Sun Java System Access Manager 7 2005Q4 Deployment Planning Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 819–2136–11
January 2006

060125@13215

Contents

Tables

Figures

Preface

The *Sun Java System Access Manager 7 2005Q4 Deployment Planning Guide* provides planning and deployment solutions for Sun Java™ System Access Manager based on the solution life cycle.

Access Manager is a component of the Sun Java™ Enterprise System (Java ES), a set of software components that provide services needed to support enterprise applications distributed across a network or Internet environment.

Who Should Use This Book

This book is intended for deployment architects and business planners responsible for the planning, analysis, and design of an Access Manager deployment. This book might also be useful for system integrators who are responsible for the design and implementation of the specific aspects of an Access Manager deployment.

Before You Read This Book

Readers should be familiar with the following components and concepts:

- Access Manager technical concepts, as described in the *Sun Java System Access Manager 7 2005Q4 Technical Overview*
- Deployment platform: Solaris™ or Linux operating system
- Web container that will run Access Manager: Sun Java System Application Server, Sun Java System Web Server, BEA WebLogic, or IBM WebSphere Application Server

- Technical concepts: Lightweight Directory Access Protocol (LDAP), Java™ technology, JavaServer Pages™ (JSP) technology, HyperText Transfer Protocol (HTTP), HyperText Markup Language (HTML), and eXtensible Markup Language (XML)

How This Book Is Organized

This book is organized as follows:

Chapter 1 introduces Sun Java System Access Manager.

Chapter 2 describes the business analysis phase of the solution life cycle, when you define business goals by analyzing a business problem and identifying the business requirements and business constraints to meet that goal.

Chapter 3 describes the technical requirements phase of the solution life cycle, when you perform a usage analysis, identify use cases, and determine quality of service requirements for the proposed deployment solution.

Chapter 4 describes the logical design phase of the solution life cycle, when you design a logical architecture showing the interrelationships of the logical components of the solution.

Chapter 5 describes the deployment design phase of the solution life cycle, when you design a high-level deployment architecture and a low-level implementation specification, and prepare a series of plans and specifications necessary to implement the solution.

Chapter 6 describes the implementation phase of the solution life cycle. For example, deploying Access Manager on multiple servers or installing and configuring Access Manager session failover.

Appendix A describes the directory layout after you install Access Manager.

Appendix B describes the password encryption key and how to change it after installation.

Related Books

Related documentation is available as follows:

Access Manager Core Documentation

The Access Manager core documentation set contains the following titles:

- The *Sun Java System Access Manager 7 2005Q4 Release Notes*will be available online after the product is released. It gathers an assortment of last-minute information, including a description of what is new in this current release, known problems and limitations, installation notes, and how to report issues with the software or the documentation.
- The *Sun Java System Access Manager 7 2005Q4 Technical Overview* provides an overview of how Access Manager components work together to consolidate access control functions, and to protect enterprise assets and web-based applications. It also explains basic Access Manager concepts and terminology.
- The *Sun Java System Access Manager 7 2005Q4 Deployment Planning Guide* (this guide) provides information for planning an Access Manager deployment within an existing information technology infrastructure.
- The *Sun Java System Access Manager 7 2005Q4 Performance Tuning Guide* provides information about how to tune Access Manager and its related components for optimal performance.
- The *Sun Java System Access Manager 7 2005Q4 Administration Guide* describes how to use the Access Manager console as well as manage user and service data via the command line interface.
- The *Sun Java System Access Manager 7 2005Q4 Federation and SAML Administration Guide* provides information about the Federation module based on the Liberty Alliance Project specifications. It includes information on the integrated services based on these specifications, instructions for enabling a Liberty-based environment, and summaries of the application programming interface (API) for extending the framework.
- The *Sun Java System Access Manager 7 2005Q4 Developer's Guide* provides information about customizing Access Manager and integrating its functionality into an organization's current technical infrastructure. It also contains details about the programmatic aspects of the product and its API.
- The *Sun Java System Access Manager 7 2005Q4 C API Reference* provides summaries of data types, structures, and functions that make up the public Access Manager C APIs.

- The *Sun Java System Access Manager 7 2005Q4 Java API Reference*provides information about the implementation of Java packages in Access Manager.
- The *Sun Java System Access Manager Policy Agent 2.2 User's Guide* provides an overview of the policy functionality and the policy agents available for Access Manager.

Updates to the *Access Manager Release Notes* and links to modifications of the core documentation can be found on the Access Manager documentation web site (`http://docs.sun.com/app/docs/coll/1292.1`).

Sun Java Enterprise System Product Documentation

For useful information for related products, see the following documentation collections on the Sun Java Enterprise System documentation web site (`http://docs.sun.com/prod/entsys.05q4`):

- Sun Java System Directory Server:

  `http://docs.sun.com/coll/1316.1`
- Sun Java System Web Server:

  `http://docs.sun.com/coll/1308.1`
- Sun Java System Application Server:

  `http://docs.sun.com/coll/1310.1`
- Sun Java System Message Queue:

  `http://docs.sun.com/coll/1307.1`
- Sun Java System Web Proxy Server:

  `http://docs.sun.com/coll/1311.1`

Related Third-Party Web Site References

Third-party URLs are referenced in this document and provide additional, related information.

Note – Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation, Support, and Training

Sun Function	URL	Description
Documentation	`http://www.sun.com/documentation/`	Download PDF and HTML documents, and order printed documents
Support and Training	`http://www.sun.com/support/`	Obtain technical support, download patches, and learn about Sun courses

Typographic Conventions

The following table describes the typographic changes that are used in this book.

TABLE P–1 Typographic Conventions

Typeface or Symbol	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`** `Password:`

TABLE P–1 Typographic Conventions *(Continued)*

Typeface or Symbol	Meaning	Example
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. Perform a *patch analysis*. Do *not* save the file. [Note that some emphasized items appear bold online.]

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P–2 Shell Prompts

Shell	Prompt
C shell prompt	`machine_name%`
C shell superuser prompt	`machine_name#`
Bourne shell and Korn shell prompt	`$`
Bourne shell and Korn shell superuser prompt	`#`

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions.

To share your comments, go to `http://docs.sun.com` and click Send Comments. In the online form, provide the document title and part number. The part number is a seven-digit or nine-digit number that can be found on the title page of the book or at the top of the document.

For example, the title of this book is *Sun Java System Access Manager 7 2005Q4 Deployment Planning Guide*, and the part number is 819-2136.

CHAPTER 1

Introduction to Deployment Planning for Access Manager

Sun Java™ System Access Manager (Access Manager) is part of the Sun Identity Management infrastructure that allows an organization to manage secure access to Web applications and other resources both within an enterprise and across business-to-business (B2B) value chains. This chapter introduces the basic Access Manager deployment planning principles, including:

About Access Manager

Access Manager is a component of Sun Java™ Enterprise System (Java ES), a set of software components that provide services that support enterprise applications distributed across a network or Internet environment. Access Manager provides these major functions:

- Centralized authentication and authorization services using both role-based and rule-based access control
- Single sign-on (SSO) for access to an organization's Web-based applications
- Federated identity support with the Liberty Alliance Project and Security Assertions Markup Language (SAML)

- Logging of critical information including administrator and user activities by Access Manager components for subsequent analysis, reporting, and auditing. Logging is based on the J2SE logging APIs (`java.util.logging`).

Access Manager is also part of the Sun Identity Management Suite, which provides the functions required to use, share, and manage identity information, including directory services, access management, provisioning, and federation. The products in the Identity Management Suite include:

- Sun Java System Access Manager
- Sun Java System Directory Server Enterprise Edition
- Sun Java System Federation Manager
- Sun Java System Identity Auditor
- Sun Java System Identity Manager
- Sun Java System Identity Manager Service Provider Edition

For more information about each component, see the Sun Software web site: `http://www.sun.com/software/`.

The following figure shows the Access Manager, Identity Manager, and Directory Server identity management components.

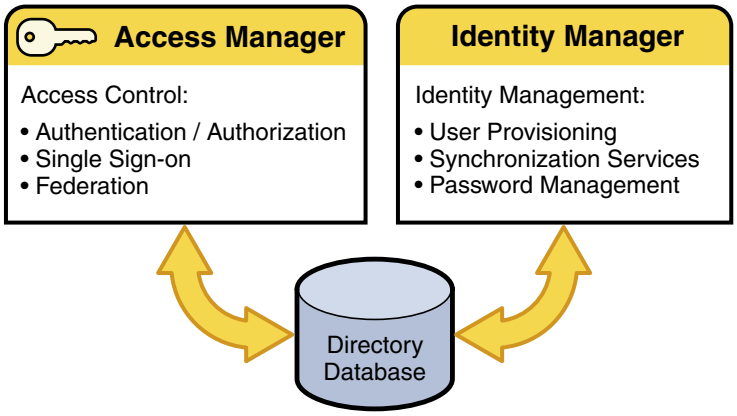

FIGURE 1–1 Sun Identity Management Components

Sun Java System Identity Manager provides user provisioning, password management, synchronization services, comprehensive audit and reporting, and delegated administration. Identity Manager is not a component of Sun Java Enterprise System. To use Identity Manager in your deployment or to obtain more information, contact your Sun Microsystems technical representative or a Sun sales office: `http://www.sun.com/sales-n-service/WWSales.html`.

