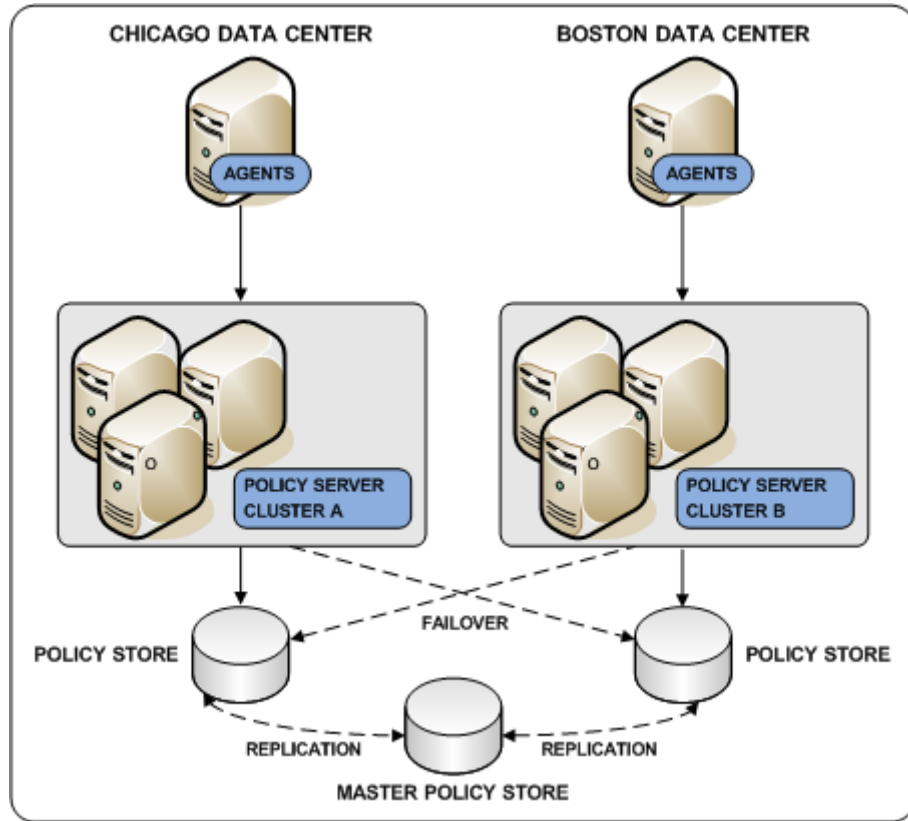
Master Policy Store

Deploying a master policy store with replicated versions is a way to achieve policy store redundancy. A single master policy store lets each Policy Server communicate with the closest replicated version. This method of communication:

- Improves the performance of geographically separated Policy Servers. Sending Policy Server requests to policy stores outside a certain locale can result in increased network communication overhead and network congestion.
- Allows for failover. If a primary policy store fails, Policy Servers failover to a secondary store.

Note: For more information about configuring replication, see your vendor-specific documentation. For more information about configuring policy store failover, see the *Policy Server Administration Guide*.

The following diagram illustrates a single master policy store environment:



Multi-Mastered Policy Stores

Deploying LDAP directories using multi-master technology is a way to achieve policy store redundancy. A multi-master policy store lets each Policy Server communicate with the closest replicated version. This method of communication:

- Improves the performance of geographically separated Policy Servers. Sending Policy Server requests to policy stores outside a certain locale can result in increased network communication overhead and network congestion.
- Allows for failover. If a primary policy store fails, Policy Servers failover to a secondary store.

The following configuration is recommended when configuring an LDAP policy store in multi-master mode:

- A single master should be used for all administration.
- A single master should be used for key storage.

This master does not need to be the same as the master used for Administration. However, we recommend that you use the same master store for both keys and administration. In this configuration, all key store nodes should point to the master rather than a replica.

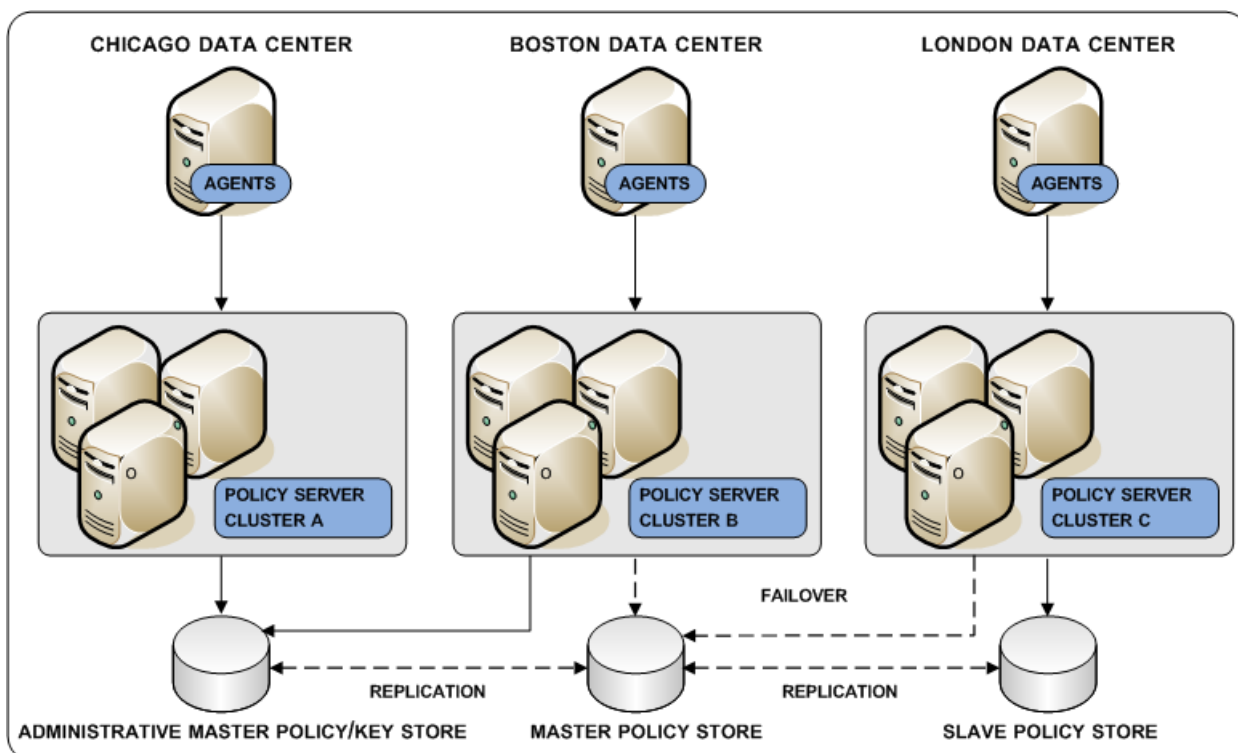
Note: If you use a master for key storage other than the master for administration, then all key stores must use the same key store value. No key store should be configured to function as both a policy store and a key store.

- All other policy store masters should be set for failover mode.

Due to possible synchronization issues, other configurations may cause inconsistent results, such as policy store corruption or Agent keys that are out of sync.

Contact CA SiteMinder® Support for assistance with other configurations.

The following diagram illustrates a multi-master policy store environment:



Policy Server to Audit Store Communication

By default, each Policy Server stores its own audit information to a text file. This text file is known as the Policy Server log. You can configure a Policy Server to log audit data directly to a database.

CA SiteMinder® audit logs are typically used for audit and compliance purposes. Consider the following:

- Having all Policy Servers write to a centralized audit store, where all data can be queried at once, may be preferable. If you deploy a centralized audit store, we recommend a highly available deployment.

Note: For more information about configuring an audit store, see the *Policy Server Installation Guide*. For more information about configuring failover, see the *Policy Server Administration Guide*.

Important! If you enable synchronous auditing, we recommend configuring failover to prevent an audit store outage from stopping all Policy Server authentications and authorizations. The Policy Server does not return the result of Agent authentication and authorization requests until the record is saved in the audit database. Users are not authenticated or authorized until the record is saved. For more information about configuring failover, see the *Policy Server Administration Guide*.

- If your deployment cannot permit Policy Servers to write to a centralized audit store, you can use the `smauditimport` utility to import individual Policy Server logs into a centralized audit store.

Note: For more information about Policy Server logging and the `smauditimport` tool, see the *Policy Server Administration Guide*.

More information:

[CA SiteMinder® Audit Database](#) (see page 17)

Policy Server to Session Store Communication


If you deploy a session store, all Policy Servers in the environment must use the same session store database.

Deploying a master session store is a way to achieve session store redundancy. A master session store lets each Policy Server communicate with the closest replicated version. This method of communication:

- Improves the performance of geographically separated Policy Servers. Sending Policy Server requests to a centralized session store outside a certain locale can result in increased network communication overhead and network congestion.
- Allows for failover. If a primary session store fails, Policy Servers failover to a secondary session store.

Note: For more information about configuring replication, see your vendor-specific documentation. For more information about configuring session store failover, see the *Policy Server Administration Guide*.

The following diagram illustrates all Policy Servers sharing a common view into a session store.



More information:

[Session Store](#) (see page 17)

Enhanced Session Assurance Architecture and Performance Considerations

The Enhanced Session Assurance feature prevents session hijacking and replay. When you log in, a DeviceDNA™ check is performed to fingerprint the end-user device. The device is validated by fingerprinting every 5 minutes by default and comparing the new fingerprint against the original fingerprint that is taken during the log in.

The work to perform the initial fingerprinting and subsequent re-checks will increase demands on the CA SiteMinder® architecture. Specifically, the CA SiteMinder® Policy Server and the instances of CA SiteMinder® SPS that run the Enhanced Session Assurance flow application are impacted. The time to authenticate a user increases, as will the time to authorize access to a resource when a re-validation of the fingerprint is required. The actual amount of additional latency during authentication depends on a number of parameters such as, but not limited to, network connection speeds, server capacity, and the end user's device. Internal testing has shown that authentication latency for an application configured with Enhanced Session Assurance may increase by 60%. This is due to additional redirections and processing time for the DeviceDNA™ collection and calculations. The actual percentage increase in each environment varies depending on the current authentication latency, the network, and speed of the computer that is accessing the resource. Additionally, resource (For example, CPU utilization) consumption for systems hosting the CA SiteMinder® family components that participate in an Enhanced Session Assurance transaction is higher if the transaction did not use Enhanced Session Assurance.

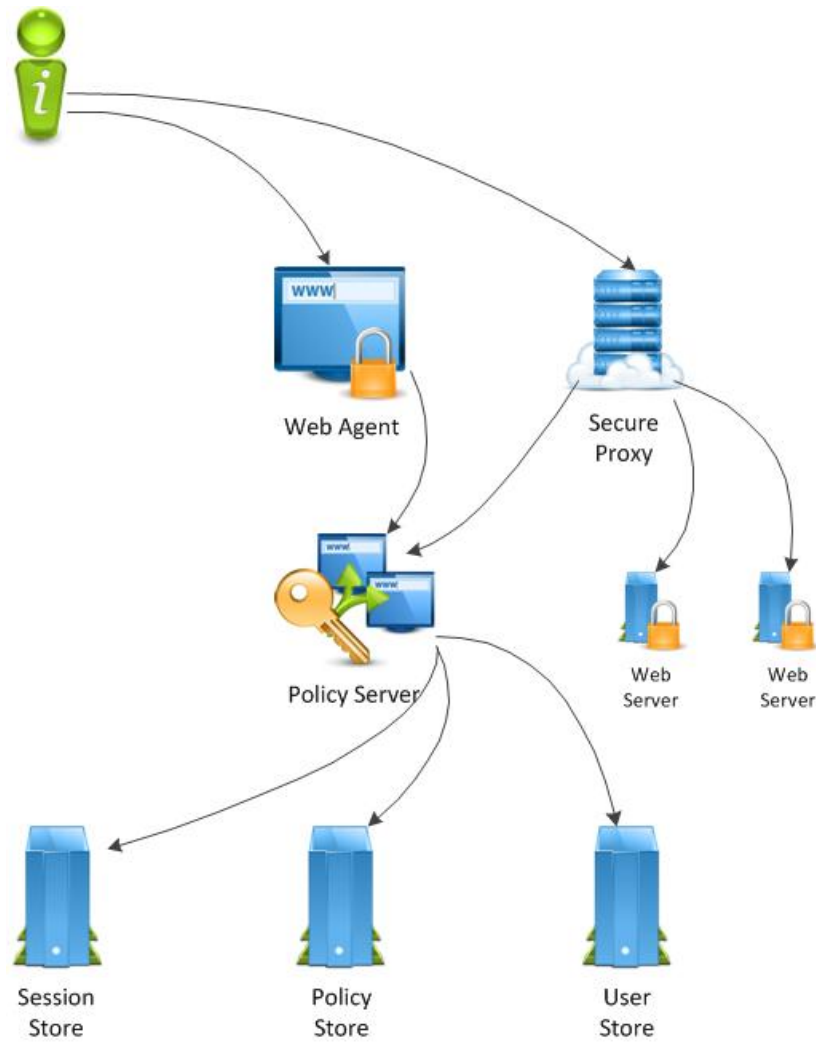
The application that drives the DeviceDNA checks is hosted on a CA SiteMinder® SPS instance. The instance can be a CA SiteMinder® SPS instance that performs the standard CA SiteMinder® SPS functions, such as web proxy or SAML federation functions or it can be a separate stand-alone CA SiteMinder® SPS instance that is dedicated to servicing the Enhanced Session Assurance transactions. Performance of the CA SiteMinder® SPS platform is also dependent on a number of parameters such as, but not limited to, authentication and authorization transactions per second, the ratio of authentications to authorizations within the environment, the length of user sessions, and the frequency of re-validations.

This section describes the different architectures that can be used for deploying CA SiteMinder® Enhanced Session Assurance and outlines choices to minimize the performance impact on existing CA SiteMinder® applications that do not use the feature in your environment.

Regardless of the architecture you choose to deploy, it is very important to introduce Enhanced Session Assurance gradually. The best practice is to install it in a development environment and test its impact on that environment by enabling it for different applications or realms gradually while measuring its impact on performance.

Basic Architecture

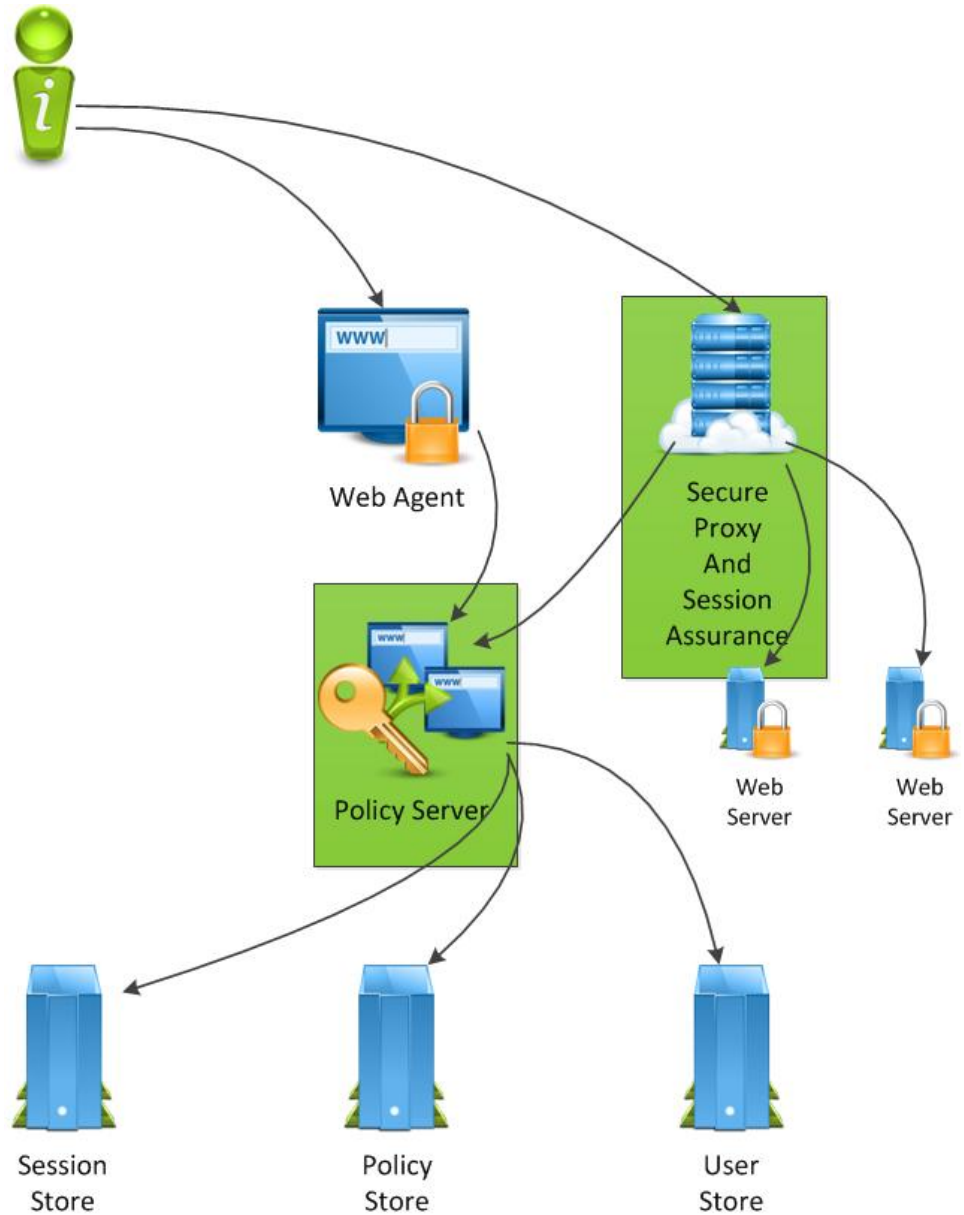
The following diagram illustrates an abbreviated, basic CA SiteMinder® architectural diagram that does not use Enhanced Session Assurance:



This architecture uses both Web Agents and CA SiteMinder® SPS for proxying to other web applications.

Possible Architecture 1–Use Existing Components

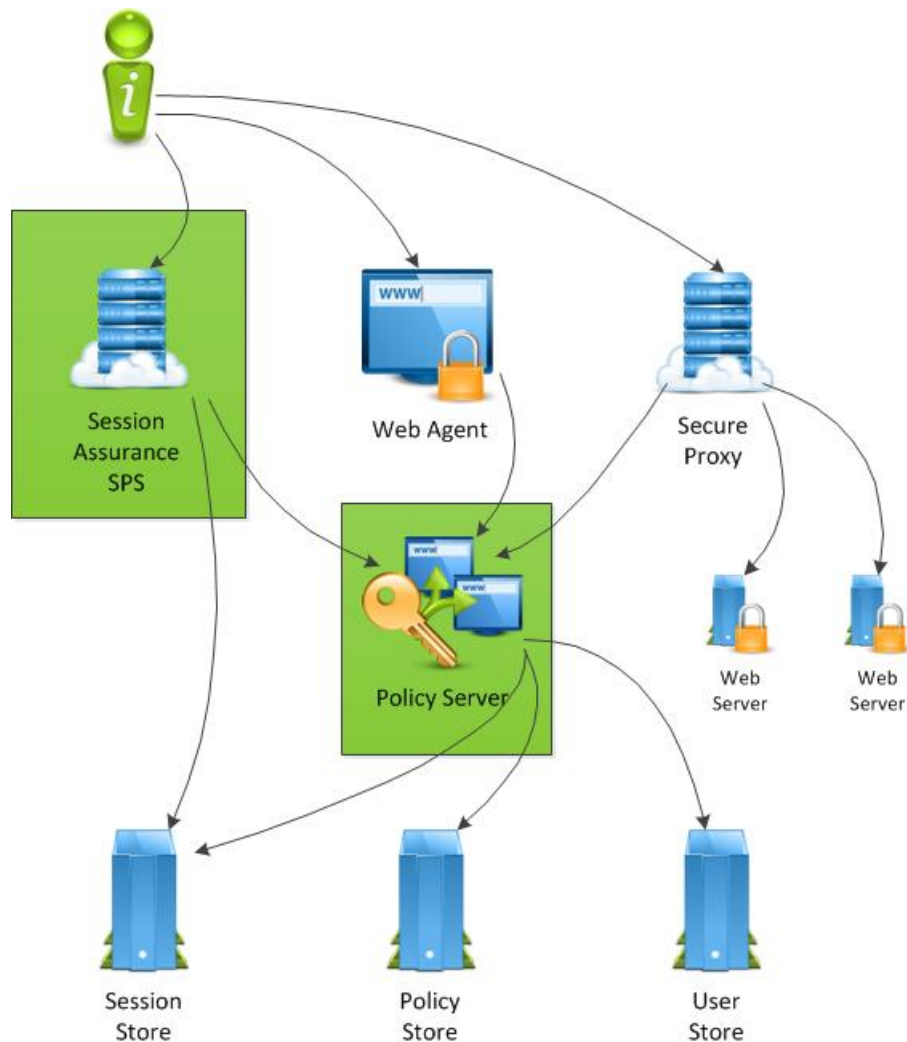
The following diagram illustrates a CA SiteMinder® architecture that uses existing components to deploy Enhanced Session Assurance:



In this architecture, Policy Server and CA SiteMinder® SPS that are highlighted in green can be used for Enhanced Session Assurance. Using the existing Policy Server and CA SiteMinder® SPS means that no additional hardware is needed to deploy the feature. However, in this architecture, as the Enhanced Session Assurance load increases, the CPU utilization of Policy Server and CA SiteMinder® SPS increases. If the load increases until either component's threads are fully utilized or the CPU cannot keep up with load, all the CA SiteMinder® transactions, regardless of whether they are configured to use Enhanced Session Assurance or not, will be negatively impacted.

Possible Architecture 2—Use Existing Policy Server


The following diagram illustrates a CA SiteMinder® architecture that uses a new CA SiteMinder® SPS instance to deploy Enhanced Session Assurance:



In this architecture, a new CA SiteMinder® SPS is introduced as highlighted in green. This CA SiteMinder® SPS fulfills all Enhanced Session Assurance tasks to avoid increased CPU utilization or a performance decrease of the other CA SiteMinder® SPS instance that is used to proxy requests to back-end web servers. However, by sharing the same Policy Server across both the CA SiteMinder® SPS instance running the Enhanced Session Assurance flow application, and the other agents and CA SiteMinder® SPS instances, if the Policy Server utilization demand is increased beyond its capacity, then all the CA SiteMinder® transactions from applications and agents will be affected.

Possible Architecture 3–Full Separation of the Session Assurance Components

The following diagram illustrates a possible CA SiteMinder® architecture that uses a new Policy Server and CA SiteMinder® SPS to deploy Enhanced Session Assurance:



In this architecture, a new CA SiteMinder® SPS instance is deployed specifically for hosting Enhanced Session Assurance. The new CA SiteMinder® SPS communicates with a new Policy Server. Though this architecture increases hardware in the environment, it keeps the existing Policy Server load and performance as close to the basic architecture as possible.

This architecture is recommended for large organizations that want to quickly rollout Enhanced Session Assurance. This helps in minimizing the chances of Enhanced Session Assurance increasing the CPU utilization or monopolizing threads that are needed for the general processing of requests.

Chapter 3: Plan a CA SiteMinder® Implementation

Implementation Planning Overview

The decisions related to how you implement CA SiteMinder® depend on:

- How you map your applications to the CA SiteMinder® access management models
- The CA SiteMinder® features you plan to use
- How you plan to manage CA SiteMinder® Policy Servers and Agents

We recommend that you consider the information in this section before you deploy and configure CA SiteMinder®.

Policy Management Models

CA SiteMinder® policy management models let you define access permissions for web resources and their respective user populations. A policy management model establishes the following:

- What resource is protected.
- Who can access the resource.
- What type of access user populations have.
- What happens when CA SiteMinder® grants access to the resource.
- What happens when CA SiteMinder® denies access to a resource.

All CA SiteMinder® functionality is available, regardless of which model you use. The primary difference between the models is the level of CA SiteMinder® knowledge required to configure each. The following Administrative UI objects represent the policy management models:

- Applications
- Policy domains and domain objects

Note: The following CA SiteMinder® core objects are required to configure an application object or domain policy:

- A host configuration object
- An agent configuration object

- An agent object
- A user directory object

Note: For more information about these objects, see the *Policy Server Configuration Guide*.

Policy Management Using Application Objects

Application objects provide an intuitive method of defining a complete security policy for a web application, website, or web service. Applications associate resources with user roles to specify entitlement policies that determine what users can access what resources.

Note: An application object defines policy information that can otherwise be configured in a policy domain and its subobjects. That is, realms, rules, rule groups, responses, and policies. The following table summarizes this relationship.

Application Dialogs and Group Boxes	Equivalent Domain Component
General settings	Policy domain and the root location of the protected resources.
Components	Realm and the location of the resources within the application that share the same security requirements.
Resource	Rule and the required authentication or authorization actions.
Application Roles	User directory lookups.

Note: For more information about policy management using applications, see the *Policy Server Configuration Guide*.

Policy Management Using Policy Domains and Domain Objects

Before SiteMinder r12.0, policy domains and domain objects (realms, rules, responses, policies, and so on) were the only way to protect resources. For Policy Server administrators already comfortable with them, policy domains and domain objects can still be used to configure security policies for resources.

A CA SiteMinder® policy is comprised of the following individual CA SiteMinder® objects:

- A domain
- At least one realm in the domain

- At least one rule or rule groups in the domain
- (Optional) One or more responses or response groups in the domain

A policy object binds these core objects to identify resources, user populations, and the required actions when CA SiteMinder® grants or denies access to the resource. As such, configuring a CA SiteMinder® policy requires an understanding of each object.

Note: For more information about each of these objects and their individual CA SiteMinder® policy roles, see the *Policy Server Configuration Guide*.

Identify the Applications to Secure

Which applications are you planning to secure? How do they map to the CA SiteMinder® access management models?

Begin thinking about the individual applications in the organization and the individual resources (URLs) within each application that require the same level of protection. We recommend identifying the following:

- Logical groupings of resources, often an individual application, that are associated with one or more user populations. These logical groupings map to either a CA SiteMinder® policy domain or the resource filter of an EPM application. A CA SiteMinder® policy domain or the resource filter of an EPM application represent the root location of the application.
- Sets of individual resources (URLs) within an application that have the same security (authentication and authorization) requirements. Sets of resources that share the same security requirements map to either a CA SiteMinder® policy realm or an EPM application component.

Grouping resources in this way helps you map applications to the CA SiteMinder® access management models.

When gathering information about each application, use a resource table similar to the following to help organization information:

Resource	Domain/Application Resource Filter	Realm/Component Resource Filter
Example: Corporate Portal	Example: Performance Management application	Example: Manager resources

Note: Identifying the applications that require protection also aids in capacity planning.

More information:

[Capacity Planning Introduced](#) (see page 87)

[Ignore Extensions Parameter](#) (see page 143)

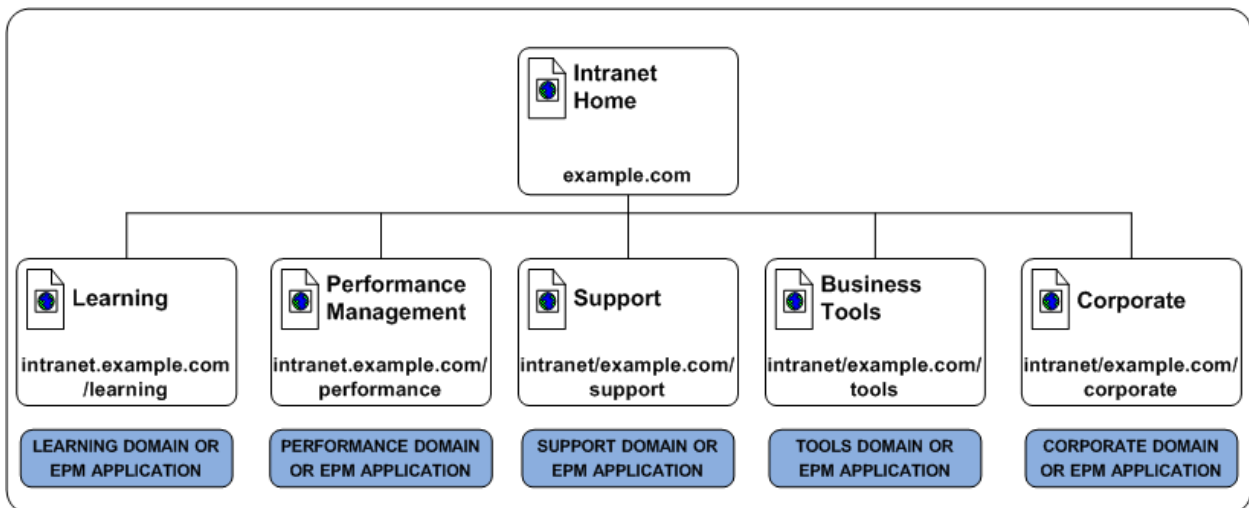
Group Resources into Domains or EPM Applications

Defining a CA SiteMinder® policy domain or an EPM application depends on identifying logical groups of resources, often an individual application, that are associated with one or more user populations. Grouping resources at this level helps you to identify the sets of individual resources (URLs) within an application that share the same security requirements.

Note: For more information about a CA SiteMinder® policy domain or an EPM application, see the *Policy Server Configuration Guide*.

A strategy for determining these requirements is to review a site map of the organization.

For example, a fictitious company, has a corporate intranet that the following site map represents:



In this example, the corporate portal is separated into the following logical groups of resources:

- Learning
- Performance Management
- Support

- Business Tools
- Corporate

The resource table for the corporate intranet looks like the following:

Resource	Domain/EPM Application Filter	Realm/Component Filter
Corporate Intranet	intranet.example.com	N/A
Learning	intranet.example.com/learning	N/A
Performance Management	intranet.example.com/performance	N/A
Support	intranet.example.com/support	N/A
Business Tools	intranet.example.com/tools	N/A
Corporate	intranet.example.com/corporate	N/A

More information:

[Domains and Authentication Performance](#) (see page 157)

Group Resources into Realms or EPM Components

Defining a CA SiteMinder® policy realm or an EPM component depends on identifying sets of individual resources (URLs) that share the same security or personalization requirements within a CA SiteMinder® policy domain or EPM application. The contents of a realm or EPM component share the same authentication scheme. As a result, identifying these resources early in the process can help you determine the authentication schemes required to meet individual security requirements.

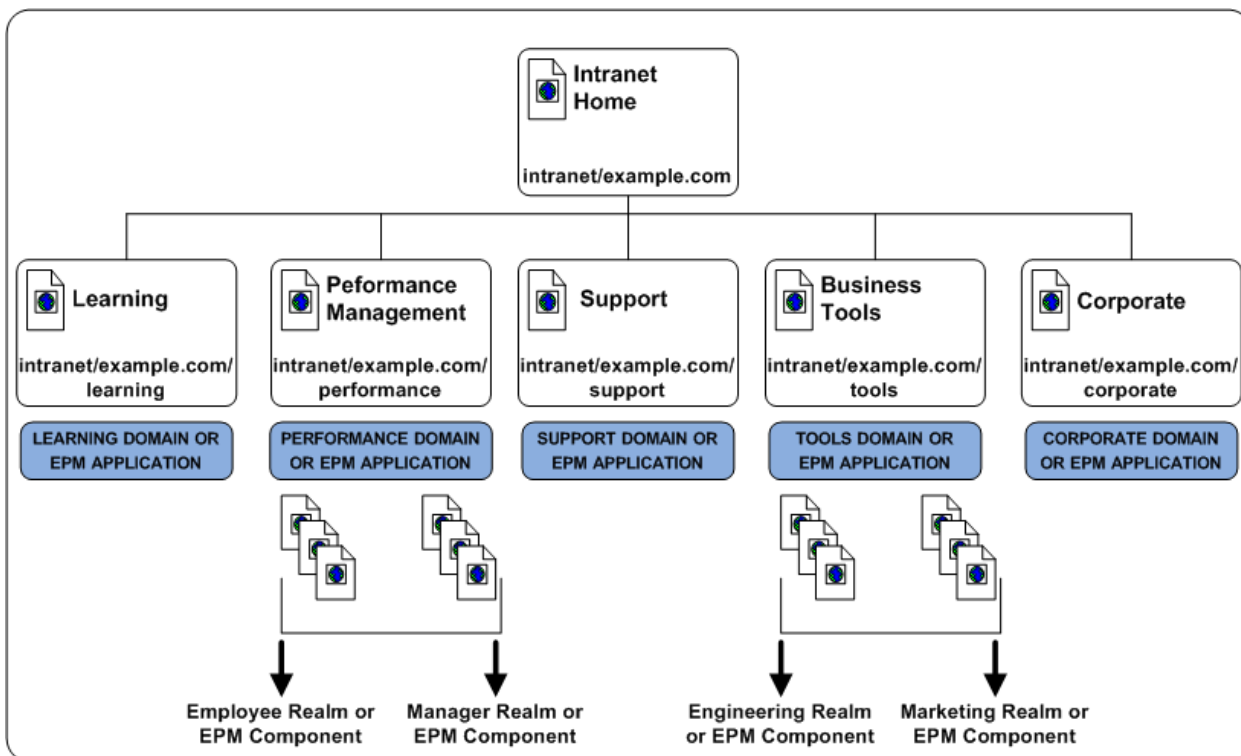
Note: For more information about CA SiteMinder® policy realms and EPM components, see the *Policy Server Configuration Guide*.

For example, although the Performance Management and Business Tools applications each let a specific user population access the root of the application, each application contains additional CA SiteMinder® policy realms or EPM components to provide a level of security or personalization appropriate for the resource:

- The Performance Management application contains resources that only full-time employees can access and resources that only managers can access.

- The Business Tools application contains resources that only Research and Development employees can access and resources that only Marketing employees can access.

Note: Although not illustrated, CA SiteMinder® policy rules and EPM resources are used to control specific Web Agent, authentication, and authorization events. For more information, see the *Policy Server Configuration Guide*.



The resource table for the applications looks like the following:

Resource	Domain/EPM Application Filter	Realm/Component Filter
Corporate Intranet	intranet.example.com	N/A
Learning	intranet.example.com/learning	N/A
Performance Management	intranet.example.com/performance	/employee /manager
Support	intranet.example.com/support	N/A
Business Tools	intranet.example.com/tools	/engineering /marketing

Resource	Domain/EPM Application Filter	Realm/Component Filter
Corporate	intranet.example.com/corporate	N/A

Identify User Stores

CA SiteMinder® can authenticate and authorize users through one or more connections to existing user stores in your enterprise network. After you identify the [applications to secure](#) (see page 55), consider the following questions:

- Do the applications use a centralized user store or use separate user stores for authentication?
- If the applications use separate stores, does this project include a task to centralize the user identities into a single store?
- Do the applications use the same store to authenticate and authorize users? Or is a separate store or stores used for authorization?

Identifying the stores each application uses helps you to:

- Identify the user store connections a CA SiteMinder® Administrator must configure in a CA SiteMinder® policy domain to protect the resource.

Note: For more information about configuring user store connections within a domain, see the *Policy Server Configuration Guide*.

- Determine if your environment requires the CA SiteMinder® directory mapping feature. By default, CA SiteMinder® assumes that users are authenticated and authorized against the same user store or stores. However, you can configure a CA SiteMinder® policy domain to authenticate against one or more stores and authorize against others.

Note: For more information about directory mapping, see the *Policy Server Configuration Guide*.

When gathering information about each application, use a table similar to the following to organize information:

User Store Name	User Store Type	Authentication?	Authorization?

Identify Authentication Methods

CA SiteMinder® supports multiple authentication methods to meet the varying levels of protection your resources require:

- Basic
- Forms-based user ID and password
- Hardware and software token-based, such as RSA® Ace/SecurID®
- Integrated Windows Authentication (IWA)
- Information Card Authentication Schemes (ICAS), such as Microsoft Windows CardSpace
- MIT Kerberos
- Server-based, such as RADIUS and SafeWord
- X.509 Certificate-based
- Custom Authentication schemes created using the CA SiteMinder® SDK

After you identify the [applications to secure](#) (see page 55), in which we recommend identifying sets of resources (URLs) that share the same security requirements, consider the following questions:

- Are their authentication guidelines, regulations, or laws your organization is required to meet for specific types of resources?
- How sensitive and valuable is the information?
- What types of users are accessing this information?
- What type of security do these users expect?

Answering these types of questions helps you to

- Identify the authentication methods your environment requires
- Identify the authentication schemes a CA SiteMinder® Administrator must configure to protect a specific resource.

Note: For more information about configuring authentication schemes, see the *Policy Server Configuration Guide*.

When gathering information about each resource, we recommend organizing your information by the applications you plan on securing. For example, the following table assumes that an application is grouped into individual domains and realms, as detailed in [applications to secure](#) (see page 55).

Resource	URL	Realm	Authentication Method

Resource	URL	Realm	Authentication Method

Identify Password Management Options

Do any security policies require your organization to manage user passwords? Do you anticipate managing user passwords in the future?

You can use CA SiteMinder® password policies to enforce the password requirements of your enterprise. A password policy can validate the user's password against any of the following types of characteristics before accepting it:

Composition

Verifies the minimum or maximum length, the types of characters allowed, and if or how often any of those characters can be repeated in a password.

Age

Verifies the time limits for how long the same password can be used, how long a password can remain inactive before it must be changed, and how long or how often before an expired password can be reused. You can specify one of the following responses for users with expired passwords:

- disable their accounts
- force them to change their passwords

Attempts

Records the number of times the user has previously entered an incorrect password, and takes one of the following actions when that number is exceeded:

- disables the account.
- waits a specified time period before allowing either one login attempt or reenabling the account.

Note: For more information, see the *CA SiteMinder® Policy Server Configuration Guide*.

Password Policy Considerations

If you plan to implement password policies in your enterprise, consider the following:

- CA SiteMinder® requires read/write access to the user directory, including exclusive use of several attributes within that directory to store passwords and password-related information.
- Password policies can affect CA SiteMinder® performance because of the additional user directory searches required to validate passwords. Password policies that are configured to search only part of a user directory, instead of the entire directory, can also affect performance.
- If your user directory has a native password policy, this policy must be:
 - Less-restrictive than the CA SiteMinder® password policy or
 - Disabled

Otherwise the native password policy accepts or rejects passwords without notifying CA SiteMinder®. Consequently, CA SiteMinder® cannot manage those passwords.