For a detailed description of Access Manager, see the *Sun Java System Access Manager 7 2005Q4 Technical Overview*.

Access Manager Deployment Planning

Deployment planning is a critical step in the successful implementation of an identity management solution. Each enterprise has its own set of goals, requirements, and priorities to consider. Successful deployment planning is the result of careful preparation, analysis, and design. Errors and missteps that occur anywhere during the planning process can result in a system that can misfire in many ways. Significant problems can arise from a poorly planned system. For example, the system could under-perform, be difficult to maintain, be too expensive to operate, could waste resources, or could be unable to scale to meet increasing needs.

Access Manager deployment planning as described in this guide follows the solution life cycle. The solution life cycle includes the process of planning, designing, and implementing an Access Manager enterprise software solution based on Java Enterprise System.

Solution Life Cycle

The solution life cycle, shown in the following figure, is a useful tool for planning and tracking a deployment project. The life cycle structures the preparation, analysis, and design necessary for successful deployment planning into a series of ordered phases. Each phase consists of related tasks that result in outputs that are carried forward as inputs to subsequent phases. The tasks within each phase are iterative, requiring thorough analysis and design before generating the outputs for that phase.

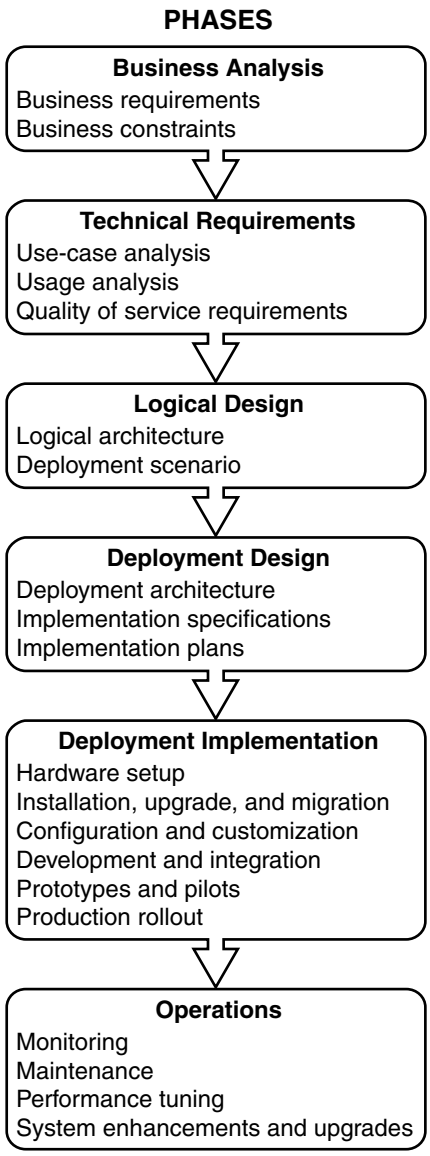

FIGURE 1–2 Solution Life Cycle

The organization of this manual is based on phases within the solution life cycle. The following sections in this chapter briefly describe each life cycle phase. For a more detailed description of these phases, see *Sun Java Enterprise System 2005Q4 Deployment Planning Guide*.

Business Analysis Phase

During business analysis, you define the business goals of a deployment project and state the business requirements that must be met to achieve those goals. When stating the business requirements, consider any business constraints that might affect the ability to achieve the business goal. Without proper business analysis, you run the risk of an incomplete solution.

During the business analysis phase you create business requirements documents that you later use as inputs to the technical requirements phase.

See Chapter 2.

Technical Requirements Phase

The technical requirements phase starts with the business requirements and business constraints defined during the business analysis phase and translates them into technical specifications that can be used to subsequently design the deployment architecture. The technical requirements specify quality of service (QoS) features, such as performance, availability, security, and others.

During the technical requirements phase, you create documents that contain the following information:

- Analysis of user tasks and usage patterns
- Use cases that model user interaction with the planned system
- Quality of service requirements derived from the business requirements, possibly taking into consideration the analysis of user tasks and usage patterns

The resulting usage analysis, use cases, and QoS requirements documents are inputs to the logical design phase of the solution life cycle. The usage analysis also plays a significant role in the deployment design phase.

See Chapter 3.

Logical Design Phase

During logical design, using use cases from the technical requirements phase as inputs, you identify the Access Manger components necessary to implement a solution. You also identify components that provide support to those Java ES components, and any additional custom-developed components necessary to meet the business requirements. You then map the components within a logical architecture that shows the interrelationships among the components. The logical architecture does not specify any hardware required to implement the solution.

The output of the logical design phase is the logical architecture. The logical architecture and the QoS requirements from the technical requirements phase form a deployment scenario, which is the input to the deployment design phase.

See Chapter 4.

Deployment Design Phase

During deployment design, you map the components specified in the logical architecture to a physical environment, producing a high-level deployment architecture. You also create an implementation specification, which provides low-level details specifying how to build the deployment architecture. Additionally, you create a series of plans and specifications that detail different aspects of implementing the software solution.

Project approval occurs during the deployment design phase. During project approval, the cost of the deployment is assessed. If approved, contracts for implementation of the deployment are signed, and resources to build the project are acquired. Often, project approval occurs after the implementation specification has been detailed. However, approval can also occur upon completion of the deployment architecture.

The outputs of the deployment design phase include the following:

- **Deployment architecture.** A high-level design document that represents the mapping of components to network hardware and software.
- **Implementation specifications.** Detailed specifications used as blueprints for building the deployment.
- **Implementation plans.** A group of plans and specifications that cover various aspects of implementing an enterprise software solution. Implementation plans include a migration plan, installation plan, user management plan, test plan, and others.

See Chapter 5

Implementation Phase

During the implementation phase, you work from specifications and plans created during deployment design to build the deployment architecture and implement the solution. Depending on the nature of your deployment project, this guide documents the following tasks:

- Installing Access Manager on multiple host servers
- Configuring an Access Manager deployment as a site
- Using a Load Balancer with Access Manager
- Implementing Access Manager session failover
- Setting session quota constraints
- Enabling session property change notifications
- Tuning a deployment

See Chapter 6.

CHAPTER 2

Business Analysis for Access Manager

Sun Java™ System Access Manager allows an organization to deploy an identity management solution for employees, contractors, customers, and suppliers. During the business analysis phase of the solution life cycle, you define business goals by analyzing a business problem and identifying the business requirements and business constraints to meet that goal. This chapter contains the following sections:

About Business Analysis

Business analysis starts with stating the business goals. You then analyze the business problems you must solve and identify the business requirements that must be met to achieve the business goals. Consider also any business constraints that limit your ability to achieve the goals. The analysis of business requirements and constraints results in a set of business requirements documents.

You use the resulting set of business requirements documents as a basis for deriving technical requirements in the technical requirements phase. Throughout the solution life cycle, you measure the success of your deployment planning and ultimately the success of your solution according to the analysis performed in the business analysis phase.

Defining Access Manager Business Requirements

This section provides specific business requirements to consider for Access Manager (that is, which business requirements imply a need for an Access Manager solution).

Sun Java™ System Access Manager is a complex, distributed identity management system that, when properly deployed, secures access to a wide variety of data and services spanning an enterprise's organizations. To ensure proper control over corporate resources, appropriate planning of the deployment process is required. This chapter offers information about how to plan the deployment, including:

Defining Resources

Because an identity management solution involves a broad variety of systems throughout an organization, proper Access Manager deployment requires a variety of resources. The following corporate resources will be involved or required in the deployment process.

Human Resources

You should consider the various business and political relationships within an organization. A team of individuals should be assembled with a direct or matrixed reporting structure. Typically, Access Manager deployments have small teams that might consist of a project manager and several dedicated System Administrators. These people report to the Team Lead and further up to an owner who has responsibility across a number of related projects and often reports directly to an executive sponsor. This group is often augmented by virtual team members consisting of Sun technical resources, and LOB Application Administrators, which are used as required.

While this structure might not meet your exact needs, it does represent a fairly typical deployment team model. Although not necessarily distinct individuals, the following abstract technical roles representing various skill sets further define a typical Access Manager deployment team.

Executive Sponsors

Successful identity management deployments traditionally cross organizational and political boundaries, which requires buy-in and support from those setting direction for the company. It is critical that executive sponsorship be in place. Planning meetings are an important process for gaining insight from those with a vested interest in the deployment. As the project plan is developed, ensure that its deliverables are inline with the goals of the company as a whole. For example, if cost reduction is a core business driver, collect statistics on current identity management costs and then determine costs such as using the help desk for password resets? Having tangible statistics available can help define a specific return on investment (ROI) as the deployment team attempts to gain executive support. Other company issues that might be relevant include:

- Who benefits from the identity management deployment?
- What organizational problems does an identity management solution solve?
- How does the company address internal issues that might slow the deployment?

Often the identity management concepts and the value of an Access Manager deployment must be related to other executives. A business and technology evangelist can sell the new infrastructure to executives, helping to drive the demand for integration and aid in the acceptance and ultimate success of the infrastructure changes.

Team Lead

A team lead should be chosen as the party responsible for the project's success. The team lead must be in charge and have the authority to make the project's goals happen. The team lead might be a logically distributed role, perhaps between a technical lead, a project manager, and an executive. However you define this role, the goal is to show continued progress and demonstrated success throughout the deployment process to maintain executive sponsorship.

Project Management

A project manager is responsible for the coordination of schedules. The project manager maintains a schedule that correlates the availability of services, support provided by the core IT group and the integration of the various line-of-business (LOB) applications. This person must have strong communication skills and understand the political aspects of the company. The project manager must also balance the needs of the internal customers with the availability of resources in order to support new applications joining the environment.

LOB applications are vital to running an organization. They are generally large programs with capabilities that tie into databases and database management systems. They can include accounting, supply chain management, and resource planning applications. Increasingly, LOB applications are being connected with network applications that have user interfaces and with personal applications such as e-mail and address books.

Systems Analyst

A systems analyst is responsible for assessment and categorization of the various data and services to be integrated into the Access Manager deployment. The systems analyst interviews the LOB application owners and gathers details on technical requirements including platform, architecture, and the deployment schedule. With this information, the systems analyst formulates a plan about how the application will be integrated into the deployment in order to meet their customer's requirements. The systems analyst must be an IT generalist, with broad knowledge of various application architectures and platforms. Detailed knowledge of Access Manager architecture, services, agents, and APIs is also required.

Line-of-Business (LOB) Application Administrators

LOB application administrators are technical specialist with intimate knowledge of, and control over, the LOB application and are responsible for integration of the Access Manager policy agents, or policy enforcement point, into their application. They must clearly communicate the LOB application's architecture, its integration points, and appropriate schedules. They are typically responsible for defining the access control model represented in Access Manager policies. They might perform custom programming to enhance the integration between Access Manager and their application (for example, session coordination). Finally, they are generally responsible for quality assurance (QA) and the regression testing of their application within the newly-deployed environment.

System Administrators

It is critical that appropriate resources are in place to deploy and maintain the availability of Access Manager. System administrators are required at the following levels. Additional administrators might also include a web container administrator who is responsible for the deployment and performance of the software container in which Access Manager is deployed.

Access Manager Administrator

The Access Manager administrator is responsible for the deployment and maintenance of Access Manager. This administrator assures the availability of the common services, provides necessary enhancements to the infrastructure in general, and configures

policies and roles in particular. This administrator also helps support integration efforts by developing guidelines, and offers technical support to the LOB application administrators. An understanding of Java, XML, LDAP, HTTP, and web application architectures is critical.

Directory Server Administrator

Corporate directory services used for authentication and authorization are often already managed by a group within the organization before the Access Manager deployment is even considered. The Directory Server administrator is responsible for the availability of the directory services, as well as for accepting and integrating additions or modifications to the currently defined LDAP schema and identity data, including changes that are required to support the identity management infrastructure.

Hardware, Datacenter, and Network Administrator

Large organizations typically find economies of scale by separating hardware, operating system, data center, and network administration from middleware administration. If this is the case in your company, it is essential that there is clear communication between these various administrators. It may be critical to the deployment's success to have access to certain machines or to establish certain network configurations; keeping these administrators aware of project milestones and requirements can facilitate a smooth rollout.

Independent Software Vendors

Sun Microsystems and other independent software vendors (ISV) are critical partners in the successful deployment of Access Manager. Purchasing packaged software allows an enterprise to diminish and distribute the cost and risk of software development across multiple organizations.

An ISV makes and sells software products that can run on one or more types of computer hardware or operating system platforms. The companies that make the platforms (for example, Sun, IBM, Hewlett-Packard, Apple, or Microsoft) encourage and lend support to the ISV.

It is in the best interest of all parties involved for ISV to develop cooperative relationships and drive successful deployments. Engage Sun technical services and other ISVs to help bootstrap the project and to convey knowledge they have gained from previous Access Manager deployments. Using technical services, as well as an open discussion with your account team (who can act as an intermediary between Access Manager engineers and your deployment team) can help insure your investment and a successful deployment.

Third Party Affiliates

If you are planning on leveraging the Federation Management capabilities of Access Manager, you will be collaborating with external partners and third party affiliates. Consider an initial deployment of this functionality in conjunction with your own internal deployment. In this case, it is important to involve the LOB application that owns the business functionality that will be delivered and to maintain communication with the technical resources of all parties. Your legal counsel can also help to establish a good relationship between involved parties.

Funding

The core IT group is often responsible for the cost of the deployment project. In fact, it is common to have internal funds transferred from an LOB application to the core group in order to fund portions of the identity management project. But, even when a single LOB application group is providing initial funding, the needs of the larger organization should be balanced with the needs of the funding group.

Setting Goals

By setting goals, an organization defines where it wants to be after the Access Manager deployment is finished. The deployment strategy is to plan a roadmap for reaching these objectives and move towards it. Goals are created by defining the expectations of all involved and getting approval early in the process.

In general, an identity management solution enhances security and improves infrastructure manageability while decreasing costs. More specifically, some common goals (and their benefits) that Access Manager allows an organization to meet include:

- Implementation of a scalable infrastructure to meet the expected increase in digital identities (employees, partners, and customers).
- Consolidation of the creation and management of identity profiles with each group controlling their own data.
- Cost reduction through vendor consolidation, user self-management and related administration costs.
- Improvement of security through immediate identity profile termination.
- Improved transparency of security models and access rights.
- Condensing time required for access to critical systems.
- Removal of user rights to critical systems as roles or affiliation within the organization change.

Ultimately, these goals, combined with an understanding of the motivation of all groups involved and information gleaned from a site survey, can be used to design an infrastructure for the deployment. In addition, they can be used throughout the deployment process to keep interested parties engaged and encourage project endorsement.

Gathering Information

A site survey can be used to gather information about the applications and data stores that will be integrated into the deployment. In addition, these departmental interviews help to forge an understanding of the motivation of the groups involved by defining their particular functions and goals. Once collected, the information can solidify buy-in from the executive sponsors as well as serve as a design blueprint. The following groups of individuals can help in a site survey:

- Users provide feedback about the applications they use on a daily basis.
- Human resources provides information about hiring and termination processes.
- Support personnel offer insight into problems that cross organizational boundaries.
- Application administrators and developers can provide technical information about the line-of-business (LOB) applications to be integrated into the deployment.
- Network administrators have knowledge of the organization's technical baseline for performance and standards.

An initial survey might include gathering information about the following items:

- "Business Processes" on page 29
- "IT Infrastructure" on page 30
- "Virtual Data" on page 31

Business Processes

The business processes are the procedures that diverse groups in the organization define to do their job. Processes can include procedures for:

- Issuing payroll
- Purchasing and accounts payable
- Authorizing employee travel
- Departmental budgeting
- Terminating employees

It is imperative to assess these processes because they are generally supported by the applications used by each business unit. Things to consider include:

- Do the current processes cause delays?
- Are there a number of different processes that perform the same function?
- Can processes be standardized across business unit boundaries?
- How complex are the processes? Can they be consolidated or simplified?
- Can the current processes handle organizational changes?

Any changes to be made to the processes should be initiated prior to the beginning of the deployment.

IT Infrastructure

The IT infrastructure includes all the hardware servers, operating systems, and integrated applications that will be integrated into the Access Manager deployment. Consider the following:

- What applications will leverage Access Manager?

  Applications might include critical internal applications such as those for human resources and accounting or less-critical employee portals. Also leveraging the functionality of Access Manager might be external business-to-business applications that deal with both confidential financial information and less confidential sales material, or business-to-consumer shopping carts that are concerned with credit card data and purchase histories.

- What systems will leverage Access Manager?

  Consider the hardware on which applications are being deployed as well as their operating systems. An Access Manager deployment, at the minimum, includes a web container to run the application, a Sun Java System Directory Server (or existing data store), and Access Manager. Additional hardware servers might run their own web containers with corporate resources and on which Access Manager policy agents can be installed for improved security purposes.

- What Access Manager services will each department leverage?

  Consider the default and custom services integrated within Access Manager. Role and policy strategies will have to be mapped and defined for each department. Authentication modules need to be assessed and custom services, if any, need to be developed.

Other technical considerations also include:

- Are there incompatibilities in the infrastructure?
- Does the current system experience slowdowns or down time?
- Are the applications sufficiently secure?
- Are there virus control procedures?
- Can applications be customized based on user entitlements?

For more information, see "Evaluating Applications" on page 32.

Virtual Data

Virtual data is a catch-all phrase for the profiles that will access, the configurations that will be accessible from, and the data that will be secured by Access Manager. Virtual data includes, but is not limited to, user profiles (such as employees or customers), data and service access rules, and other types of corporate data.