- By default, if a user enters incorrect information when changing a password, CA SiteMinder® returns a generic failure message. This message does not specify the failure reason. Create and enable the `DisallowForceLogin` registry key to change the default behavior and explicitly tell users why the change failed.
- If you use password policies on multiple Policy Servers, synchronize the system times of all servers. Synchronizing times helps to avoid the disabling of accounts or forcing password changes prematurely.

Note: For more information, see the *CA SiteMinder® Policy Server Configuration Guide*.

Identify Who Will Manage Your Web Agents

Web Agents connect to a Policy Server upon startup. The Policy Server contains an Agent Configuration Object (ACO), which directs the associated Web Agent to the location of its configuration parameters.

How your applications are deployed throughout your organization can help you determine the most efficient method of storing the configuration parameters for your CA SiteMinder® Web Agents. Consider the following questions:

1. Are most of your web applications deployed on a large server farm with the same security requirements?
2. Are most of your web applications managed by a centralized person or group?
3. Are most of your web applications deployed on separate web servers with different security requirements?
4. Are most of your web applications managed by different personnel in different departments or physical locations?

If you answered yes to questions one or two in the previous list, try the following configuration method:

- Central Configuration

Manages the parameters for one or more agents from an agent configuration object (ACO) that resides in the Policy Server. With a central agent configuration, you can update the parameter settings of several agents at once. Generally, each distinct web application uses a separate ACO, whose settings are shared among all the agents that protect the web application. For example, if you have five agents protecting one accounting application, you can create one ACO with the settings that you want for the application. All five agents would use the parameter settings from the same ACO.

For different applications, we recommend using separate agent configuration objects. For example, if you want to protect a human resources application with stricter security requirements, create a separate ACO for the human resources application.

When an agent starts, it reads the AllowLocalConfig parameter values of its associated ACO. When the value is set to no, then the agent uses the parameter settings from the ACO (except the agent log and trace file settings). Agent log files and trace files can always be controlled locally, regardless of ACO settings.

Note: We recommend using central agent configuration (wherever possible) because it simplifies agent configuration and maintenance.

If you answered yes to questions three or four in the previous list, try the following configuration method:

- Local Configuration

Manages each Web Agent individually using a file installed on the web server itself. When a Web Agent starts, it reads the value of the AllowLocalConfig parameter of its associated Agent Configuration Object (ACO). If the value is set to yes, then the Web Agent uses the parameter settings from LocalConfig.conf file on the web server. The parameter settings from the LocalConfig.conf file override any settings stored in an ACO on the Policy Server.

The following questions can help you identify other situations where local agent configuration better serves the needs of your enterprise:

- Will your enterprise deploy some of your Web Agents on reverse proxy servers?

For example, you want to protect your internal resources with a large group of Web Agents, while implementing reverse-proxy servers in a few locations. You can use local configuration to manage the reverse-proxy Web Agents.

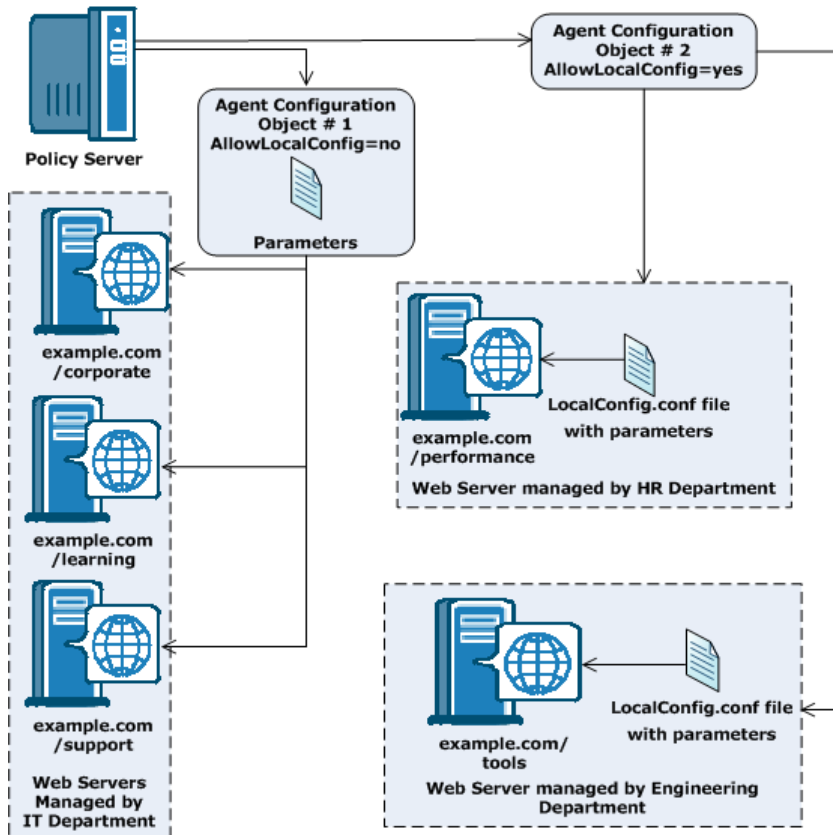
- Do you want to allow local web server administrators to change some Web Agent configuration settings but not others?

For example, your organization uses CA SiteMinder® to manage and enforce security policies, but allows web server administrators in remote offices to customize their log on and log off pages. You can add individual parameters to the value of the AllowLocalConfig parameter of the ACO to allow the administrators to change only those settings for the customized pages but no others.

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

Central and Local Configurations Together

You can also use a combination of central and local configuration to meet your needs. For example, you can manage three similar web servers with central configuration, while managing the other two servers with local configuration. See the following illustration for an example:



Identify Data Centers

Multiple factors, which are discussed later, can influence how you decide to implement CA SiteMinder® components across multiple data centers. Identifying the data centers and the purpose each is to serve in your CA SiteMinder® environment prepares you to make informed decisions when determining how to implement CA SiteMinder® components. Consider the following questions:

- How many data centers does your deployment include and where is each center located?
- If you have multiple data centers:
 - Will they all be active or are some only intended for disaster recover or backup?
 - Will each protected application reside in a single data center or across multiple centers?
 - Will you configure failover on the data center-level or across data centers?
 - What is the bandwidth and throughput between the data centers?

When gathering information about each data center, use a resource table similar to the following to organize your results:

Data Center Name	Location	Purpose

More information:

[Multiple Data Centers](#) (see page 113)

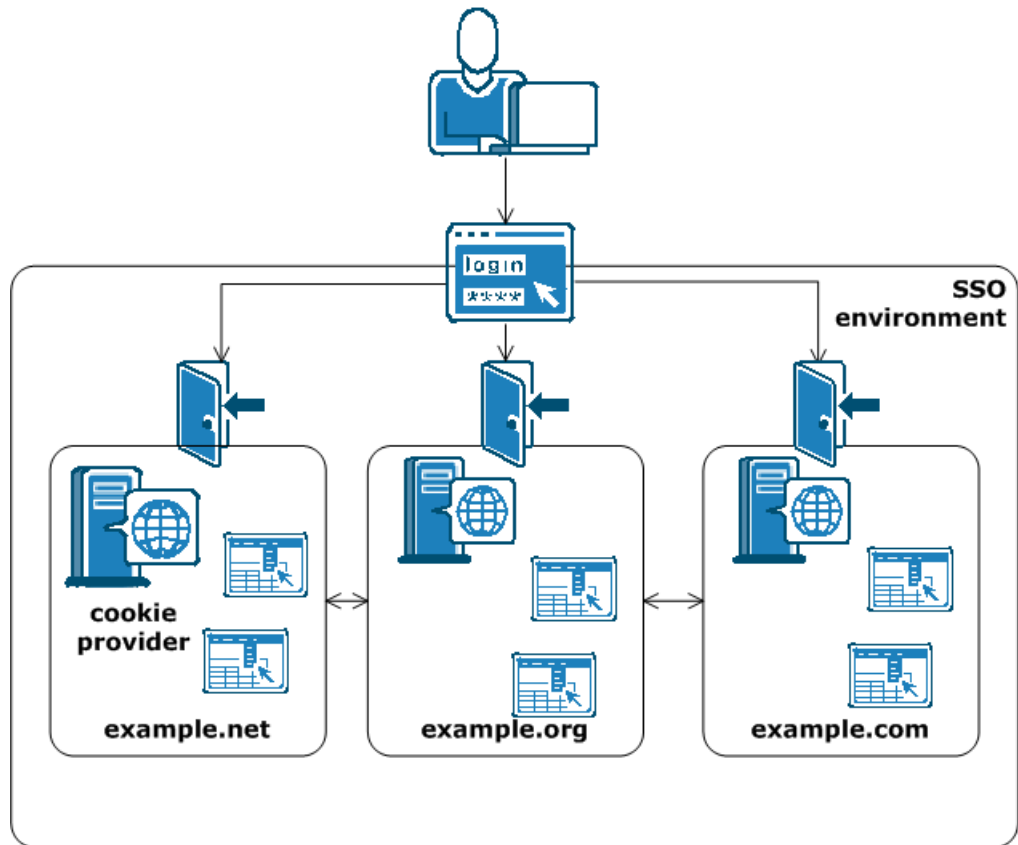
Identify Resources to be Secured with Multiple Cookie Domains

Will the single-sign on environment in your enterprise extend across multiple cookie domains?

CA SiteMinder® implements single sign-on across multiple cookie domains using a CA SiteMinder® Web Agent configured as a cookie provider.

The cookie domain where the cookie provider Web Agent resides is named the cookie provider domain. All the other Web Agents from the other cookie domains within the single sign-on environment, point to one cookie provider.

The following illustration shows an example of an SSO environment using multiple cookie domains:

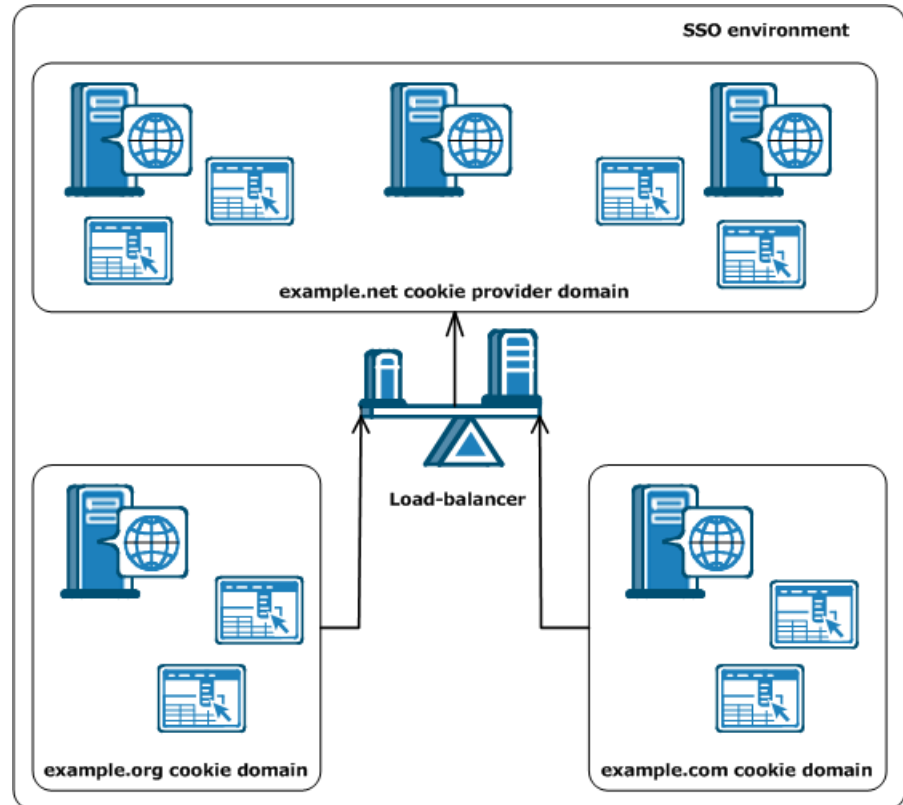


Note: For more information about cookie providers, see the *Web Agent Configuration Guide*.

Load-balancing for SSO between Cookie Provider Domains and Other Cookie Domains

Will the Agents in your single-sign on environment use load-balancing?

All agents in an SSO environment must refer to a single cookie provider domain. Add a load-balancer between the web servers in your cookie provider domain and the other cookie domains in your SSO environment. The following illustration shows an example:



The Web Agent in the example.org cookie domain points and the Web Agent in the example.com cookie domain both point to the same cookie provider domain of example.net. A load-balancer distributes the traffic evenly between all the web servers in the example.net cookie provider domain.

Determine if Partnerships Require CA SiteMinder® Federation

Do existing or planned business-to-business (B2B) partnerships require your organization to share identity information securely with partners?

CA SiteMinder® Federation lets you extend CA SiteMinder® functionality to partner sites by enabling identity federation. CA SiteMinder® Federation offers two deployment options: legacy federation and partnership federation.

Federated transactions between partner organizations let your enterprise:

- Exchange user identity information between partners in a secure fashion.
- Establish a link between a user identity at a partner and a user identity in your company.
- Enable single sign-on across partner web sites in multiple domains.
- Handle different user session models between partner sites, such as single-logout across all partner web sites or separate sessions for each partner web site.
- Control access to resources based on user information that is received from a partner.
- Allow interoperability across heterogeneous environments.

CA SiteMinder® Federation lets your enterprise generate, consume, or generate and consume assertions. CA SiteMinder® Federation supports the following standards and protocols:

- SAML 1.0 (legacy federation only)
- SAML 1.1 and 2.0
- Microsoft ADFS/WS-Federation (legacy federation only)
- SAML browser artifact protocol
- SAML POST protocol
- WS-Federation Passive Requestor Profile protocol (legacy federation only)

Note: CA SiteMinder® Federation is separately licensed from CA SiteMinder®. Contact your CA account representative for more information about licensing. For more information about federation, see the *CA SiteMinder® Federation Legacy Federation Guide* or the *CA SiteMinder® Federation Partnership Federation Guide*.

If your organization plans on implementing federation, use a table similar to the following table to identify partners and the possible methods for enabling identity federation.

Partner	Standard	Protocol

Determine if Advanced Encryption Standards are Required

Does your organization require the use of Federal Information Processing Standard (FIPS) 140–2 compliant algorithms?

The CA SiteMinder® implementation of the Advanced Encryption Standard (AES) supports the FIPS 140–2 standard. FIPS is a US government computer security standard used to accredit cryptographic modules that meet the AES.

The Policy Server uses certified FIPS 140–2 compliant cryptographic libraries. These cryptographic libraries provide a FIPS mode of operation when a CA SiteMinder® environment only uses AES–compliant algorithms to encrypt sensitive data. A CA SiteMinder® environment can operate in one of the following FIPS modes of operation.

- FIPS–compatibility
- FIPS–migration
- FIPS–only

Note: For more information about the cryptographic libraries CA SiteMinder® uses and the AES algorithms used to encrypt sensitive data in FIPS–only mode, see the *Policy Server Administration Guide*. For more information about the FIPS modes of operation and which to use when installing the Policy Server, see the *Policy Server Installation Guide*.

If you are implementing AES encryption through FIPS–only mode, consider the following:

- All third–party components, including directory servers, databases, and drivers must be configured to support FIPS–compliant algorithms.

Note: For more information about your vendors ability to support the FIPS 140–2 standard, see the vendor–specific documentation.

- If the environment uses X.509 Client Certificate authentication schemes, be sure that the user certificates are generated using only FIPS–compliant algorithms.

- If the Policy Servers are to connect to policy stores or user stores using SSL, be sure that the Policy Servers and directory stores use certificates that are FIPS-compliant.
- All Web Agents that ship with CA SiteMinder® r12.x are FIPS-compliant. To determine if other agents are FIPS-compliant, see the agent-specific documentation.

Important! An environment that is running in FIPS-only mode cannot operate with and is not backward compatible to earlier versions of CA SiteMinder®. This requirement includes all agents, custom software using older versions of the Agent API, and custom software using PM APIs or any other API that the Policy Server exposes. Re-link all such software with the current versions of the respective SDKs to achieve the required support for FIPS-only mode.

Determine if Virtualization is to be Used

Will CA SiteMinder® be implemented to a virtual environment?

Consider the following before implementing CA SiteMinder® to a virtual environment:

- Be sure to review the [CA policy on virtualization](#).
- Be sure to:
 - Understand the virtual environment and the performance overhead the host system can impose on applications.
 - Tune the virtual environment to eliminate as much as the performance overhead as possible.

Note: For more information about performance tuning the virtual environment, see the vendor-specific documentation.
- Be sure to size the CPU, disk space, and memory available to the virtual environment. Use the system requirements detailed in each CA SiteMinder® installation guide to determine how many components to deploy to the entire system.
- Be aware of issues associated with clock synchronization and multiple operating systems. Unsynchronized clocks can result in unexpected CA SiteMinder® behavior.

- When considering where to deploy components:
 - We recommend deploying Policy Servers to the virtual environment. We recommend that Policy Servers have their own Ethernet port. A dedicated port helps to prevent CA SiteMinder® from missing requests because it is competing with other virtual hosts for available bandwidth.
 - We recommend deploying Web Agents to virtualized web servers.
 - We recommend deploying all CA SiteMinder® data stores to physical hardware and operating systems. Directory servers and databases can become very resource dependent. If deployed to the virtualized environment, this dependency can result in performance degradation.

Determine how to Manage Policy Servers

Should individual business units be responsible for managing Policy Servers? Or can a single business unit manage all Policy Servers centrally?

Local Policy Server Management

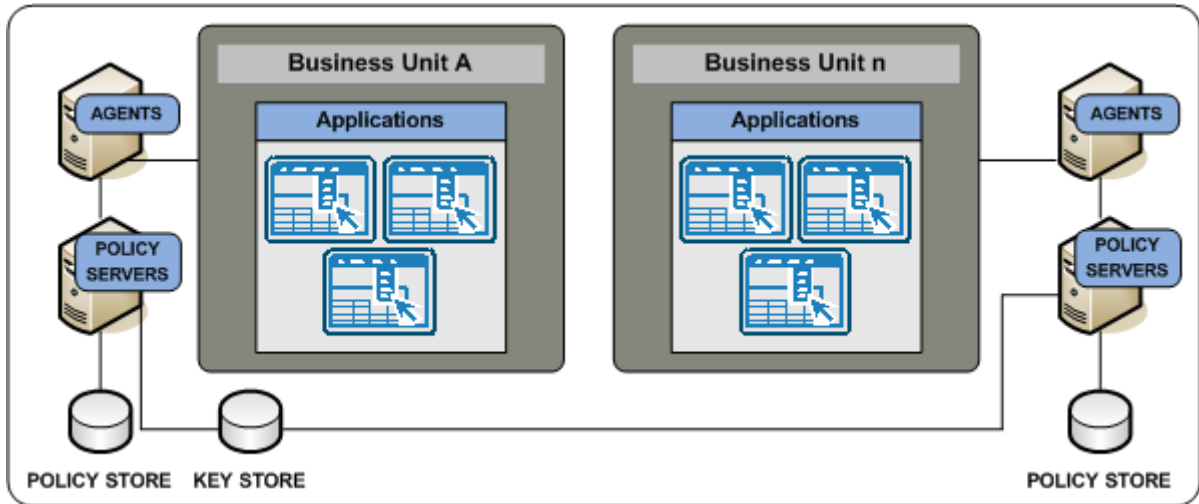
If individual business units manage Policy Servers and policy stores locally, consider that local Policy Server management:

- Lets each business unit manage their security requirements based on their individual needs.
- Can increase the complexity of the CA SiteMinder® infrastructure:
 - Local Policy Server management can result in more Policy Server and policy stores to manage and upgrade.
 - If single sign-on is a requirement, local Policy Server management results in additional CA SiteMinder® configuration. As illustrated, Policy Servers in both business units must share a key store to let all CA SiteMinder® Agents share the same keys.

Note: The illustration details a shared key store to depict a single sign-on requirement. A shared key store is not the only way to implement single sign-on and additional requirements exist. For more information about key management scenarios to facilitate single sign-on, see the *Policy Server Administration Guide*.

- Can make a consistent implementation and management of CA SiteMinder® core objects, policies, and EPM applications more challenging because CA SiteMinder® administrators are located in disparate business units.

The following illustration details two business units managing Policy Servers locally:



Central Policy Server Management

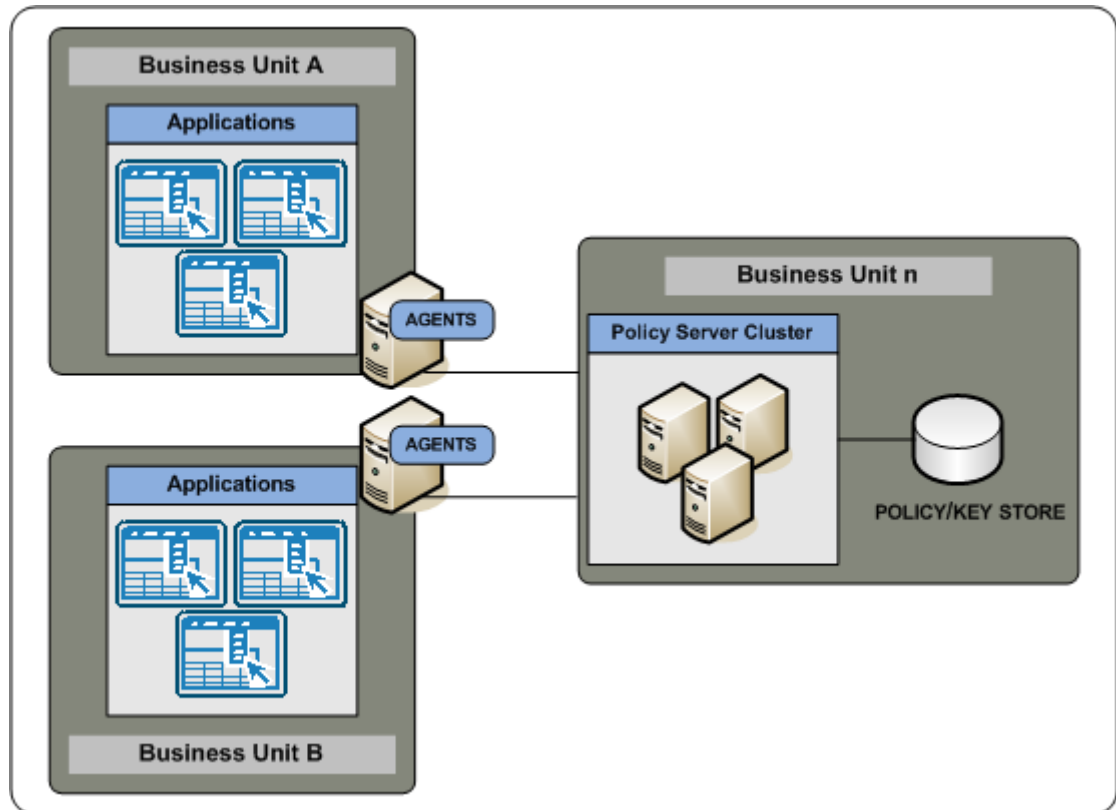
If a single business unit is to manage Policy Servers centrally, consider that central Policy Server management:

- Can facilitate a consistent implementation of CA SiteMinder® core objects, policies, and EPM applications because all CA SiteMinder® administrators are located in the same business unit.
- Can make the management of these objects easier because all CA SiteMinder® administrators are located in the same business unit.

Note: As illustrated, individual business units can continue to manage the CA SiteMinder® Agents protecting their applications.

- Can simplify the CA SiteMinder® infrastructure. Central management can result in fewer Policy Server and policy stores to manage and upgrade.
- Lets administrators monitor CA SiteMinder® performance centrally.

The following illustration details a single business unit managing all Policy Servers:



Determine how to Manage Web Agents

If you have several Web Agents which will all be configured identically, then using an Agent Configuration object on the Policy Server will make managing your Web Agents easier. A single Agent configuration object can be shared among an unlimited number of Web Agents. Configuration changes made on the Policy Server are automatically applied to any Web Agents which use the configuration object.

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

Chapter 4: Plan a CA SiteMinder® Web Services Security Implementation

Policy Management Models

CA SiteMinder® Web Services Security access management models let you define access permissions for applications and their respective user populations. An access management model establishes the following:

- What resource is protected.
- Who can access the resource.
- What type of access user populations have.
- What happens when CA SiteMinder® grants access to the resource.
- What happens when CA SiteMinder® denies access to a resource.

Almost all CA SiteMinder® Web Services Security functionality is available, regardless of which model you use. The primary difference between the models is the level of CA SiteMinder® knowledge required to configure each. The following Administrative UI objects represent the policy management models:

- Application objects
- Policy domains and policy objects

Note: The following CA SiteMinder® core objects are required to configure an application object or CA SiteMinder® domain policy:

- A host configuration object
- An agent configuration object
- An agent object
- A user directory object

For more information about these objects, see the *Policy Server Configuration Guide*.

Policy Management Using Application Objects

The recommended method for creating and managing new security policies for your CA SiteMinder® Web Services Security environment is to define application objects that represent one or more related web services and then generate the component and resource settings that define what to protect from associated WSDL files.

Note: Application objects do not support policy expressions using variable objects. Content-based authorization using variables must be implemented using policy domains and policies.

Policy Management Using Policy Domains and Policies

For Policy Server administrators already comfortable with CA SOA Security Manager or CA SiteMinder®, policy management using policy domains and domain objects (realms, rules, responses, policies, and so on) — can still be used to perform manual configuration of security policies for web service resources.

Domains and domain objects must also be used in the following situations:

- To modify policies created traditionally and migrated from a previous CA SOA Security Manager deployment.
- To implement content-based authorization using variables.

Identify the Web Services to Secure

Which web services are you planning to secure? How do they map to the CA SiteMinder® Web Services Security policy management methods?

Begin thinking about the individual web services in your organization and the sets of operations within each web service that require the same level of protection. We recommend identifying the following:

- Logical groupings of resources, often offered by individual web services, that are associated with one or more user populations.
- Sets of individual operations within a web service that have the same security (authentication and authorization) requirements.

Note: Identifying the web services that require protection also aids in capacity planning.

More information:

[CA SiteMinder® Web Services Security Capacity Planning Introduced](#) (see page 103)

Identify User Stores

CA SiteMinder® Web Services Security can authenticate and authorize users through one or more connections to existing user stores in your enterprise network. After you identify the web services to secure, consider the following questions:

- Do the web services use a centralized user store or use separate user stores for authentication?
- If the web services use separate stores, does this project include a task to centralize the user identities into a single store?
- Do the web services use the same store to authenticate and authorize users? Or is a separate store or stores used for authorization?

Identifying the stores each web service uses helps you to:

- Identify the user store connections a CA SiteMinder® Administrator must configure in a CA SiteMinder® policy domain to protect the resource.

Note: For more information about configuring user store connections within a domain, see the *Policy Server Configuration Guide*.

- Determine if your environment requires the CA SiteMinder® directory mapping feature. By default, CA SiteMinder® assumes that users are authenticated and authorized against the same user store or stores. However, you can configure a CA SiteMinder® policy domain to authenticate against one or more stores and authorize against others.

Note: For more information about directory mapping, see the *Policy Server Configuration Guide*.

When gathering information about each web service, use a table similar to the following to organize information:

User Store Name	User Store Type	Authentication?	Authorization?

More information:

[Identify the Web Services to Secure](#) (see page 78)

Identify Authentication Methods

CA SiteMinder® Web Services Security supports four authentication methods to meet the varying levels and types of protection your resources require:

XML Document Credential Collector

Validates XML messages using credentials gathered from the message itself by mapping fields within the document to fields within a user directory.

XML Digital Signature

Validates XML documents digitally signed with valid X.509 certificates.

WS-Security

Validates XML messages using credentials gathered from WS-Security headers in the SOAP envelope of an incoming message.

CA SiteMinder® Web Services Security can produce and consume WS-Security tokens, enabling you to use the WS-Security authentication scheme to deploy a multiple-web service implementation across federated sites.

SAML Session Ticket

Validates XML messages using credentials obtained from CA SiteMinder® Web Services Security synchronized-sessioning SAML assertions (which contain an encrypted combination of a CA SiteMinder session ticket and a CA SiteMinder user public key) placed in the message HTTP header, SOAP envelope, or cookie.

CA SiteMinder® Web Services Security can generate and consume SAML Session Ticket assertions. This enables you to use the SAML Session Ticket authentication scheme to deploy a multiple-web service implementation within a single Policy Server domain.

After you identify the web services to secure, in which we recommend identifying web service operations that share the same security requirements, consider the following questions:

- Are their authentication guidelines, regulations, or laws your organization is required to meet for specific types of resources?
- How sensitive and valuable is the information?
- What types of users are accessing this information?
- What type of security do these users expect?

Answering these types of questions helps you to

- Identify the authentication methods your environment requires
- Identify the authentication schemes a CA SiteMinder® Administrator must configure to protect a specific resource.

Note: For more information about configuring authentication schemes, see the *Policy Server Configuration Guide*.

More information:

[Identify the Web Services to Secure](#) (see page 78)

Identify Who Will Manage Your SiteMinder WSS Agents

SiteMinder WSS Agents connect to a Policy Server upon startup. The Policy Server contains an Agent Configuration Object (ACO), which directs the associated SiteMinder WSS Agent to the location of its configuration parameters.

How your web services are deployed throughout your organization can help you determine the most efficient method of storing the configuration parameters for your SiteMinder WSS Agents. Consider the following questions:

1. Are most of your web services deployed on a large server farm with the same security requirements?
2. Are most of your web services managed by a centralized person or group?
3. Are most of your web services deployed on separate servers with different security requirements?
4. Are most of your web services managed by different personnel in different departments or physical locations?

CA SiteMinder® Web Services Security offers the following configuration methods:

Central Configuration

If you answered yes to questions one or two in the previous list, try central configuration in which one or more SiteMinder WSS Agents are managed from an Agent Configuration Object (ACO) that resides in the Policy Server. With central configuration, you can update the parameter settings of several SiteMinder WSS Agents at once. Generally, each distinct web service uses a separate ACO, whose settings are shared among all the SiteMinder WSS Agents that protect the web service. For example, if you have five SiteMinder WSS Agents protecting one accounting web service, you can create one ACO with the settings for the web service. All five SiteMinder WSS Agents would use the parameter settings from the same ACO.

For different applications, we recommend using separate Agent Configuration Objects. For example, if you want to protect a human resources web service with stricter security requirements, create a separate ACO for the human resources web service.

When a SiteMinder WSS Agent starts, it reads the value of the AllowLocalConfig parameter of its associated ACO. If the value is set to no, then the SiteMinder WSS Agent uses the parameter settings from the ACO.

Note: We recommend using central agent configuration (wherever possible) because it simplifies agent configuration and maintenance.

Local Configuration

If you answered yes to questions three or four in the previous list, try local configuration in which each SiteMinder WSS Agent is managed individually using a file installed on the server itself. When a SiteMinder WSS Agent starts, it reads the value of the AllowLocalConfig parameter of its associated Agent Configuration Object (ACO). If the value is set to yes, then the SiteMinder WSS Agent uses the parameter settings from LocalConfig.conf file on the application or web server. The parameter settings from the LocalConfig.conf file override any settings stored in an ACO on the Policy Server. If you answered yes to questions three or four in the previous list, try the following configuration method:

Note: For more information about the location of the LocalConfig.conf file on your application or web server, see the corresponding SiteMinder WSS Agent Guide.

You can also use a combination of central and local configuration to meet your needs. For example, you can manage three similar web servers with central configuration, while managing the other two servers with local configuration.

The following questions can help you identify other situations where local agent configuration better serves the needs of your enterprise:

- Will your enterprise deploy custom SiteMinder WSS Agents on XML gateways?
For example, you want to protect your internal resources with a large group of SiteMinder WSS Agents, while implementing XML gateways in a few locations. You can use local configuration to manage the custom SiteMinder WSS Agents on XML gateways.
- Do you want to allow local server administrators to change some SiteMinder WSS Agent configuration settings but not others?
For example, your organization uses CA SiteMinder® to manage and enforce security policies, but allows application and web server administrators in remote offices to customize their log on and log off pages. You can add individual parameters to the value of the AllowLocalConfig parameter of the ACO to allow the administrators to change only those settings for the customized pages but no others.

Identify Data Centers

Multiple factors, which are discussed later, can influence how you decide to implement CA SiteMinder® components across multiple data centers. Identifying the data centers and the purpose each is to serve in your CA SiteMinder® environment prepares you to make informed decisions when determining how to implement CA SiteMinder® components. Consider the following questions:

- How many data centers does your deployment include and where is each center located?
- If you have multiple data centers:
 - Will they all be active or are some only intended for disaster recover or backup?
 - Will each protected application reside in a single data center or across multiple centers?
 - Will you configure failover on the data center-level or across data centers?
 - What is the bandwidth and throughput between the data centers?

When gathering information about each data center, use a resource table similar to the following to organize your results:

Data Center Name	Location	Purpose

More information:

[Multiple Data Centers](#) (see page 113)

Determine if Advanced Encryption Standards are Required

Does your organization require the use of Federal Information Processing Standard (FIPS) 140–2 compliant algorithms?

The CA SiteMinder® implementation of the Advanced Encryption Standard (AES) supports the FIPS 140–2 standard. FIPS is a US government computer security standard used to accredit cryptographic modules that meet the AES.

The Policy Server uses certified FIPS 140–2 compliant cryptographic libraries. These cryptographic libraries provide a FIPS mode of operation when a CA SiteMinder® environment only uses AES–compliant algorithms to encrypt sensitive data. A CA SiteMinder® environment can operate in one of the following FIPS modes of operation.

- FIPS–compatibility
- FIPS–only

Note: For more information about the cryptographic libraries CA SiteMinder® uses and the AES algorithms used to encrypt sensitive data in FIPS–only mode, see the *Policy Server Administration Guide*. For more information about the FIPS modes of operation and which to use when installing the Policy Server, see the *Policy Server Installation Guide*.

If you are implementing AES encryption through FIPS–only mode, consider the following:

- All third–party components, including directory servers, databases, and drivers must be configured to support FIPS–compliant algorithms.
Note: For more information about your vendors ability to support the FIPS 140–2 standard, see the vendor-specific documentation.
- If the environment uses X.509 Client Certificate authentication schemes, be sure that the user certificates are generated using only FIPS–compliant algorithms.
- If the Policy Servers are to connect to policy stores or user stores using SSL, be sure that the Policy Servers and directory stores use certificates that are FIPS–compliant.
- All SiteMinder WSS Agents that ship with CA SiteMinder® 12.x are FIPS–compliant. To determine if other agents are FIPS–compliant, see the agent–specific documentation.

Important! An environment that is running in FIPS–only mode cannot operate with and is not backward compatible to earlier versions of CA SiteMinder®. This requirement includes all agents, custom software using older versions of the CA SiteMinder® Web Services Security SDK. Re–link all such software with the current versions of the SDK to achieve the required support for FIPS–only mode.

Determine if Virtualization is to be Used

Will CA SiteMinder® be implemented to a virtual environment?

Consider the following before implementing CA SiteMinder® to a virtual environment:

- Be sure to review the [CA policy on virtualization](#).
- Be sure to:
 - Understand the virtual environment and the performance overhead the host system can impose on applications.
 - Tune the virtual environment to eliminate as much as the performance overhead as possible.

Note: For more information about performance tuning the virtual environment, see the vendor–specific documentation.
- Be sure to size the CPU, disk space, and memory available to the virtual environment. Use the system requirements detailed in each CA SiteMinder® installation guide to determine how many components to deploy to the entire system.
- Be aware of issues associated with clock synchronization and multiple operating systems. Unsynchronized clocks can result in unexpected CA SiteMinder® behavior.
- When considering where to deploy components:
 - We recommend deploying Policy Servers to the virtual environment. We recommend that Policy Servers have their own Ethernet port. A dedicated port helps to prevent CA SiteMinder® from missing requests because it is competing with other virtual hosts for available bandwidth.
 - We recommend deploying Web Agents to virtualized web servers.
 - We recommend deploying all CA SiteMinder® data stores to physical hardware and operating systems. Directory servers and databases can become very resource dependent. If deployed to the virtualized environment, this dependency can result in performance degradation.

Determine how to Manage Policy Servers

Should individual business units be responsible for managing Policy Servers? Or can a single business unit manage all Policy Servers centrally?

Local Policy Server Management

If individual business units manage Policy Servers and policy stores locally, consider that local Policy Server management:

- Lets each business unit manage their security requirements based on their individual needs.
- Can increase the complexity of the CA SiteMinder® infrastructure; local Policy Server management can result in more Policy Server and policy stores to manage and upgrade.
- Can make a consistent implementation and management of CA SiteMinder® core objects, policies, and application objects more challenging because CA SiteMinder® administrators are located in disparate business units.

Centralized Policy Server Management

If a single business unit is to manage Policy Servers centrally, consider that central Policy Server management:

- Can facilitate a consistent implementation of CA SiteMinder® core objects, policies, and application objects because all CA SiteMinder® administrators are located in the same business unit.
- Can make the management of these objects easier because all CA SiteMinder® administrators are located in the same business unit.

Note: As illustrated, individual business units can continue to manage the CA SiteMinder® Agents protecting their applications.

- Can simplify the CA SiteMinder® infrastructure. Central management can result in fewer Policy Server and policy stores to manage and upgrade.
- Lets administrators monitor CA SiteMinder® performance centrally.

Determine how to Manage SiteMinder WSS Agents

If you have several SiteMinder WSS Agents which will be configured identically, then using an Agent Configuration object on the Policy Server will make managing those Agents easier. A single Agent configuration object can be shared among an unlimited number of SiteMinder WSS Agents. Configuration changes made on the Policy Server are automatically applied to any SiteMinder WSS Agents which use the configuration object.