- What assets will Access Manager be protecting?

  Access Manager secures access to all types of data and services. An administrator can regulate who can view or configure Access Manager data as well as control access to applications, portals, and services.

- What users will leverage Access Manager?

  Users might include employees, business partners, suppliers, and current or potential customers. Each user will have a profile that includes, at a minimum, their user ID and password. Employees will undoubtedly have larger and more confidential profiles than customers who access external sales information.

- What data will be accessible?

  Data might include public information, internal information, confidential information, and restricted data. Data might also include sales information on an external web site, confidential employee profiles, access rules that protect corporate resources, server configuration information, and federated customer profiles.

- What is the authoritative source of the data?

  Often multiple schemas that define different types of data are used. These definitions need to be reconciled within your deployment. Be aware of data ownership issues, allowing the various LOB applications to maintain control over their data, where appropriate. It is imperative to balance the demands of the satellite groups in order to provide service that is representative of the overall enterprise as all services are critical to the larger organization.

Other technical considerations also include:

- Is the same information defined in multiple attributes?
- Do users have multiple cross-organizational profiles?
- Are the data stores located in front of the firewall?
- Is the data consistent across different data stores?
- How often is new data added or existing data modified?

For more information, see "Categorizing Data" on page 34.

Evaluating Applications

Identity management services are generally provided as a centralized IT function with corporate and business unit applications forming the extended system. Upkeep of this system hierarchy involves a core IT group that manages and maintains the server infrastructure and a satellite group of employees to maintain the LOB applications.

As large organizations often have hundreds (or even thousands) of deployed internal applications, evaluating all of them would be time-intensive and cost-prohibitive. When conducting an application survey, focus on applications that meet the following criteria:

- Are of particularly high value to the organization.
- Would naturally benefit from integration into a single sign-on infrastructure.
- Are indicative of standard programming and deployment platforms within your organization.
- Are generally receptive to the identity management infrastructure.
- Are currently in the early process of deployment and might logically have time lines that coincide with the Access Manager deployment.

You might develop a spreadsheet that can be used to organize the information from the most promising applications. An overall metric can be developed to compare the value of the application to the complexity of its integration. This metric might be considered an application's degree of fitness for deployment. An example of a highly fit application might be a web application that delegates authentication to an application server on which an Access Manager policy agent is installed for security. All user information would be stored in an LDAP directory.

An example of an unfit application might have a text-based interface, running on a mainframe computer. In this case, it would be advantageous to integrate other applications while waiting for a new version of the mainframe application.

The following sections describe types of information that can be gathered when evaluating your organization's applications. This step also helps in determining the resources that will be protected.

Platform Information

General platform information, based on your existing technologies and hardware, can be used to assess the appropriateness of an application as a candidate for integration. Collected platform information might include the following:

- What operating system (including version) do the applications run on?
- Which web containers (including version) do the applications run in?

- What programming model was used to develop the applications ? (such as Java, ASP/.NET, or C)
- Are there plans to upgrade the applications? If so, what is the time line?

LOB applications might also be running third party applications (such as portals, content management databases, or human resource systems). These applications do not always deploy on platforms supported by Access Manager policy agents. If a policy agent is required, determine the deployment criteria of these applications and schedule their deployment based on the availability of a policy agent.

Security Models

It is important to document the existing security models used within the LOB applications. Typically, applications that use external authentication or authorization are candidates for deployment as well as applications that rely on external directory services. Security information might include the following:

- What authentication mechanisms are currently being used?
- Are their special authentication requirements (such as 2-factor authentication)?
- Is there a pluggable interface for external authentication mechanisms?
- What authorization mechanisms are currently being used?
- Can (or should) authorization be externalized?
- What user data repositories are being used? Can these be externalized?
- Who can access the application? Are there existing roles or groups in place? Under what special conditions are they granted access?

Lifecycle of a Session

An identity's session lifecycle is an important topic to consider when evaluating authentication applications. Make sure you have a clear picture of how a user session is created, managed, and destroyed. Clearly document this process because it will be needed during the application's integration.

Customization and Branding

Consider any specific branding or look and feel requirements for the application. Often times, it is important to maintain an individual look and feel or to simply maintain consistency of user experience. Ensure that any customization and branding requirements are noted with your application assessment because time must be scheduled for this activity.

Categorizing Data

Having analyzed your applications and categorized them into fitness levels, you must now begin categorizing the data and services offered by those applications. This information will be used to build a security model. The categorization process itself is a procedure of data and service categorization, followed by cataloguing the existing authentication and authorization systems.

The information collected in Evaluating Applications is used for the former portion of the process. A good methodology might be to organize the collected information into various tiers of security. These tiers would indicate the amount of risk associated with data loss, application compromise, abuse, or other illicit types of access. Using well-defined categories can help to simplify the mapping of resources into a security model incorporating authentication and authorization requirements.

The data or service is separated into four levels of security. The X axis is the data or service and the Y axis is the security level associated with it. Tier 1 is illustrated with a minimal amount of security and might be data applicable to a public web site. Tier 4, on the other hand, requires maximum security and might be financial or human resources (HR) data. Your organization's categorization might have more or less tiers, but this chart shows how typical it is for large amounts of data to have low associated risk, and thus, low security requirements. As risk associated goes up, security requirements also go up. (Usually, there is very little data with high security requirements, and a lot of data with limited or no security requirements.)

The following figure shows the security requirements for data and services within a typical organization.

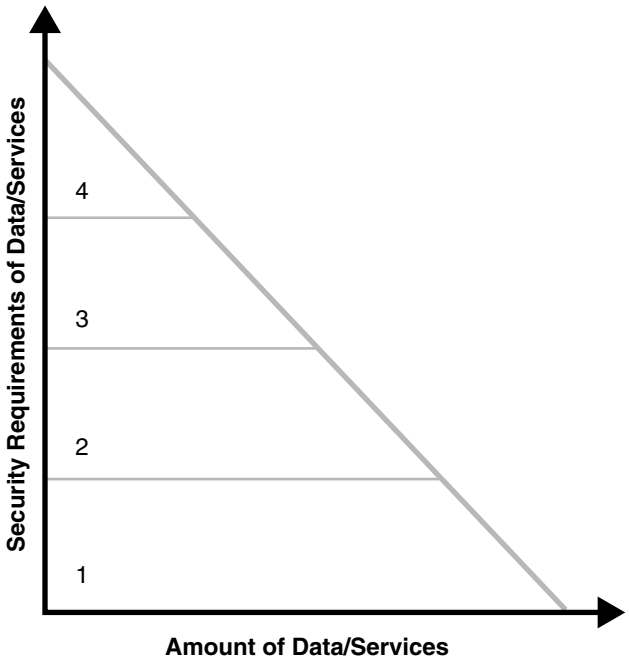

FIGURE 2–1 Security Requirements of Data and Services

Keep in mind that you are planning to build functional groupings of data and service types so that authentication and authorization functions can be mapped to them. Too many tiers can inject extra complexity into your process, while too few tiers might not offer enough flexibility. It is also important to note that there might be data with too much risk to place on the network at all. If relevant, make sure distinctions are made between internal and externally available data. Keep authentication and authorization requirements in mind as you build out these tiers, as well as conditional qualifiers such as such as access time of data and network location.

Mapping To Authentication

With the data categorized according to security level, the next step is to inventory authentication and authorization mechanisms. Using a current list of available authentication mechanisms, associate those mechanisms with the security tiers defined. For example, the following association might be appropriate for the data categorized in the previous figure.

- Tier 1 data might be appropriate for anonymous authentication with no access control.
- Tier 2 data might require password protection only.
- Tier 3 data might require hard token or certificate authentication.

- Tier 4 data might require multi-factor authentication (or might not be placed on the network at all).

You should ensure a clean mapping between authentication requirements and the data and services categorization. If there is none, look for common criteria between those items that do not match. Don't hesitate to make multiple charts if logical distinctions occur.

For example, separate charts can be made for intranet and extranet applications. You might also categorize data based upon a functional security domain such as human resources (HR) or finance. While not a universally applicable tool, categorizing your data in this manner can help you to understand your security requirements and to map them into logically manageable groups.

Mapping To Authorization

Using the data available from your application assessment, examine each of the applications to determine a scalable authorization model. Typically, it is best to look for common groups and roles used across applications. Ideally, these groups and roles will map to functional roles within the organization. You should also determine the source of those groups and roles (where does the membership data live and how is it modeled). For example, the data might be in Sun Java System Directory Server.

If not, custom plug-ins might be required. If a robust grouping model is in place, begin associating each application with existing groups or roles. If not, begin planning a group or role mechanism, finding common relationships between functional user types and access to specific applications. When completed, you should have the following items:

- A clear map of existing groups and roles.
- A clear understanding of where that data lives and who is the authority over its quality and management.
- A clear understanding of new groups or roles that need to be created to facilitate your deployment or to reduce cost and complexity of the deployment.
- A mapping of existing and future grouping mechanisms to your categorized applications.
- Notes on additional conditions required by the applications to allow access to a certain group or role.

With this basic security model (categorization of data, with correlation to authentication and authorization mechanisms), you can now put together a time line to drive your deployment.

Building a Time Line

From the information you have gathered, you should build a preliminary time line. The following sections describes the steps to build a time line for a generic schedule of deployment.

Deployment Design

This phase of the time line is where the concepts, business needs, and user requirements are put in their proper context. A total view of the deployment takes shape. Components are described, technological requirements are defined, and a complete architecture is mapped out. Storyboarding login screens or creating data flow charts are two ways of initiating this design phase.

Proof-of-Concept

A proof-of-concept enables the design to be tested in a business environment. Organizations often have a test case database, a set of pre-configured test cases coupled with their expected results. The proof-of-concept can be applied to this test database, and, if all goes well, the documented results will be equivalent to the new results. A proof-of-concept aims to answer all question posed by the Deployment Design, proving that it meets all needs efficiently and with minimal risk.

It is generally fast allowing for ample time to refine the design based on a limited set of data. There are usually several rounds of proof-of-concept, followed by design refinement. The last round in the proof-of-concept should be integration of some internal applications. Integration of a corporation's shared services often adheres to a standard model of sign-on by early adopters, followed by general participation and, lastly, the stragglers. Demonstrated success with early adopters makes it easier to use those applications as references for general adoption.

Early Adoption

Mission critical or revenue-building applications should not be chosen as your first application. A less risky strategy is to choose an important application that will not completely disrupt business operations if there are issues during roll-out. For example, a divisional portal serves as a natural staging ground for a single sign-on (SSO) roll out, rather than an accounting system at the close of a fiscal period.

Also, limit the number of applications roll-outs in the early phases so process flaws can be driven out, results demonstrated, and immediate success recognized. Minimizing organizational risk while maximizing visibility is the optimal roll-out strategy. This plan positions the deployment team with the appropriate product experience to take on critical applications.

General Participation

Although the deployment project begins with a single application, the requirements of other internal customers should be assessed at the same time so that a general purpose system can be built. The central IT group should be able to accommodate the diverse criteria and schedules of the satellite groups in order to provide service that is representative of the larger organization. Schedules must have sufficiently large windows, allowing the satellite groups time to build changes and upgrades into their application's deployment and quality assurance (QA) cycle.

Production Environment

Following the proof-of-concept, the refined design can be replicated into a production environment. The purpose of a production environment is to demonstrate the function of the design in a non-artificial environment, ensuring its proper behavior. The environment is compared to the behavior as observed in the proof-of-concept, and as defined in the deployment design. It is also tested for stability.

An assessment is made and reports are generated. Early adoption applications go live in the production environment as they are ready. Incrementally phase new applications through the test phase and into production. Other applications are incrementally added to the production environment by working them through the proof-of-concept cycle as the early adopters have been.

Sample time lines are not available, because they vary based upon project complexity. However, this process typically takes place in a span of two to three months.

Deployment Road Map

Mapping out your Access Manager integration is imperative to ensuring its success. This process include collecting information concerning hardware, currently deployed applications, identity data, and access hierarchy. Access Manager deployment can be broken down into the following phases:

1. Identify business objectives such as:
    - Increase operational efficiency.
    - Assure data security.
    - Assure continued productivity by understanding the scope and relationships within the organization and analyzing the behavioral changes needed to support the business objectives.

2. Develop a high-level technology analysis and map it to the business objectives by listing technology services and tools needed to meet business objectives.
3. Define initiatives for each technology service such as:
   - Storing employee history and data accumulated through personalization.
   - Accomplishing password synchronization and identity administration through identity management.
   - Realizing enterprise security through the development of role strategies.
4. Prioritize initiatives based on items such as statistical accuracy, predictability, scope, cost, impact, complexity, behavior, infrastructure, benefit, support, and dependencies.

CHAPTER 3

Technical Requirements

During the technical requirements phase of the solution life cycle you perform a usage analysis, identify use cases, and determine quality of service requirements for the proposed deployment solution. This chapter provides a high-level technical overview of the requirements related to this process for Sun Java™ System Access Manager 7 2005Q4, including:

Deployment Options

There are several key factors that an organization should consider when planning for an Access Manager deployment. These considerations generally deal with risk assessment and a growth strategy. For example:

- How many users is your deployment expected to support, and what is your projected growth rate?

  It is critical that user growth and system usage are monitored and that this data is compared with the projected data to ensure that the current capacity is capable of handling the projected growth.
- Do you have plans to add additional services that might impact the current design?

  The architecture being developed now is optimized for the current service. Your future plans should also be examined.

In addition, the architecture should provide a foundation for the objectives detailed in the following sections.

Security

Consider the following options when you are planning for a secure internal and external networking environment:

- Server-based firewalls provide an additional layer of security by locking down port-level access to the servers. As with standard firewalls, server-based firewalls lock down incoming and outgoing TCP/IP traffic.
- Minimization refers to removing all unnecessary software and services from the server in order to minimize the opportunity for exploitation of the vulnerabilities of a system.
- A Split-DNS infrastructure has two zones that are created in one domain. One zone is used by an organization's internal network clients, and the other is used by external network clients. This approach is recommended to ensure a higher level of security. The DNS servers can also use load balancers to improved performance.

High Availability

Deployments strive for no single point of failure (SPOF) as well as continuos availability to its users. Different products achieve availability in different ways; for example, clustering or multi-master replication. The desired high availability refers to a system or component that is continuously operational for a specified length of time. It is generally accomplished with multiple host servers that appear to the user as a single highly available system. In a deployment that meets the minimal requirements (all applications on a single server), the SPOFs might include:

- Access manager web container
- Directory Server
- Java™ Virtual Machine (JVM)
- Directory Server hard disk
- Access Manager hard disk
- Policy agents

Planning for high availability centers around backup and failover processing as well as data storage and access. For storage, a redundant array of independent disks (RAID) is one approach. For any system to be highly available, the parts of the system should be well-designed and thoroughly tested before they are used. For example, a new application program that has not been thoroughly tested is likely to become a frequent point-of-breakdown in a production system.

Clustering

Clustering is the use of multiple computers to form a single, highly available system. Clustering is often crucial for the Sun Java System Directory Server data store. For example, a clustered multi-master replication (MMR) server pair can increase the availability of each master instance by ensuring availability.

Scalability

Horizontal scaling is achieved by connecting multiple host servers so they work as one unit. A load balanced service is considered horizontally scaled because it increases the speed and availability of the service. Vertical scaling, on the other hand, is increasing the capacity of existing hardware by adding resources within a single host server. The types of resources that can be scaled include CPUs, memory, and storage. Horizontal scaling and vertical scaling are not mutually exclusive; they can work together for a deployment solution. Typically, servers in an environment are not installed at full capacity, so vertical scaling is used to improve performance. When a server approaches full capacity, horizontal scaling can be used to distribute the load among other servers.

Hardware Requirements

The minimum configuration for an Access Manager deployment is a single host server running Access Manager and a web container such as Sun Java System Web Server. Directory Server can be running on the same server or on a different server. In a multiple server deployment, Access Manager instances and their respective web containers are installed on a different host servers, with a load balancer distributing client requests to the various Access Manager instances. Usually, Directory Server and Access Manager are installed on different servers.

For optimum performance, run Access Manager on a 100 Mbytes or greater Ethernet network. A minimum configuration Access Manager deployment (a single server running Access Manager and a web container) should have one or more CPUs, with greatly diminishing returns on processor performance after four CPUs. Two to four CPUs per host server are recommended. A minimum of 512 Mbytes of RAM is necessary for basic testing of the software.

For an actual deployment, 1 Gbytes of RAM is recommended for threads, the Access Manager SDK, the HTTP server, and other internals; 2 Gbytes for basic operation and object allocation space, and 100 Mbytes per 10,000 concurrent sessions. Each Access Manager is recommended to cap out at 100,000 concurrent sessions, after which horizontal load balancing should be used (assuming the 4 Gbytes memory limitation of 32–bit applications).

Software Requirements

Access Manager has specific minimum software requirements, including:

For the latest information about the software requirements, including supported releases, any required patches, and known limitations, see the *Sun Java System Access Manager 7 2005Q4 Release Notes*.

Operating System Requirements

Access Manager 7 2005Q4 is supported on these platforms:

- Solaris™ Operating System (OS), SPARC® based systems, versions 8, 9, and 10
- Solaris™ OS on x86 platforms, versions 9 and 10
- Red Hat™ Linux, Advanced Server and Enterprise Server

For information about downloading OS patches and patch clusters, see SunSolve Online at `http://sunsolve.sun.com/`.

To list the patches currently installed on a Solaris system, use the `showrev -p` command.

Web Container Requirements

Access Manager 7 2005Q4 supports the following web containers for either a full installation or an SDK-only installation:

- Sun Java System Web Server
- Sun Java System Application Server
- BEA WebLogic
- IBM WebSphere Application Server

For the supported versions of these web containers, see the *Sun Java System Access Manager 7 2005Q4 Release Notes*.