Note: For more detailed information about how to configure Agent-related objects, see the *CA SiteMinder® Web Services Security Policy Configuration Guide* and the *SiteMinder WSS Agent Guide* for your SiteMinder WSS Agent type.

Chapter 5: CA SiteMinder® Capacity Planning

This section contains the following topics:

[Capacity Planning Introduced](#) (see page 87)

[Use Case: Capacity Planning](#) (see page 88)

[How to Estimate a Sustained Authentication Rate](#) (see page 89)

[Estimate a Peak Authentication Rate](#) (see page 93)

[How to Estimate a Sustained Authorization Rate](#) (see page 94)

[Estimate a Peak Authorization Rate](#) (see page 99)

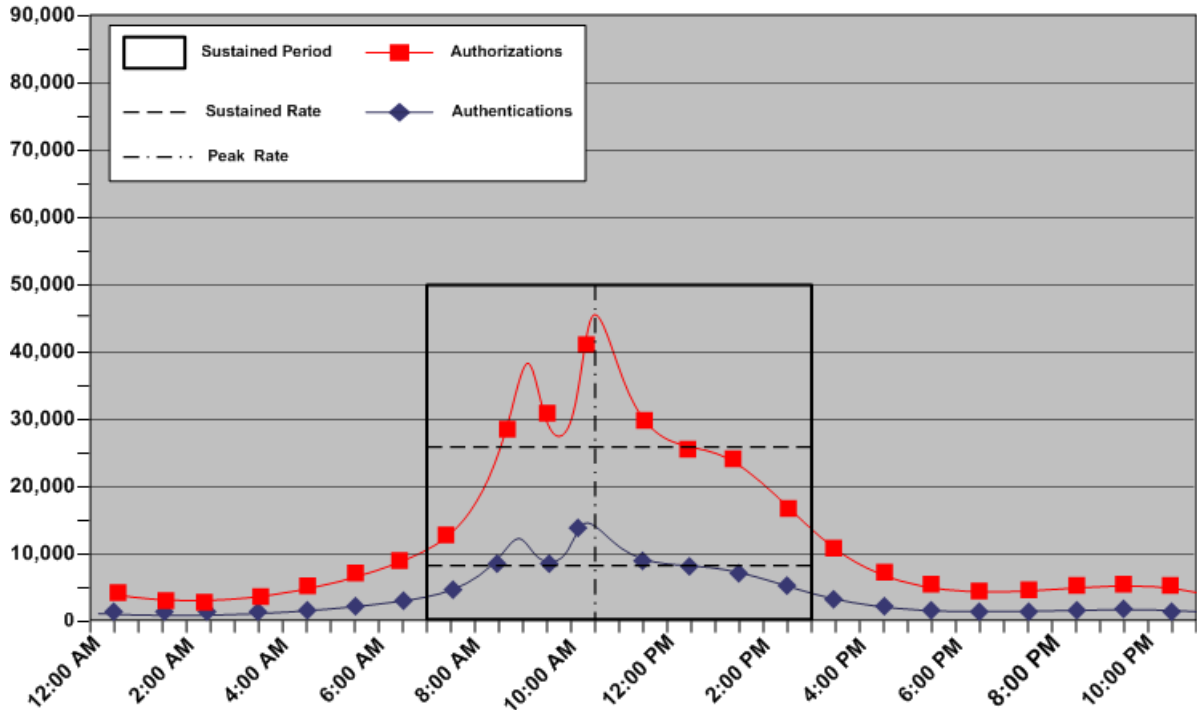
Capacity Planning Introduced

Planning a CA SiteMinder® deployment with performance in mind is the first step to maintaining high enterprise availability and performance standards. A good approach is to estimate the number of expected authentications and authorizations CA SiteMinder® must handle per application. The following general factors influence CA SiteMinder® performance:

- Sustained authentication and authorization rates. The rate at which users authenticate to an application and request protected resources fluctuates throughout your business day. Some periods can generate relatively few authentication requests, and therefore relatively few authorization requests, while others generate more. The sustained authentication and authorization rates represent a sustained period during which CA SiteMinder® must service an average number of authentication and authorization requests.
- Peak authentication and authorization rates. During sustained periods of activity, user activity may spike. The peak authentication and authorization rates represent a period during which CA SiteMinder® must service the highest number of authentication and authorization requests.

Note: Although a number of other factors can influence CA SiteMinder® performance, such as performance tuning and network bandwidth, the previous factors can help you make informed decisions when implementing Policy Servers and Agents, and when determining if existing user stores can handle the anticipated CA SiteMinder® workload.

The following graphic illustrates how authentication and authorization rates fluctuate throughout the day, are sustained for a specific period, and peak within that period:



Note: Authenticating and authorizing users results in a number of reads, and if Password Policies are enabled, writes, to a user store. Determining sustained and peak rates helps you determine the load under which your user stores must operate to service Policy Server requests.

More information:

[Performance Tuning Introduced](#) (see page 129)

Use Case: Capacity Planning

The purpose of the following use case is to illustrate how a fictitious organization approaches capacity planning by modeling the usage of their application. The use case is referenced throughout this chapter for examples.

The company is planning to deploy CA SiteMinder®. The company has 100,000 users in a single user store. Password Services is enabled for this store.

Some users log into the portal application once a day, while other users login as much as three times per day.

How to Estimate a Sustained Authentication Rate

Estimating the sustained authentication rate of an application is the process of determining:

- How the total number of authentication requests fluctuate throughout your business day
- How the authentication requests translate into requests per second.

Complete the following steps to estimate the sustained authentication rate for an application:

1. Estimate daily authentications.
2. Estimate the sustained authentication rate.

Estimate Daily Authentications

What is the estimated number of daily authentications for the application?

The number of users directly affect daily authentications (authentication load). When users log into the application, CA SiteMinder® authenticates them. Therefore, think of the authentication load of the application as the total logins per day.

Note: When determining the authentication load, we recommend beginning with an evaluation interval of 24 hours. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding of usage throughout the year.

All users logging into the application each day is unlikely, so estimating total logins begins with determining the percentage of users that log in once a day, which the following represents:

$$(total_users * percentage_users) * (number_of_logins) = daily_logins$$

total_users

Represents the total number of users with access to the application.

percentage_users

Represents the percentage of users who log in the same number of times per day.

number_of_logins

Represents the number of times the particular set of users login.

daily_logins

Represents the number of logins the particular set of users creates.

Example 1: The company has 100,000 users, 75 percent of which log in once a day.

$$(100,000 * 0.75) \times (1) = 75,000 \text{ logins}$$

However, some users logging into the application two or more times a day is more likely.

Example 2: The company has 100,000 users, 5 percent of which log in twice a day and 1 percent of which log in three times a day.

$$(100,000 * 0.05) \times (2) = 10,000 \text{ logins}$$

$$(100,000 * 0.01) \times (3) = 3,000 \text{ logins}$$

The total logins per day are the sum of each of the login calculations.

Example 3: The company has 100,000 users:

- 75 percent of which log in once a day, creating 75,000 logins.
- Five percent of which log in twice a day, creating 10,000 logins.
- One percent of which log in three times a day, creating 3,000 logins.

The authentication load for the portal application is 88,000 logins.

Note: The percentage of users logging in does not have to equal 100 percent because all users will not log into the application each day.

The following table illustrates each of the previous examples:

Total Users	Percent of Total Users	Logins Per Day	Logins
100,000	75	1	75,000
100,000	5	2	10,000
100,000	1	3	3,000
Authentication Load			88,000

The company uses the authentication load to estimate the sustained authentication rate.

Estimate a Sustained Authentication Rate

What is the sustained authentication rate for the application?

The sustained authentication rate is based on the authentication load. Specifically, when and at what rate the authentications occur. The chance that the authentication load is uniformly spread across your business day is unlikely. Rather, the rate at which requests occur fluctuates, remaining between the lowest and highest (peak) levels for a sustained period. Estimating the sustained authentication rate is the process of identifying a sustained period during which the system is servicing an average amount of authentication requests.


When estimating a sustained authentication rate, we recommend using the daily authentication load to determine:

- The rate at which the authentication requests occur throughout your business day.

Note: We recommend beginning with an evaluation period of 24 hours, broken down into one-hour increments. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding of usage throughout the year.

- The sustained period during which the system is servicing an average number of authentication requests.
- The approximate number of authentication requests that occur during the sustained period.

The following figure is an example of these metrics:



Identifying these metrics helps you to estimate the number of authentication requests, per second, that CA SiteMinder® must service to maintain the average rate at which users authenticate, which the following represents:

$$\frac{(\text{authentication_load} * \text{percentage_of_authentication_requests})}{\text{number_of_sustained_hours} / 3600} = \text{sustained_authentication_rate}$$

authentication_load

Represents the number of daily authentications for the application.

percentage_of_authentication_requests

Represents the percentage of authentication requests that occur when the system is operating at sustained levels.

Example: If the authentication load is 50,000 logins, and 32,000 logins occur during the sustained period, then the value is 64percent (0.64)

number_of_sustained_hours

Represents the number of hours in which the system is operating at the sustained level.

Note: 3,600 represents the number of seconds in an hour.

sustained_authentication_rate

Represents the number of authentication requests, per second, that CA SiteMinder® must service during the period of sustained activity.

Example: Estimate the Sustained Authentication Rate

The company has determined that their application portal has an authentication load of 88,000 logins. The application portal is available to customers 24 hours a day, seven days a week. Using system activity reports to break down a typical day results in the following metrics:

- The system is operating at sustained levels for approximately five hours (9:00 AM - 2:00 PM).
- During sustained levels, approximately 9,000 authentications requests occur per hour.
- Approximately 45,000 (9,000 * 5) authentication requests, or 51 percent (45,000 / 88,000) of the daily authentication load, occur during these hours.

$(88,000 * 0.51) / 5 / 3600 = 2.49$ authentications per second.

The portal application has a sustained authentication rate of 2.49 authentications per second.

Estimate a Peak Authentication Rate


What is the peak authentication rate for the application?

The peak authentication rate is based on the sustained authentication rate, specifically, when and at what rate the system is operating at peak levels. Estimating the peak authentication rate is the process of identifying when the system is servicing the highest level of authentication requests.

When estimating the peak authentication rate, we recommend using the metrics you gathered when determining the sustained authentication rate to determine:

- The hour when the system is servicing the highest number of authentication requests
- The approximate number of authentication requests that occur during this period

The following figure is an example of these metrics:



Identifying these metrics helps you to estimate the number of authentication requests, per second, that CA SiteMinder® must service to maintain the peak rate at which users authenticate, which the following represents:

$$\frac{(\text{authentication_load} \times \text{percentage_of_transactions})}{\text{number_of_hours} / 3600} = \text{peak_authentication_rate}$$

Note: This rate is based on the single busiest hour. There can be periods when the peak authentication rate exceeds the hourly calculation.

authentication_load

Represents the number of daily authentications for the application.

percentage_of_transactions

Represents the percentage of transactions that occur when the system is operating at peak levels.

number_of_hours

Represents the number of hours in which the system operates at peak levels.

Note: 3,600 represents the number of seconds in an hour.

peak_authentication_rate

Represents the peak authentication rate for the application.

Example: Estimate the Peak Authentication Rate

The company has determined that their portal application has a daily authentication load of 88,000 logins. System activity reports detail that during the single busiest hour of the day 18,000 authentication requests occur. This number represents approximately 20 percent of the authentication load:

$$18,000 / 1 / 3600 = 5 \text{ authentications per second}$$

The portal application has a peak authentication rate of five authentications per second.

Note: This example is based on the single busiest hour. There can be periods when the peak authentication rate during the hour exceeds five authentications per second.

More information:

[Increase the Amount of Available Sockets for the Agent](#) (see page 135)

How to Estimate a Sustained Authorization Rate

Estimating the sustained authorization rate for the application is the process of determining:

- How the total number of authorization requests fluctuate throughout your business day.
- How the authorization requests translate in to requests per second.

Complete the following steps to estimate the peak authorization rate for an application:

1. Estimate daily authorizations.
2. Estimate the sustained authorization rate.

Estimate Daily Authorizations

What is the estimated number of daily authorizations for the application?

The number of total logins (authentication load) and the number of page "hits" each authenticated user makes directly affects the number of daily authorizations (authorization load). A web page "hit" usually requires an authorization. Therefore, think of the authorization load of an application as total authorizations per day.

Note: When estimating the authorization load, we recommend that you begin with an evaluation interval of 24 hours. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding usage throughout the year.

All users requesting the same number of pages per login is unlikely, so calculating total authorizations begins with determining the percentage of logins that generate one page hit, which the following represents:

$$\text{authentication_load} * \text{percentage_of_authenticated_users} * \text{page_visits} = \text{daily_authorizations}$$

authentication_load

Represents the estimated number of daily authentications for the application.

percent_of_authenticated_users

Represents the percentage of authenticated users that visit the same number of pages after login.

page_visits

Represents the number of pages a particular set of authenticated users visits after login.

Note: A page can result in multiple GET/POST because it contains multiple objects. The total number of authorizations per page is the number of GET requests, plus the number of POST requests, minus the number of extensions the Web Agent ignores. For the purpose of this guide, each of the following examples assume that a page visit generates one GET/POST. For more information about configuring a Web Agent to allow access to specific resources types without checking policies, see the *Web Agent Configuration Guide*.

daily_authorizations

Represents the number of authorizations a particular set of authenticated users require.

Example 1: Estimate Daily Authorizations

As detailed in [Estimate Daily Authentications](#) (see page 89), the portal application has an authentication load of 88,000 logins. Twenty-five percent of which visit one page after login:

$$88,000 * 0.25 * 1 = 22,000 \text{ authorizations}$$

However, some logins generating more than one page hit is more likely.

Example 2: Estimate Daily Authorizations

The portal application has an authentication load of 88,000 logins:

- 50 percent of which visit 10 pages after login.
- 25 percent of which visit 15 pages after login.

$$88,000 * 0.5 * 10 = 440,000 \text{ authorizations}$$

$$88,000 * 0.25 * 15 = 330,000 \text{ authorizations}$$

The total authorizations per day (authorization load) is the sum of each of the authorization calculations.

Example 3: Estimate Daily Authorizations

The portal application has an authentication load of 88,000 logins:

- 25 percent of which generate one page hit after login, creating 22,000 authorizations.
- 50 percent of which generate 10 page hits after login, creating 440,000 authorizations.
- 25 percent of which generate 15 page hits after login, creating 330,000 authorizations.

Note: The percentage of authenticated users must equal 100 percent because each authenticated user generates at least one page hit.

Therefore, the authorization load for the portal application is 792,000.

The following table illustrates each of the previous examples:

Page Hits	Percent of Total Logins	Authentication Load	Authorizations
1	25	88,000	22,000
10	50	88,000	440,000

Page Hits	Percent of Total Logins	Authentication Load	Authorizations
15	25	88,000	330,000
Authorization Load			792,000

The company uses the authorization load to estimate the sustained authorization rate.

Estimate a Sustained Authorization Rate

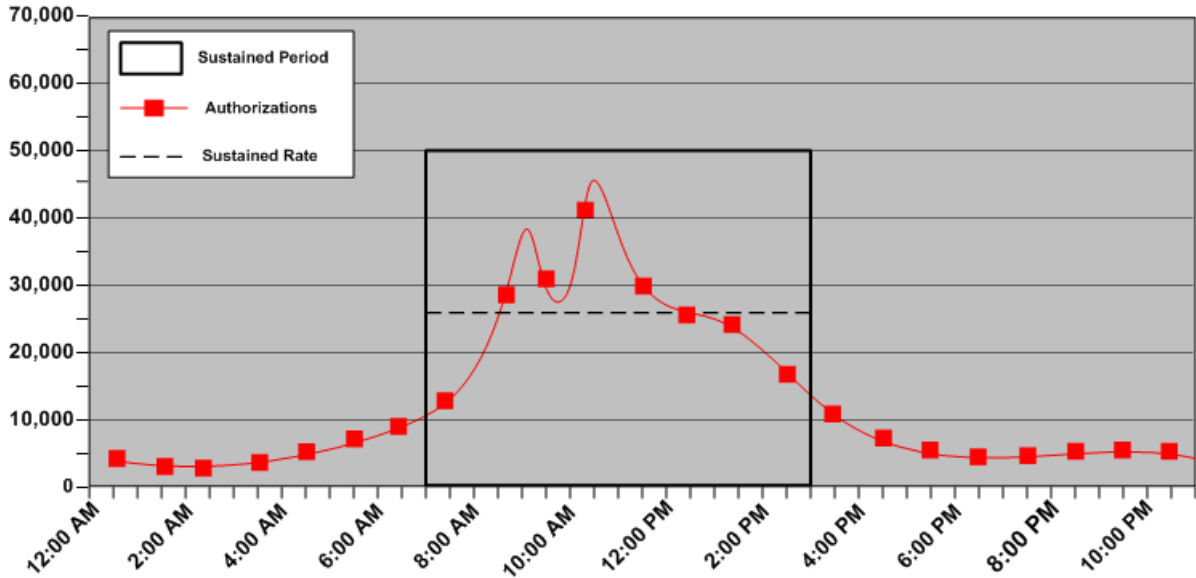
What is the sustained authorization rate for the application?

The sustained authorization rate is based on the authorization load, specifically, when and at what rate the authorizations occur. The chance that the authorization load is uniformly spread across your business day is unlikely. Rather, the rate at which requests occur fluctuates, remaining between the lowest and highest (peak) levels for a sustained period. Estimating the sustained authorization rate is the process of identifying a sustained period during which the system is servicing an average amount of authorization requests.

When estimating a sustained authorization rate, we recommend that you use the daily authorization load to determine:

- The rate at which the authorization requests occur throughout your business day.
Note: We recommend beginning with an evaluation period of 24 hours, broken down into one-hour increments. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding of usage throughout the year.
- The sustained period during which the system is servicing an average number of authorization requests.
- The approximate number of authorization requests that occur during the sustained period.

The following figure is an example of these metrics:



Identifying these metrics helps you to estimate the number of authorization requests, per second, that CA SiteMinder® must service to maintain the average rate at which authorization requests occur, which the following represents:

$$\frac{(\text{authorization_load} * \text{percentage_of_authorization_requests})}{\text{number_sustained_hours} / 3600} = \text{sustained_authorization_rate}$$

authorization_load

Represents the number of daily authorizations for the application.

percentage_of_authorization_requests

Represents the percentage of authorization requests that occur when the system is operating at sustained levels.

Example: If the authorization load is 500,000 requests, and 320,000 requests occur during the sustained period, then the value is 64 percent (0.64)

number_of_sustained_hours

Represents the number of hours in which the system is operating at the sustained level.

Note: 3,600 represents the number of seconds in an hour.

sustained_authentication_rate

Represents the number of authorization requests, per second, that CA SiteMinder® must service during the period of sustained activity.

Example: Estimate a Sustained Authorization Rate

As detailed in [Estimate Daily Authorizations](#) (see page 95), the portal application has an authorization load of 792,000. The application portal is available to customers 24 hours a day, seven days a week. Using system activity reports to break down a typical day results in the following metrics:

- The system is operating at sustained levels for approximately five hours (9:00 AM - 2:00 PM)
- During sustained levels, approximately 75,000 authorization requests occur per hour.
- Approximately 375,000 (75,000 * 5) authorization requests, or 47 percent (375,000 / 792,000) of the daily authorization load, occur during these hours.

$$(762,000 * 0.47) / 5 / 3600 = 19.90 \text{ authorizations per second}$$

The portal application has a sustained authorization rate of 19.90 authorizations per second.

Estimate a Peak Authorization Rate


What is the peak authentication rate for the application?

The peak authorization rate is based on the sustained authorization rate, specifically, when and at what rate the system is operating at peak levels. Estimating the peak authorization rate is the process of identifying when the system is servicing the highest level of authorization requests.

When estimating the peak authorization rate, we recommend using the metrics that you gathered when determining the sustained authorization rate to determine:

- The hour the system is servicing the highest number of authorization requests
- The approximate number of authorization requests that occur during this period

The following figure is an example of these metrics:



Identifying these metrics helps you to estimate the number of authentication requests, per second, that CA SiteMinder® must service to maintain the peak rate at which users authenticate, which the following represents:

$$(authorization_load * percentage_of_transactions) / number_of_hours / 3600 = peak_authorization_rate$$

Note: This rate is based on the single busiest hour. There can be times when the peak authorization rate exceeds the hourly calculation.

authorization_load

Represents the number of daily authorizations for the application.

percentage_of_transactions

Represents the percentage of transactions that occur when the system is operating at peak levels.

number_of_hours

Represents the number of hours in which the system is operating at peak levels.

peak_authorization_rate

Represents the peak authorization rate for the application.

Example: Estimate a Peak Authorization Rate

As detailed in [Estimate Daily Authorizations](#) (see page 95), the portal application has an authorization load of 792,000. System activity reports detail that during the single busiest hour of the day, 260,000 authorization requests occur. This number represents approximately 33 percent of the authorization load.

$$(792,000 * 0.33) / 1 / 3600 = 72.6 \text{ authorizations per second}$$

The portal application has a peak authentication rate of 72.6 authorizations per second.

Chapter 6: CA SiteMinder® Web Services Security Capacity Planning

This section contains the following topics:

[CA SiteMinder® Web Services Security Capacity Planning Introduced](#) (see page 103)

[Use Case: Capacity Planning](#) (see page 104)

[How to Estimate a Sustained Request Rate](#) (see page 104)

[Estimate a Peak Request Rate](#) (see page 108)

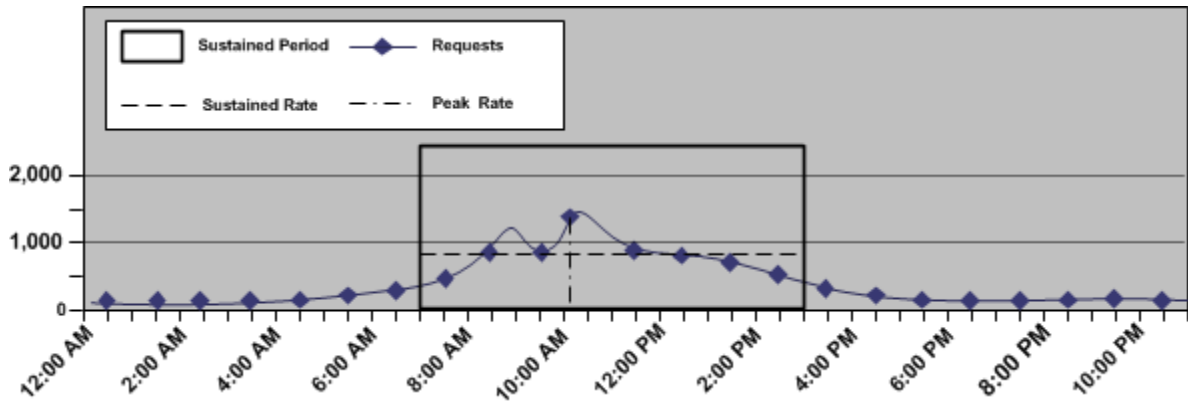
[Other Factors to Consider When Capacity Planning](#) (see page 110)

CA SiteMinder® Web Services Security Capacity Planning Introduced

Planning a CA SiteMinder® Web Services Security deployment with performance in mind is the first step to maintaining high enterprise availability and performance standards. A good approach is to estimate the number of expected requests CA SiteMinder® must handle per web service. The following are the most significant factors that influence CA SiteMinder® performance:

- Sustained request rates. The rate at which web service clients send requests to protected web service resources fluctuates throughout your business day. Some periods can generate relatively few requests, and therefore require relatively few authentications and authorizations, while others generate more. The sustained request rates represent a sustained period during which CA SiteMinder® must service an average number of authentication and authorization requests.
Note: Each web service request triggers one authentication and one authorization event.
- Peak request rates. During sustained periods of activity, web service client activity may spike. The peak request rates represent a period during which CA SiteMinder® must service the highest number of authentication and authorization requests.

The following graphic illustrates how request rates fluctuate throughout the day, are sustained for a specific period, and peak within that period:



Note: Authenticating and authorizing requests results in a number of reads from a user store. Determining sustained and peak rates helps you determine the load under which your user stores must operate to service Policy Server requests.

More information:

[Performance Tuning Introduced](#) (see page 129)

Use Case: Capacity Planning

The purpose of the following use case is to illustrate how example.com, a fulfillment service organization approaches capacity planning by modeling the usage of their order fulfillment web service. The use case is referenced throughout this chapter for examples.

The company is planning to deploy CA SiteMinder® Web Services Security to protect its web service. The company has 10,000 users in a single user store.

Some web service clients send a single status request to the inventory operation of the web service once a day while others may send as many as four batched order requests per day to the fulfillment operation.

How to Estimate a Sustained Request Rate

Estimating the sustained request rate for a web service is the process of determining:

- How the total number of requests fluctuate throughout your business day
- How the requests translate into requests per second.

Complete the following steps to estimate the sustained request rate for a web service:

1. Estimate daily requests
2. Estimate the sustained request rate

Estimate Daily Requests

What is the estimated number of daily requests for the web service?

The number of web service clients directly affect daily requests (request load). When web service clients send a request to the web service, CA SiteMinder® authenticates them. Therefore, think of the request load of the web service as the total requests per day.

Note: When determining the request load, we recommend beginning with an evaluation interval of 24 hours. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding of usage throughout the year.

All web service clients sending requests to the web service each day is unlikely, so estimating total requests begins with determining the percentage of web service clients that send a request once a day, which the following represents:

$$(total_clients * percentage_clients) * (number_of_requests) = daily_logins$$

total_clients

Represents the total number of clients with access to the application.

percentage_clients

Represents the percentage of clients that send requests the same number of times per day.

number_of_requests

Represents the number of times the particular set of clients send requests.

daily_logins

Represents the number of logins the particular set of clients creates.

Example

The company has 10,000 users, 60 percent of which send an inventory status request once a day.

$$(10,000 * 0.6) \times (1) = 6,000 \text{ logins}$$

Additionally, 30 percent of users send one order fulfillment request per day, 20 percent of users send two order fulfillment requests a day, 10 percent send three order fulfillment requests a day, and 10 percent send four order fulfillment requests a day.

$$(10,000 * 0.3) \times (1) = 3,000 \text{ logins}$$

$$(10,000 * 0.2) \times (2) = 4,000 \text{ logins}$$

$$(10,000 * 0.1) \times (3) = 3,000 \text{ logins}$$

$$(10,000 * 0.1) \times (4) = 4,000 \text{ logins}$$

The total requests per day are the sum of each of the request calculations. The request load for the fulfillment web service is therefore 20,000 logins.

Note: The percentage of clients making requests is not necessarily equal to 100 percent because not all clients will necessarily send a request to the service each day.

The company uses the request load to estimate the sustained request rate.

Estimate a Sustained Request Rate


What is the sustained request rate for the web service?

The sustained request rate is based on the request load. Specifically, when and at what rate the requests occur. The chance that the request load is uniformly spread across your business day is unlikely. Rather, the rate at which requests occur fluctuates, remaining between the lowest and highest (peak) levels for a sustained period. Estimating the sustained request rate is the process of identifying a sustained period during which the system is servicing an average number of requests.

When estimating a sustained request rate, we recommend using the daily request load to determine:

- The rate at which the requests occur throughout your business day.
Note: We recommend beginning with an evaluation period of 24 hours, broken down into one-hour increments. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding of usage throughout the year.
- The sustained period during which the system is servicing an average number of requests.
- The approximate number of requests that occur during the sustained period.

The following figure is an example of these metrics:



Identifying these metrics helps you to estimate the number of requests, per second, that CA SiteMinder® must service to maintain the average rate at which users authenticate, which the following represents:

$$\frac{(\text{request_load} * \text{percentage_of_requests})}{\text{number_of_sustained_hours} / 3600} = \text{sustained_request_rate}$$

request_load

Represents the number of daily requests for the application.

percentage_of_requests

Represents the percentage of requests that occur when the system is operating at sustained levels.

Example: If the request load is 5,000 logins, and 3,000 logins occur during the sustained period, then the value is 64percent (0.64)

number_of_sustained_hours

Represents the number of hours in which the system is operating at the sustained level.

Note: 3,600 represents the number of seconds in an hour.

sustained_request_rate

Represents the number of requests, per second, that CA SiteMinder® must service during the period of sustained activity.

Example: Estimate the Sustained Request Rate

The company has determined that their web service has a request load of 2,000 logins. The web service is available to customers 24 hours a day, seven days a week. Using system activity reports to break down a typical day results in the following metrics:

- The system is operating at sustained levels for approximately five hours (9:00 AM - 2:00 PM).
- During sustained levels, approximately 2,500 requests occur per hour.
- Approximately 1,250 (250 * 5) requests, or 62.5 percent (1,250 / 2,000) of the daily request load, occur during these hours.

$(2,000 * 0.625) / 5 / 3600 = 0.0694$ requests per second.

The fulfillment web service has a sustained request rate of 0.694 requests per second.

Estimate a Peak Request Rate


What is the peak request rate for the web service?

The peak request rate is based on the sustained request rate, specifically, when and at what rate the system is operating at peak levels. Estimating the peak request rate is the process of identifying when the system is servicing the highest level of requests.

When estimating the peak request rate, we recommend that you use the metrics you gathered when determining the sustained request rate to determine:

- The hour when the system is servicing the highest number of requests
- The approximate number of requests that occur during this period

The following figure is an example of these metrics:



Identifying these metrics helps you to estimate the number of requests, per second, that CA SiteMinder® must service to maintain the peak rate at which web service clients authenticate, which the following represents:

$$\frac{(\text{request_load} \times \text{percentage_of_transactions})}{\text{number_of_hours} / 3600} = \text{peak_request_rate}$$

Note: This rate is based on the single busiest hour. There can be periods when the peak request rate exceeds the hourly calculation.

request_load

Represents the number of daily requests for the web service.

percentage_of_transactions

Represents the percentage of transactions that occur when the system is operating at peak levels.

number_of_hours

Represents the number of hours in which the system operates at peak levels.

Note: 3,600 represents the number of seconds in an hour.

peak_request_rate

Represents the peak request rate for the application.

Example: Estimate the Peak Request Rate

The company has determined that their web service has a daily request load of 8,800. System activity reports detail that during the single busiest hour of the day 1,800 requests occur. This number represents approximately 20 percent of the request load:

$$1,800 / 1 / 3600 = 0.5 \text{ requests per second}$$

The fulfillment web service has a peak request rate of five requests per second.

Note: This example is based on the single busiest hour. There can be periods when the peak request rate during the hour exceeds five requests per second.

More information:

[Increase the Amount of Available Sockets for the Agent](#) (see page 135)

Other Factors to Consider When Capacity Planning

Although request rates are the most significant factors in CA SiteMinder® Web Services Security capacity planning decisions other factors can also influence CA SiteMinder® performance, in particular the authentication schemes used to protect your web services.

Also consider performance tuning and network bandwidth in your capacity planning process.

Chapter 7: Configuration Considerations

This section contains the following topics:

[Security Zones](#) (see page 111)

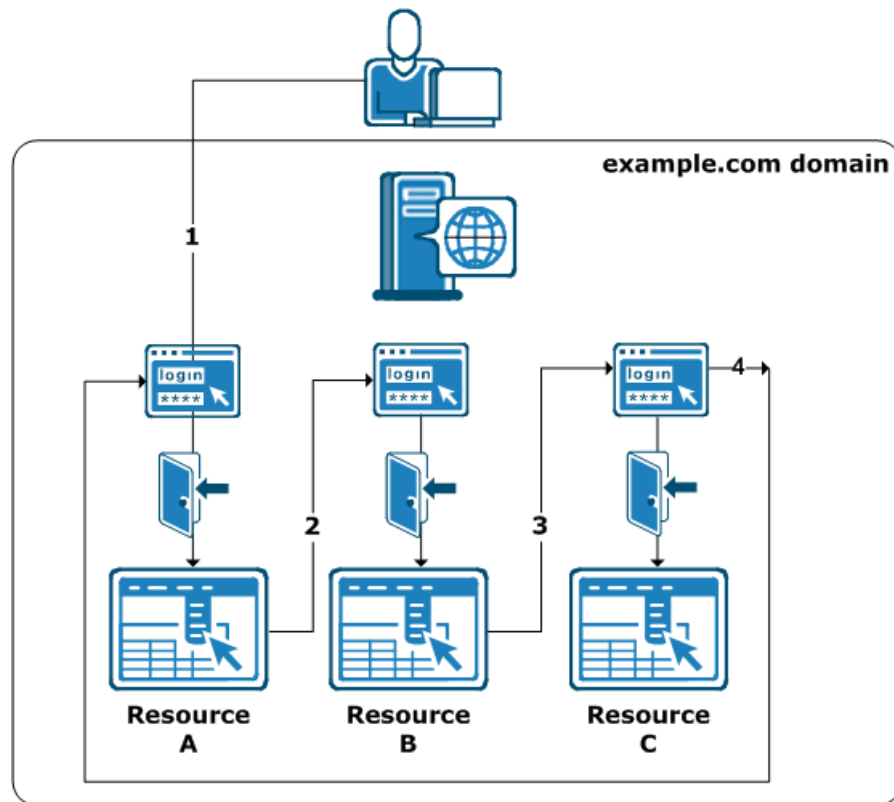
[Multiple Data Centers](#) (see page 113)

[Authentication and a Centralized Login Server](#) (see page 121)

Security Zones

Security Zones are groups of resources in a single cookie domain that a CA SiteMinder® Web Agent protects. Users authenticate once, and can then access other resources in the zones (for which they are authorized) without being rechallenged.


Without Security Zones, users could possibly be challenged each time they access a protected resource in the same cookie domain; even if they have previously been authenticated by CA SiteMinder® for another resource in the cookie domain. The following illustration shows an example:



Consider implementing Security Zones in the following situations:

- You have several resources in a cookie domain, but you want to apply different access restrictions to those resources.
- You want to enable SSO between different resources in the same cookie domain.
- You want to create groups of resources that span several cookie domains and allow SSO between them.
- You have a large organization with a single cookie domain, and you use multiple instances of CA SiteMinder® to protect resources in your organization. Security Zones let you separate the resources to control access within the single cookie domain. Without Security Zones, the cookies used by one CA SiteMinder® instance could possibly overwrite the cookies of another CA SiteMinder® instance (cookie stomping) because the cookie domain name is the same for both instances.

The following illustration shows how Security Zones can be used so that only a single log in allows a user access to resources in Security Zones 1 and 3, but prevents access to unauthorized resources in Security Zone 2:



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

Multiple Data Centers

CA SiteMinder® treats a global deployment the same as multiple data centers in the same continent. As such, factors outside of CA SiteMinder® affect the performance of a multi–data center deployment. The following key factors include:

- Network latency
- Resiliency

We recommend that you consider the following outside factors as you plan for a multi–data center deployment:

- Network infrastructure
- Application locations
- User locations
- User store vendors and their restrictions, such as the number of masters allowed

Best Practices

Consider the following when configuring data centers:

- Collocating the following components in each data center helps to reduce the effect network latency and resiliency has on CA SiteMinder® performance:
 - CA SiteMinder® Agents
 - Policy Servers
 - User stores

Note: If a CA SiteMinder® feature, such as Password Services, requires a write–enabled store, we recommend having a write–enabled store in each data center.
- If all components cannot be in the same data center, we recommend at least collocating Policy Servers and user stores in the same data center.

Architectural Considerations

Consider the following architectural factors when planning for a CA SiteMinder® data center:

- CA SiteMinder® Password Services attempts to perform an LDAP write to the user account on every authentication.
Note: For more information about Password Services, see the *Policy Server Configuration Guide*.
- CA SiteMinder® follows LDAP write referrals when communicating with a read-only consumer directory.
- If you deploy a master policy store with replicated versions, consider using a local host file on the Policy Server host system (LDAP) or the ODBC data source to point Policy Servers to the local policy store. Using this method lets all Policy Servers share the same policy store and avoids the latency that can occur when all Policy Servers must communicate with the policy store over the wide area network (WAN).
- If you deploy master/consumer user stores, consider using a local host file on the Policy Server host system (LDAP) or the ODBC data source name (DSN) to point Policy Servers to the local consumer. Using this method lets all Policy Servers read the same user store and avoids the latency that can occur when all Policy Servers must read user account information over the WAN.

Example: Local Host Files Pointing Policy Servers to the Local Consumer User Store

Two geographically separated data centers include Policy Servers pointing to a consumer user store named myusers.

- The local consumer in data center one is available at 111.11.111.1
- The local consumer in data center two is available at 222.22.222.2

To point Policy Server to the local consumer

1. From the Policy Server host systems in data center one, use a local host file to map myusers to 111.11.111.1.
2. From the Policy Server host systems in data center two, use a local host file to map myusers to 222.22.222.2.

Multiple Data Center Use Cases

The purpose of the following use cases is to get you thinking about your CA SiteMinder® data centers in terms of network latency and resiliency. The use cases begin with a simple deployment and progress into more complex scenarios.

These use cases are intended to identify techniques that you can use as part of a global architecture and are not intended as a final architecture. Extrapolate the necessary infrastructure from these cases to configure data centers that best meet the needs of your organization.

All Components in One Data Center

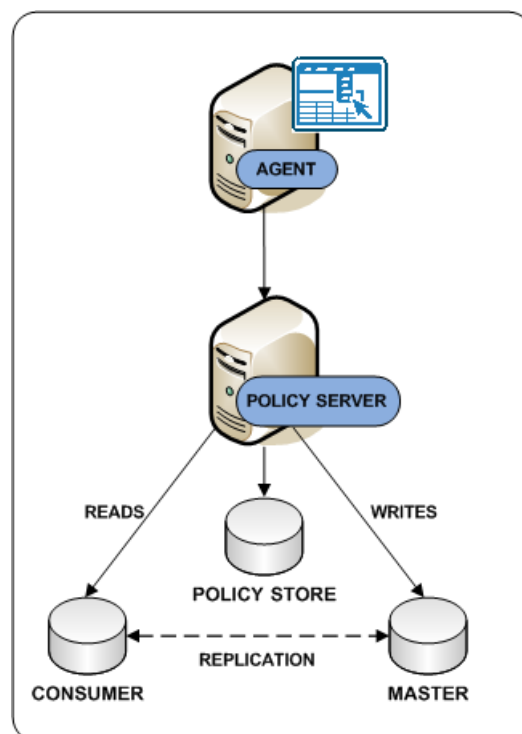
The simplest deployment includes all required CA SiteMinder® components in a single data center.

The following diagram illustrates:

- All applications in a single data center.
- A Policy Server writing to a master user store. CA SiteMinder® Password Services attempts to perform an LDAP write to the user account on every authentication.

Important! For more information about multi-mastered LDAP user store support limitations, see the *Policy Server Release Notes*.

- CA SiteMinder® reading a consumer user store.



Consider the following:

- Although not illustrated, CA SiteMinder® supports database clusters that are configured for write and read-only transactions.
- You can configure multiple components in a data center for operational continuity, redundancy, and high availability.

More information:

[Redundancy and High Availability](#) (see page 32)

All Components in Multiple Data Centers

You extend the CA SiteMinder® environment by deploying multiple data centers. The following factors can influence your decision to implement multiple data centers:


- The network infrastructure
- The location of applications
- The location of users

The following diagram illustrates applications in multiple data centers:

- Each data center uses its own policy store.
- Data center one contains the primary policy store. Data center two contains the replicated version, as the dotted line details.

Note: Every Policy Server in the deployment must share a common view into the same policy store. For more information about policy store redundancy, see [Policy Server to Policy Store Communication](#) (see page 41).

- Each data center uses its own master/consumer user stores.
- Important!** For more information about multi-mastered LDAP user store support limitations, see the *Policy Server Release Notes*.
- A centralized replicated session store enables single sign-on between all applications.



More information:

[Policy Server to Policy Store Communication](#) (see page 41)

[All Components in One Data Center](#) (see page 115)


CA SiteMinder® Agent Communicating Across a Data Center

If all components cannot be in the same data center, we recommend collocating Policy Servers and user stores in the same data center.

The following diagram illustrates applications in multiple data centers:

- Data center one only contains a web server with a CA SiteMinder® agent. The agent communicates across the wide area network to a Policy Server in data center two.
- Data centers two and three share a common view into the policy store through a [master/replicated policy store](#) (see page 41).

- Data centers two and three use their own [master/consumer user stores](#) (see page 115).
- A centralized replicated session store enables single sign-on between all applications.



More information:

[Policy Server to Policy Store Communication](#) (see page 41)

Policy Server Communicating Across a Data Center

If all components cannot be in the same data center, we recommend collocating Policy Servers and user stores in the same data center.

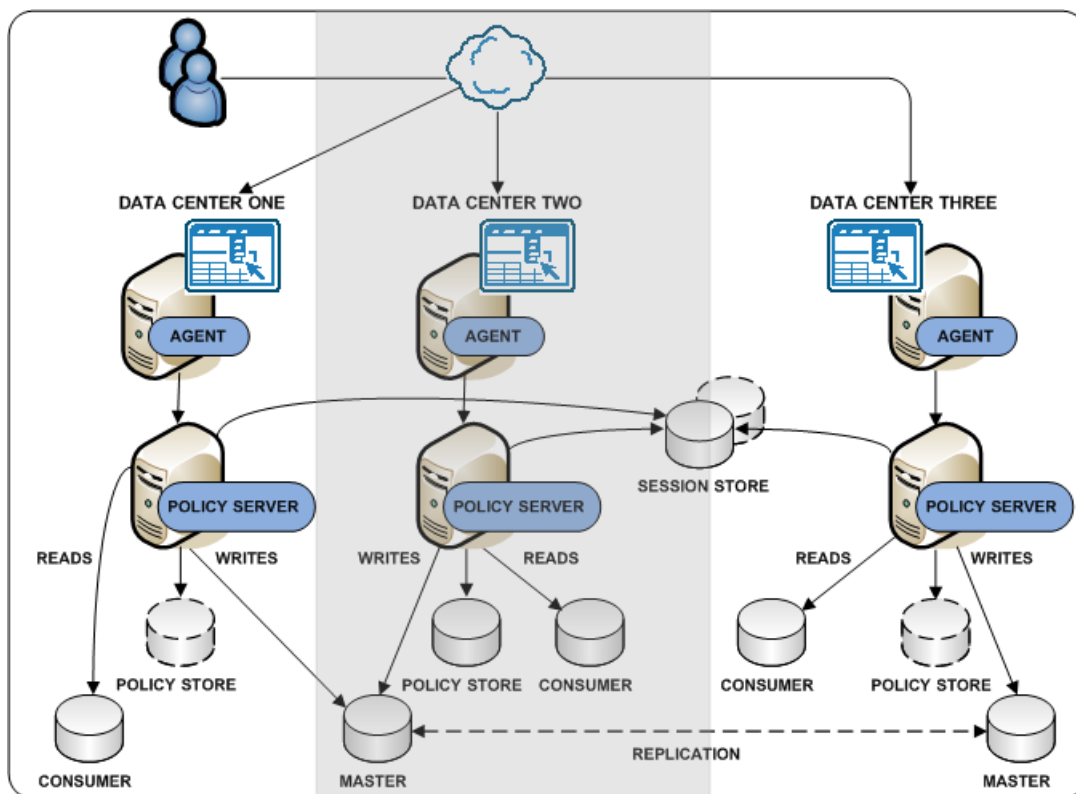
The following diagram illustrates applications in multiple data centers:

- Data center one contains an agent and a Policy Server. The Policy Server only communicates across the wide area network to perform LDAP writes to the master user store in data center two.

Important! We do not recommend configuring a Policy Server to communicate across the wide area network to perform LDAP reads and writes.

- All data centers share a common view into the policy store through a [master/replicated policy store](#) (see page 41).
- A centralized replicated session store enables single sign-on between all applications.
- Data centers two and three use their own [master/consumer user stores](#) (see page 115).

Important! For more information about multi-mastered LDAP user store support limitations, see the *Policy Server Release Notes*.



More information:

[Policy Server to Policy Store Communication](#) (see page 41)

[Master Policy Store](#) (see page 41)

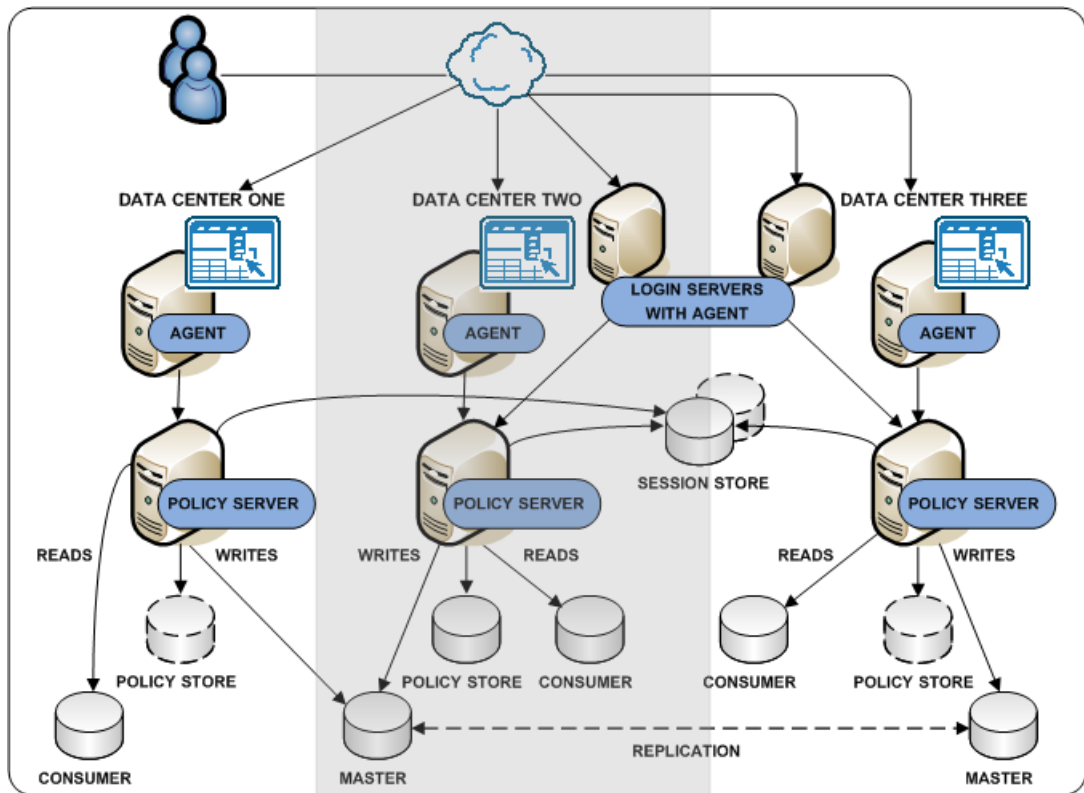
[All Components in One Data Center](#) (see page 115)

Login Server Controlling User Store Writes

The location of LDAP writable masters can constrain a CA SiteMinder® deployment. To eliminate requirements for writable masters in each data center, consider using one or more centralized login servers.

The following diagram illustrates a multiple data center deployment:

- The Policy Server in data center one is [communicating across the WAN to perform an LDAP write](#) (see page 119).
- The remaining data centers include [all components](#) (see page 116).
- A login server is deployed in data centers two and three.



When a user requests access to a protected URL in data center one:

1. The agent redirects the request to the login server in data center two. The redirect is based on the authentication scheme that is protecting the resource.

Note: For more information about authentication schemes, see the *Policy Server Configuration Guide*.

2. The Policy Server in data center two authenticates the user and writes to the master user store.
3. The Policy Server creates a CA SiteMinder® session ticket and passes it back to the original protected URL.

Note: For more information about user sessions, see the *Policy Server Configuration Guide*.

4. The agent places the CA SiteMinder® session ticket into a cookie. The agent uses the cookie to handle subsequent authentication and authorization requests in the data center, until one of the following occurs:
 - The user requests another resource that requires re-authentication.
 - The session expires.

Authentication and a Centralized Login Server

A CA SiteMinder® deployment typically includes applications for which different authentication (login) requirements exist. These requirements can result in numerous login pages that the individual application owners must manage. Managing these login pages locally can introduce inconsistencies, such as page design and the presentation of error messages, that can affect the overall authentication experience.

We recommend managing login pages centrally to help:

- Create consistency across your applications. If a single CA SiteMinder® team owns all login pages, the team can implement them consistently and manage them easier.
- Minimize the number of login pages. Minimizing the number of entry points into applications creates the impression that users are logging into a centralized infrastructure, rather than individual applications.

Consider the following when configuring login pages:

- Identify applications that share the same authentication requirements and reuse the same login page.
- Use a centralized login server to host all login pages
- Configure login pages to inform users when:
 - They have failed to provide valid credentials.
 - Too many attempts have resulted in a failed authentication.

Centralize Login Pages

Application login requirements can range from basic user name/password authentication to forms-based authentication to digital certificates. If possible, we recommend:

- Managing all login pages from a central login server to avoid duplication on every web application.
- Managing all other system-wide resources, such as password services pages, error pages, and terms and conditions pages from a central server.

Note: For more information about authentication schemes, see the *Policy Server Configuration Guide*.

Managing login pages centrally is the process of identifying applications that share the same login requirements. Consider the following when configuring authentication:

- Try to avoid creating separate login pages for each application. As CA SiteMinder® adoption increases, managing a login page for every application may not be sustainable.
- Identify applications that share the same authentication requirements. If possible, use a single login page as an entry point into these applications.

Use a table similar to the following to group applications by authentication requirements:

Auth Scheme Name	Type	Login Page Server	Login Page URL

Example: Grouping applications by authentication requirements

A CA SiteMinder® environment protects ten applications:

- Five of the applications require forms-based authentication.
- Three of the applications require Windows-based authentication.
- Two of the applications require basic user name/password authentication.

By identifying applications that share the same authentication requirements, three login pages replace the need for eight, as detailed by the following table:

Auth Scheme Name	Type	Login Page Server	Login Page URL
Auth1	Forms	login.acme.com	/login.asp
Auth2	Windows	login.acme.com	/smgetcrd.ntc
Auth3	Basic	login.acme.com	n/a

Best Practices

Consider the following when configuring login pages:

- Display an error message when a user fails to authenticate properly.
- Redirect users to a page that displays a message that the number of login attempts has been exceeded.
- We recommend using forms-based authentication to redirect users. If you are unable to use forms-based authentication, you can use the CA SiteMinder® OnAuthAttempt and OnAuthReject responses to redirect users.

Note: For more information about responses, see the *Policy Server Configuration Guide*.

- If you configure forms-based authentication, consider creating a dynamic page, such as login.asp, to create a tighter integration with your existing infrastructure.
- If creating a dynamic page is not possible, use the sample login FCC file (login.fcc) that is included as part of the Web Agent installation to configure a login FCC file. The default location for the sample file is `web_agent_home\samples_default\forms`. The forms directory is the default location for files that the Forms Credential Collector (FCC) processes.

web_agent_home

Specifies the Web Agent installation path.

Note: For more information about the login FCC as it applies to forms-based authentication, see the *Policy Server Configuration Guide*. For more information about configuring the login FCC with a Web Agent and how the FCC process requests, see the *Web Agent Configuration Guide*.

- We recommend creating a separate directory on the Web Agent host system for all login pages. Using a location other than the forms directory helps to prevent the sample files from being accidentally overwritten.
- Display a custom logoff page after a user logs out successfully.

Note: For more information about configuring a logoff page, see the *Web Agent Configuration Guide*.

Login Page Use Cases

The purpose of the following use cases is to get you thinking about configuring CA SiteMinder® authentication.

These use cases reflect best practices and are intended to identify techniques that you can use as part of a global architecture. These use cases are not intended as a final architecture. Extrapolate the necessary infrastructure from these cases to configure login pages that best meet the needs of your organization.

Stand-Alone Login Page

In this use case, CA SiteMinder® directs users to a stand-alone login page when they request a protected resource. Specifically:

- A dynamic login page (login.asp) is deployed to the Web Agent host system.
- The dynamic login page is coded to:
 - Post to a login FCC file (login.fcc).
 - Display an error message when the SMTRYNO cookie is present in the web browser of the user.

Note: For more information about the SMTRYNO cookie, see the *Web Agent Configuration Guide*.


- The login FCC file is configured with an @directive (@smretries) to redirect users to a failed authentication page (login.unauth) after two failed authentication attempts.

Note: For more information about configuring an FCC file with @directives, see the *Policy Server Configuration Guide*.

- A CA SiteMinder® administrator has configured a form-based authentication scheme named Auth1. The target of Auth1 is login.asp.

Note: For more information about configuring authentication schemes, see the *Policy Server Configuration Guide*.

The following diagram illustrates the authentication process for this use case:



1. A user requests a protected resource.
2. The Web Agent contacts the Policy Server, which determines that the resource is protected.
3. The Web Agent redirects the user request to login.asp.
4. The user submits invalid credentials. The credentials are posted to the login.fcc file and processed by the FCC.
5. The FCC forwards the credentials to the Policy Server.
6. The Policy Server determines that the credentials are invalid and notifies the FCC.
7. The FCC inserts the SMTRYNO cookie into the web browser of the user and redirects the user to the login page.
8. The login page refreshes with an error message. The error message states that invalid credentials were supplied and to try again.
9. The user submits invalid credentials. The credentials are posted to the login.fcc file and processed by the FCC.
10. The FCC forwards the credentials to the Policy Server.

11. The Policy Server determines that the credentials continue to be invalid and notifies the FCC.
12. The user has exceeded the maximum number of failed authentication attempts and is redirected to a page that displays a failed authentication message.

Embedded Form on a Web Portal

In this use case, a form is embedded on a web portal home page. Users enter credentials in the form and are redirected to the protected resource upon authentication. Specifically:


- A web portal home page (portal.asp) includes an embedded form that prompts users for credentials. The home page:
 - Contains a target variable that points to the protected resource.
 - Posts to a login FCC file (login.fcc).
- A stand-alone login page (login.asp) is deployed to the Web Agent host system. If users try to access the protected resource directly, this page prompts users for credentials. The login page:
 - Posts to the login FCC file.
 - Displays an error message when the SMTRYNO cookie is present in the web browser of the user.

Note: For more information about the SMTRYNO cookie, see the *Web Agent Configuration Guide*.
- The login FCC file is configured with an @directive (@smretries) to redirect users to a failed authentication page (login.unauth) after two failed authentication attempts.

Note: For more information about configuring an FCC file with @directives, see the *Policy Server Configuration Guide*.
- A CA SiteMinder® administrator has configured a form-based authentication scheme named Auth1. The target of Auth1 is login.asp.

Note: For more information about configuring authentication schemes, see the *Policy Server Configuration Guide*.

The following diagram illustrates the authentication process for this use case:



1. A user navigates to the web portal home page.
 2. The Web Agent contacts the Policy Server, which determines that the resource is unprotected.
 3. The user submits invalid credentials. The credentials are posted to the login.fcc file and processed by the FCC.
 4. The FCC forwards the credentials to the Policy Server.
 5. The Policy Server determines that the credentials are invalid and notifies the FCC.
 6. The FCC inserts the SMTRYNO cookie into the web browser of the user and redirects the user to the login page. The login page appears with an error message. The error message states that invalid credentials were supplied and to try again.
- Note:** Although not illustrated, if the user accessed the protected resource directly, the login page would appear without an error message because the web browser would not contain the SMTRYNO cookie.
7. The user submits invalid credentials. The credentials are posted to the login.fcc file and processed by the FCC.
 8. The FCC forwards the credentials to the Policy Server.

9. The Policy Sever determines that the credentials continue to be invalid and notifies the FCC.
10. The user has exceeded the maximum number of failed authentication attempts and is redirected to a page that displays a failed authentication message.

Chapter 8: Performance Tuning

This section contains the following topics:

[Performance Tuning Introduced](#) (see page 129)

[Performance Tuning Roadmap](#) (see page 130)

[Web Tier Performance](#) (see page 131)

[Application Tier Performance](#) (see page 149)

[Data Tier Performance](#) (see page 165)

[Periodic Maintenance Tasks](#) (see page 188)

Performance Tuning Introduced

The Policy Server evaluates and enforces access control policies by servicing three basic requests:

- **IsProtected**—is the requested resource protected?
- **IsAuthenticated**—did the user requesting the resource present credentials to establish an identity?
- **IsAuthorized**—is the authenticated user authorized to view the protected resource?

Servicing each of these requests creates transactions between CA SiteMinder® components. CA SiteMinder® performance tuning is the iterative process of increasing throughput and reducing latency by:


- Understanding where and when these transactions occur
- Identifying the CA SiteMinder® settings and features that affect performance
- Using third-party and CA SiteMinder® tools to measure performance and identify infrastructure bottlenecks

A good strategy is to examine performance factors in the web, application, and data tiers.

Note: CA SiteMinder® is middleware and is not deployed independently. The following sections focus on tuning CA SiteMinder® components in the Web and Application tiers, but not how to tune the actual Web, Application, or Data tiers themselves. See your vendor–specific documentation for more information about tuning the web servers, directory servers, and databases in your environment.

Performance Tuning Roadmap

Performance tuning is an iterative process, and as such, it is important to address the Web, Application, and Data tiers on an individual basis to understand how each can affect overall performance. You can often achieve better performance by changing configuration settings in CA SiteMinder® Agents, Policy Servers, or the CA SiteMinder® policy objects themselves. The following diagram represents a standard deployment and details the individual components that are central to performance.



1. The types of web and applications servers deployed in your environment can affect how a CA SiteMinder® Agent and Policy Server communicate.
2. The number of available sockets can affect the efficiency in which an agent and Policy Server communicate.
3. CA SiteMinder® policy design can affect the efficiency in which the Policy Server services authentication and authorization requests.
4. The Policy Server performs a series of services to authenticate and authorize users. These services result in number of reads and writes, collectively known as requests, to a user directory. A contributing factor to CA SiteMinder® performance is determining whether your user directories can handle this workload during sustained and peak periods of operation.
5. The user directory itself can affect CA SiteMinder® performance.

More information:

[User Store Capacity Planning](#) (see page 168)

[CA SiteMinder® Policy Design and Performance](#) (see page 149)

[Server Performance](#) (see page 132)

[Reduce Traffic between Your Agents and the Policy Server](#) (see page 137)

[Data Tier Guidelines](#) (see page 165)

[Web Tier Socket Usage](#) (see page 134)

Web Tier Performance

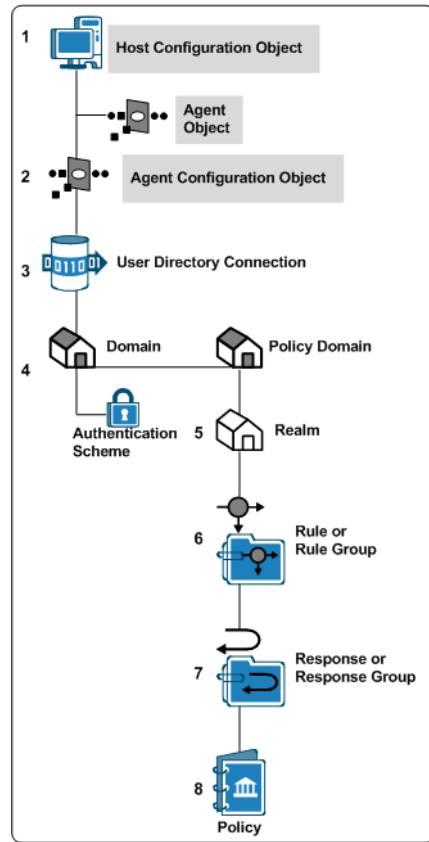
When a CA SiteMinder® Agent intercepts a request sent to a web or application server, the agent makes the following calls to the CA SiteMinder® Policy Server:

- isProtected
- isAuthenticated
- isAuthorized

Each of the previous calls generates traffic between the agent in the Web Tier, and the Policy Server in the Application Tier. The following settings can help you adjust the performance of the Web Tier:

- Change the timeout interval for Policy Server requests.
- Change the number of sockets that are available for an agent to use for Policy Server connections.
- Use the agent caches to reduce the number of calls an agent makes to the Policy Server.

The shaded items shown in the following illustration contain settings that affect the performance of your Web Tier:



Server Performance

CA SiteMinder® agents can be installed on a number of supported web and application servers. The performance of the hosting server determines the performance of the CA SiteMinder® web tier. The following items affect how your web server performs with CA SiteMinder®:

- The processor speed of your web server
- The amount of memory in your web server

CA SiteMinder® Agent Performance

The following factors influence CA SiteMinder® Web Agent performance:

- Web or application server CPU and available memory.
- Policy Server latency (how quickly the Policy Server responds to agent requests).

If too few web servers are available to handle the number of requests, the following types of problems can occur:

- Delays of or inability of users to log in.
- Delays in users receiving the resources they requested.
- CPU usage at or near maximum capacity.

Anticipating the number of requests serviced by each web or application server during peak periods can help you determine the ideal number of web servers for your CA SiteMinder® environment.

Use any of the following methods to estimate the number of requests:

- Complete a capacity planning effort.
- Generate a CA SiteMinder® Activity report for each agent in your environment.
- Generate a performance report for your web server.

Note: For more information, see the documentation provided by your web server vendor.

More information:

[Estimate a Peak Authentication Rate](#) (see page 93)

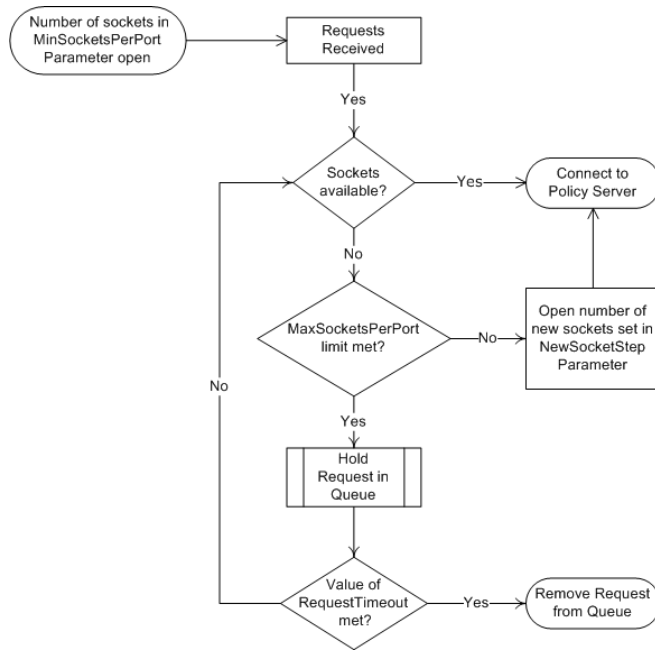
[Estimate a Peak Authorization Rate](#) (see page 99)

Web Tier Socket Usage

When a CA SiteMinder® Agent starts, it opens the number of sockets specified by the MinSocketsPerPort parameter in the Host Configuration Object on the Policy Server. If more requests are received, the Agent adds a specified number of new sockets to the connection pool until the maximum number of sockets is reached. When all sockets used, any additional requests (up to 300) are held in a queue, until one of the following events occurs:

- A socket pair becomes available and the request is sent to the Policy Server.
- The request times out, and the user must try again to access the resource.

The following illustration describes this process:



The Host Configuration Object on the Policy Server contains the parameters that control the number of sockets used.

Increase Request Timeout Interval during Heavy Loads

Consider increasing the length of time that requests from CA SiteMinder® Agents are held in the Policy Server queue if your network has any of the following conditions:

- Heavy traffic
- Slow connections

The RequestTimeout parameter in the Host Configuration Object on the Policy Server controls how long the Agents wait for responses from the Policy Server. If the interval is too short, the requests time out and the user receives an error message.

Note: For more information, see the *CA SiteMinder® Policy Server Configuration Guide*.

Increase the Amount of Available Sockets for the Agent

If your capacity planning estimates reveal that the number of user requests per CA SiteMinder® agent exceeds 60 at any given moment (20 requests in process and 40 in the queue), increase the value of the MaxSocketsPerPort parameter.

After increasing the value of the MaxSocketsPerPort parameter in the Administrative UI, verify that the Max Connections setting in the Policy Server Management Console is high enough to accommodate all the Agent processes in your CA SiteMinder® environment. This setting determines the maximum number of connections available to a specific Policy Server.

Note: For multiprocess web servers (such as an Apache-based server in pre-fork mode), you can reduce this number of sockets to one. Because each process uses only a single thread to communicate with the CA SiteMinder® Policy Server, only one socket is required.

More information:

[Estimate a Peak Request Rate](#) (see page 108)

Increase NewSocketStep Setting

When the CA SiteMinder® Agent requires additional sockets from the connection pool during peak loads, the NewSocketStep parameter determines the number of sockets obtained each time.

If the value of the NewSocketStep parameter is set too low, response time during peak periods suffers because the Agent takes extra time to create socket connections.

To help avoid slow response times, use your capacity planning estimates to determine how many requests your Agents handle, and then increase the value of the NewSocketStep parameter accordingly.

The ideal number for this parameter is one large enough to prevent the Agent from spending too much time creating sockets for requests as the load on the web or application server increases.

We recommend experimenting with different settings until you find what works best in your CA SiteMinder® environment.

Note: For multiprocess web servers (such as an Apache-based server in pre-fork mode), you can reduce this number of sockets to one. Because each process uses only a single thread to communicate with the CA SiteMinder® Policy Server, only one socket is required.

More information:

[Estimate a Peak Authentication Rate](#) (see page 93)

[Estimate a Peak Request Rate](#) (see page 108)

Minimum Sockets per Port Setting

When a CA SiteMinder® Agent starts, it opens the number of sockets specified by the MinSocketsPerPort parameter in the Host Configuration Object on the Policy Server. These sockets maintain a constant connection to the Policy Server.

For most types of web and application servers (including Apache-based servers in worker mode), we recommend leaving this parameter at its default setting. Increasing this parameter occupies additional sockets unnecessarily by leaving them open even when the Agent is not receiving any requests for resources.

Note: For multiprocess web servers (such as an Apache-based server in pre-fork mode), you can reduce this number of sockets to one. Because each process uses only a single thread to communicate with the CA SiteMinder® Policy Server, only one socket is required.

Examples of Relationships between Socket Settings

The type web server that you are using determines the relationship between the socket-allocation parameters on the Policy Server.

Since single-process multiple-threaded web servers operate differently than multiple-process single-threaded web servers, the allocation of sockets on your Policy Servers differs for each type of web server.

Note: See the documentation from the vendor of your web server to determine its type.

The following illustration describes the formula for single-process, multiple-threaded web servers:

Socket Settings Formula for Single-Process/Multiple-Threaded Web Servers

$$\begin{array}{|c|} \hline \text{MaxSocketsPerPort} \\ \hline 20 \\ \hline \end{array} \times \begin{array}{|c|} \hline \text{Number of Ports on which} \\ \text{Service Listens} \\ \hline 1 \\ \hline \end{array} = \begin{array}{|c|} \hline \text{MaxSocketsPerPort} \\ \hline 20 \\ \hline \end{array}$$

The following illustration describes the formula for multiple-process, single-threaded web servers:

Socket Settings Formula for Multiple-Process/Single-Threaded Web Servers

$$\begin{array}{|c|} \hline \text{Max Processes} \\ \hline 150 \\ \hline \end{array} \times \begin{array}{|c|} \hline \text{MinSocketsPerPort} \\ \hline 1 \\ \hline \end{array} \times \begin{array}{|c|} \hline \text{Number of Ports on which} \\ \text{Service Listens} \\ \hline 1 \\ \hline \end{array} = \begin{array}{|c|} \hline \text{MaxSocketsPerPort} \\ \hline 150 \\ \hline \end{array}$$

The following illustration describes the formula for multiple-process, multiple-threaded web servers:

Socket Settings Formula for Multiple-Process/Multiple-Threaded Web Servers

$$\begin{array}{|c|} \hline \text{MaxSocketsPerPort} \\ \hline 20 \\ \hline \end{array} \times \begin{array}{|c|} \hline \text{Number of Ports on which} \\ \text{Service Listens} \\ \hline 1 \\ \hline \end{array} \times \begin{array}{|c|} \hline \text{Max Processes} \\ \hline 150 \\ \hline \end{array} = \begin{array}{|c|} \hline \text{MaxSockets} \\ \hline 3000 \\ \hline \end{array}$$

Use any of the previous formulas as guides when adjusting your socket settings.

Reduce Traffic between Your Agents and the Policy Server

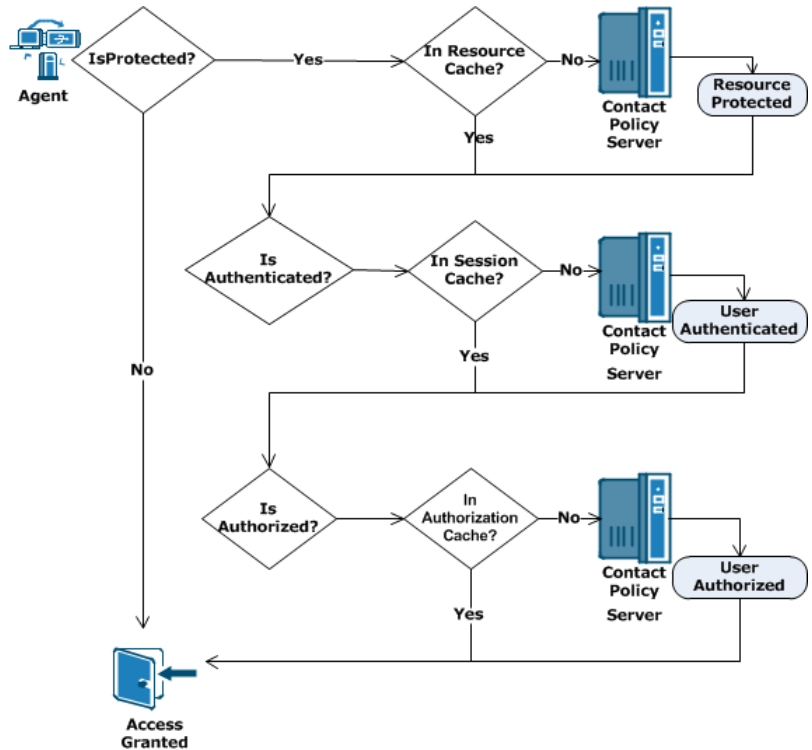
CA SiteMinder® Agents multiple caches and configuration parameters that you can use together to reduce the amount of traffic between your Agents and Policy Servers. Generally, these settings are most efficient in CA SiteMinder® environments where the policies and URIs usually remain static.

How Agent Caches Work

The CA SiteMinder® Agent searches the following caches for the information it needs before contacting the CA SiteMinder® Policy Server:

- Resource Cache
- Session Cache
- Authorization Cache

Because retrieving information from a cache is quicker than contacting the Policy Server, performance improves. The following illustration describes this process:



Resource Cache

Each CA SiteMinder® Agent uses a resource cache to store the following information it receives from the Policy Server temporarily:

- Whether a resource is protected
- Any additional response attributes included in the policy

The Agent searches the resource cache to determine if a resource is protected before contacting the Policy Server. If the resource exists in the cache, traffic to the Policy Server is reduced because the Agent does not make an IsProtected call to the Policy Server.

Two Agent configuration parameters affect the resource cache. Consider the following as you plan your CA SiteMinder® deployment:

Resource Cache Timeout

We recommend basing the timeout interval of the Agent resource cache on the results of your capacity planning tests. A timeout interval that is too small limits the effectiveness of the resource cache. The value of the ResourceCacheTimeout parameter in your Agent configuration determines the timeout interval of the resource cache.

Resource Cache Size

We recommend using a resource cache that is 10 percent larger than the largest number of URIs that you expect users to request. If you are protecting an application that uses dynamic URLs (such as URLs with query strings) consider using the IgnoreQueryData parameter instead of adjusting the size of the resource cache. The value of the MaxResourceCacheSize agent configuration parameter determines the size of the resource cache.

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

Resource Cache and URL Query Strings

If you want to protect applications that use URL query strings, you can still take advantage of the resource cache by configuring the Web Agent to ignore the data in the query string. When the query string data is ignored, the truncated URL is stored in the resource cache. Query strings are ignored by setting the value of the IgnoreQueryData parameter in your Web Agent configuration.

Important! Do not enable this setting if you have policies which depend on URL query data.

The following table shows how ignoring the query strings in a URL determines whether the items from resource cache are used, or if the Web Agent contacts the Policy Server instead:

Requested URL with query string	Truncated URL stored in cache	Cached item used	Policy Server Contacted
/exampleapplication/page1.html ?user=firstuser	/exampleapplication/ page1.html	No	Yes
/exampleapplication/page1.html ?user=seconduser		Yes	No
/exampleapplication/page2.html ?user=seconduser	/exampleapplication/ page2.html	No	Yes

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

Session Cache (authentication)

Each CA SiteMinder® Agent uses a session cache to store the authentication information of users whom the Policy Server has already authenticated.

The Agent searches the session cache to determine whether a user is authenticated before contacting the Policy Server for authentication. The session cache improves performance by reducing the number of authentication calls to the Policy Server.

The authentication for a user ends when any of the following events occur:

- The user logs out.
- The session associated with a user expires.
- The age of an item in the cache exceeds 60 minutes.

Authentication information is removed from the session cache and discarded.

Authorization Cache

Each CA SiteMinder® Agent uses an authorization cache to store the authorization identification of users whom the Policy Server has already authorized.

The Agent searches the authorization cache to determine whether a user is authorized before contacting the Policy Server for authorization. The authorization cache improves performance by reducing the number of authorization calls to the Policy Server.

The authorization for a user ends when any of the following events occur:

- The user logs out.
- The session associated with a user expires.

The authorization identification is removed from the cache and discarded.

Session and Authorization Cache Settings

A combination of Policy Server settings and agent configuration parameters control the session cache and authorization cache. Use the results of your capacity planning as a guide to determine the best values for the following settings in your CA SiteMinder® deployment:

Session Timeouts

We recommend setting the session timeouts as follows:

- Set the maximum session timeout to match the sustained amount of time that the largest number of users are accessing the protected applications.
- Set the idle session timeout to an interval that meets all the following criteria:
 - Long enough to prevent the user from being logged out while working.
 - Short enough to log the user out automatically when the application is not being used (such as when a user leaves the computer without logging out).

Policy Server settings determine the timeout intervals.

Session Cache Size

Base the size of this cache on the number of users that you expect to access a resource for a sustained period during the session timeout interval. Include users who logout and log back in during the session timeout period in your sizing estimate. Do not include users whom you expect to make relatively few requests in your sizing estimate (because these users have a small effect on the session cache and authorization cache). A Web Agent configuration parameter named `MaxSessionCacheSize` determines the size of both the session cache and the authorization cache.

More information:

[How to Estimate a Sustained Authentication Rate](#) (see page 89)

Caching and Anonymous Users

The anonymous authentication schemes offered by CA SiteMinder® do *not* provide access control to resources that they protect. Anonymous authentication schemes allow the following for unidentified users on your network:

- Track how often a user returns to your sites.
- Track what a particular user does while visiting your sites (such as the pages the user viewed during a visit).
- Display personalized content for a particular user.

When users request resources protected by an anonymous authentication scheme, the Policy Server assigns a Global Unique Identifier (GUID), and stores it in the browser of the associated user. CA SiteMinder® uses this GUID to identify the user.

If you plan to use an anonymous authentication scheme, implementing the following items can improve performance in your CA SiteMinder® environment:

- Separate web servers to handle the anonymous requests.
- Configure the Web Agent on each separate web server to cache the anonymous requests by setting the CacheAnonymous parameter.

Using separate web servers and Web Agents for anonymous users keeps the caches on the other web servers that service requests for the protected resources from being flushed too often.

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

Other Parameters That Affect Web Agent Performance

The following parameters also affect Web Agent performance:

- PSPollInterval
- IgnoreExt
- IgnoreURL

Policy Server Poll Interval Parameter

CA SiteMinder® Agents contact the Policy Server regularly to receive any updated policies or encryption keys. The time interval for contacting the Policy Server can be adjusted by changing the PSPollInterval agent configuration parameter.

Increasing the time interval can reduce unnecessary traffic between the Agents and the Policy Server. Consider increasing the interval when your CA SiteMinder® environment has any of the following characteristics:

- You have many Agents.
- Most of the CA SiteMinder® policies are static, and do not change often.