When a policy agent is installed on an Access Manager web container, it uses approximately 10 Mbytes of disk space. This additional space must be considered when configuring the web container. For more information, see the *Sun Java System Access Manager Policy Agent 2.2 User's Guide*.

Directory Server Requirements

Access Manager 7 2005Q4 has the following requirements for an LDAP directory server:

- The Access Manager information tree, which contains the following information, is stored in Sun Java System Directory Server:
  - How users authenticate
  - Which resources users can access
  - What information is available to applications after users are given access to resources
- The Access Manager identity repository is used to store user data such as users and groups. Access Manager 7 2005Q4 can use Sun Java System Directory Server or an LDAP version 3 (LDAP v3) compliant directory server as the identity repository.

For more information about the Access Manager information tree and identity repository, see the *Sun Java System Access Manager 7 2005Q4 Technical Overview*.

Java Development Kit (JDK) Software Requirements

For the specific version of the JDK software required by Access Manager 7 2005Q4, see the *Sun Java System Access Manager 7 2005Q4 Release Notes*.

Access Manager Session Failover Requirements

If you are planning to implement Access Manager session failover, these components are required:

- Web container to run Access Manager: Sun Java System Web Server, Sun Java System Application Server, IBM WebSphere Application Server, or BEA WebLogic.
- Sun Java System Directory Server. All Access Manager instances must access the same Directory Server.
- Sun Java System Message Queue. The Message Queue broker cluster manages the session messages between Access Manager instances and the session store database.

- Berkeley DB by Sleepycat Software, Inc. (`http://www.sleepycat.com/`) is the default session store database. Use the version that is distributed with the Sun Java Enterprise System 2005Q4 release.

Access Manager session failover is supported on the following platforms:

- Solaris™ OS, SPARC® based systems, versions 8, 9, and 10
- Solaris™ OS on x86 platforms, versions 9 and 10
- Red Hat™ Linux, Advanced Server and Enterprise Server

For the latest information about the supported versions of these platforms and components, see the *Sun Java System Access Manager 7 2005Q4 Release Notes*.

For more information, see "Implementing Access Manager Session Failover" on page 89.

Web Browser Requirements

Access Manager administrators and end users use web browsers to perform administrative and user management tasks. For information about the supported web browsers for this release, see the *Sun Java System Access Manager 7 2005Q4 Release Notes*.

Access Manager Schema

In general, a schema is a set of rules imposed on data that defines how it is stored. Sun Java System Directory Server uses the Lightweight Directory Access Protocol (LDAP) schema to define how its data is stored. Object classes define attributes in a LDAP schema. In Directory Server, each data entry must have one or more object class(es) to specify the type of object the entry describes and define the set of attributes it contains. Each entry then is basically a set of attributes and their corresponding values and the list of object classes to which these attributes correspond.

Access Manager uses Sun Java System Directory Server as its data repository, which includes the Access Manager schema that extends the Directory Server schema. When Access Manager is installed, the Access Manager schema, from the `ds_remote_schema.ldif` and `sunone_schema2.ldif` files, is integrated with the Directory Server schema. The `ds_remote_schema.ldif` file defines the LDAP object classes and attributes that are specifically used by Access Manager. The `sunone_schema2.ldif` file loads the Access Manager LDAP schema object classes and attributes.

You can view the `ds_remote_schema.ldif`, `sunone_schema2.ldif`, and other Access Manager LDIF files in the following directories:

- Solaris systems: `/etc/opt/SUNWam/config/ldif`
- Linux systems: `/etc/opt/sun/identity/config/ldif`

Marker Object Classes

Identity entries created using the Access Manager console and stored in Directory Server are appended with marker object classes. Marker object classes define the designated entries as those which Access Manager will manage. The object classes will not interfere with other aspects of the directory tree, such as servers or hardware. As well, existing identity entries can not be managed using Access Manager until they are modified to include these object classes. More detailed information about the marker object classes can be found in the Access Manager Developer's Guide. For information about migrating existing Directory Server data for use with Access Manager, see the *Sun Java Access Manager 6 2005Q1 Migration Guide*.

Administrative Roles

Delegated administration of the LDAP entries (mapped to each identity-related object in Access Manager) are implemented through the use of pre-defined roles and access control instructions (ACIs). Default administrative roles and their defined ACIs are created during Access Manager installation and can be viewed and managed using the Access Manager Console. In Access Manager 7 2005Q4 in Realm mode, roles depend on policies rather then ACIs.

When an Access Manager identity-related object is created, the appropriate administrative roles (and thus, corresponding ACIs) are also created and assigned to the LDAP entry for that object. The role can then be assigned to an individual user who maintains control of that object's node. For example, when Access Manager creates a new organization, several roles are automatically created for it and stored in Directory Server:

- Organization Administrator has read and write access to all entries in the configured organization.
- Organization Help Desk Administrator has read access to all entries in the configured organization and write access to the `userPassword` attribute in those entries.

- Organization Policy Administrator has read and write access to all policies in the organization.

The assignation of any of these roles to a user gives that user all the permissions accorded that role.

The following table summarizes the Access Manager administrator roles and the permissions that apply to each one.

TABLE 3–1 Default and Dynamic Roles and Their Permissions

Role	Administrative Suffix	Permissions
Top-level Organization Admin (amadmin)	Root level	Read and write access to all entries (such as roles, policy, and groups) under top-level organization.
Top-level Organization Help Desk Admin	Root level	Read and write access to all passwords under top-level organization.
Top-level Organization Policy Admin	Root level	Read and write access to policies at all levels. Used by sub-organizations to delegate referral policy creation.
Organization Admin	Organization only	Read and write access to all entries (such as roles, policy, and groups) under the created sub-organization only.
Organization Help Desk Admin	Organization only	Read and write access to all passwords under the created sub-organization only.
Organization Policy Admin	Organization only	Read and write access to all policies under the created sub-organization only.
Container Admin	Container only	Read and write access to all entries (such as roles, policy, and groups) under the created container only.
Container Help Desk Admin	Container only	Read and write access to all passwords under the created container only.
Group Admin	Group only	Read and write access to all entries (such as roles, policy, and groups) under the created group only.
People Container Admin	People Container only	Read and write access to all entries (such as roles, policy, and groups) under the created people container only.
User (self-administrator)	User only	Read and write access to attributes in the user entry only (except for user attributes such as `nsroledn` and `inetuserstatus`).

Using roles instead of group-based ACIs is more efficient and requires less maintenance. Filtered roles are simpler for LDAP clients, because they can just ask for the `nsRole` attribute of a user. Roles do suffer though from scope limitations, where a role must be a peer of a parent of a member of that role.

For more information about default ACIs, see the Access Manager Console Online Help.

Access Manager Administrative Accounts

During the installation of Access Manager, the following administrative accounts are created:

- Administrator user ID (`amadmin`) is the Access Manager top-level administrator that has unlimited access to all entries managed by Access Manager. You cannot change the default name, `amadmin`.

  During installation, you must provide a password for `amadmin`. To change the `amadmin` password after installation, use the Directory Server Console or the `ldapmodify` utility.
- Bind DN user for LDAP, Membership, and Policy services (`amldapuser`) is the administrative user that has read and search access to all Directory Server entries. You cannot change the default name, `amldapuser`.

  During installation, you must provide a password for `amldapuser`. Do not use the same password that you used for `amadmin`. To change the `amldapuser` password after installation, use the Directory Server Console or the `ldapmodify` utility.

  If you change the `amldapuser` password, you must also modify the LDAP authentication service and policy configuration services to reflect the change (`amldapuser` is the default user used in these services). You must make changes in each organization where these services are registered.
- Proxy user (`puser`) can take on any user's privileges (for example, an organization administrator or end user).
- Admin user (`dsameuser`) is used for binding purposes when the Access Manager SDK performs operations on Directory Server that are not linked to a particular user (for example, retrieving service configuration information).

Both `puser` and `dsameuser` have an associated password that is stored in encrypted format in the `serverconfig.xml` file, in the following directories:

- Solaris systems: `/etc/opt/SUNWam/config`
- Linux systems: `/etc/opt/sun/identity/config`

After installation, it is recommended that you change the password for `puser` and `dsameuser`, but do not use the same password that you used for `amadmin` or `amldapuser`. To change the `puser` or `dsameuser` password, use the `ampassword` utility:

- The `ampassword --admin` (or `-a`) option changes the password for `dsameuser`. (This option does not change the `amadmin` password.)

- The `ampassword --proxy` (or `-p`) option changes the password for `puser`.

Changing the `puser` or `dsameuser` password depends on your deployment.

If Access Manager is deployed on a single host server:

1. Use the `ampassword` utility to change the respective password in Directory Server and in the local `serverconfig.xml` file.
2. Restart the Access Manager web container.