Note: For more information, see the *Agent Configuration Guide* or *Agent Guide* for your Agent.

Important! Increasing the PSPollInterval parameter also affects how quickly the Agents enforce CA SiteMinder® policy changes. For example, suppose you change a Policy to revoke access for a terminated employee at 10:30, and your PSPollInterval parameter has a value of 3600 (the number of seconds in an hour). The Web Agents would not enforce the changed policy until as late as 11:30.

Ignore Extensions Parameter

If the resources you want to protect with CA SiteMinder® contain many images or files that you do *not* want to protect, you can reduce traffic between your Web Agents and Policy Servers by configuring the Web Agent to ignore certain file extensions.

Performance improves because the Web Agent does *not* make the following calls to the Policy Server:

- IsProtected
- IsAuthenticated
- IsAuthorized
- Login

Requests for the associated resources are passed directly to the web server and the user is granted access.

Identifying the resources you want to protect first can help you determine which file extensions, if any, you want your Web Agents to ignore.

Add any file extensions you want to ignore are to the IgnoreExt parameter of your Web Agent configuration.

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

More information:

[Identify the Applications to Secure](#) (see page 55)

Ignore URL Parameter

If you want to leave the resources in certain subdirectories unprotected, you can configure the Web Agent to ignore certain uniform resource identifiers (URI).

For example, if each of your web servers has a subdirectory named pictures, and you want to leave those directories on protected, you can set the IgnoreURL parameter in your Web Agent configuration.

Performance improves because the Web Agent does *not* make the following calls to the Policy Server:

- IsProtected
- IsAuthenticated
- IsAuthorized
- Login

Requests for the associated resources are passed directly to the web server and the user is granted access.

Improve Agent Performance through Load Balancing

When you have multiple CA SiteMinder® Agents and Policy Servers, dynamic load balancing reduces latency and improves throughput because the Agents distribute requests among all the Policy Servers. Dynamic load balancing gives the Agents faster access to Policy Servers and more efficient authentication and authorization.

CA SiteMinder® provides software-based failover and load-balancing of their communication with multiple Policy Servers. The EnableFailover parameter of the Host Configuration Object uses one of the following values to determine how Web Agent connections are handled:

- When the value is set to yes, the Agent always tries to connect to the first Policy Server listed (from left to right) in the Host Configuration Object. If you have multiple Policy Servers, all the Agents try to connect to the first one. The other servers in the list are not contacted unless the first server in the list is not available. In high-volume environments, this configuration is less efficient than load balancing because some Policy Servers handle many connections while others handle fewer connections, if any.

- When the value is set to no, load-balancing is enabled. The Agents balance their requests among all the Policy Servers in listed in the Host Configuration Object in a round-robin fashion. We recommend this setting because it produces better throughput when using multiple Policy Servers. Failover still occurs if one of the load balancing Policy Servers is not available.

Note: For more information, see the *CA SiteMinder® Policy Server Configuration Guide*.

CA SiteMinder® also supports the use of hardware load balancers to provide high performance dynamic load balancing of connections between CA SiteMinder® Agents and Policy Servers. When configured to expose multiple Policy Servers through a virtual IP address, hardware load balancers handle distribution of load between all Policy Servers associated with that virtual address. Because the Agent does not need to handle failover or load balancing, set the EnableFailover parameter to yes to disable CA SiteMinder® load balancing. Configure only the VIP or VIPs that expose groups of Policy Servers in the Host Configuration Object.

CA SiteMinder® Failover and Load Balancing with Multi-Threaded Web and Application Servers

CA SiteMinder® Agents running on multi-threaded web and application servers (such as Sun Java System, IIS, an Apache-based server in worker mode, or WebSphere Application Server), open the minimum number of sockets to a Policy Server at startup.

If you configure your environment for failover or load-balancing between Policy Servers, then the Agent opens the minimum number of sockets to each Policy Server at startup. Connections to a load-balanced Policy Server occur in the same way, although fewer sockets are opened to each Policy Server, because each is getting only half of the total requests.

If configured for failover, and an error occurs between the Agent and the primary Policy Server, then connections to the failover Policy Server are used. Failover occurs per service, so there could be active connections to both the primary and the failover Policy Servers at once. Once the primary Policy Server comes back up, the sockets opened to the failover server remain. All new sockets are opened to the primary Policy Server.

More information:

[Web Agent and Policy Server Interaction using Apache-based Web Server Worker Mode](#) (see page 147)

CA SiteMinder® Failover and Load Balancing with Multi-Process Web and Application Servers

A CA SiteMinder® Agent running on a multi-process web or application server (such as an Apache-based server running in pre-fork mode) opens the same number of connections to *all* configured Policy Servers, regardless of whether failover has occurred or not.

When failover occurs, it happens independently for each child, because each child process has its own connections to the Policy Server. This results in a 500 error for each socket as failover takes place. After the primary Policy Server comes back up, the sockets opened to the failover server remain open. All new sockets are opened to the primary Policy Server.

More information:

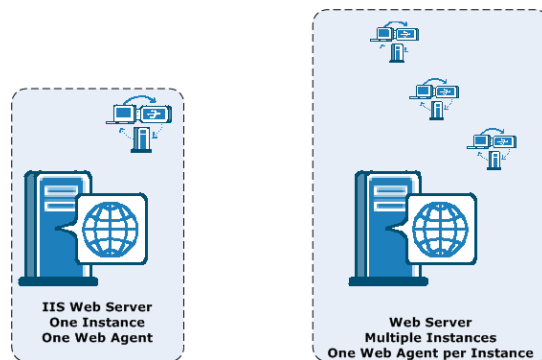
[Web Agent and Policy Server Interaction using Apache-based Web Server Pre-Fork Mode](#) (see page 147)

Web Servers, Web Agents, and Web Server Processes

Each CA SiteMinder® agent requires its own web server instance. IIS web servers, for example, operate using a single instance on the computer on which is it installed. The number of IIS agents equals the number of IIS web servers.

For other web servers that support multiple instances per computer, you can install and configure one CA SiteMinder® agent for each instance. For example, you could have one computer that runs three separate web server instances. Each instance has its own agent. Therefore, one computer operates three CA SiteMinder® agents.

The following illustration shows an example:



For Apache web servers, the following multi-processing modules (MRMs) affect how the CA SiteMinder® agent processes connect to the Policy Server:

Pre-fork mode

Creates child processes to handle additional requests.


Worker mode

Obtains additional threads from the connection pool to handle additional requests.

Web Agent and Policy Server Interaction using Apache-based Web Server Worker Mode

Apache-based web servers in worker mode use threads to handle connections to the CA SiteMinder® Policy Server. Threads are obtained from a connection pool as needed to create additional connections to the Policy Server during heavy loads.

The following illustration describes this process:



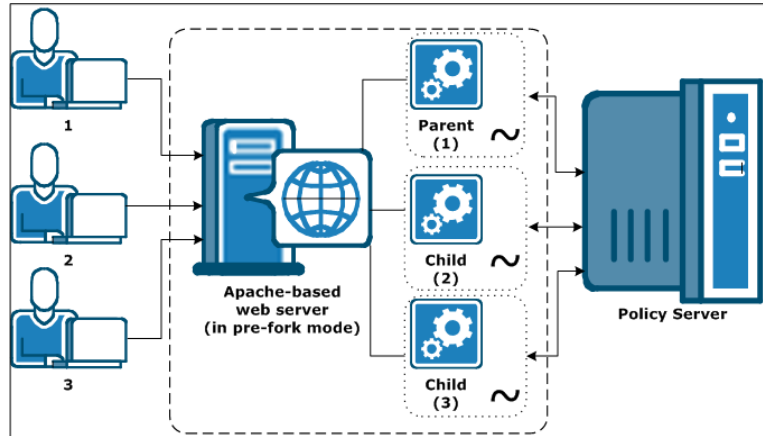
More information:

[CA SiteMinder® Failover and Load Balancing with Multi-Threaded Web and Application Servers](#) (see page 145)

Web Agent and Policy Server Interaction using Apache-based Web Server Pre-Fork Mode

When an Apache-based web server in pre-fork mode receives a request, the web server spawns a child process to communicate with the CA SiteMinder® Policy Server. When more requests are received, more child-processes are spawned to handle them. Each child process spawned by the Apache-based web server has its own independent connections to the CA SiteMinder® Policy Server.

The following illustration describes this process:



For Apache-based web servers, the value of the MaxClients parameter (in the httpd.conf file) determines the number of child processes spawned the web server. When a parent process from an Apache-based web server spawns a child process, the child process opens an initial connection to the CA SiteMinder® Policy Server.

An important distinction exists between the number of Web Agents, and the number of Web Agent processes. Each Web Agent requires its own web server instance. IIS web servers, for example, only operate as a single instance, so the number of IIS Web Agents equals the number of IIS web servers. For other types of servers, it is possible to have multiple server instances listening on different ports within one physical web server.

The maximum number of sockets opened from an Apache-based web server to a CA SiteMinder® Policy Server equals the value of the MaxClients parameter multiplied by the number of Web Agent processes. For example, if the value of the MaxClients parameter of your server is set to 150, and you have five Web Agent processes, then the maximum number of possible sockets opened is 750.

Using a multiprocess web server affects the ratio of Web Agent processes to Policy Servers in your CA SiteMinder® environment. The limiting factor often becomes the number of connections between the Web Agent processes and the Policy Server, not the number of transactions per second.

Before deploying Web Agents, verify that the CA SiteMinder® Policy Servers receiving the requests can handle the maximum number of connections that the related web servers could open.

More information:

[CA SiteMinder® Failover and Load Balancing with Multi-Process Web and Application Servers](#) (see page 145)

Application Tier Performance

Policy Servers evaluate policies in the application tier and user credentials and attributes in the data tier to protect resources. Consider the following guidelines to performance tune the application tier:

- The amount of system resources required to authenticate users affects performance.
- The amount of system resources required to authorize users affects performance.
- The number of Policy Server requests to CA SiteMinder® user directories during authentication and authorization affects performance.

CA SiteMinder® Policy Design and Performance

CA SiteMinder® policies define how users interact with resources. When you create CA SiteMinder® policies in the Administrative UI, you link together (bind) objects that identify users, resources, and actions associated with the resources.

You can improve or degrade performance in the way you configure specific CA SiteMinder® components or by choosing to enable optional features. A performance strategy includes:

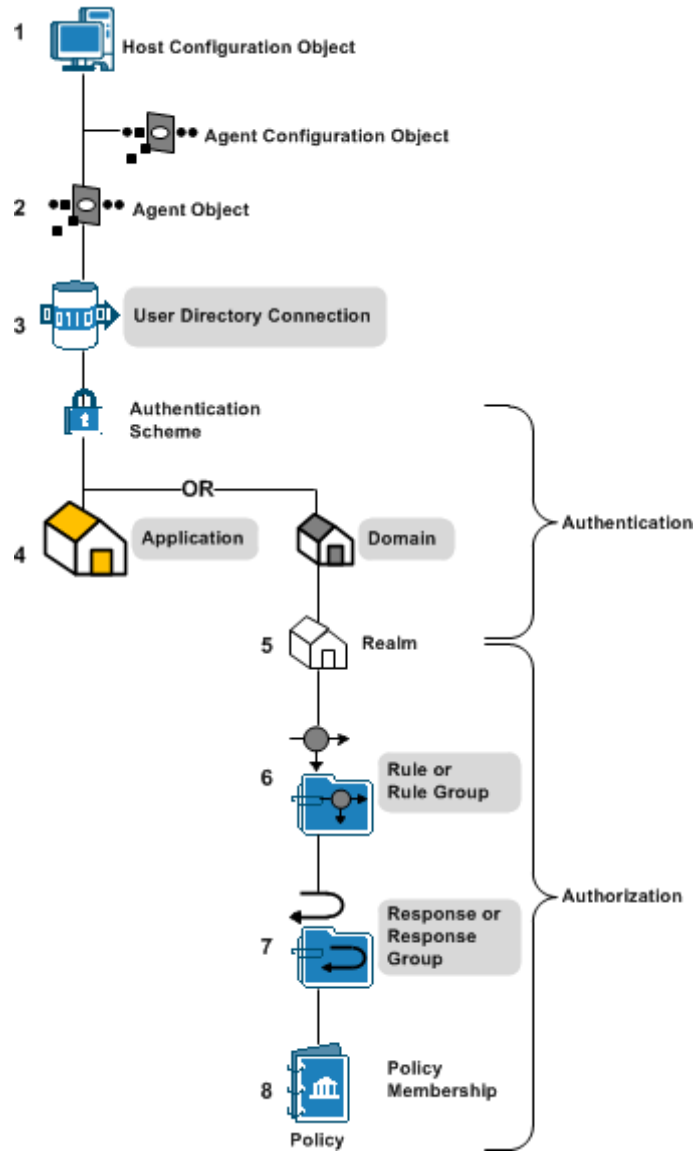
- Identifying the CA SiteMinder® policy objects that can affect performance
- Identifying the CA SiteMinder® parameters and features that affect user authentication
- Identifying the CA SiteMinder® parameters and features that affect user authorization

The business rules and security requirements of your enterprise should ultimately dictate your CA SiteMinder® policy design. The following guidelines are available to help you balance CA SiteMinder® performance, while meeting these requirements.

CA SiteMinder® Policy Objects and Performance Roadmap

CA SiteMinder® requires that you configure core CA SiteMinder® policy objects in a specific order. The following diagram lists this order, where shaded items represents objects that affect performance during user authentication or authorization.

Note: The Host Configuration Object (HCO) and Agent Configuration Object (ACO) affect the performance of your Web tier.



More information:

[Web Tier Performance](#) (see page 131)

Applications

You can improve or degrade performance during authentication and authorization in the way you configure applications.

An application is a Policy Server object that defines a complete security policy for one or more related web services. Applications associate web service resources with user roles to specify entitlement policies that determine what web service users can access what web service application resources.

When you create an application, you bind it to one or more user directory connections against which the Policy Server attempts to authenticate users. Therefore, the number of directory connections, and order in which they are listed, directly affects CA SiteMinder® performance during authentication.

The number of web service ports and operations that are defined as protected resources in an application correlates to CA SiteMinder® performance during authorization.

Resources can be bound to one or more responses. When a resource is accessed, the associated response returns information to an agent, such as user attributes, DN attributes, static text, or customized active responses.

The types of responses you bind to web service resources directly correlate to CA SiteMinder® performance during authorization.

Domains

You can improve or degrade performance during authentication in the way you configure domains.

A CA SiteMinder® policy domain is a logical grouping of resources associated with one or more user directories. When you create a domain, you bind one or more user directory connections to the domain.

The Policy Server attempts to authenticate users using these directory connections. Therefore, the number of directory connections, and order in which they are listed, directly correlates to CA SiteMinder® performance during authentication.

Note: For more information about configuring domains, see the *Policy Server Configuration Guide*.

More information:

[Group Resources into Domains or EPM Applications](#) (see page 56)

[Domains and Authentication Performance](#) (see page 157)

Realms

You can improve or degrade performance during authentication in the way you configure realms.

You group the resources in a domain into one or more realms. A realm is a set of resources (URLs) with a common security (authentication) requirement. The resource filter you define and the authentication scheme you select directly correlate to performance during authentication:

- The resource filter functions as the root of the protected resources. The Policy Server must evaluate the resource filter to determine if the requested resource is protected (IsProtected?).
- The authentication scheme associated with the realm determines the type of credentials users must present to gain access to the resources in the realm (IsAuthenticated?).

Realm settings also determine:

- How CA SiteMinder® handles user sessions. CA SiteMinder® creates a user session in the context of the realm to which the user authenticated against.
- If the realm can be used to control actions during authentication.

Note: For more information about realms, see the *Policy Server Configuration Guide*. For more information about authentication schemes, see the *Policy Server Configuration Guide*.

More information:

[Group Resources into Realms or EPM Components](#) (see page 57)
[Realms and Authentication Performance](#) (see page 157)

Rules and Rule Groups

You can improve or degrade performance during authorization in the way you configure realms.

You create rules or rule groups in the context of a realm. Rules:

- Identify the specific resources within a realm that require protection
- May be used to either allow or deny access to the resource based on specific authentication or authorization events.

The resource filter you define in the rule, which is prefixed with the realm filter, identifies the resource that requires protection.

The Policy Server evaluates rules to determine which resource filter best matches the requested resource. Upon a match, the Policy Server fires the policies to which the rule is bound to determine if the user is authorized to access the resource.

The number of rules within a realm and how you define each of the resource filters directly correlates to CA SiteMinder® performance during authorization.

Note: For more information about rules, see the *Policy Server Configuration Guide*.

More information:

[Rules and Authorization Performance](#) (see page 159)

Responses

You can improve or degrade performance during authorization in the way you configure responses.

Responses or response groups are bound to specific rule or rule groups. When a rule fires, a response can:

- Customize the amount of time user sessions remain valid.
- Redirect the user to other resources.
- Customize the content the user receives based on attributes contained in a user directory.
- Pass static text, user attributes, DN attributes, customized active responses, or the runtime values of defined variables from the Policy Server to a CA SiteMinder® Agent.
- Instruct a SiteMinder WSS Agent to generate WS-Security headers and SAML Session Tickets

Policies rules can be bound to one or more responses. The types of responses you bind to CA SiteMinder® policy rules directly correlates to CA SiteMinder® performance during authorization.

Note: For more information about responses, see the *Policy Server Configuration Guide*.

More information:

[Responses and Authorization Performance](#) (see page 160)

Authentication Guidelines

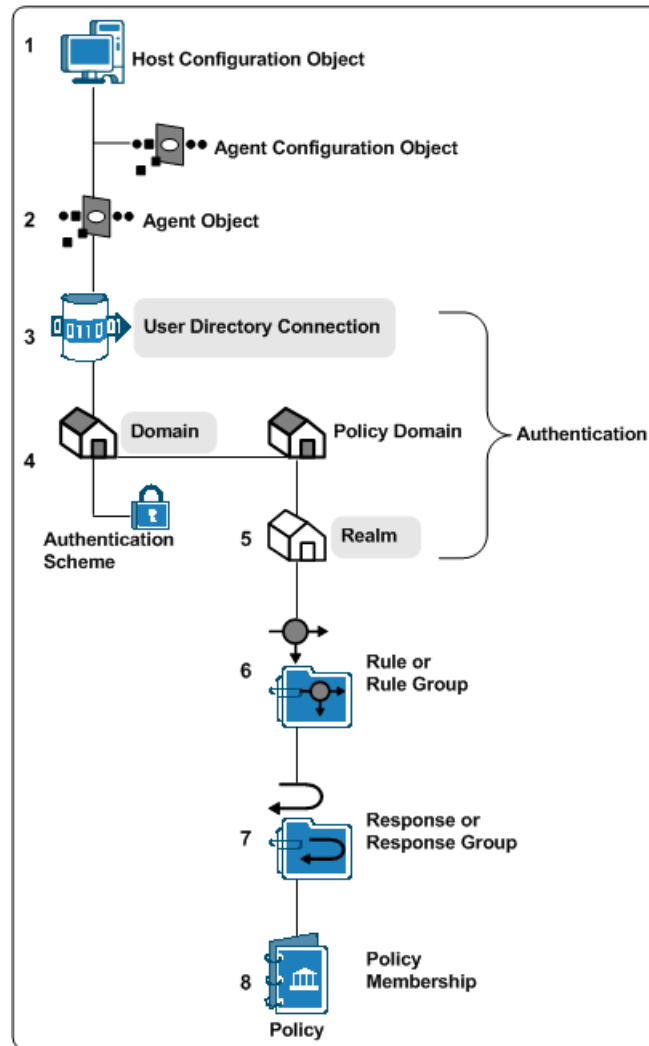
CA SiteMinder® performance during the authentication (IsAuthenticated?) step typically correlates with:

- The system resources used to service an authentication request
- The number of reads/writes, collectively known as requests, that the Policy Server makes to CA SiteMinder® user directories to service an authentication request

CA SiteMinder® Policy Objects and Performance Roadmap

Authentication performance can improve or degrade depending on how you configure specific CA SiteMinder® policy objects or by choosing to enable optional features associated with those objects.

CA SiteMinder® requires that you configure core CA SiteMinder® policy objects in a specific order. The following diagram lists this order, where shaded items represent objects that affect performance during user authentication.



User Directories and Authentication Performance

Configuring a domain requires that you bind one or more user directory connections to the domain. The Policy Server uses the search criteria you specify in the user directory connection to verify user credentials during the authentication step.

Note: For more information about configuring user directory connections, see the *Policy Server Configuration Guide*.

The following factors affect user authentication performance at the directory level:

- Search expressions and queries—The more complex the LDAP expression or ODBC query, the longer it takes the Policy Server to resolve the criteria to authenticate the user.
- Password Services—You can apply password policies to CA SiteMinder® user directories. Consider the following before implementing password policies:
 - The Policy Server reads attributes related to the password policy and may need to update them. Updating an attribute requires the Policy Server to write to the user directory.
 - If the password policy is configured to track login details, an additional user directory write is required for every authentication.
 - The Policy Server takes longer to resolve password policies that only apply to a specific group of users within the directory, instead of the entire directory.

CA SiteMinder® Web Services Security Authentication Schemes and Authentication Performance

Different CA SiteMinder® Web Services Security authentication schemes impose different level of WSS Agent processing overhead, which can also vary between WSS Agent types.

In general, authentication throughput is greater for authentication schemes that do not require digital signature verification or payload confidentiality.

Digital signature verification is more CPU- and data-intensive on WSS Agent for Web Servers, but also slightly impacts WSS Agents for application servers.

Domains and Authentication Performance

The following factors affect user authentication performance at the domain (or application object general) level:

- The number of directory connections in the domain—The Policy Server searches each user directory in the domain until it is able to validate the user credentials. The greater the number of user directory connections, the longer it can take the Policy Server to authenticate the user.

Evaluate ways to reduce the number of directory connections in a domain to prevent unnecessary Policy Server requests. Consider:

- Who is requesting the resources within the domain and in which directories their information is stored
 - Combining user directories when you add an organization to your CA SiteMinder® deployment
- The order in which user directory connections are listed—The Policy Server searches user directories in the order in which the domain lists them. Evaluate authentication priorities when determining the order of connections. Consider:
 - If a larger percentage of users access the application from a specific directory or directories
 - If a smaller group of users exists that are higher priority for authentication

Realms and Authentication Performance

The following factors affect user authentication performance at the realm (or application object component) level. Consider each as you configure realms:

- Credential collection—Realms are associated with a specific authentication scheme, some of which require the use of credential collectors. Agents protecting resources with these types of authentication schemes redirect users to the credential collector to gather the credentials. Gathering credentials adds an additional step to the authentication process.

Note: For more information about configuring authentication schemes, see the *Policy Server Configuration Guide*. For more information about using credential collectors, see the *Web Agent Configuration Guide*.

- Persistent Sessions—When CA SiteMinder® authenticates a user, the Policy Server issues a session ticket. A session ticket contains basic information about the user and the authentication context of the user. By default, CA SiteMinder® implements session management through non-persistent sessions, for which the Agent writes the session ticket to a cookie in the web browser of the user.

Some CA SiteMinder® features require persistent sessions. You can configure a realm for Persistent Sessions. Agents protecting resources in this realm write the session ticket a CA SiteMinder® session store, which results in additional requests to the session store for each authentication.

Important! Persistent sessions can have a significant impact on performance.

Note: For more information about user sessions, see the *Policy Server Configuration Guide*.

- Authentication Events—By default, a realm is configured to Process Authentication Events. This setting lets you define rules that fire when a user authenticates or fails to authenticate. Policy evaluation logic applies to all realms configured to process Authentication Events. This logic consumes system resources and can result in user directory requests.

Evaluate the need for event actions that occur when users authenticate to gain access to a resource. If you do not require authentication actions, disable Authentication Events for the realm to speed the authentication step.

Note: For more information about realms, see the *Policy Server Configuration Guide*.

Authorization Guidelines

CA SiteMinder® performance during the authorization step typically correlates with:

- The system resources used to service an authorization request.
- The number of reads/writes, collectively known as requests, the Policy Server makes to CA SiteMinder® user directories to service an authorization request.

The complexity of your CA SiteMinder® policy design affects each of these areas.

Policy Objects and Performance

You can improve or degrade authentication performance in the way you configure specific CA SiteMinder® policy objects or by choosing to enable optional features associated with those objects. The following policy objects can affect performance during user authorization:

- [Rules](#) (see page 159)
- [Responses](#) (see page 160)
- [Policy membership](#) (see page 160)

Rules and Authorization Performance

The following factors affect user authorization performance at the rule (or application object resource) level:

- Large numbers of rules in a single realm can slow authorization decisions. If a user is authenticated for a particular realm, the Policy Server must evaluate all rules within the realm to determine which of the resource filters best matches the specific resource (URL) the user is requesting.
- The type of resource filter affects how quickly the Policy Server can evaluate the resource match.

Note: For more information about rules, see the *Policy Server Configuration Guide*.

The following filters are listed in the order in which they have the smallest affect on performance:

- Exact match—Defining a resource filter with a specific resource has the smallest affect on performance. The Policy Server only has to compare the resource filter to the URL of the requested resource.

Example: A company creates a customer realm (/customer) and specifies a rule with a specific page of their portal application (lending_home.html). The resulting resource filter is /customer/lending_home.html. Evaluating a match between the requested resource and the rule only requires the Policy Server to compare the requested resource with the resource filter to determine if it is a match.

- Exact prefix—Defining a resource filter with a prefix has a greater affect on performance than an exact match. The Policy Server must determine if the requested resource is contained within the root (realm) of the resource.

Example: A company creates an employee realm (/employee) and specifies a rule with "*.html". The * prefix specifies that all html files in the employee realm are protected. The resulting resource filter is /employee/*.html. Evaluating a match between the requested resource and the resource filter requires the Policy Server to evaluate if the requested resource is part of the employee directory and is an HTML file.

- Regular expression—Defining a resource filter with a regular expression has the greatest affect on performance. The Policy Server must evaluate the expression and compare the result to the requested resource. The complexity of the expression further affects performance.

Responses and Authorization Performance

The type of response attributes bound to rules in a CA SiteMinder® policy affect performance. The following response types are listed in the order in which they have the smallest affect on performance:

- Static—Defining a static attribute returns data that is constant.
- User attribute—Defining a user attribute returns profile information from a user's entry in a user directory.

Note: This type of response requires the Policy Server to search the user directory.

- DN attribute—Defining a DN attribute returns information associated with directory objects to which the user is related. Groups to which a user belongs, and organizational units (ou) that are part of a user DN, are examples of directory objects whose attributes can be treated as DN attributes.

Note: This type of response requires the Policy Server to search the user directory.

CA SiteMinder® Policy Membership and Authorization Performance

Policy membership is the part of a CA SiteMinder® policy that specifies which users apply to the policy. CA SiteMinder® policies are stored in domains, and as a result, you use filters to apply CA SiteMinder® policy membership to any or all users stored in the user directories bound to the domain. The type of filter you define determines how the Policy Server evaluates CA SiteMinder® policy membership.

Note: For more information about adding users to a CA SiteMinder® policy, see the *Policy Server Configuration Guide*.

The following filters are listed in the order in which they have the smallest affect on performance:

- All—"All" has the smallest affect on performance.

When CA SiteMinder® authenticates a user, the Policy Server issues a session ticket. The session ticket identifies the user directory in which the user is stored. The Policy Server only has to compare the session ticket with the directory bound to the CA SiteMinder® policy to determine that the policy applies to the user.

Note: For more information about user sessions, see the *Policy Server Configuration Guide*.

- Distinguished name—A distinguished name (dn) has a greater affect on performance than "All".

The organization or organizational unit, which contains the dn of the authenticated user, is stored in the session ticket. The Policy Server has to compare the session ticket information with the CA SiteMinder® policy membership filter to determine if the policy applies to the user.

- Group membership or search expressions—These types of filters have a greater affect on performance than distinguished names. Group membership and search expressions consume additional system resources and result in a user directory search. The Policy Server must:
 - a. Resolve the group membership or search expression
 - b. Search the user directory to determine if the CA SiteMinder® policy applies to the user.
- Nested groups—Defining CA SiteMinder® policy membership with a nested group has the greatest affect on performance.

The Policy Server must search each user group and all sub-groups in the directory to determine if the CA SiteMinder® policy applies to the user.

Important! Directories with deep group hierarchies can have a significant effect on the time it takes the Policy Server to evaluate policy membership.

Note: You can enable the User Authorization cache to reduce the number of requests the Policy Server makes to user directories to resolve policy membership.

More information:

[User Authorization Cache](#) (see page 161)

User Authorization Cache

The user authorization cache reduces the number of user directory requests to determine CA SiteMinder® policy membership by storing the relationship between users and policies.

Note: The user authorization cache does not store data about the user, store user attribute values, or cache user entries.

For example, three policies are configured to apply to an "Administrator" group, to which user A belongs. The first-time the Policy Server evaluates CA SiteMinder® policy membership, it must resolve the group membership and make three requests (one for each policy) to the user directory to determine that each CA SiteMinder® policy applies.

The Policy Server writes these results to the user authorization cache. Subsequent policy evaluation does not require the Policy Server to make user directory requests. Rather, the Policy Server uses the cached authorization information to determine policy membership.

Note: The Policy Server polls for policy updates periodically. The default interval is 60 seconds. If the policy membership changes, the Policy Server reloads the policy and removes the cache entries that are related to the updated policy.

More information:

[CA SiteMinder® Policy Membership and Authorization Performance](#) (see page 160)

User Authorization Cache Efficiency

The user authorization cache is most efficient when:

- All user requests during a session are consistently sent (persisted) to the same server.
- All CA SiteMinder® agents are configured for Policy Server failover, not round–robin load balancing.

If these factors are not met, the efficiency of the User Authorization cache is reduced.

Example: the user authorization cache and agents configured to round–robin load balance

The more Policy Servers that are in the CA SiteMinder® agent round–robin pool, the greater the chance that the efficiency of the user authorization cache is reduced.

If a single CA SiteMinder® Agent is configured to round–robin between two Policy Servers, the first request for a protected resource results in a user authorization cache entry on one of the Policy Servers. There is approximately a 50 percent chance that the Policy Server that does not have the cache entry must service the second request. Moving forward, however, both Policy Servers have cached the data for subsequent requests.

Consider now, the effect of a single Agent configured to round–robin between 10 Policy Servers. After a Policy Server authorizes a user and enters the result in to the authorization cache, there is only a 10 percent chance that the same Policy Server services the next request. In this configuration, 5 cache misses must occur before there is a 50 percent chance of a cache hit.

Note: Policy Server clusters can reduce the effect round–robin load balancing has on the user authorization cache.

Estimate the Size of the User Authorization Cache

The default size of the user authorization cache is 10 MB. You can estimate the amount of space the user authorization cache requires and use the Policy Server Management Console to adjust the default size.

To estimate the size of the user authorization cache

1. Use the following formula to estimate the number of cache entries:

$$\text{expected_users} * \text{number_of_policies_per_session} = \text{entries}$$

expected_users

Specifies the total number of users authenticating to the applications CA SiteMinder® is protecting.

number_of_policies_per_session

Specifies the average number of CA SiteMinder® policies that apply to a user during the session.

Note: Each CA SiteMinder® policy has the potential to enter a unique entry into the user authorization cache.

entries

Specifies the number of cache entries authorizations can create.

2. Use the following formula to estimate the size of the cache:

$$(\text{entries} * .000062) + 1$$

Note: .000062 represents the approximate size of a cache entry in MB.

Auditing and Performance

By default, the Policy Server writes audit events to a text file, which is known as the Policy Server log. Optionally, you can configure the Policy Server to log events to an audit database.

Note: For more information about configuring the Policy Server to log events to an audit database, see the *Policy Server Administration Guide*. For more information about configuring an audit database, see the *Policy Server Installation Guide*.

Consider the following factors if you decide to log events to an audit database:

- Performance associated with authentication and authorization is affected because CA SiteMinder® is logging all authentication and authorization decisions to the database.
- (Optional) Synchronous logging—You can configure synchronous logging at the realm level. If configured, the Policy Server prevents the result of each authentication and authorization request until the record is saved in the audit database. Users are not authenticated or authorized until the record is saved.

Load Balancing the Application Tier

Tuning the various CA SiteMinder® Agent parameters and following the CA SiteMinder® policy design guidelines may not significantly improve the amount of time it takes the Policy Server to service authentication and authorization requests.

When you have multiple Agents and Policy Servers, dynamic load balancing reduces latency and improves throughput because the Agents distribute requests among all of the Policy Servers.

More information:

[Redundancy and High Availability](#) (see page 32)

Data Tier Performance

Poor performance associated with CA SiteMinder® data stores, especially user directories, is one of the most common reasons for poor CA SiteMinder® performance. Data tier performance typically correlates with two general areas:

- The data tier itself. A user directory that is not properly tuned or lacks sufficient system resources can degrade CA SiteMinder® performance.
- The capacity under which your user directories have to operate. CA SiteMinder® authentication and authorization services result in a number of reads and writes, collectively known as requests, to a user directory. Conduct a capacity planning effort on the user directory itself to be sure it can handle the CA SiteMinder® workload.

A performance strategy includes:

- Determining that the data tier itself is not the primary reason for poor performance.
- Identifying the number of authentications and authorizations CA SiteMinder® must service in a given period.

Note: The sustained and peak rates at which user authentication and authorization occur can be calculated.

- Estimating how many user directory requests each user authentication, and the subsequent authorizations, create.

More information:

[Capacity Planning Introduced](#) (see page 87)

[CA SiteMinder® Web Services Security Capacity Planning Introduced](#) (see page 103)

Data Tier Guidelines

The Policy Server interacts with the data tier using standard protocols. If your directory servers and databases are tuned to maximize performance with their normal clients, then these modifications can translate into improved CA SiteMinder® performance.

Note: See your vendor-specific documentation for tuning guidance.

There are several general considerations to improving CA SiteMinder® performance as it relates to the performance of your user directories. Examine the following areas:

- The system resources available to the user directory and any external resources that may contend for those resources
- The use of Secure Socket Layer

- The efficiency in which CA SiteMinder® can search the user directory
- The use of static IP addresses
- The use of replication

System Resources

The system resources available to the user directory directly correlates to CA SiteMinder® performance. If the user directory is operating at a high level of utilization, then no amount of CA SiteMinder® tuning can improve performance.

Be sure that the system hosting the user directory is not degrading performance due to:

- A slow CPU or I/O system
- Insufficient memory
- An incorrectly configured buffer cache
- Insufficient or fragmented disk space

Secure Socket Layer and User Directories

Consider the following if you are planning to implement SSL in your CA SiteMinder® environment:

- Configuring the Policy Server and an LDAP user directory to communicate over SSL reduces performance. Review your security requirements to determine if SSL is mandatory.
- If you decide to configure SSL, do not place an SSL accelerator between the Policy Server and the directory server or the Policy Server assumes a single instance of the directory. This can cause inconsistent writes across multiple user directories behind the accelerator.

Static IP Addresses and User Directories

When you configure user directory connections in the Administrative UI, consider using static IP addresses rather than hostnames. Although the time the Policy Server takes to resolve hostnames is negligible, using static IP addresses removes Domain Naming Services (DNS) dependencies.

User Directory Searches

Making sure that CA SiteMinder® can efficiently search users directories directly correlates with performance. Consider the following:

- Use directory indexing to enhance search results for CA SiteMinder®:
 - LDAP—the objectClass attribute, in addition to all other attributes used in searches, should be indexed.
Note: Microsoft recommends using the objectCategory attribute instead of objectClass. Failing to index the objectClass attribute in Active Directory can result in significant performance degradation.
 - ODBC—All fields defined as search criteria in CA SiteMinder® schema queries should be indexed.
Note: See your vendor-specific documentation for more information about indexing.
- Design queries to return manageable sets of user groups.
Note: If you are unable to optimize the query, set the maximum search results parameter to limit large result sets from degrading overall performance.
- Optimize SQL query schemes for ODBC with any standard SQL analyzer.

Replication

Replication can degrade performance in the following situations:

- When the master replica, in a master–slave replication, only allows write requests. Password Services usually requires updates to the password blob attribute for each authentication. If only the master replica can process writes, then each write request is redirected to the master.
The redirection results in additional time spent on the authentication step, and the master-replica may not be able to accommodate the rate at which writes occur.
- When LDAP referrals are enabled. LDAP referrals can degrade performance because each request may involve more than one request to a directory.

User Store Capacity Planning

The Policy Server performs a series of services to authenticate and authorize users. These services result in number of reads and writes, collectively known as requests, to a user directory. A significant contributing factor to CA SiteMinder® performance is determining whether your user directories can handle this workload during sustained and peak periods of operation.

The following general factors influence CA SiteMinder® performance:

- Total operations and sustained user directory search rates—Total operations is the combined number of requests the Policy Server must service when handling authentication and authorization requests. The rate at which these operations occur fluctuate throughout your business day.


In turn, the rate at which the Policy Server makes user directory requests to process the operations fluctuates. Some periods generate relatively few user directory requests, while others generate more.

The sustained user directory search rate represents a period during which the Policy Server makes an average number of user directory requests to service an average number of operations.

- Total operations and peak user directory search rates—During sustained periods of activity, user activity can spike. The peak user directory search rate represents a period during which the Policy Server makes the highest number of user directory requests to process peak numbers of operations.

The following graphic illustrates:

- The relationship between total operations and the user directory search rate.
- How each rate fluctuates throughout the day, is sustained for a specific period, and peaks within that period.



We recommend using the following guidelines to estimate the load under which your user directories have to operate. Once you have estimated the load, you can use any standard tool to create the load on the directory and track the results.

Note: Many factors can contribute to failing to achieve the required numbers. See your vendor-specific documentation for tuning guidance.

More information:

[Policy Server](#) (see page 12)

[How to Estimate a Sustained Authentication Rate](#) (see page 89)

[How to Estimate a Sustained Authorization Rate](#) (see page 94)

User Store Capacity Planning Checklist

Estimating the number of user directory requests that the Policy Server must make to service authentication and authorization requests requires specific information. Gather the following before beginning a user store capacity plan:

- The total number of daily authentications (authentication load) for the application.
- The total number of daily authorizations (authorization load) for the application.
- The sustained and peak periods during which users are authenticating to the application and requesting protected resources.

Note: A capacity planning effort can help you identify metrics related to authentication load, authorization load, and sustained and peak levels of user activity.

- The total number of enabled policies. For each CA SiteMinder® policy determine:
 - If the CA SiteMinder® policy membership filter results in one or more user directory searches.
 - If the responses bound to the CA SiteMinder® policy results in one or more user directory searches.

More information:

[Capacity Planning Introduced](#) (see page 87)

[CA SiteMinder® Policy Membership and Authorization Performance](#) (see page 160)

[Responses and Authorization Performance](#) (see page 160)

How to Estimate a Sustained User Directory Search Rate

Estimating a sustained user directory search rate is the process of determining:

- How the total number of user directory requests fluctuate throughout your business day
- How the user directory requests translate into requests per second over a sustained period.

Complete the following steps to estimate the sustained user directory search rate:

1. Use the authentication guidelines to estimate the number of user directory requests that the authentication load creates.
2. Use the authorization guidelines to estimate the number of user directory requests that the authorization load creates.
3. Estimate the sustained user directory search rate.

Use Authentication Guidelines to Estimate Directory Searches

A Policy Server makes a number of user directory requests to service each authentication request. Some of the user directory requests are required, while others can be avoided.

Estimate the number of Policy Server requests that each authentication creates using the following guidelines:

(Required) Two searches to authenticate each user:

- One search/query, per store, to identify the user
- One search/query to verify the user credentials

(Optional) Additional searches may be required depending on how you design policies and if you decide to enable Password Services:

- One search/query for each CA SiteMinder® policy that is bound to a rule that fires when a user is authenticated (OnAuth rule).

Note: For more information about configuring rules, see the *Policy Server Configuration Guide*. For more information about the relationship a rule has to a CA SiteMinder® policy, see the *Policy Server Configuration Guide*.

- One search/query for each CA SiteMinder® policy that is bound to a response that returns user attributes.

Note: For more information about responses and their relationship to rules, see the *Policy Server Configuration Guide*.

- One write/update per user store enabled for Password Services. If Password Services does not apply to the user directories in the CA SiteMinder® policy domain, a write/update is not required.

Note: For more information about Password Services, see the *Policy Server Configuration Guide*.

The following use cases detail how you can use each guideline to determine the total number of user directory searches the authentication load creates.

Case 1: User Authentication and Directory Requests

A company has:

- Deployed one user directory for their banking application.
- Completed a capacity planning effort. The results of which show that users create an authentication load of 88,000 logins.

The company uses the following formula to begin estimating the number of requests the Policy Server sends to the user directory to service the authentication load:

$$\text{authentication_load} * 2 * \text{number_of_user_stores} = \text{requests_for_authentication}$$

authentication_load

Specifies the number of daily authentications for the application.

Note: Two (2) is a constant. Authenticating a users results in two requests. One search to identify the user and one bind to verify credentials.

number_of_user_stores

Specifies the number of user stores in the implementation.

requests_for_authentication

Specifies the number of user directory requests that the authentication load creates.

Result: $88,000 * 2 * 1 = 176,000$ requests.

The company uses this estimate to determine the total number of user directory requests required to service the daily authentication load.

Case 2: Policy Design and User Directory Requests

A company has configured four policies to protect the application portal, one of which is bound to a rule that fires upon a successful authentication.

The company uses the following formula to continue estimating the number of requests the Policy Server sends to the user directory to service the authentication load:

$$\text{authentication_load} * (\text{percent_of_policies} * \text{number_of_searches}) = \text{requests_for_authentication}$$

authentication_load

Specifies the number of daily authentications for the application.

percent_of_policies

Specifies the total number of enabled policies, represented as a percentage, that are:

- bound to an onAuth rule
- create the same number of user directory searches

Example: Four enabled CA SiteMinder® policies exist. One is bound to an OnAuth rule. This policy generates one user directory search to determine policy membership. Twenty-five percent of the enabled policies fire on authentication and generate one user store search. The remaining policies do not fire during authentication.

number_of_searches

Specifies the number of requests that the Policy Server makes to determine if the CA SiteMinder® policy applies to each authenticated user.

requests_for_authentication

Specifies the number of user directory requests that the authentication load creates.

Result: $88,000 * 0.25 * 1 = 22,000$ requests

The company uses this estimate to determine the total number of user directory requests required to service the daily authentication load.

Case 3: Responses and User Directory Requests

A company has defined one CA SiteMinder® policy with an OnAuth rule. This policy requires that a common name (cn) attribute response be returned when the policy fires. The company defines a Web Agent response to return this value and binds it to the CA SiteMinder® policy rule.

The company uses the following formula to continue estimating the number of requests the Policy Server sends to the user directory to service the authentication load:

*authentication_load * percent_of_policies * number_of_responses_per_policy = requests_for_authentication*

authentication_load

Specifies the number of daily authentications for the application.

percent_of_policies

Specifies the total number of enabled policies, represented as a percentage, that are bound to a specific number of responses that return user attributes.

Example: If there are four enabled policies, and one uses a response to return a user attribute, then twenty-five percent of the policies require a user directory search.

number_of_responses_per_policy

Specifies the number of responses bound to the CA SiteMinder® policy.

requests_for_authentication

Specifies the number of user directory requests that the authentication load creates.

Result: $88,000 * 0.25 * 1 = 22,000$ requests

The company uses this estimate to determine the total number of user directory requests required to service the daily authentication load.

Case 4: Password Services and Directory Requests

A company has enabled Password Services for their user store. The company uses the following formula to continue estimating the number of requests the Policy Server sends to the user directory to service the authentication load:

$$\text{authentication_load} * 1 = \text{requests_for_authentication}$$

authentication_load

Represents the number of daily authentications for the application.

Note: One (1) is a constant. Tracking user login details requires one write to the user directory for each authentication.

requests_for_authentication

Represents the number of user directory requests that the authentication load creates.

Result: $88,000 * 1 = 88,000$ requests.

The company uses this estimate to determine the total number of user directory requests required to service the daily authentication load.

Case 5: Total Directory Requests for Authentication

A company uses the individual totals from each use case to determine the total number of requests the Policy Server sends to the user store to service the authentication load:

- 176,000 requests to identify 88,000 unique users and their credentials
- 22,000 requests to determine if the OnAuth CA SiteMinder® policy applies to those users

- 22,000 requests to return the common name attribute upon authentication
- 88,000 requests for the password policy

Result: 176,000 + 22,000 + 22,000 + 88,000 = 322,080 requests

The company uses this result and the results based on the authorization load to estimate the sustained rate at which the user store must service Policy Server requests.

Use Authorization Guidelines to Estimate Directory Searches

A Policy Server makes a number of user directory requests to authorize a user. Some of the user directory requests are required to determine CA SiteMinder® policy membership, while others are dependent on CA SiteMinder® policy design. You can estimate the number of Policy Server requests that each authorization creates using the following guidelines.

- One search/query for each CA SiteMinder® policy in the policy domain.
Note: This guideline only applies to policies whose membership filter results in one or more user directory requests. For more information about the relationship between CA SiteMinder® policy membership and user directory requests, see Policy Membership and Authorization Requests.
- One search/query for each policy that is bound to a response that returns user attributes.
Note: For more information about the relationship between responses and user directory requests, see Responses and Authorization Performance.

The following use cases detail how you can use each guideline to determine the total number of user directory searches the authorization load creates.

Note: The user authorization cache can significantly reduce the number of authorization-related requests to user directories.

More information:

[CA SiteMinder® Policy Membership and Authorization Performance](#) (see page 160)

[Responses and Authorization Performance](#) (see page 160)

[User Authorization Cache](#) (see page 161)

Case 1: Policy Membership and User Directory Requests

A company has enabled three policies protect their portal application:

- Policy A requires one user directory request to determine CA SiteMinder® policy membership.
- Policies B may require up to two user directory requests to determine CA SiteMinder® policy membership.
- Policies C may require up to three user directory requests to determine CA SiteMinder® policy membership.

Additionally, the results of a capacity planning effort show that the application has an authorization load of 726,000.

The company uses the following formula to begin estimating the number of requests that the Policy Server sends to the user directory to service the authorization load:

$$\text{authorization_load} \times \text{percent_of_policies} * \text{number_of_searches} = \text{daily_authorization_requests}$$

authorization_load

Specifies the number of daily authorizations for the application.

percent_of_policies

Specifies the number of enabled policies, represented as a percentage, that may result in the same number of user directory requests to determine CA SiteMinder® policy membership.

Note: The total percentage must equal 100 percent.

number_of_searches

Specifies the number of user directory requests that the Policy Server may make to determine CA SiteMinder® policy membership.

daily_authorization_requests

Specifies the number of user directory requests to service the authorization request.

Result:

- Policy A— $792,000 * 0.33 * 1 = 261,360$ requests
- Policies B and C— $792,000 * 0.66 * 2 = 1,045,440$ requests
- Total user directory requests— $261,360 + 1,045,440 = 1,306,800$ requests

The company uses this estimate to determine the total number of user directory requests required to service the daily authorization load.

More information:

[User Authorization Cache](#) (see page 161)

Case 2: Responses and User Directory Searches

A company has enabled three policies to protect their portal application, two of which are bound to responses that return user attributes:

- Policy A returns one user attribute when it fires.
- Policy B returns two user attributes when it fires.
- Policies C is not bound to responses that return user attributes.

The company uses the following to estimate the number of user directory requests that the Policy Server makes to resolve responses that return user attributes:

*authorization_load * percent_of_policies * number_of_responses =
daily_authorization_requests*

authorization_load

Specifies the number of daily authorizations for the application.

percent_of_policies

Specifies the number of enabled policies, represented as a percentage, that result in the same number of user directory requests because of responses returning user attributes.

Note: The total percentage must equal 100 percent.

number_of_responses

Specifies the number of responses bound to the CA SiteMinder® policy.

daily_authorization_requests

Specifies the number of user directory requests to service the authorization request.

Result:

- Policy A— $792,000 * 0.2 * 1 = 158,000$
- Policy B— $792,000 * 0.2 * 2 = 316,800$
- Policies C— $792,000 * 0.6 * 0 = 0$
- Total user directory request— $158,000 + 316,800 + 0 = 526,000$

The company uses this estimate to determine the total number of user directory requests required to service the daily authorization load.

Case 3: Total Directory Requests for Authorization

The company uses the individual totals from each use case to determine the total number of requests the Policy Server sends to the user directory to service the authorization load:

- 1,203,440 requests to resolve CA SiteMinder® policy membership.
- 526,000 requests to return user attributes associated with responses.

Result: $1,203,440 + 526,000 = 1,729,440$ requests

The company uses these result and the results based on the authentication load to estimate the sustained rate at which the user store must service Policy Server requests.

Estimate the Sustained User Directory Search Rate

The sustained user directory search rate is based on the total number of operations (authentication load plus authorization load), specifically, when and at what rate these requests occur. The chance that these requests are uniformly spread across your business day is unlikely. Rather, the rate at which these requests occur fluctuates, remaining between the lowest and highest (peak) levels for a sustained period.


Estimating the sustained user directory search rate is the process of identifying:

- A sustained period during which the system is servicing an average number of operations.
- How these requests translate into user directory searches.

When estimating the sustained user directory search rate, we recommend using the daily authentication load and authorization load to identify:

- The rate at which total operations occur throughout the day
Note: We recommend beginning with an evaluation period of 24 hours, broken down into one-hour increments. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding of usage throughout the year.
- The sustained period during which the system is servicing an average number of requests
- The approximate number of requests that occur during the sustained period.

The following figure is an example of these metrics.



Case: Estimate the Sustained User Directory Search Rate

The company has determined that:

- The daily authentication load and authorization load for the application result in approximately 888,000 total operations.
- The total operations result in approximately 2,051,520 user directory requests.
- The system is operating at sustained levels for approximately five hours (9:00 AM - 2:00 PM).
- During sustained levels, approximately 84,000 operations occur, per hour.
- Approximately 420,000 (84,000 * 5) operations, or 48 percent (420,000 / 888,000) of the total operations, occur during these hours.

The company uses the following formula to estimate the sustained user store search rate:

$$\frac{(total_user_directory_requests * percentage_of_requests) / number_of_hours}{3600} = sustained_user_directory_search_rate$$

total_user_directory_requests

Represents the daily number of requests the Policy Server makes to the user directory to service authentication and authorization requests.

percentage_of_requests

Represents the percentage of total operations that occur when the system is operating at sustained levels.

number_of_hours

Represents the number of hours when the system is operating at a sustained rate.

sustained_user_directory_search_rate

Represents the number of requests, per second, the Policy Server makes to the user directory to maintain the sustained rate of operation.

Result: $(2,051,520 * 0.48) / 5 / 3600 = 54.7$ user directory requests per second.

The Policy Server makes 54.7 requests, per second, to the user directory when servicing authentication and authorization requests during sustained levels of operation.


Estimate the Peak User Directory Search Rate

The peak user directory search rate is based on the total number of operations (authentication load plus authorization load), specifically, when and at what rate the system is operating at peak levels. Estimating the peak user directory search rate is the process of identifying when the system is servicing the highest level of operations and how these requests translate into user directory searches.

When estimating the peak authorization rate, we recommend using the metrics that you gathered when determining the sustained authorization rate to determine:

- The hour the system is servicing the highest number operations.
- The approximate number of operations that occur during this period.

The following figure is an example of these metrics:



Case: Estimate the Peak User Directory Search Rate

A company has determined the application results in a total of 888,000 operations per day. These operations result in approximately 2,051,520 user directory searches. Using metrics gathered during a capacity planning exercise, the company has determined that during the single busiest hour, approximately 278,000 operations, or 31 percent of the total operations, occurred.

The company uses the following formula to estimate the peak user store search rate.

$$(total_user_directory_requests * percentage_of_requests) / number_of_hours / 3600 = peak_authentication_request_rate$$

total_authentication_requests

Represents the total number of requests the Policy Server sends to the user store.

percentage_of_requests

Represents the percentage of operations that occur when the system is operating at peak levels.

number_of_hours

Represents the number of hours in which the system operates at peak levels.

peak_user_directory_request_rate

Represents the number of requests, per second, that the Policy Server makes to the user store to maintain the peak authentication rate.

Result: $(2,051,520 * 0.31) / 1 / 3600 = 176.6$ requests per second.

The Policy Server makes 176.6 requests, per second, to the user directory when servicing authentication and authorization requests during peak levels of operation.

User Store Capacity Planning

The Policy Server performs a series of services to authenticate and authorize web service request messages. These services result in number of reads and writes, collectively known as requests, to a user directory. A significant contributing factor to CA SiteMinder® Web Services Security performance is determining whether your user directories can handle this workload during sustained and peak periods of operation.

The following general factors influence CA SiteMinder® Web Services Security performance:

- Total web service requests and sustained user directory search rates—The Policy Server must handle an authentication and authorization operation for each incoming web service request. The rate at which these requests occur fluctuate throughout your business day.


In turn, the rate at which the Policy Server makes user directory requests to process the operations fluctuates. Some periods generate relatively few user directory requests, while others generate more.

The sustained user directory search rate represents a period during which the Policy Server makes an average number of user directory requests to service an average number of operations.

- Total requests and peak user directory search rates—During sustained periods of activity, web service request activity can spike. The peak user directory search rate represents a period during which the Policy Server makes the highest number of user directory requests to process peak numbers of authentication and authorization operations.

The following graphic illustrates:

- The relationship between total requests and the user directory search rate.
- How each rate fluctuates throughout the day, is sustained for a specific period, and peaks within that period.



We recommend using the following guidelines to estimate the load under which your user directories have to operate. Once you have estimated the load, you can use any standard tool to create the load on the directory and track the results.

Note: Many factors can contribute to failing to achieve the required numbers. See your vendor–specific documentation for tuning guidance.

More information:

[Policy Server](#) (see page 12)

[How to Estimate a Sustained Request Rate](#) (see page 104)

User Store Capacity Planning Checklist

Estimating the number of user directory requests that the Policy Server must make to service web service requests requires specific information. Gather the following before beginning a user store capacity plan:

- The total number of daily web service requests (request load) for the web service.
- The sustained and peak periods during which web service clients are sending requests to the web service.
- The total number of enabled policies. For each CA SiteMinder® policy determine:
 - If the policy membership filter results in one or more user directory searches.
 - If the responses bound to the policy result in one or more user directory searches.

More information:

[CA SiteMinder® Policy Membership and Authorization Performance](#) (see page 160)

[CA SiteMinder® Web Services Security Capacity Planning Introduced](#) (see page 103)

How to Estimate a Sustained User Directory Search Rate

Estimating a sustained user directory search rate is the process of determining:

- How the total number of user directory requests fluctuate throughout your business day
- How the user directory requests translate into requests per second over a sustained period.

Complete the following steps to estimate the sustained user directory search rate:

1. Use the authentication guidelines to estimate the number of user directory requests that the authentication load creates.
2. Use the authorization guidelines to estimate the number of user directory requests that the authorization load creates.
3. Estimate the sustained user directory search rate.

Use Authentication Guidelines to Estimate Directory Searches

A Policy Server makes a number of user directory requests to service each authentication request. Some of the user directory requests are required, while others can be avoided.

Estimate the number of Policy Server requests that each authentication creates using the following guidelines:

(Required) Two searches to authenticate each user:

- One search/query, per store, to identify the user
- One search/query to verify the user credentials

(Optional) Additional searches may be required depending on how you design policies:

- One search/query for each CA SiteMinder® policy that is bound to a rule that fires when a user is authenticated (OnAuth rule).

Note: For more information about configuring rules, see the *Policy Server Configuration Guide*. For more information about the relationship a rule has to a CA SiteMinder® policy, see the *Policy Server Configuration Guide*.

- One search/query for each CA SiteMinder® policy that is bound to a response that returns user attributes.

Note: For more information about responses and their relationship to rules, see the *Policy Server Configuration Guide*.

Use Authorization Guidelines to Estimate Directory Searches

A Policy Server makes a number of user directory requests to authorize a user. Some of the user directory requests are required to determine CA SiteMinder® policy membership, while others are dependent on CA SiteMinder® policy design. You can estimate the number of Policy Server requests that each authorization creates using the following guidelines.

- One search/query for each CA SiteMinder® policy in the application or policy domain.

Note: This guideline only applies to policies whose membership filter results in one or more user directory requests. For more information about the relationship between CA SiteMinder® policy membership and user directory requests, see Policy Membership and Authorization Requests.

- One search/query for each policy that is bound to a response that returns user attributes.

Note: For more information about the relationship between responses and user directory requests, see Responses and Authorization Performance.

Note: The user authorization cache can significantly reduce the number of authorization-related requests to user directories.

More information:

[User Authorization Cache](#) (see page 161)

Estimate the Sustained User Directory Search Rate

The sustained user directory search rate is based on the total number of operations (authentication load plus authorization load), specifically, when and at what rate these requests occur. The chance that these requests are uniformly spread across your business day is unlikely. Rather, the rate at which these requests occur fluctuates, remaining between the lowest and highest (peak) levels for a sustained period.

Estimating the sustained user directory search rate is the process of identifying:

- A sustained period during which the system is servicing an average number of operations.
- How these requests translate into user directory searches.

When estimating the sustained user directory search rate, we recommend using the daily authentication load and authorization load to identify:

- The rate at which total operations occur throughout the day
Note: We recommend beginning with an evaluation period of 24 hours, broken down into one-hour increments. However, depending on the requirements of your enterprise, you can compare your daily results over a period of weeks or months to gain a better understanding of usage throughout the year.
- The sustained period during which the system is servicing an average number of requests
- The approximate number of requests that occur during the sustained period.

Case: Estimate the Sustained User Directory Search Rate

The company has determined that:

- The daily authentication load and authorization load for the application result in approximately 888,000 total operations.
- The total operations result in approximately 2,051,520 user directory requests.
- The system is operating at sustained levels for approximately five hours (9:00 AM - 2:00 PM).
- During sustained levels, approximately 84,000 operations occur, per hour.
- Approximately 420,000 (84,000 * 5) operations, or 48 percent (420,000 / 880,000) of the total operations, occur during these hours.

The company uses the following formula to estimate the sustained user store search rate:

$$\frac{(total_user_directory_requests * percentage_of_requests) / number_of_hours}{3600} = sustained_user_directory_search_rate$$

total_user_directory_requests

Represents the daily number of requests the Policy Server makes to the user directory to service authentication and authorization requests.

percentage_of_requests

Represents the percentage of total operations that occur when the system is operating at sustained levels.

number_of_hours

Represents the number of hours when the system is operating at a sustained rate.

sustained_user_directory_search_rate

Represents the number of requests, per second, the Policy Server makes to the user directory to maintain the sustained rate of operation.

Result: $(2,051,520 * 0.48) / 5 / 3600 = 54.7$ user directory requests per second.

The Policy Server makes 54.7 requests, per second, to the user directory when servicing authentication and authorization requests during sustained levels of operation.

Estimate the Peak User Directory Search Rate

The peak user directory search rate is based on the total number of operations (authentication load plus authorization load), specifically, when and at what rate the system is operating at peak levels. Estimating the peak user directory search rate is the process of identifying when the system is servicing the highest level of operations and how these requests translate into user directory searches.

When estimating the peak authorization rate, we recommend using the metrics that you gathered when determining the sustained authorization rate to determine:

- The hour the system is servicing the highest number operations.
- The approximate number of operations that occur during this period.

Case: Estimate the Peak User Directory Search Rate

A company has determined the application results in a total of 888,000 operations per day. These operations result in approximately 2,051,520 user directory searches. Using metrics gathered during a capacity planning exercise, the company has determined that during the single busiest hour, approximately 278,000 operations, or 31 percent of the total operations, occurred.

The company uses the following formula to estimate the peak user store search rate.

$(total_user_directory_requests * percentage_of_requests) / number_of_hours / 3600 = peak_authentication_request_rate$

total_authentication_requests

Represents the total number of requests the Policy Server sends to the user store.

percentage_of_requests

Represents the percentage of operations that occur when the system is operating at peak levels.

number_of_hours

Represents the number of hours in which the system operates at peak levels.

peak_user_directory_request_rate

Represents the number of requests, per second, that the Policy Server makes to the user store to maintain the peak authentication rate.

Result: $(2,051,520 * 0.31) / 1 / 3600 = 176.6$ requests per second.

The Policy Server makes 176.6 requests, per second, to the user directory when servicing authentication and authorization requests during peak levels of operation.

Periodic Maintenance Tasks

The following lists details the tasks you can perform for general CA SiteMinder® maintenance. The CA Services implementation team typically covers the details for these tasks, which are based on the specific environment:

- Apply operating system patches.
Frequency: monthly or as required
- Apply CA SiteMinder® cumulative patches.
Frequency: monthly or as required
- Monitor the performance of CA SiteMinder® using the CA SiteMinder® OneView Monitor, CA Wily (or an equivalent tool).
Frequency: continuous
- Monitor the performance of the backend repositories.
Frequency: continuous
- Back up the backend repositories using native or CA SiteMinder® tools.
Frequency: dependent on the requirements of your organization
- Maintain the backend repositories using native tools. Examples of this maintenance include the following items:
 - Indexing
 - Backing up the transaction logs to reduce the consumption of disk space.**Frequency:** dependent on the requirements of your organization
- Remove the tombstones of deleted objects from the policy store by running the XPSSweeper utility.
Frequency: Every 24 hours. This schedule helps reduce the size of the policy store.
- Archive CA SiteMinder® log files.
Frequency: dependent on the requirements of your organization

- Audit CA SiteMinder® policies and adjust/optimize as required.
Frequency: dependent on the requirements of your organization
- Audit authentication and authorization failures. Escalate the events as required.
Frequency: continuous

Chapter 9: Diagnose Implementation Issues

This section contains the following topics:

[Diagnose Issues Introduced](#) (see page 191)

[Policy Server/Policy Store Connection Issues](#) (see page 192)

[Work with Support](#) (see page 193)

[Locate Knowledge Base Articles](#) (see page 202)

[Measure CA SiteMinder® Performance](#) (see page 202)

Diagnose Issues Introduced

The problems you can encounter during a CA SiteMinder® implementation vary and are unique to each environment. Problems can be related to the deployment of individual components to the overall performance of the environment.

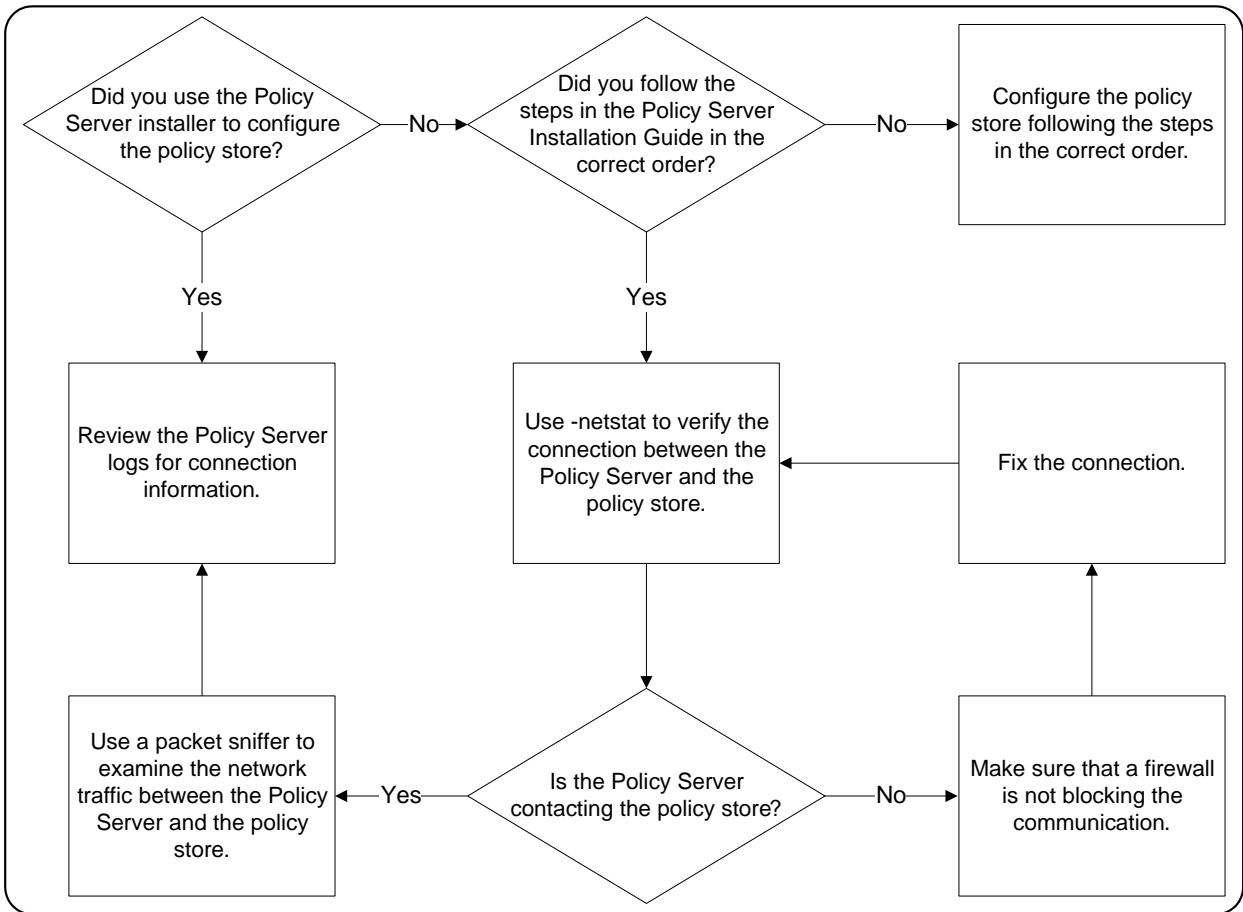
The following sections detail:

- How to diagnose common implementation issues
- How to work with Support to resolve issues efficiently
- Where to locate additional CA SiteMinder® documentation to help troubleshoot issues
- Some tools that you can use to measure CA SiteMinder® performance

Policy Server/Policy Store Connection Issues

Various problems are associated with connecting a Policy Server with a properly configured policy store. These problems can range from an incorrectly configured policy store to network and database connections.

Use the following flowchart to diagnose problems:



Consider the following:

- Using a packet sniffer on the Policy Server host system lets you record error messages that the policy store sends to the Policy Server. If a connection request contains an error message stating that the connection was refused, the database or directory server functioning as the policy store is preventing the connection.

- Reviewing the Policy Server logs lets you identify information about the connections the Policy Server is attempting to make. Common reasons the connections fail include:
 - The Policy Server is using invalid administrator credentials to access the policy store.
 - The administrator account that the Policy Server is using does not have read access.

Note: Policy Server logs are located in *siteminder_home/log*.

siteminder_home

Specifies the Policy Server installation location.

Work with Support

If you require assistance from the CA SiteMinder® Support team, there is specific information you can gather and include when opening a Support ticket. Including as much information as possible helps to reduce the amount of time it takes the Support team to resolve the issue.

Environment Information

Gather as much of the following information as possible and include it when you open a Support ticket:

- The operating system on which the Policy Server is installed, including service pack level.

Example: Windows 2008 SP2
- The web server on which the CA SiteMinder® Agent is installed.

Example: Windows 2008 SP2, IIS 7.0
- The version, including the service pack and the cumulative release (CR), of the Policy Server.

Example: r12.0 SP2 CR1
- The version, including the service pack and the CR, of the CA SiteMinder® Agents communicating with the Policy Server.

Example: r12.0 SP2 CR1
- The policy store type (LDAP/ODBC) and the specific vendor and version.

Example: Oracle 10g R2
- The specific vendor and version of other CA SiteMinder® data stores.

- If applicable, any other CA products, or third-party products integrated with CA SiteMinder®.
- Any custom code or third-party authentication schemes that are deployed in the environment. Custom code includes code provided by Global Solutions Engineering (GSE) or code developed by your organization.
- Any changes that were recently made to the environment, such as an upgraded CA SiteMinder® component or new hardware.
- When the problem started.

Note: You can use the CA SiteMinder® Platform Support Matrix to verify that the issue is not related to an operating system or third-party product that CA SiteMinder® does not support. For more information, see the CA SiteMinder® Platform Support Matrix.

Log Files

Depending on the problem you are experiencing, Support may request one or more of the following log files:

Component	Files
Policy Server	<ul style="list-style-type: none">■ The Policy Server log (smpls.log)■ The Policy Server profiler log (smtracedefault.log)■ The audit log (smaccess.log)
Web Agent	<ul style="list-style-type: none">■ The Web Agent log■ The Web Agent trace log■ The web server error log■ The web server access log
WSS Agent	<ul style="list-style-type: none">■ WSS Agent log■ XML Processing Message Log■ Web Agent trace log (WSS Agent for Web Servers only)■ Application server or web server error log■ Application server or web server access log

Consider the following:

- All Policy Server logs are located in *ps_home*\log.

ps_home

Specifies the Policy Server installation path.

Note: For more information about configuring the Policy Server profiler, see the *Policy Server Administration Guide*. For more information about auditing, see the *Policy Server Administration Guide*.

- Web and WSS Agent logs have no default location or default names.

Note: For more information about configuring Web Agent logging, see the *Web Agent Configuration Guide*. For more information about configuring WSS Agent logging, see the appropriate WSS Agent Guide.

Policy Server Crash

If the Policy Server has crashed, the following lists the information that helps Support probe for additional details. This information is not required to open a Support ticket, but is information that Support is likely to request. If you provide this information initially, it can reduce the amount of time it takes Support to resolve the issue.

1. Provide environment information.
2. Describe the problem in as much detail as possible. For example:
 - How often the process is crashing.
 - The number of times the crash has occurred.
 - A description of what was happening on the server when it crashed.
 - The steps to reproduce the crash.
3. Attach the UNIX core file or Windows dump file. If you are attaching these files, consider the following:
 - (UNIX) If possible, provide a packaged core.
 - (Windows) Be sure that this file is a full dump file and not mini dumps produced by a program error debugging tool.
4. Attach the policy store data.
5. Attach the Policy Server log and the Policy Server audit log.
6. Modify the Policy Server trace log output.
7. Attach the Policy Server profiler log.

Note: If you are attaching log files, be sure that the set of files matches. Also ensure that all the files are from the same time as when the issue occurred.

More information:

[Environment Information](#) (see page 193)

[Log Files](#) (see page 194)

Attach the Policy Store Data

Support is better able to identify the problem by examining the policy store data. Export the policy store and attach the CA SiteMinder® data information file (smdif) to the ticket.

Note: For more information about exporting the policy store, see the *Policy Server Administration Guide*.

Modify the Policy Server Trace Log

Support is better able to identify the problem by examining the Policy Server trace log. If the Policy Server is crashing, you can use the Policy Server profiler to capture the problem.

Note: If the Policy Server is hung, it may not be possible to capture the problem. Instead of using the Policy Server profiler, force a core dump.

The Policy Server profiler uses a default configuration file to log Policy Server actions to a trace log. The default settings include information about components and data:

- Components represent logical groups of actions that the Policy Server executes.
- Data represents the actual pieces of data that the Policy Server must trace.

CA SiteMinder® Support uses component and data settings that are not included in the default configuration file to begin the troubleshooting process. Modify the default settings before submitting the Policy Server trace log.

Note: For more information about configuring the Policy Server profiler, see the *Policy Server Administration Guide*.

Example: Modified Components

Modify the default trace configuration to include the following components:

- Server

The Server component includes additional subcomponents. After you add the Server component, remove the following subcomponents:

- Policy_Object
- Policy_Object_Cache

- Administration
- Audit_Logging
- Tunnel_Service
- JavaAPI

Example: Modified Data Types

Modify the default trace configuration to include the following data types.

Important! The order in which the data types are listed determine the order in which the data is logged. Be sure that the data types are listed in the following order.

- Date
- Time
- Precise Time
- Pid
- Tid
- SrcFile
- Function
- AgentName
- TransactionName
- TransactionID
- Resource
- Realm
- Rule
- Domain
- Group
- Policy
- User
- Directory
- AgentType
- ReturnValue
- ErrorString
- ErrorValue

- AuthStatus
- AuthReason
- AuthScheme
- ClusterID
- RequestIPAddr
- Returns
- Result
- Message

Agent Crash

If an agent has crashed, the following lists the information that helps Support probe for more details. This information is not required to open a Support ticket, but is information that Support is likely to request. If you provide this information initially, it can reduce the amount of time it takes Support to resolve the issue.

1. Gather environment information.
2. Describe the problem in as much detail as possible. For example:
 - How often the process is crashing.
 - The number of times the crash has occurred.
 - A description of what was happening on the server when it crashed.
 - The steps to reproduce the crash.
3. Attach the UNIX core file or Windows dump file. If you are attaching these files, consider the following items:
 - (UNIX) If possible, provide a packaged core.
 - (Windows) Be sure that this file is a full dump file and not a mini dump that has been produced by a program error debugging tool.
4. Attach the agent log and the web or application server error log.

Note: If you are attaching log files, be sure that the set of files matches. Also ensure that all the files are from the same time as when the issue occurred.
5. Attach a tar or zip of the web server binary directory.

Note: This step does not apply to agents running on an IIS web server.
6. For Web Agents or WSS Agents for Web Servers, attach the Web Agent trace log and the web server access log.

More information:

[Environment Information](#) (see page 193)

[Log Files](#) (see page 194)

Resource Leaks

If a system resource, such as memory, file handles, network connections, sockets, or disk space, is not being released, the following lists the information that helps Support probe for additional details. This information is not required to open a Support ticket, but is information that Support is likely to request. If you provide this information initially, it can reduce the amount of time it takes Support to resolve the issue.

1. Gather environment information.
2. Describe the problem in as much detail as possible. Include at least the following:
 - The frequency of the resource leak.
 - The size of the resource leak. Measure the resource leak over a time period with a tool that can show the resource allocation, such as prstat.
 - The tool that you used to measure the resource leak.
 - The effect the resource leak has on the system.
Example: The system crashes or hangs.
 - The steps to reproduce the resource leak or a reproduction test based on the application traffic.
3. Attach logs:
 - (Policy Server) If you are experiencing a Policy Server issue, attach the Policy Server log and the Policy Server audit log.
 - (CA SiteMinder® Agent) If you are experiencing an Agent issue, attach the Agent log and the web server or application server error log.

Note: If you are attaching log files, be sure that the set of files matches. Also ensure that all the files are from the same time as when the issue occurred.

More information:

[Environment Information](#) (see page 193)

[Log Files](#) (see page 194)

Functional Issues

A functional issue is defined as an issue where CA SiteMinder® is not performing as detailed by the documentation. If you are experiencing a functional issue, the following lists the information that helps Support probe for additional details. This information is not required to open a Support ticket, but is information that Support is likely to request. If you provide this information initially, it can reduce the amount of time it takes Support to resolve the issue.

1. Gather environment information.
2. Describe the problem in as much detail as possible, including the steps to reproduce the issue.
3. Attach logs:
 - (Policy Server) If you are experiencing a Policy Server issue, attach the Policy Server log and the Policy Server audit log.
 - (CA SiteMinder® Agent) If you are experiencing an Agent issue, attach the Agent log and the corresponding web server or application server error log.

Note: If you are attaching log files, be sure that the set of files matches. Also ensure that all the files are from the same time as when the issue occurred.

4. Export the policy store to a CA SiteMinder® data information file (smdif) and attach the file.

Note: For more information about exporting the policy store, see the *Policy Server Administration Guide*.

5. Attach logs:
 - (Policy Server) If you are experiencing a Policy Server issue, attach the Policy Server profiler log.
 - (CA SiteMinder® Agent) If you are experiencing an Agent issue, attach all Agent logs and the web server or application server access log.

More information:

[Environment Information](#) (see page 193)

[Log Files](#) (see page 194)

Random Issues

A random issue is defined as an issue that occurs sporadically, and although functional in nature, does not have a pattern that can be reproduced. If you are experiencing a random issue, the following lists the information that helps Support probe for additional details. This information is not required to open a Support ticket, but is information that Support is likely to request. If you provide this information initially, it can reduce the amount of time it takes Support to resolve the issue.