If Access Manager is deployed on multiple host servers:

1. On the first server, use the `ampassword` utility to change the respective password in Directory Server and in the local `serverconfig.xml` file.
2. Encrypt the new password using the `ampassword --encrypt` (or `-e`) option.
3. On each additional server where Access Manager is deployed, change the password manually in the `serverconfig.xml` file, using the new encrypted password from Step 2.
4. On each server where you changed the password, including the first server, restart the Access Manager web container.

For information about the `ampassword` utility, see the *Sun Java System Access Manager 7 2005Q4 Administration Guide*.

Schema Limitations

Access Manager abstractly represents the entries it manages. This means, for example, that an organization in Access Manager is not necessarily the same as an organization in Directory Server. Whether a specific Directory Information Tree (DIT) can be managed or not depends on how the you choose to represent or manage your directory entries and whether your DIT fits into the limitations of each Access Manager type.

The following sections describe these limitations of the Access Manager schema. At the end of this section, several "Examples of Unsupported DITs" on page 53 are included.

Only One Type of Entry Can be Marked as an Organization

By adding the Access Manager `sunISManagedOrganization` auxiliary class to any entry, Access Manager can manage this entry as if it is an organization. However, only one type of entry may be marked as an organization in Access Manager. For example, if you have an entry `o=sun` and another entry `dc=ibm` in your DIT, you cannot mark them both as organizations.

In the following example, if you want both the `dc` and `o` entries to be organizations, the DIT structure will not be manageable using Access Manager:

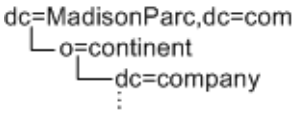

The entry at the Access Manager root suffix, however, does not count as one entry. Therefore, in the following example, the DIT structure can be managed by Access Manager:

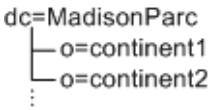

If you were able to add `dc=company1` below `o=continent1`, then this DIT would be manageable only if `dc` is marked as a container. Container is another abstract type in Access Manager that typically maps to an `OrganizationalUnit`. In most DITs, you would add the `iplanet-am-managed-container` entry to all `OrganizationlUnits`.

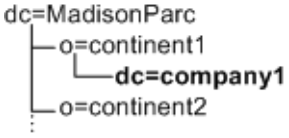

However, you could add this marker object class to any entry type. The DIT structure in the following example is allowed:

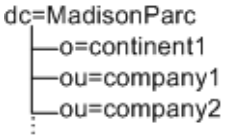

In this example, because you cannot mark both `o=` and `ou=` entries as organizations, you could mark the `o=` entries as `organization` and the `ou=` entries as `containers`. When exposed in the console, both organizations and containers have the same options. You can create subordination or subcontinents, people containers, groups, roles, and users under both of them.

People Containers Must be Parent Entries for Users

Another abstract entry type is the people container. The Access Manager type assumes that this entry is a parent entry for users. When you mark an entry as a people container with `iplanet-am-managed-people-container`, the UI will assume it can only contain sub-people containers or users. The attribute `OrganizationUnit` is typically used as a people container, but any entry can be this type in Access Manager as long as it has the `iplanet-am-managed-people-container` object class and it has a Access Manager manageable parent of type `organization` or `container`.

Only One Organization Description is Allowed in the Access Manager XML

The Access Manager organization is defined in `amEntrySpecific.xml`. Only one organization description is allowed in this file. As a result, when you customize directory entry properties, or create administration pages or search pages in the UI, your custom attributes apply globally to the entire Access Manager configuration. This Access Manager requirement may not meet the needs of some companies, especially hosting companies, that require different display attributes for each organization in the deployment.

In the following example, Edison-Watson is a hosting company that provides internet services to a number of companies. CompanyA wants to display fields for capturing a user's name First Name, Surname, and Badge Number. CompanyB wants to display fields for capturing a user's First Name, Last Name, and Employee Number.

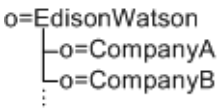

The organization description is defined at the root level (`o=EdisonWatson`), and not at the organization level. By default, the UI for both CompanyA and CompanyB must be identical. Also, all services globally define attributes to be of the subschema type user. So if CompanyA has attributes for its users in the auxiliary class `CompanyA-user`, and CompanyB has attributes in `CompanyB-user` then CompanyB's attributes will be overridden and will not be displayed.

As a workaround, you can modify the ACIs to work for user display. However, this workaround will not address the attributes in Search and Create windows.

Examples of Unsupported DITs

In the following example, you would need three types of organization makers: `o`, `ou`, and `l`. Assuming that `l=california` and `l=alabama` are not a people containers, this DIT would not work with Access Manager:

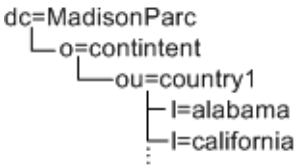

In the following example, you would need three types of Access Manager markers (`dc`,`o`,`ou`) plus the people container type (`ou=people`). Under these assumptions, the DIT would not work with Access Manager:

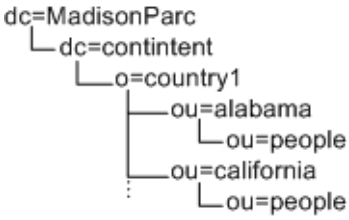

CHAPTER 4

Logical Design with Access Manager

During the logical design phase of the solution life cycle, you design a logical architecture showing the interrelationships of the logical components of the solution. The logical architecture and the usage analysis from the technical requirements phase form a deployment scenario, which is the input to the deployment design phase. This chapter contains the following sections about logical design for Sun Java™ System Access Manager:

About Logical Architectures

A logical architecture identifies the software components needed to implement a solution, showing the interrelationships among the components. The logical architecture and the quality of service requirements determined during the technical requirements phase form a deployment scenario. The deployment scenario is the basis for designing the deployment architecture, which occurs in the next phase, deployment design.

Designing a Logical Architecture

When you design a logical architecture, use the use cases identified during the technical requirements phase to determine the Java Enterprise System (Java ES) components that provide the services necessary for the solution. You must also identify any components providing services to the components you initially identify.

You place the Java ES components within the context of a multi-tiered architecture according to the type of services that they provide. Understanding the components as part of a multi-tiered architecture helps you later determine how to distribute the services provided by the components and also helps determine a strategy for implementing quality of service (such as scalability, availability, and others.)

For more detailed information about logical architectures and the solution life cycle, see the *Sun Java Enterprise System 2005Q4 Deployment Planning Guide*.

Access Manager Components

An Access Manager deployment includes the following products and components:

- "Web Container" on page 56
- "Directory Server" on page 56
- "Message Queue and Berkeley DB for Session Failover" on page 57

Web Container

Access Manager must run in one of the following web containers:

- Sun Java System Web Server
- Sun Java System Application Server
- BEA WebLogic Server
- IBM WebSphere Application Server

For the specific versions of each web container that are supported, see the *Sun Java System Access Manager 7 2005Q4 Release Notes*.

Directory Server

Access Manager requires an LDAP directory server for these entities:

- "Access Manager Information Tree" on page 57
- "Identity Repository" on page 57

Access Manager Information Tree

Access Manager 7 2005Q4 requires Sun Java System Directory Server to store the Access Manager information tree. Access Manager creates and maintains the Access Manager information tree, which includes the following information pertinent to system access:

- How users authenticate
- Which resources users can access
- What information is available to applications after users are given access to resources

Identity Repository

Access Manager requires an identity repository to store user data such as users and groups. Previous versions of Access Manager required Sun Java System Directory Server as the identity repository. However, in addition to Sun Java System Directory Server, Access Manager 7 2005Q4 also supports an LDAP version 3 (LDAP v3) compliant directory server.

Message Queue and Berkeley DB for Session Failover

If you are planning to implement session failover, Access Manager requires these additional components:

- Sun Java System Message Queue. The Message Queue broker cluster manages the session messages between Access Manager instances and the session store database.

- Berkeley DB by Sleepycat Software, Inc. (`http://www.sleepycat.com/`) is the default session store database.

Java ES Components That Use Access Manager

Access Manager is usually deployed with other Java ES component products, including:

- Sun Java System Portal Server, Sun Java System Messaging Server, Sun Java System Calendar Server, Sun Java System Instant Messaging, and Sun Java System Communications Express: To provide single sign-on (SSO)
- Sun Java System Web Server: To provide optional access control.

Example Access Manager Logical Architectures

This section provides the following scenarios as examples of logical architectures for Access Manager solutions, including:

Access Manager Web Deployment

A common Access Manager deployment has a web browser accessing an application or resource deployed on a web server. The application or resource is protected by Access Manager and communicates with it using a policy agent also installed on the web server. The web server might also have the Access Manager SDK deployed. This scenario does not restrict the number of web servers on a machine or the instances of Access Manager deployed on multiple machines. For example, a machine might have multiple web servers, each deploying an instance of Access Manager. Similarly, multiple web servers might also be running on different machines, each deploying an instance of Access Manager.

The following figure shows an Access Manager web deployment scenario.