1. Gather environment information.
2. Describe the problem in as much detail as possible. For example:
 - When the issue started.
 - The frequency of the issue.
 - The effect the issue has on the system.
Example: Transactions are taking more time than usual.
3. Attach logs:
 - (Policy Server) If you are experiencing a Policy Server issue:
 - attach the Policy Server log with the point of failure, the Policy Server audit log with the point of failure, and the Policy Server profiler log with the point of failure.
 - attach the Policy Server profiler log with the system functioning correctly.
 - (CA SiteMinder® Agent) If you are experiencing an Agent issue:
 - attach all Agent logs with the points of failure.
 - attach all Agent log with the system functioning correctly.

Note: If you are attaching log files, be sure that the set of files matches. Also ensure that all the files are from the same time as when the issue occurred.

More information:

[Environment Information](#) (see page 193)

[Log Files](#) (see page 194)

Locate Knowledge Base Articles

The CA SiteMinder® bookshelf is only one resource that is available to you. CA SiteMinder® knowledge base (KB) articles are available on the CA Technical Support site. These articles address various topics related to managing and troubleshooting a CA SiteMinder® environment.

To locate CA SiteMinder® KB articles

1. Log into the [Technical Support site](#).
2. Click Support by Product.
The Support by Product page appears.
3. Locate CA SiteMinder® in the product list and click the link.
The CA SiteMinder® product page appears.
4. Enter search criteria under Search Support. Search Support is located on the right side of the screen.
Information matching the search criteria appears.

Measure CA SiteMinder® Performance

Measuring CA SiteMinder® performance is an iterative process that involves gathering metrics that reflect how different components of your deployment are performing. We recommend measuring round-trip times between each pair of components to determine if performance standards are being met and to identify potential bottlenecks.

Note: Avoid using traditional performance metrics, such as CPU usage, as the sole determining factor for tuning a CA SiteMinder® deployment. For example, the system hosting the Policy Server can be running at low CPU usage under load, but this factor does ensure that the Policy Server has reached optimal performance.

Tools you can use to measure CA SiteMinder® performance include:


- Network sniffers
- The CA SiteMinder® OneView Monitor
- The CA SiteMinder® Test Tool
- Directory server utilities and SQL analyzers

Network Sniffers

You can use third-party network sniffers to gather insight into the size and content of a request for unencrypted data without affecting test results. Sniffers can also provide alerts to extra packets being sent, long delays between load balancing, and redirection technologies that logs alone cannot capture.

Note: If the network is set up on a switched hub configuration, place the sniffer between the client and the server on the client-side hub.

The following diagram illustrates a network sniffer in a standard CA SiteMinder® deployment.



CA SiteMinder® OneView Monitor

You can use the CA SiteMinder® OneView Monitor to identify performance bottlenecks and gather metrics about resource usage in a CA SiteMinder® deployment. The OneView Monitor also displays alerts when certain events, such as component failure, occur by collecting operational data from the following CA SiteMinder® components:

- Policy Server
- CA SiteMinder® Agent

The OneView Monitor can identify performance bottlenecks between CA SiteMinder® Agents and the Policy Server by providing metrics such as:

- The average number of authentication attempts and the average time it takes to authenticate a user.
- The average number of authorization attempts and the average time it take to authorize a user.
- The number of cache hits and misses.

Note: For a complete list of the data types the OneView Monitor can provide, see the *Policy Server Administration Guide*. For more information about installing the OneView Monitor, see the *Policy Server Installation Guide*.

CA SiteMinder® Test Tool

You can use the CA SiteMinder® Test Tool utility to test the interaction between CA SiteMinder® Agents and a Policy Server. The Test Tool emulates a Web Agent, which lets you isolate Policy Server performance.

The Test Tool can perform three types of tests:

Functionality

Tests policies to be sure that they are configured correctly.

Regression

Tests whether changes, such as migrating a policy store or implementing a new feature, affect the deployment.

Stress

Test the performance of a Policy Server as it receives multiple requests.

Note: For more information about the Test Tool utility, see the *Test Tool Help*.

Directory Server Utilities and SQL Analyzers

You can use directory server utilities to simulate Policy Server requests to the directory server or database to isolate query lags. You can also use SQL analyzers to analyze response times between the Policy Server and user directories.

Chapter 10: Product Integrations

This section contains the following topics:

[CA Arcot WebFort and RiskFort Integration](#) (see page 205)

[CA Arcot A-OK Integration](#) (see page 215)

[CA DataMinder Content Classification Service Integration](#) (see page 223)

[CA Identity Manager Roles and Access Control](#) (see page 237)

CA Arcot WebFort and RiskFort Integration

You use the CA Arcot Adapter™ (Adapter) to integrate CA SiteMinder® with an on-premise implementation of the CA Arcot WebFort strong authentication solution and the CA Arcot RiskFort adaptive authentication solution.

Consider the following before you begin:

- The integration requires a minimum version of the Adapter and CA Arcot RiskFort.
- The integration requires a minimum version of CA Arcot WebFort.

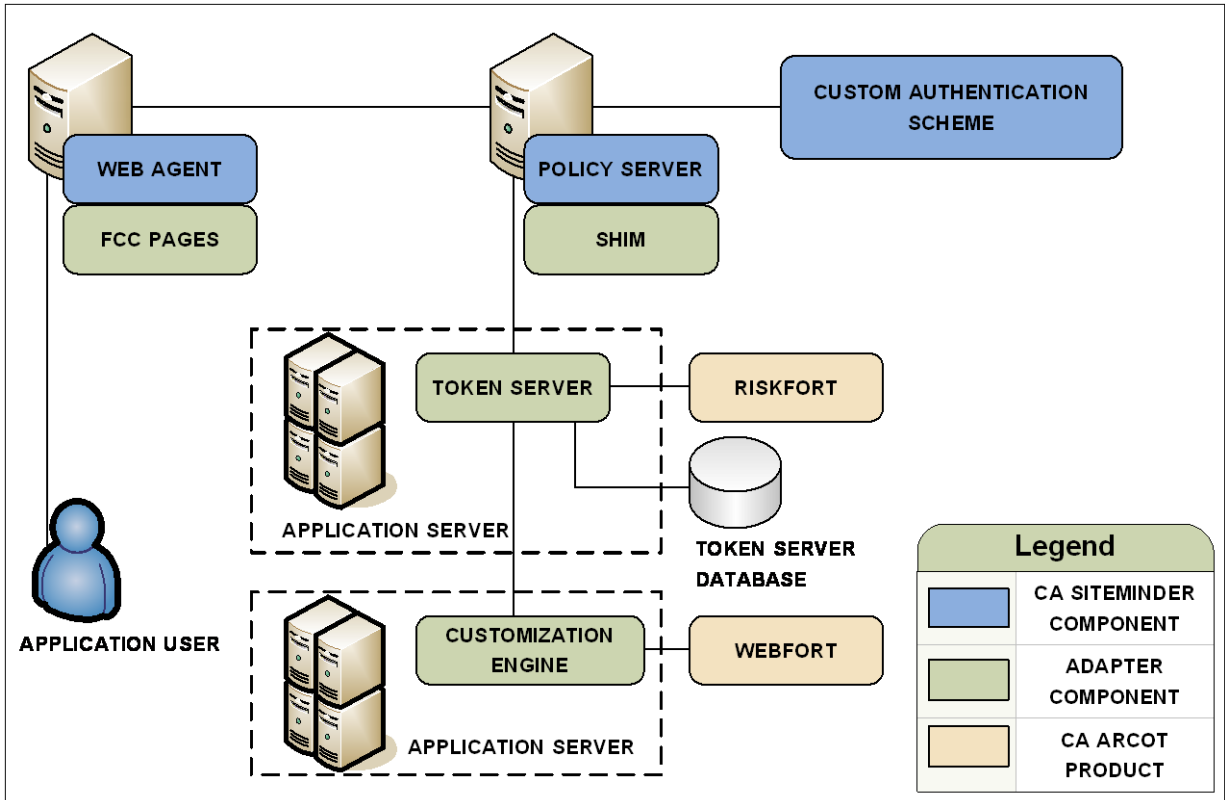
Note: For more information about the supported versions, see the CA SiteMinder® Platform Support Matrix.

The purpose of the following diagram is to:

- Illustrate how the Adapter and its components, CA Arcot RiskFort, and CA Arcot WebFort integrate in a CA SiteMinder® environment.
- Detail the major components and their general relationships. This is not a workflow diagram.

Note: For more information about installing and configuring all CA Arcot components, see the CA Arcot documentation.

Figure 3: CA SiteMinder and CA Arcot integration architecture



Authentication in an On-Premise Arcot Integration

CA Arcot assumes authentication services in an integrated environment by guiding users through the authentication (CA Arcot WebFort) and risk evaluation (CA Arcot RiskFort) processes. During the authentication process:

- CA Arcot WebFort provides the strong authentication, which helps to ensure that the identities of the users requesting the CA SiteMinder®-protected resources are legitimate.

Note: For more information about strong authentication, see the *CA Arcot WebFort Installation and Deployment Guide*. For more information about configuring the supported authentication methods, see the *CA Arcot WebFort Administration Guide*.

- CA Arcot RiskFort collects a range of data to complete a risk evaluation, which determines the level of risk associated with each transaction.

Note: For more information about risk evaluation and risk scores, see the *CA Arcot RiskFort Installation and Deployment Guide*. For more information about configuring risk scoring, see the *CA Arcot RiskFort Administration Guide*.

The result of the risk evaluation is a risk score and corresponding advice, which is a recommend action, such as allow or deny the authentication.

CA Arcot forwards the advice to the Policy Server, which if necessary, continues with its authorization services.

Note: For more information about the Adapter workflow and the role of each CA Arcot component during authentication, see the *CA Arcot Adapter for CA SiteMinder® Installation and Configuration Guide*.

Confidence Levels and CA SiteMinder® Authorization

The Policy Server maintains authorization services in an integrated environment and can apply the risk score to authorization decisions. The risk score is created during the [authentication process](#) (see page 206).

The Policy Server applies the risk score as a CA SiteMinder® confidence level (confidence level). A confidence level is based on a risk score, and as such, is also an integer that represents the likelihood that the transaction is safe.

You can apply a confidence level to both access management models:

- If you are protecting resources with policies, you can apply a confidence level to the following objects:
 - a policy realm
 - an active policy expression
- If you are protecting resources with EPM applications, you can apply a confidence level to the following objects:
 - an application component
 - an application role that is comprised of a named expression that references the SM_USER_CONFIDENCE_LEVEL CA SiteMinder® generated attribute.

Note: Applying a confidence level to a policy realm or an application component requires that you [enable confidence level support](#) (see page 210). Using an active policy expression or an application role to apply a confidence level remains supported from previous releases and is enabled by default. For more information about applying a confidence level to policies and applications, see the *Policy Server Configuration Guide*.

The following example workflow details the relationship between both values and explains how the Policy Server applies a confidence level to authorization decisions:

1. After the user is successfully authenticated, the Adapter converts the risk score to a confidence level using the following algebraic formula:

$$(100 - \text{risk score}) * 10 = \text{confidence level}$$

2. The Adapter inserts the confidence level into the CA SiteMinder® session ticket.

Note: For more information about session tickets, see the *Policy Server Configuration Guide*.

3. As the user requests protected resources, the Policy Server compares the confidence level in the session ticket to the confidence level configured in the policy or application.

4. The following actions can occur:

- If the policy rule is configured to allow access and the confidence level of the user is equal to or greater than the confidence level configured in the policy realm or the active policy expression, the policy rule is triggered.

Note: If the confidence level of the user is less than the confidence level configured in the policy, CA SiteMinder® denies access.

- If the policy rule is configured to reject access and the confidence level of the user is less than the value configured in the policy realm or the active policy expression, the policy rule is triggered.

- If the confidence level of the user is less than the confidence level configured in the application role, the user is excluded from the role membership and CA SiteMinder® denies access.

- If the confidence level of the user is equal to or greater than the confidence level configured in the application component, CA SiteMinder® grants access.

More information:

[Policy Management Models](#) (see page 53)

Risk Scores and Confidence Levels Compared

Although a risk score and a confidence level both help ensure that the transaction is safe, there are differences between both values. Consider the following differences when planning for authorization decisions:

CA Arcot Risk Score	CA SiteMinder® Confidence Level
A numeric scale (0–100) represents a risk score.	A numeric scale (0–1000) represents a confidence level.

CA Arcot Risk Score	CA SiteMinder® Confidence Level
The lower the risk score, the greater the chance that the transaction is safe.	The higher the confidence level, the greater the chance that the transaction is safe. Note: A value of zero (0) represents no confidence. No confidence results in CA SiteMinder® denying access to the requested resource.

The following example workflow details the inverse relationship between a risk score and a confidence level:

1. A user requests a CA SiteMinder® protected resource and is forwarded to CA Arcot for authentication.
2. The Adapter guides the user through authentication and risk analysis. Based on the CA Arcot evaluation and scoring rules, the user is authenticated with a risk score of 30. The lower risk score is representative of a safe transaction.

Note: For more information about risk evaluation and scoring rules, see the *CA Arcot RiskFort Administration Guide*.

3. The Adapter:
 - a. Forwards the authentication decision to the Policy Server
 - b. Converts the risk score to a confidence level using the following algebraic formula:

$$(100 - \text{risk score}) * 10 = \text{confidence level}$$

In this example, the Adapter converts the risk score to a confidence level using the following algebraic formula:

$$(100 - 30) * 10 = 700$$

The higher confidence level is representative of a safe transaction.

4. The Adapter inserts the confidence level into the session ticket of the user.
5. The user requests a resource protected by a policy or an application that requires a confidence level of at least 700.
6. The Policy Server grants access to the resource.

Enable Confidence Level Support for Authorization Decisions

You can optionally apply a confidence level to authorization decisions. Consider the following items:

- You can apply a confidence level to the following objects:
 - A policy realm
 - An active policy expression
 - An application component
 - An application role that includes a named expression, which references the SM_USER_CONFIDENCE_LEVEL CA SiteMinder® generated attribute.

Note: For more information about applying a confidence level to policies and applications, see the *Policy Server Configuration Guide*.

- You only need enable confidence level support to apply a confidence level to a realm or an application component. Using an active policy expression or an application role to apply a confidence level remains supported from previous releases and is enabled by default.

Follow these steps:

1. Log in to any Policy Server host system in the CA SiteMinder® environment.
2. Start the XPSConfig utility.
XPSConfig prompts for an option.
3. Enter SM and press Enter.
XPSConfig prompts for an option.
4. Enter 15 and press Enter.
The ConfidenceLevelSupportEnabled parameter appears.
5. Enter C and press Enter.
The pending value of the parameter appears as True.
6. Quit the XPSConfig utility.
7. Restart the Policy Server.
Confidence level support is enabled.

CA Arcot Integration Use Cases

The following use cases detail how you can integrate CA SiteMinder® with CA Arcot strong authentication and risk evaluation. The use cases begin with a simple integration and progress into more complex scenarios.

CA Arcot Authentication and Risk Analysis

The simplest deployment includes integrating the Adapter and all related components with CA SiteMinder®.

The Adapter guides users through the authentication (CA Arcot WebFort) and risk evaluation (CA Arcot RiskFort) processes to apply a [risk score during authentication](#) (see page 206).

Follow these steps:

1. Be sure that CA Arcot RiskFort and CA Arcot WebFort are installed and configured.

Note: For more information, see the respective CA Arcot installation and deployment guide.

2. Install and deploy the CA Arcot Adapter and all related components. These components include a set of Forms Credential Collector files. These files let you use the Adapter HTML forms authentication scheme to gather user credentials.

Note: For more information about installing and configuring the Adapter and all related components, see the *CA Arcot Adapter for CA CA SiteMinder® Installation and Configuration Guide*.

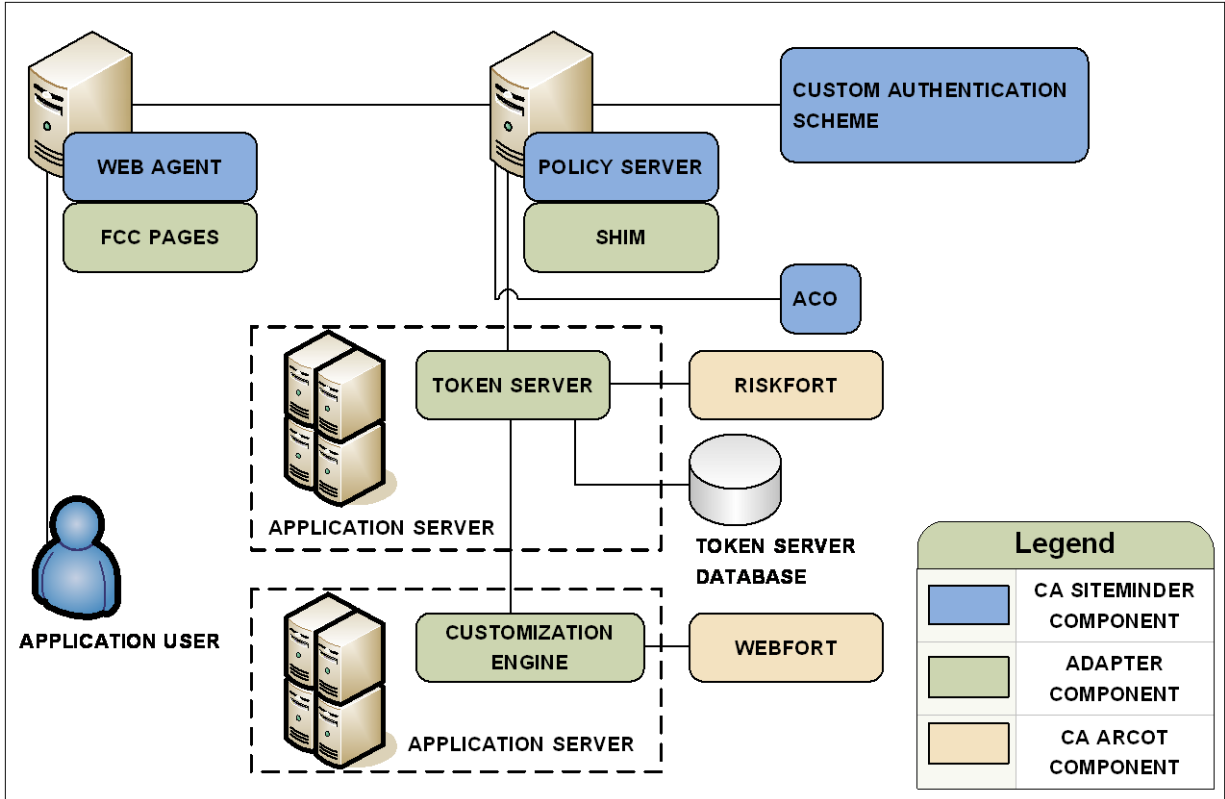
3. Do the following steps:

- a. Configure a CA SiteMinder® Custom authentication scheme to call the Adapter library.
- b. Determine which Web Agents are included in the CA Arcot integration. Configure the respective Agent Configuration Objects (ACO) to support the integration.

Note: For more information about the required custom authentication scheme and ACO settings, see the *CA Arcot Adapter for CA CA SiteMinder® Installation and Configuration Guide*. For more information about configuring an authentication scheme and ACO parameters, see the *Policy Server Configuration Guide*.

The following diagram illustrates this deployment scenario:

Figure 4: CA Arcot authentication and risk analysis



CA SiteMinder® Authentication and CA Arcot Risk Analysis

You can configure the Adapter for risk evaluation only by integrating a CA SiteMinder® authentication scheme. A CA SiteMinder® authentication scheme that is part of the integration is known as backing authentication.

If you use a CA SiteMinder® authentication scheme as backing authentication, the Shim acts as an interface between CA SiteMinder® and the CA SiteMinder® authentication scheme.

Note: For more information about backing authentication, see the *CA Arcot Adapter for CA SiteMinder® Installation and Configuration Guide*. Not all CA SiteMinder® authentication schemes are supported for backing authentication. For more information, see the CA SiteMinder® Platform Support Matrix.

Follow these steps:

1. Complete the steps listed in [CA Arcot Authentication and Risk Analysis](#) (see page 211).

Important! The integration requires that a CA SiteMinder® Custom authentication scheme is configured. The CA SiteMinder® Custom authentication scheme calls the required Adapter library. This library is required even if you are deploying backing authentication.

2. Be sure that you configure the CA SiteMinder® Custom authentication scheme with a valid CA Arcot parameter. This parameter must represent a user flow that supports the CA SiteMinder® authentication scheme that is functioning as backing authentication. You enter this value in in the Parameter field.

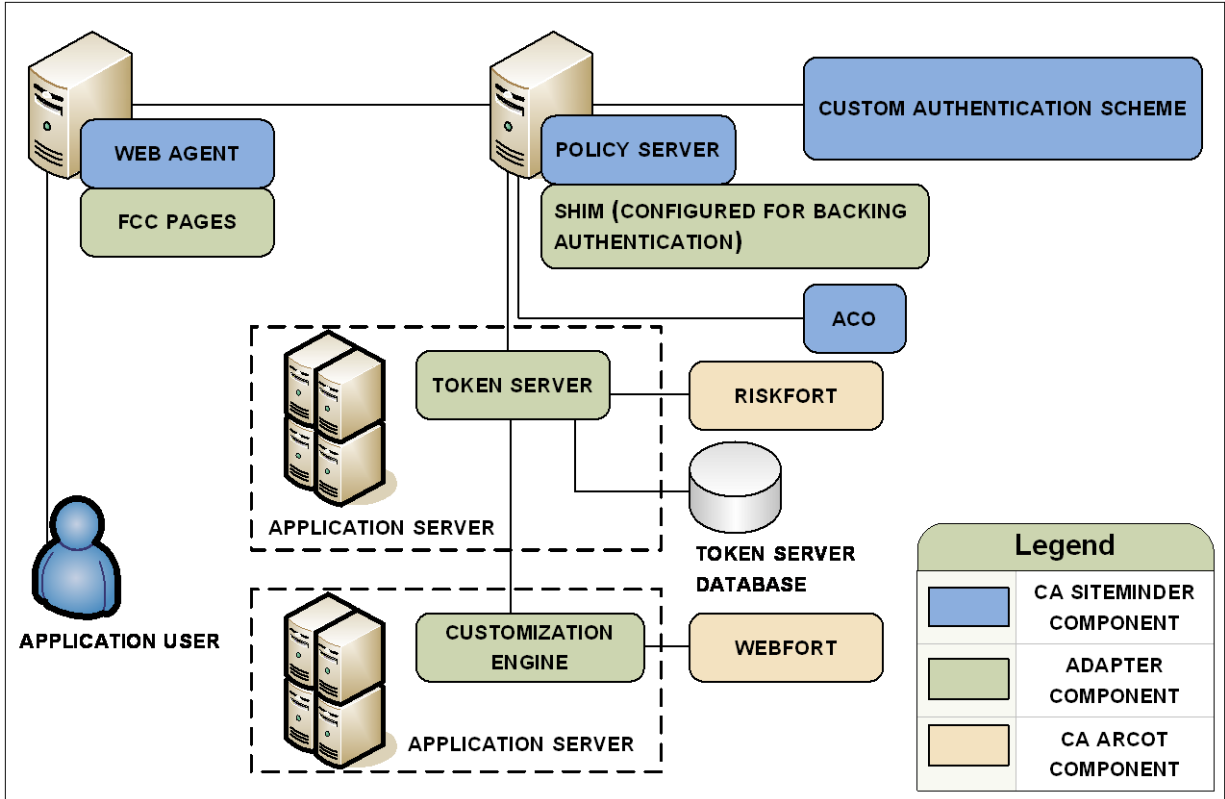
Note: For more information about user flows and the corresponding parameter values, see the *CA Arcot Adapter for CA CA SiteMinder® Installation and Configuration Guide*. For more information about configuring a CA SiteMinder® Custom authentication scheme, see the *Policy Server Configuration Guide*.

3. Configure the Shim to use the CA SiteMinder® authentication scheme as a backing authentication.

Note: For more information about configuring a backing authentication scheme, see the *CA Arcot Adapter for CA CA SiteMinder® Installation and Configuration Guide*.

The following diagram illustrates this deployment scenario:

Figure 5: CA SiteMinder authentication and CA Arcot risk analysis



CA SiteMinder® Authorization and Confidence Levels

You can extend the Policy Server authorization services by adding a [confidence level](#) (see page 207) to both access management models.

Adding a confidence level lets you apply the CA Arcot risk analysis results to authorization decisions.

Follow these steps:

1. Complete the steps in [CA Arcot Authentication and Risk Analysis](#) (see page 211) or [CA SiteMinder® Authentication and CA Arcot Risk Analysis](#) (see page 212).
2. (Optional) If you plan on applying a confidence level to a policy realm or an application component, [enable confidence level support](#) (see page 210). Using an active policy expression or an application role to apply a confidence level remains supported from previous releases and is enabled by default.

3. Do one of the following steps:
 - If you are using policies to protect resources, add a confidence level to one or more policy realms or active policy expressions.
 - If you are using applications to protect resources, add a confidence level to one or more application components or application roles.

Note: For more information about applying a confidence level to policies and applications, see the *Policy Server Configuration Guide*.

More information:

[Policy Management Models](#) (see page 53)

User Store Consideration

All CA SiteMinder® users to which the integration applies must be made available to the CA Arcot WebFort database.

Contact CA Arcot Support for assistance.

Note: For contact information, see the *CA Arcot Adapter for CA SiteMinder® Installation and Configuration Guide*.

CA Arcot A-OK Integration

You use the CA Arcot A–OK Adapter™ (A–OK Adapter) to integrate CA SiteMinder® with the hosted CA Arcot A–OK service.

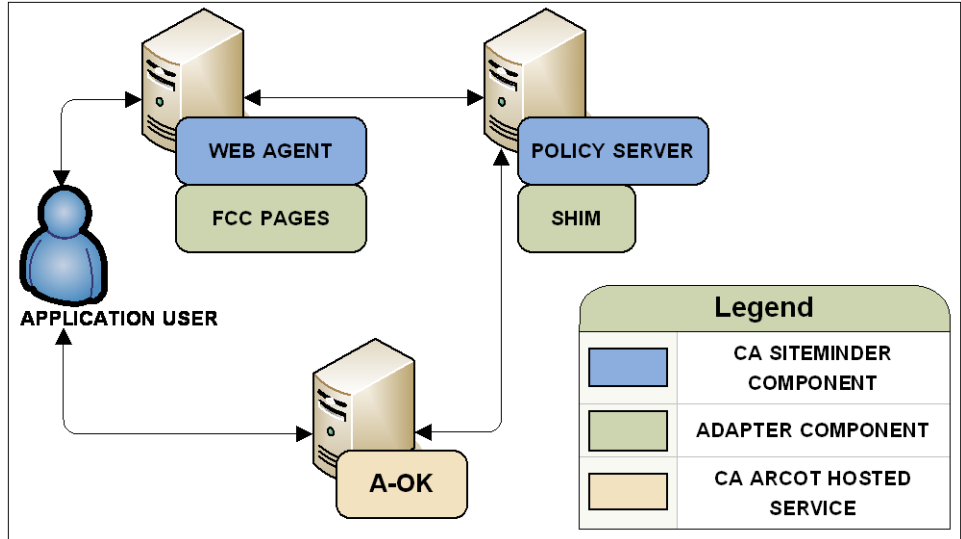
Note: The integration requires a minimum version of the A–OK Adapter. For more information about the supported version, see the CA SiteMinder® Platform Support Matrix.

The purpose of the following diagram is to:

- Illustrate how the A–OK Adapter and its components integrate in a CA SiteMinder® environment.
- Detail the major components and their general relationships. This is not a workflow diagram.

Note: For more information about installing and configuring the A-OK Adapter, see the *CA Arcot A-OK Adapter for CA SiteMinder® Installation and Configuration Guide*.

Figure 6: CA SiteMinder and CA Arcot A-OK integration architecture



Authentication in a Hosted CA Arcot Integration

CA Arcot A-OK assumes authentication services in an integrated environment by guiding users through the authentication and risk evaluation processes. CA Arcot A-OK uses a series of SAML requests and responses to step through the authentication workflow.

Note: For more information about the authentication workflow, see the *CA Arcot A-OK Adapter for CA SiteMinder® Installation and Configuration Guide*.

The result of the risk evaluation is a risk score and corresponding advice, which is a recommend action, such as allow or deny the authentication.

CA Arcot A-OK forwards the advice to the Policy Server, which if necessary, continues with authorization services.

Note: For more information about managing user credentials and configuring the rules associated with the risk evaluation process, see the *CA Arcot A-OK User Administration Guide*.

Confidence Levels and CA SiteMinder® Authorization

The Policy Server maintains authorization services in an integrated environment and can apply the risk score to authorization decisions. The risk score is created during the [authentication process](#) (see page 216).

The Policy Server applies the risk score as a CA SiteMinder® confidence level. A confidence level is based on a risk score, and as such, is also an integer that represents the likelihood that the transaction is safe.

You can apply a confidence level to both access management models:

- If you are protecting resources with policies, you can apply a confidence level to the following objects:
 - a policy realm
 - an active policy expression
- If you are protecting resources with EPM applications, you can apply a confidence level to the following objects:
 - an application component
 - an application role that is comprised of a named expression that references the SM_USER_CONFIDENCE_LEVEL CA SiteMinder® generated attribute.

Note: Applying a confidence level to a policy realm or an application component requires that you [enable confidence level support](#) (see page 219). Using an active policy expression or an application role to apply a confidence level remains supported from previous releases and is enabled by default. For more information about applying a confidence level to policies and applications, see the *Policy Server Configuration Guide*.

The following example workflow details the relationship between both values and explains how the Policy Server applies a confidence level to authorization decisions:

1. After the user is successfully authenticated, the A-OK Adapter converts the risk score to a confidence level using the following algebraic formula:

$$(100 - \text{risk score}) * 10 = \text{confidence level}$$

2. The A-OK Adapter inserts the confidence level into the CA SiteMinder® session ticket.

Note: For more information about session tickets, see the *Policy Server Configuration Guide*.

3. As the user requests protected resources, the Policy Server compares the confidence level in the session ticket to the confidence level configured in the policy or application.

4. The following actions can occur:

- If the policy rule is configured to allow access and the confidence level of the user is equal to or greater than the confidence level configured in the policy realm or the active policy expression, the policy rule is triggered.

Note: If the confidence level of the user is less than the confidence level configured in the policy, CA SiteMinder® denies access.

- If the policy rule is configured to reject access and the confidence level of the user is less than the value configured in the policy realm or the active policy expression, the policy rule is triggered.
- If the confidence level of the user is less than the confidence level configured in the application role, the user is excluded from the role membership and CA SiteMinder® denies access.
- If the confidence level of the user is equal to or greater than the confidence level configured in the application component, CA SiteMinder® grants access.

Risk Scores and Confidence Levels Compared

Although a risk score and a confidence level both help ensure that the transaction is safe, there are differences between both values. Consider the following differences when planning for authorization decisions:

CA Arcot Risk Score	CA SiteMinder® Confidence Level
A numeric scale (0–100) represents a risk score.	A numeric scale (0–1000) represents a confidence level.
The lower the risk score, the greater the chance that the transaction is safe.	The higher the confidence level, the greater the chance that the transaction is safe. Note: A value of zero (0) represents no confidence. No confidence results in CA SiteMinder® denying access to the requested resource.

The following example workflow details the inverse relationship between a risk score and a confidence level:

1. A user requests a CA SiteMinder® protected resource and is forwarded to CA Arcot A-OK for authentication.
2. The A-OK Adapter guides the user through authentication and risk analysis. Based on the CA Arcot A-OK evaluation and scoring rules, the user is authenticated with a risk score of 30. The lower risk score is representative of a safe transaction.

Note: For more information about managing user credentials and configuring the rules that are associated with the risk evaluation process, see the *CA Arcot A-OK User Administration Guide*.

3. The A-OK Adapter:
 - a. Forwards the authentication decision to the Policy Server.
 - b. Converts the risk score to a confidence level using the following algebraic formula:

$$(100 - \text{risk score}) * 10 = \text{confidence level}$$

In this example, the A-OK Adapter converts the risk score to a confidence level using the following algebraic formula:

$$(100 - 30) * 10 = 700$$

The higher confidence level is representative of a safe transaction.

4. The A-OK Adapter inserts the confidence level into the session ticket of the user.
5. The user requests a resource protected by a policy or an application that requires a confidence level of at least 700.
6. The Policy Server grants access to the resource.

Enable Confidence Level Support

You can optionally apply a confidence level to authorization decisions. Consider the following items:

- You can apply a confidence level to the following objects:
 - A policy realm
 - An active policy expression
 - An application component
 - An application role that includes a named expression, which references the SM_USER_CONFIDENCE_LEVEL CA SiteMinder® generated attribute.

Note: For more information about applying a confidence level to policies and applications, see the *Policy Server Configuration Guide*.

- You only need enable confidence level support to apply a confidence level to a realm or an application component. Using an active policy expression or an application role to apply a confidence level remains supported from previous releases and is enabled by default.

Follow these steps:

1. Log in to any Policy Server host system in the CA SiteMinder® environment.
2. Start the XPSConfig utility.
XPSConfig prompts for an option.
3. Enter SM and press Enter.
XPSConfig prompts for an option.
4. Enter 15 and press Enter.
The ConfidenceLevelSupportEnabled parameter appears.
5. Enter C and press Enter.
The pending value of the parameter appears as True.
6. Quit the XPSConfig utility.
7. Restart the Policy Server.
Confidence level support is enabled.

CA Arcot A-OK Integration Use Cases

The following use cases detail how you can integrate CA SiteMinder® with CA Arcot A-OK strong authentication and risk evaluation. The use cases begin with a simple integration and progress into more complex scenarios.

CA Arcot A-OK Authentication and Risk Analysis

The simplest deployment includes integrating the A-OK Adapter and all related components with CA SiteMinder®.

The A-OK Adapter guides users through the authentication and risk evaluation processes to apply a risk score during the [authentication process](#) (see page 216).

Follow these steps:

1. Be sure that the CA Arcot A-OK service is available.

2. Install and deploy the A-OK Adapter and all related components. These components include a set of Forms Credential Collector files. These files let you use the A-OK Adapter HTML forms authentication scheme to gather user credentials.

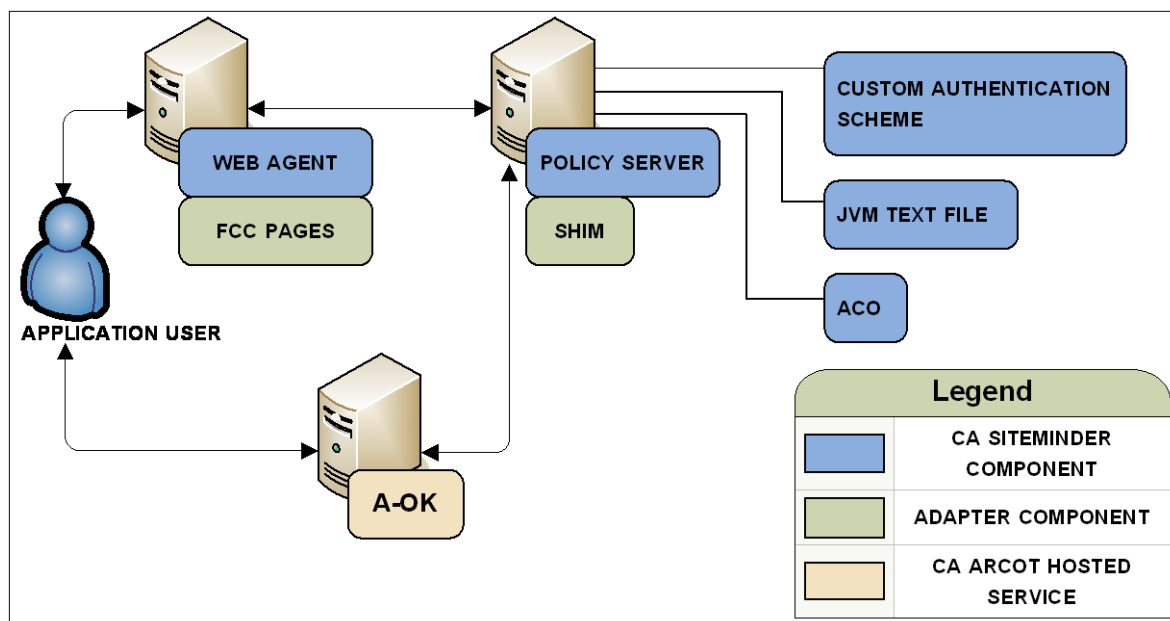
Note: For more information about installing and configuring the A-OK Adapter and all related components, see the *CA Arcot A-OK Adapter for CA SiteMinder® Installation and Configuration Guide*.

3. Complete the following steps:
 - a. Configure a CA SiteMinder® Custom authentication scheme to call the A-OK Adapter library.
 - b. Determine which Web Agents are included in the CA Arcot A-OK integration. Configure the respective Agent Configuration Objects (ACO) to support the integration.
 - c. Add the A-OK Adapter JAR files, certificates, and properties files to the Java Virtual Machine (JVM) file (JVMOptions.txt) of the Policy Server.

Note: For more information about the required custom authentication scheme, ACO settings, and edits to the Policy Server JVM file, see the *CA Arcot A-OK Adapter for CA SiteMinder® Installation and Configuration Guide*. For more information about configuring an authentication scheme and ACO parameters, see the *Policy Server Configuration Guide*.

The following diagram illustrates this deployment scenario:

Figure 7: CA Arcot A-OK authentication and risk analysis



CA SiteMinder® Authorization and Confidence Levels

You can extend the Policy Server authorization services by adding a [confidence level](#) (see page 217) to both access management models.

Adding a confidence level lets you apply the CA Arcot A-OK risk analysis results to authorization decisions.

Follow these steps:

1. Complete the steps in [CA Arcot A-OK Authentication and Risk Analysis](#) (see page 220).
2. (Optional) If you plan on applying a confidence level to a policy realm or an application component, [enable confidence level support](#) (see page 219). Using an active policy expression or an application role to apply a confidence level remains supported from previous releases and is enabled by default.
3. Complete one of the following steps:
 - If you are using policies to protect resources, add a confidence level to one or more policy realms or active policy expressions.
 - If you are using applications to protect resources, add a confidence level to one or more application components or application roles.

Note: For more information about applying a confidence level to policies and applications, see the *Policy Server Configuration Guide*.

More information:

[Policy Management Models](#) (see page 53)

User Store Consideration

All CA SiteMinder® users to which the integration applies must be made available to the CA Arcot A-OK hosted service.

Contact CA Arcot Support for assistance.

Note: For contact information, see the *CA Arcot A-OK Adapter for CA SiteMinder® Installation and Configuration Guide*.

CA DataMinder Content Classification Service Integration

A CA SiteMinder® integration with the CA DataMinder Content Classification Service (CCS) lets the Policy Server use CCS content assessments to make content-aware authorization decisions.

Consider the following items before you begin:

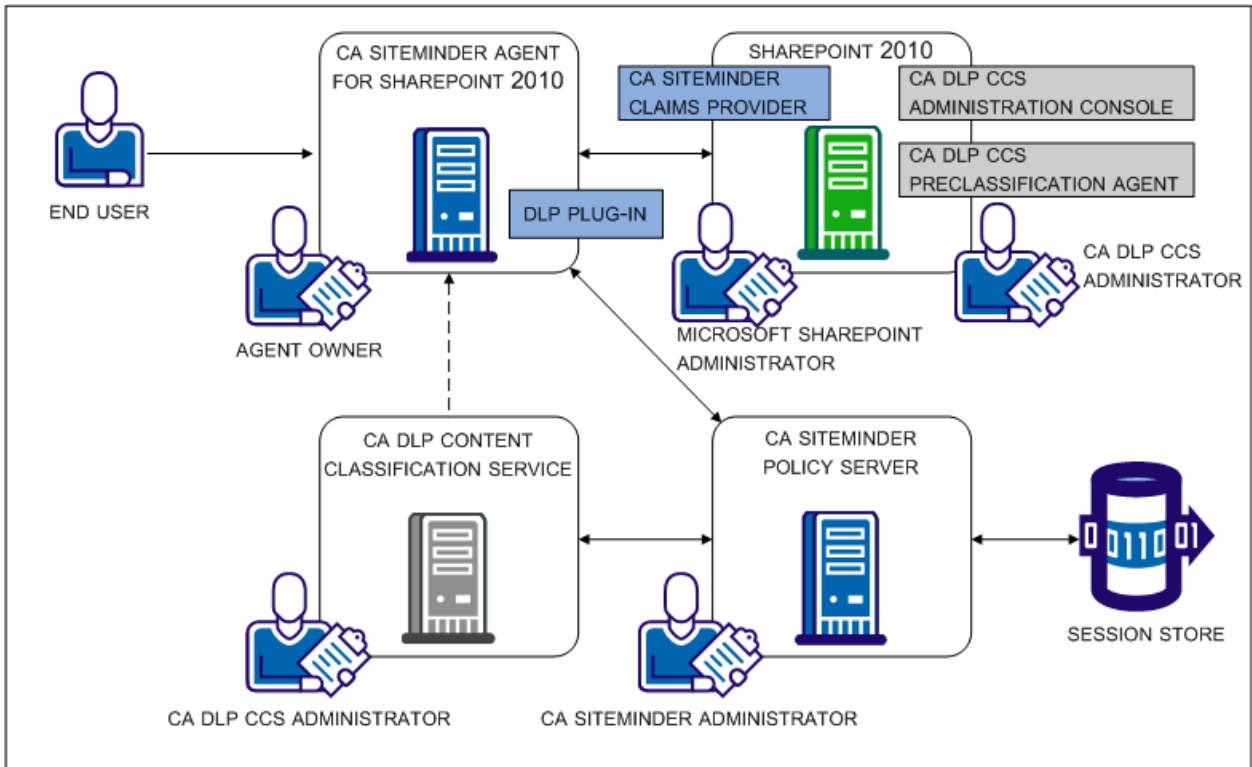
- The integration requires a minimum version of CA SiteMinder®, the CCS, and the CA SiteMinder® Agent for SharePoint.

Note: For more information, see the CA SiteMinder® Platform Support Matrix.

- Multiple organizational roles are required to enable the integration. Coordinate the integration with the following people:
 - A CCS administrator
 - A CA SiteMinder® administrator
 - The owner of the CA SiteMinder® agent for SharePoint

The purpose of the following diagram is to:

- Illustrate the general relationship between the CCS and CA SiteMinder® components in an integrated environment. The diagram is not intended to represent workflow or represent every component deployed in the integrated environment.
- Associate the individuals responsible for installing or configuring a required component.



CA DataMinder Content Classification Service

The role of the CCS in the integration is to make available predefined content classifications to the CA SiteMinder® Policy Server. The classifications correspond to document types commonly found in a corporate environment. The Policy Server uses the classifications to make content-aware authorization decisions.

As the dotted line in [CA DataMinder Content Classification Service](#) (see page 223) illustrates, if a content classification is unavailable at the time of the Policy Server authorization decision, the CCS can request the resource directly to classify or re-classify it. The CCS:

- Passes the result to the Policy Server to make the authorization decision.
- Adds the result to the CCS classification cache for future authorization decisions.

Note: For more information about the CCS and content classifications, see the *CA DataMinder Content Classification Service Integration Guide*. The guide is included in the CA DataMinder Content Classification Service bookshelf.

CA DataMinder Content Classification Service Preclassification Agent

The role of the CA DataMinder CCS preclassification agent in the integration is to scan and classify SharePoint documents offline. Classifying documents offline avoids the need to retrieve a document classification as part of the Policy Server authorization decision.

Note: For more information about the preclassification agent and classification service scans, see the *CA DataMinder Content Classification Service Integration Guide*. The guide is included in the CA DataMinder Content Classification Service bookshelf.

CA SiteMinder® Policy Server

The role of the CA SiteMinder® Policy Server in the integration is to act as the Policy Decision Point (PDP). The Policy Server:

- Maintains all authentication and authorization services in the integrated environment.
- Communicates with the CA SiteMinder® agent for SharePoint to retrieve the resource information of the protected document.
- Communicates with the CA DataMinder CCS to retrieve the content classification of the protected document. The Policy Server uses the results to make a content-aware authorization decision.

If configured to do so, the Policy Server can create a single use security token for the CA DataMinder CCS. The CA DataMinder CCS uses the token to request the resource directly. The CCS requests a resource when it must classify or re-classify it as part of the authorization decision.

Note: For more information about applying content classifications to an Enterprise Policy Management application, see the *Policy Server Configuration Guide*.

CA SiteMinder® Agent for SharePoint

The role of the CA SiteMinder® agent for SharePoint in the integration is to act as the Policy Enforcement Point (PEP). The agent for SharePoint:

- Intercepts the request for the SharePoint document.
- Extracts the resource information of the document.
- Passes the resource information to the Policy Server.

CA SiteMinder® Session Store

The role of the CA SiteMinder® session store is to make available single use security tokens to all Policy Servers in a clustered environment. If configured to do so, a Policy Server creates a security token for the CA DataMinder CCS. The token serves as credentials for the CA DataMinder CCS when it requires access to the protected document.

The CA DataMinder CCS requires access to a protected document when it cannot provide the content classification to the Policy Server. Requesting the resource lets the CCS:

- Classify or re-classify the document.
- Provide the content classification to the Policy Server.
- Add the content classification to the CA DataMinder classification cache for future Policy Server authorization requests.

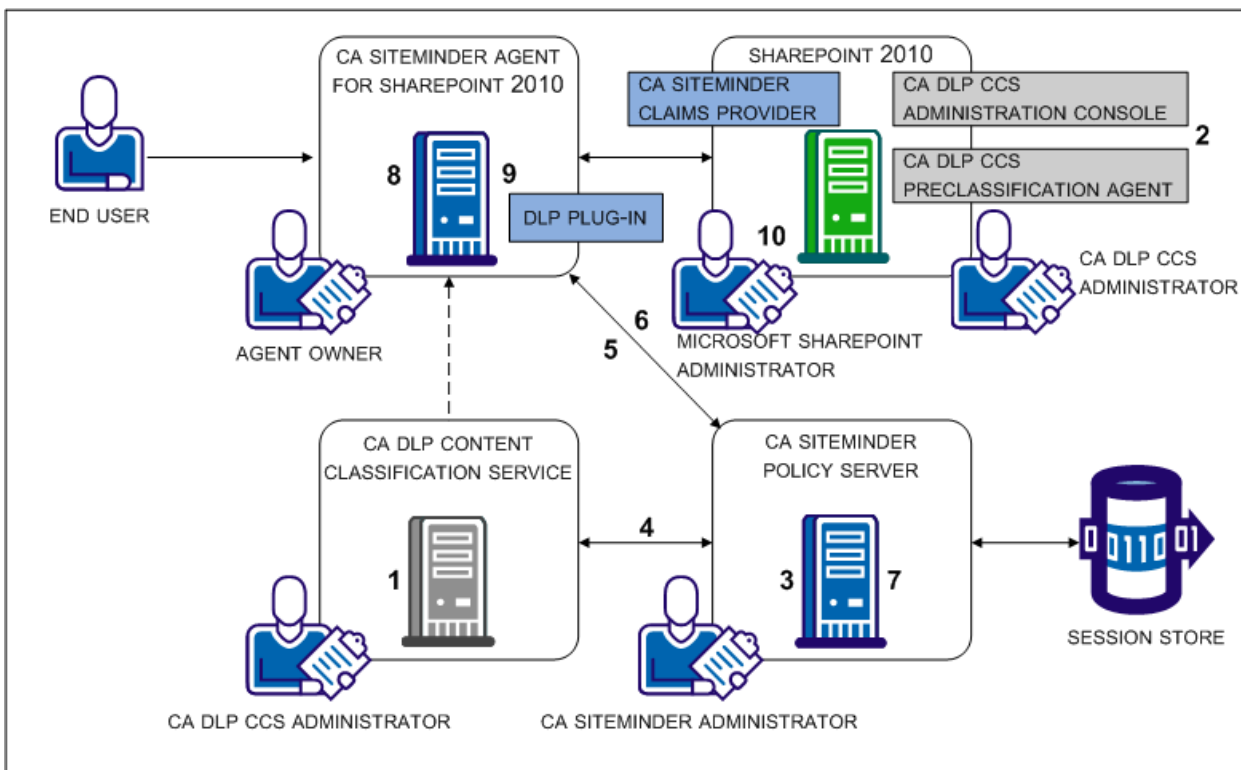
As part of the process, the agent for SharePoint returns the token to a Policy Server to validate authenticity. If the agent for SharePoint sends the validation request to a Policy Server that did not create the token and the environment:

- Includes a session store, the Policy Server retrieves the token, validates it, and authorizes the CA DataMinder CCS.
- Does not include a session store, the Policy Server cannot validate the token and denies the authorization request.

CA DataMinder Content Classification Service Integration Roadmap

The following diagram:

- Illustrates a sample CA DataMinder and CA SiteMinder® integration.
- Lists the order in which each component is installed and configured.



The following table includes each step in the figure and lists the individual responsible for the task.

Step	Action	Responsibility
1	Install and configure the CA DataMinder CCS to communicate over SSL (see page 228).	CA DataMinder CCS administrator
2	Install and configure the CA DataMinder preclassification agent (see page 228).	CA DataMinder CCS administrator
3	Enable SSL for the integration (see page 229).	CA SiteMinder® administrator
4	Configure a connection to the CA DataMinder CCS (see page 229).	CA SiteMinder® administrator
5	Modify the agent for SharePoint agent configuration object (see page 230).	CA SiteMinder® administrator

Step	Action	Responsibility
6	Enable the DLP exclusion list parameter (see page 231).	CA SiteMinder® administrator
7	Enable an authorization failure message (see page 232).	CA SiteMinder® administrator
8	Modify the proxy rules for SharePoint multi-authentication (see page 233).	SharePoint agent owner
9	Enable the DLP plug-in (see page 235).	SharePoint agent owner
10	Provide the CA DataMinder CCS with read access to SharePoint applications (see page 236).	SharePoint administrator

CA DataMinder CCS Administrator Tasks

The CA DataMinder CCS administrator is responsible for:

- Installing one or more CA DataMinder Content Classification Services and configuring each instance to communicate over SSL. The integration requires that the CA DataMinder Content Classification Service and the CA SiteMinder® Policy Server communicate securely.

A CA SiteMinder® administrator requires the CCS server certificate file to enable SSL on Policy Server host systems.

Important! Use the same certificate and password for all CCS instances when configuring them to communicate securely.

- Installing a CA DataMinder CCS preclassification agent to the SharePoint environment and scheduling classification service scans. The CA DataMinder CCS Administration console is installed with the preclassification agent.

Note: For more information, see the *CA DataMinder Content Classification Service Integration Guide*. The guide is included in the CA DataMinder Content Classification Service bookshelf.

CA SiteMinder® Administrator Tasks

The CA SiteMinder® administrator is responsible for enabling the CA SiteMinder® environment for the integration. Complete the integration steps in the following order:

1. Enable SSL for the integration.
2. Configure the connection to the CA DataMinder CCS.
3. Modify the SharePoint agent configuration object.
4. Enable the DLP exclusion list parameter.
5. Enable an authorization failure message.

Enable SSL for the Integration

The integration requires that the CA DataMinder CCS and the CA SiteMinder® Policy Server communicate securely.

- A CA DataMinder CCS administrator is required to configure all CA DataMinder CCS instances to communicate securely. Request the CCS server certificate from the CA DataMinder administrator before you begin. The server certificate is required to enable SSL for the integration.
- Enabling SSL is a local setting. Complete the following procedure for each Policy Server that is protecting SharePoint documents.

Follow these steps:

1. Create a client certificate chain file. A chain file is a single file that contains the certificate file and the respective private key.

Important! The file must be in PEM format.

2. Log in to the Policy Server host system.
3. Deploy the CCS server certificate and client certificate chain file.
4. Navigate to *siteminder_home\bin\thirdparty\axis2c*.
5. Open the following file:
`axis2.xml`
6. Locate the `SERVER_CERT` parameter. Replace the sample value with the path to the CCS server certificate file.
7. Locate the `KEY_FILE` parameter. Replace the sample value with the path to the client certificate chain file.
8. Locate the `SSL_PASSPHRASE` parameter. Replace the sample value with the passphrase used to encrypt the private key in the client certificate chain file.
9. Save the file.

Configure a Connection to a CA DataMinder Content Classification Service

The Policy Server requires a connection to a CA DataMinder CCS to:

- Retrieve the content classification of a protected document.
- Use the content classification to make a content-aware authorization decision.

Configuring the connection is a local setting. Complete the following procedure for every Policy Server that is protecting the SharePoint documents.

Follow these steps:

1. Log in to the Administrative UI with a superuser administrator account.
2. Click Policies, Configure DLP.
3. Select True from the CA SiteMinder® DLP Integration Enabled list.
4. Enter the IP address or fully qualified domain name of the primary CA DataMinder CCS.
5. (Optional) Enter additional configuration parameters.
Note: For more information about the parameters, click Help.
6. Click Save.
7. Restart the Policy Server to enable the Policy Server for the integration and to configure the connection to the CA DataMinder CCS.
8. Restart any Administrative UI that is registered with the Policy Server that has been restarted.

Modify the SharePoint Agent Configuration Object

Modifying the SharePoint agent configuration object configures the agent to extract resource information from the protected document. The agent passes the information to the Policy Server as part of the authorization process.

Follow these steps:

1. Log in to the Administrative UI.
2. Click Infrastructure, Agent Configuration Objects.
3. Locate the agent configuration object for your SharePoint 2010 agents.
4. Click the edit icon to open the object.
5. Enter the following value for the DLPSupportEnabled parameter:
SHAREPOINT
6. Click Submit.
The agent configuration object is enabled for the integration.
7. Contact the agent for SharePoint owner. The agent configuration object is the Policy Server counterpart to the web agent configuration file. A separate procedure is required on the web tier to complete the integration for the agent for SharePoint. The agent for SharePoint owner is responsible for completing the task.

Enable the DLP Exclusion List Parameter

The SharePoint 2010 agent configuration object includes the DLP exclusion list parameter. This parameter contains a set of default resources that the Policy Server excludes from CA DataMinder CCS content classifications. Excluding resources from content classifications indicates to SharePoint agents that the resource can be automatically authorized.

The integration requires that you enable the parameter.

Follow these steps:

1. Log in to the Administrative UI.
2. Click Infrastructure, Agent Configuration Objects.
3. Locate the agent configuration object for your SharePoint 2010 agents.
4. Click the edit icon to open the object.
5. Locate the following parameter:
`#DlpExclusionList`
6. Click the edit icon to open the parameter.
7. Remove the pound sign from the parameter name.
8. If you want to exclude additional resources from content classifications, add the extension to the default set.
Note: Separate the values with a comma.
9. Click OK.
10. Click Submit.

The agent configuration object is enabled.

Enable an Authorization Failure Message

By default, when users fail a DLP content check during authorization, they are redirected to a standard HTTP 403 error message.

Enable authorization failure messages to return an alternate, user-friendly message.

Follow these steps:

1. Create the custom error page using either a text file or an HTML file. Consider the following items:

- You can only redirect users to a custom error page. Applications are not supported.
- If your environment uses Internet Explorer and you are deploying a custom HTML file, include:
 - A style element in the head element.
 - A trailing line before you close the body element.

The HTML file requires these items to prevent Internet Explorer from displaying the standard error message, instead of your custom page.

2. Log in to the Administrative UI.

3. Click Infrastructure, Agent Configuration Objects.

4. Locate the agent configuration object for your SharePoint 2010 agents.

5. Click the edit icon to open the object.

6. Locate the following parameter:

#DlpErrorFile

7. Click the edit icon to open the parameter.

8. Remove the pound sign from the parameter name.

9. Enter the location of the custom error page in the Value field.

Example:

C:\custompages\dlperror.txt

10. Click OK.

11. Click Submit.

The user-friendly message is enabled.

CA Agent for SharePoint Owner Tasks

The CA Agent for SharePoint administrator is responsible for enabling the SharePoint agent environment for the integration. Complete the integration steps in the following order:

1. If SharePoint is configured for multi-authentication mode, modify the proxy rules.
2. Enable the DLP plug-in.

Modify the Proxy Rules for SharePoint Multi-Authentication

If SharePoint is configured for multi-authentication, specific CA SiteMinder Agent for SharePoint proxy rules are required to ensure that the CA DataMinder CCS classifies your SharePoint resources properly.

Contact the Sharepoint administrator to determine if multi-authentication is configured. If multi-authentication is configured, complete the following procedure.

Important! Do not use any other proxy rule settings when the SharePoint environment is configured for multi-authentication. The CA DataMinder CCS request for resources uses an HTTP header for proper forwarding by the CA SiteMinder Agent for SharePoint. If the CA SiteMinder Agent for SharePoint does not properly forward these requests using the following proxy rules, unauthorized access and disclosure of your protected information is possible.

Follow these steps:

1. Locate the following file on your CA SiteMinder Agent for SharePoint:
Agent - for - SharePoint_home\proxy-engine\conf\proxyrules.xml
2. Rename the previous file using a name similar to the following example:
proxyrules_xml_default.txt
3. Open the following file on your CA SiteMinder Agent for SharePoint with a text editor:
Agent - for - SharePoint_home\proxy-engine\examples\proxyrules\proxyrules_example2.xml
4. Save the previous file as a new file in the following location:
Agent - for - SharePoint_home\proxy-engine\conf\proxyrules.xml
5. Locate the following text in the updated proxyrules.xml file:
`://$PROXY_RULES_DTD$"`
6. Replace the previous text with the following text:
`://C:\Program
Files\CA\Agent-for-SharePoint\proxy-engine\conf\dtd\proxyrules.dtd"`

7. Locate the following text:

```
http://www.company.com
```

8. Change the previous text to the domain of your organization. Use the following example as a guide:

```
http:www.example.com
```

9. Locate the following line:

```
<nete:cond type="header" criteria="equals" headername="HEADER">
```

10. Edit the previous line to match the following line:

```
<nete:cond type="header" headername="SMSERVICETOKEN">
```

11. Locate the following line:

```
<nete:case value="value1">
```

12. Edit the previous line to match the following line:

```
<nete:case value="DLP">
```

13. Add a line after the previous line.

14. Copy and paste the following xml syntax onto the new line:

```
<nete:xprcond>
<nete:xpr>
<nete:rule>^/_login/default.aspx\?ReturnUrl=(.*)</nete:rule>
<nete:result>http://sharepoint.example.com:port_number/_trust/default.aspx?tr
ust=siteminder_trusted_identity_provider&ReturnUrl=$1</nete:result>
</nete:xpr>
<nete:xpr-default>
<nete:forward>http://sharepoint.example:port_number$0</nete:forward>
</nete:xpr-default>
</nete:xprcond>
```

15. Replace both instances of the **sharepoint.example:port_number** in the previous section with *one* of the following values:

- The host name, domain and port number of your hardware load balancer. This hardware load balancer operates between your CA SiteMinder Agent for SharePoint server and the SharePoint servers.
- host name, domain and port number of your single web front end. In this context, this web front end (WFE) refers a web server that operates in front of your "back end" SharePoint servers.

16. Replace the instance of *siteminder_trusted_identity_provider* in the previous section with the name of your CA SiteMinder® trusted identity provider.

17. Locate the following line in the file:

```
<nete:forward>http://home.company.com</nete:forward>
```

18. Replace the **home.company.com** in the previous line with *one* of the following values:

- The host name, domain and port number of your hardware load balancer. This hardware load balancer operates between your CA SiteMinder Agent for SharePoint server and the SharePoint servers.
- host name, domain and port number of your single web front end. In this context, this web front end (WFE) refers a web server that operates in front of your "back end" SharePoint servers.

19. Save the file and close your text editor.

The proxy rules are set.

Enable the DLP Plug-in

Enabling the DLP plug-in configures the agent to extract the resource information from the protected document. The agent passes the information to the Policy Server as part of the authorization process.

Important! A separate procedure is required in the application tier to enable the integration. Do not modify the web agent configuration file before the SharePoint agent configuration object is modified. The CA SiteMinder® administrator is responsible for completing the task.

Follow these steps:

1. Log in to the system hosting your CA SiteMinder Agent for SharePoint.
2. Go to the following location:

```
Agent-for-SharePoint_Home\proxy-engine\conf\defaultagent
```

Agent-for-SharePoint_Home

Indicates the directory where the CA SiteMinder Agent for SharePoint is installed.

Default: (Windows) [32-bit] C:\Program Files\CA\Agent-for-SharePoint

Default: (Windows) [64-bit] C:\CA\Agent-for-SharePoint

Default: (UNIX/Linux) /opt/CA/Agent-for-SharePoint

3. Open the following file:

```
WebAgent.conf
```

4. Uncomment (remove the # sign to the left of) the line that loads the disambiguation plug-in.

Example: (Windows [32-bit]) LoadPlugin="C:\Program Files\CA\Agent-for-SharePoint\agentframework\bin\DisambiguatePlugin.dll"

Example: (Windows [64-bit])
LoadPlugin="C:\CA\Agent-for-SharePoint\agentframework\bin\DisambiguatePlugin.dll"

Example: (UNIX/Linux)
LoadPlugin="/opt/CA/Agent-for-SharePoint/agentframework/bin/DisambiguatePlugin.so"

5. Save the file.
6. Restart the web server.

The CA SiteMinder Agent for SharePoint is configured for the CA DataMinder integration.

Microsoft SharePoint Administrator Task

The SharePoint Administrator is responsible for providing the CA DataMinder CCS with read access to the SharePoint applications that CA SiteMinder® is protecting. The CA DataMinder CCS requires read access to determine the types of content that protected documents contain.

Providing read access to the CA DataMinder CCS is local to each application. Complete the following procedure for every application that CA SiteMinder® is protecting.

Follow these steps:

1. If the CA CA SiteMinder® Claims provider is configured, the SharePoint loopback search feature is required. If the feature is not enabled, follow these steps:
 - a. Click Start, All Programs, Microsoft SharePoint 2010 Products, SharePoint 2010 Management Shell.
 - b. Use the management shell to go to the following directory:
C:\Program Files\CA\SharePointClaimsProvider\scripts
 - c. Enter the following command:
. \Set-SMClaimProviderConfiguration.ps1 -EnableLoopBackSearch
 - d. Loopback search is enabled.
2. Log in to SharePoint Central Administration.
3. Locate the Application Management section and click Manage web applications.
A list of applications appears.

4. Select an application and click User Policy in the Web Applications ribbon.
The Policy for Web Application dialog appears.
5. Click Add Users.
The Add Users wizard appears.
6. Select a Time Zone and click Next.
7. Locate the Users field and click the browse icon.
The Select People and Groups – Web Page dialog appears.
8. Locate the CA SiteMinder® trusted identity provider. Under the trusted identity provider, click the associated identifier claim.
9. Enter the following value in the Find field and click the search icon:
caservice
10. Double-click the following user icon and click OK.
caservice
The Add Users dialog appears.
11. Select the following permission and click Finish:
Full Read – Has full read-only access.
The Policy for Web Application dialog appears.
12. Click OK.
The CA DataMinder CCS has read access to the application.

CA Identity Manager Roles and Access Control

Integrating with CA Identity Manager lets you can implement policy-based access control using CA Identity Manager roles. These roles enable centralized management of user privileges in external applications.

Note: For more information about configuring the integration, see the CA Identity Manager documentation.

The integration requires:

- The Policy Server installation includes the required data definitions for the CA Identity Manager integration at the following location.

siteminder_home\mps\dd

siteminder_home

Specifies the Policy Server installation path.

- The name of the file is:

IdmSmObjects.xdd

Important! Do not import this file in to the policy store until you have completed the CA Identity Manager integration. If you import the data definitions before completing the integration, the Policy Server can reach an indeterminate state. Coordinate the integration with your CA Identity Manager administrator.

- CA Identity Manager administrators manage environments and roles in CA Identity Manager to determine user access to the applications that CA SiteMinder® secures. The CA SiteMinder® Administrative UI refers to these environments as an IDM environment.

Note: For more information about environments and roles, see the CA Identity Manager documentation.


- CA SiteMinder® administrators use the Administrative UI to associate one or more IDM environments to a policy domain and user directory. A CA SiteMinder® administrator cannot create or manage an IDM environment.
- CA SiteMinder® administrators use the Administrative UI to create a policy and associate one or more roles that are available to the IDM environments. A CA SiteMinder® administrator cannot create or manage a CA Identity Manager role.

Note: You cannot apply a CA Identity Manager role to an enterprise management application.

CA SiteMinder® can also provide details about entitlements that a CA Identity Manager user has in protected applications. As the following figure illustrates, a CA SiteMinder® administrator associates a response with an access rule in the policy. The response contains a response attribute that specifies a CA SiteMinder®-generated user attribute.

The CA SiteMinder®-generated user attribute retrieves task information from CA Identity Manager. The Policy Server passes this information to the web agent as an HTTP header variable or a cookie. The web agent makes the header variable or cookie available to the protected application for fine-grained access control.

Figure 8: CA Identity Manager and fine-grained access control



Index

A

- Agent Crash • 198
- Agent to Policy Server Communication • 32
- All Components in Multiple Data Centers • 116
- All Components in One Data Center • 115
- Application Server Vendors • 21
- Application Tier Performance • 149
- Applications • 150
- Architectural Considerations • 19, 114
- Architectural Use Cases • 22
- Attach the Policy Store Data • 196
- Auditing and Performance • 164
- Authentication and a Centralized Login Server • 121
- Authentication Guidelines • 154
- Authentication in a Hosted CA Arcot Integration • 216
- Authentication in an On–Premise Arcot Integration • 206
- Authorization Cache • 140
- Authorization Guidelines • 158

B

- Basic Architecture • 47
- Best Practices • 113, 123

C

- CA Agent for SharePoint Owner Tasks • 233
- CA Arcot A–OK Authentication and Risk Analysis • 220
- CA Arcot A-OK Integration • 215
- CA Arcot A-OK Integration Use Cases • 220
- CA Arcot Authentication and Risk Analysis • 211
- CA Arcot Integration Use Cases • 210
- CA Arcot WebFort and RiskFort Integration • 205
- CA Business Intelligence • 13
- CA DataMinder CCS Administrator Tasks • 228
- CA DataMinder Content Classification Service • 224
- CA DataMinder Content Classification Service Integration • 223
- CA DataMinder Content Classification Service Integration Roadmap • 227
- CA DataMinder Content Classification Service Preclassification Agent • 225
- CA Identity Manager Roles and Access Control • 237

- CA SiteMinder® Administrative UI • 18
- CA SiteMinder® Administrator Tasks • 228
- CA SiteMinder® Agent Communicating Across a Data Center • 117
- CA SiteMinder® Agent for SharePoint • 225
- CA SiteMinder® Agent Performance • 133
- CA SiteMinder® Agents • 12
- CA SiteMinder® Audit Database • 17
- CA SiteMinder® Authentication and CA Arcot Risk Analysis • 212
- CA SiteMinder® Authorization and Confidence Levels • 214, 222
- CA SiteMinder® Capacity Planning • 87
- CA SiteMinder® Components and Stores • 11
- CA SiteMinder® Failover and Load Balancing with Multi-Process Web and Application Servers • 145
- CA SiteMinder® Failover and Load Balancing with Multi-Threaded Web and Application Servers • 145
- CA SiteMinder® OneView Monitor • 203
- CA SiteMinder® Policy Design and Performance • 149
- CA SiteMinder® Policy Membership and Authorization Performance • 160
- CA SiteMinder® Policy Objects and Performance Roadmap • 149, 155
- CA SiteMinder® Policy Server • 225
- CA SiteMinder® Session Store • 226
- CA SiteMinder® Test Tool • 204
- CA SiteMinder® Web Services Security Agents • 13
- CA SiteMinder® Web Services Security Authentication Schemes and Authentication Performance • 156
- CA SiteMinder® Web Services Security Capacity Planning • 103
- CA SiteMinder® Web Services Security Capacity Planning Introduced • 103
- CA Technologies Product References • 3
- Caching and Anonymous Users • 141
- Capacity Planning Introduced • 87
- Case
 - Estimate the Peak User Directory Search Rate • 181, 187
 - Estimate the Sustained User Directory Search Rate • 179, 186

Case 1
Policy Membership and User Directory Requests • 176
User Authentication and Directory Requests • 172

Case 2
Policy Design and User Directory Requests • 172
Responses and User Directory Searches • 177

Case 3
Responses and User Directory Requests • 173
Total Directory Requests for Authorization • 178

Case 4
Password Services and Directory Requests • 174

Case 5
Total Directory Requests for Authentication • 174

Central and Local Configurations Together • 66

Central Policy Server Management • 74

Centralize Login Pages • 122

Certificate Data Store • 16

Clustered Components for Scale • 30

Components and Stores • 11

Confidence Levels and CA SiteMinder® Authorization • 207, 217

Configuration Considerations • 111

Configure a Connection to a CA DataMinder Content Classification Service • 229

Contact CA Technologies • 3

D

Data Stores • 14

Data Tier Guidelines • 165

Data Tier Performance • 165

Determine how to Manage Policy Servers • 73, 85

Determine how to Manage SiteMinder WSS Agents • 86

Determine how to Manage Web Agents • 75

Determine if Advanced Encryption Standards are Required • 71, 83

Determine if Partnerships Require CA SiteMinder® Federation • 70

Determine if Virtualization is to be Used • 72, 85

Diagnose Implementation Issues • 191

Diagnose Issues Introduced • 191

Directory Server Utilities and SQL Analyzers • 204

Directory Servers and Databases • 21

Documentation Changes • 4

Domains • 151

Domains and Authentication Performance • 157

E

Embedded Form on a Web Portal • 126

Enable an Authorization Failure Message • 232

Enable Confidence Level Support • 219

Enable Confidence Level Support for Authorization Decisions • 210

Enable SSL for the Integration • 229

Enable the DLP Exclusion List Parameter • 231

Enable the DLP Plug-in • 235

Enhanced Session Assurance Architecture and Performance Considerations • 46

Enterprise Resource Planning Systems • 21

Environment Information • 193

Estimate a Peak Authentication Rate • 93

Estimate a Peak Authorization Rate • 99

Estimate a Peak Request Rate • 108

Estimate a Sustained Authentication Rate • 91

Estimate a Sustained Authorization Rate • 97

Estimate a Sustained Request Rate • 106

Estimate Daily Authentications • 89

Estimate Daily Authorizations • 95

Estimate Daily Requests • 105

Estimate the Peak User Directory Search Rate • 180, 187

Estimate the Size of the User Authorization Cache • 163

Estimate the Sustained User Directory Search Rate • 178, 185

Examples of Relationships between Socket Settings • 137

External Administrative User Store • 15

F

Failover • 33, 39

Functional Issues • 200

G

Group Resources into Domains or EPM Applications • 56

Group Resources into Realms or EPM Components • 57

H

Hardware Load Balancing • 37

How Agent Caches Work • 138

How to Estimate a Sustained Authentication Rate • 89

How to Estimate a Sustained Authorization Rate • 94
How to Estimate a Sustained Request Rate • 104
How to Estimate a Sustained User Directory Search Rate • 170, 183

I

Identify Authentication Methods • 60, 80
Identify Data Centers • 67, 83
Identify Password Management Options • 61
Identify Resources to be Secured with Multiple Cookie Domains • 68
Identify the Applications to Secure • 55
Identify the Web Services to Secure • 78
Identify User Stores • 59, 79
Identify Who Will Manage Your SiteMinder WSS Agents • 81
Identify Who Will Manage Your Web Agents • 63
Ignore Extensions Parameter • 143
Ignore URL Parameter • 144
Implementation Planning Overview • 53
Improve Agent Performance through Load Balancing • 144
Increase NewSocketStep Setting • 136
Increase Request Timeout Interval during Heavy Loads • 135
Increase the Amount of Available Sockets for the Agent • 135

K

Key Store • 16

L

Load Balancing the Application Tier • 164
Load-balancing for SSO between Cookie Provider Domains and Other Cookie Domains • 69
Local Policy Server Management • 73
Locate Knowledge Base Articles • 202
Log Files • 194
Login Page Use Cases • 124
Login Server Controlling User Store Writes • 120

M

Master Policy Store • 41
Measure CA SiteMinder® Performance • 202
Microsoft SharePoint Administrator Task • 236
Minimum Sockets per Port Setting • 136
Modify the Policy Server Trace Log • 196

Modify the Proxy Rules for SharePoint Multi-Authentication • 233
Modify the SharePoint Agent Configuration Object • 230
Multi-Mastered Policy Stores • 42
Multiple Components for Operational Continuity • 26
Multiple Components for Operational Continuity Using Hardware Load Balancing • 29
Multiple Components for Operational Continuity Using SiteMinder Load Balancing • 27
Multiple Data Center Use Cases • 115
Multiple Data Centers • 113

N

Network Sniffers • 203

O

Operating Systems • 19
Other Factors to Consider When Capacity Planning • 110
Other Parameters That Affect Web Agent Performance • 142

P

Password Policy Considerations • 62
Performance Tuning • 129
Performance Tuning Introduced • 129
Performance Tuning Roadmap • 130
Periodic Maintenance Tasks • 188
Plan a CA SiteMinder® Implementation • 53
Plan a CA SiteMinder® Web Services Security Implementation • 77
Policy Management Models • 53, 77
Policy Management Using Application Objects • 54, 78
Policy Management Using Policy Domains and Domain Objects • 54
Policy Management Using Policy Domains and Policies • 78
Policy Objects and Performance • 158
Policy Server • 12
Policy Server Clusters • 35
Policy Server Communicating Across a Data Center • 119
Policy Server Crash • 195
Policy Server Poll Interval Parameter • 142
Policy Server to Audit Store Communication • 44

- Policy Server to Policy Store Communication • 41
- Policy Server to Session Store Communication • 44
- Policy Server to User Store Communication • 38
- Policy Server/Policy Store Connection Issues • 192
- Policy Store • 14
- Possible Architecture 1—Use Existing Components • 48
- Possible Architecture 2—Use Existing Policy Server • 49
- Possible Architecture 3—Full Separation of the Session Assurance Components • 50
- Product Integrations • 205

R

- Random Issues • 201
- Realms • 152
- Realms and Authentication Performance • 157
- Reduce Traffic between Your Agents and the Policy Server • 137
- Redundancy and High Availability • 32
- Replication • 167
- Resource Cache • 138
- Resource Cache and URL Query Strings • 139
- Resource Leaks • 199
- Responses • 153
- Responses and Authorization Performance • 160
- Risk Scores and Confidence Levels Compared • 208, 218
- Round Robin Load Balancing • 34, 40
- Rules and Authorization Performance • 159
- Rules and Rule Groups • 152

S

- Secure Socket Layer and User Directories • 166
- Security Zones • 111
- Server Performance • 132
- Session and Authorization Cache Settings • 141
- Session Cache (authentication) • 140
- Session Store • 17
- Simple Deployment • 22
- Simple Deployment with Optional Agents • 25
- Simple Deployment with Optional Components • 24
- Stand-Alone Login Page • 124
- Static IP Addresses and User Directories • 166
- System Resources • 166

U

- Use Authentication Guidelines to Estimate Directory Searches • 171, 184
- Use Authorization Guidelines to Estimate Directory Searches • 175, 185
- Use Case
 - Capacity Planning • 88, 104
- User Authorization Cache • 161
- User Authorization Cache Efficiency • 162
- User Directories and Authentication Performance • 156
- User Directory Searches • 167
- User Store • 15
- User Store Capacity Planning • 168, 181
- User Store Capacity Planning Checklist • 170, 183
- User Store Consideration • 215, 222

W

- Web Agent and Policy Server Interaction using Apache-based Web Server Pre-Fork Mode • 147
- Web Agent and Policy Server Interaction using Apache-based Web Server Worker Mode • 147
- Web Server Vendors • 20
- Web Servers, Web Agents, and Web Server Processes • 146
- Web Tier Performance • 131
- Web Tier Socket Usage • 134
- Work with Support • 193

Y

- Your Enterprise Environment • 19