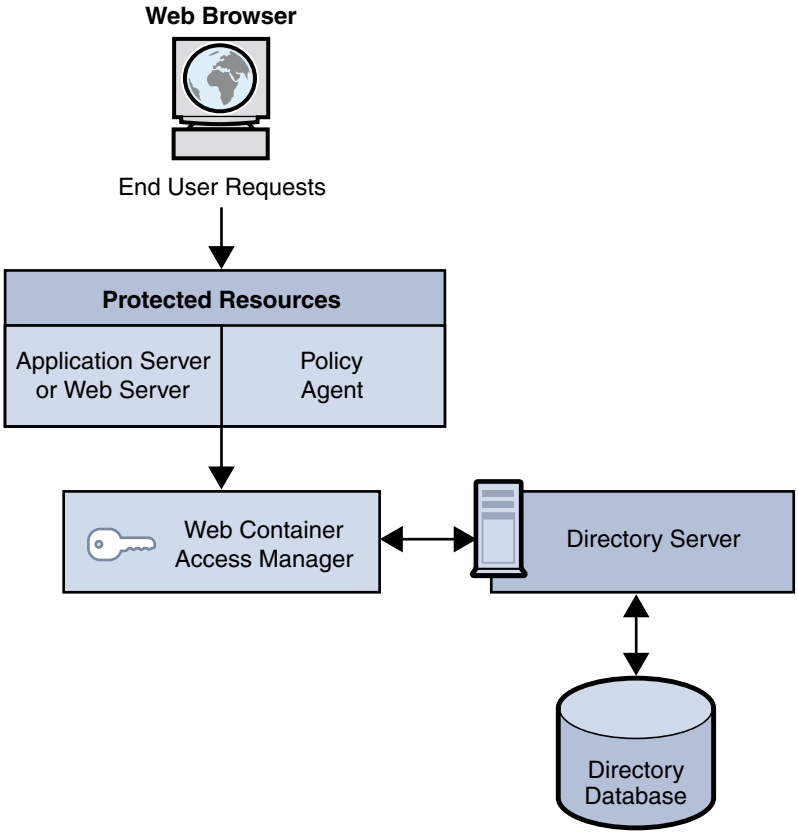

FIGURE 4–1 Access Manager Web Deployment

Access Manager Multiple Server Deployment

An Access Manager multiple server deployment has two or more host servers, with one or more instances of Access manager installed on each host server. Each Access Manager instance accesses the same Directory Server. You can configure the Directory Server instances in a multiple master replication (MMR) configuration, if required for your deployment.

The Access Manager instance installed on the first host server points to an instance of Directory Server. During installation using the Java ES installer, you can choose an existing Directory Server with or without an existing directory information tree (DIT), depending on your deployment.

You install subsequent instances of Access Manager on other host servers by running the Java ES installer, with the Access Manager instance pointing to a Directory Server with an existing DIT. Access Manager then does not write any information to Directory Server because it recognizes the Directory Server as already existing.

The following figure shows multiple Access Manager instances on different host servers with one Directory Server.

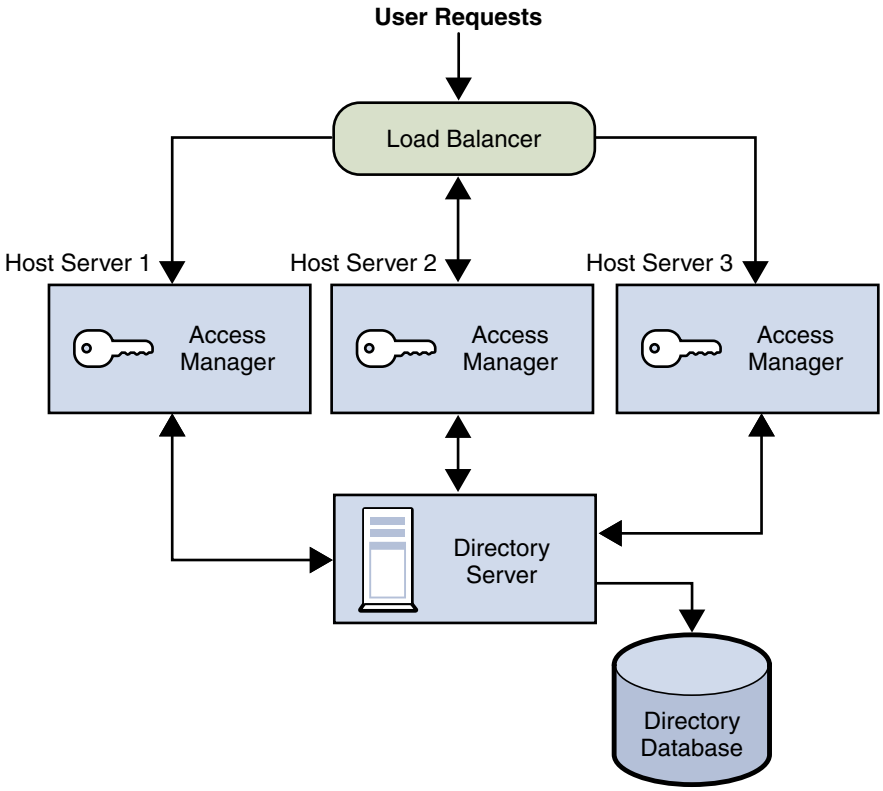

FIGURE 4–2 Multiple Access Manager Instances With One Directory Server

For more information, see "Installing Access Manager on Multiple Host Servers" on page 77.

Java Application Deployment

Another common scenario for Access Manager allows Java applications to access an Access Manager SDK installed directly on the server where they are deployed. This scenario requires an additional server with an instance of a web container (such as Sun Java System Web Server or Sun Java System Application Server) running at least one instance of Access Manager. This server also maintains the information to provide single sign-on (SSO). The following figure shows a Java application deployment scenario.

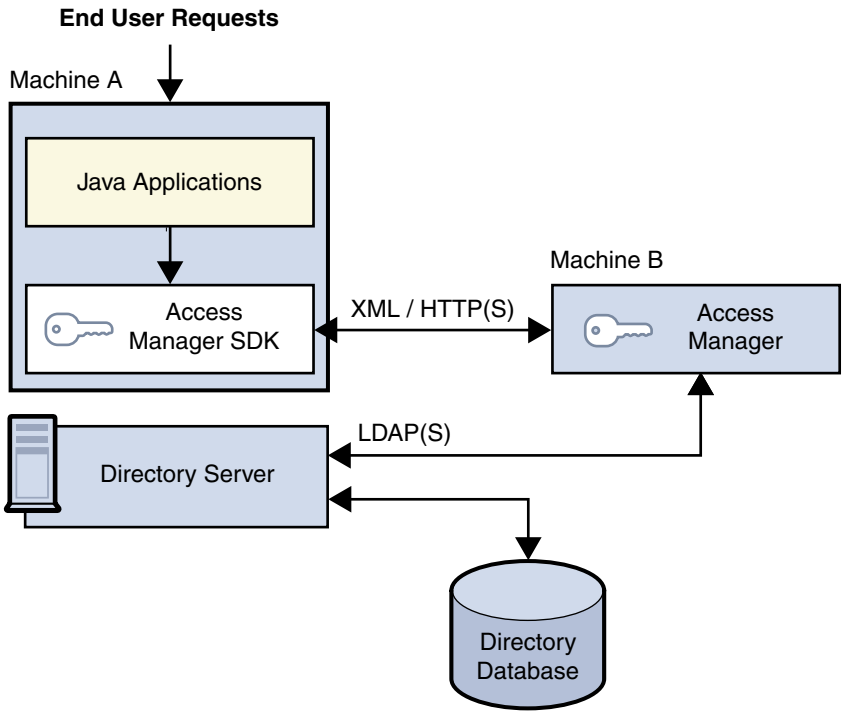

FIGURE 4–3 Java Application Deployment

Access Manager Session Failover Deployment

Access Manager provides a web container independent session failover implementation using Sun Java System Message Queue (Message Queue) as the communications broker and the Berkeley DB by Sleepycat Software, Inc. as the session store database. Access Manager session failover retains a user's authenticated session state in the event of a single hardware or software failure, which allows the user's session to fail over to a secondary Access Manager instance without losing any session information or requiring the user to login again.

Overview of Access Manager Session Failover

Access Manager 7 2005Q4 session failover includes these components:

- Two or more instances of Access Manager 7 2005Q4, with each instance running on a supported web container on two or more host servers.
- Message Queue broker cluster, which manages the session messages between the Access Manager instances and the session store database.
- Berkeley DB by Sleepycat Software, Inc. (`http://www.sleepycat.com/`), as the session store database. The Berkeley DB client daemon is `amsessiondb`.

Access Manager session failover follows the Message Queue publish/subscribe (topic destinations) delivery model:

1. When a user initiates, updates, or ends a session, Access Manager publishes a session creation, update, or deletion message to the Message Queue broker cluster.
2. The Berkeley DB client (`amsessiondb`) subscribes to the Message Queue broker cluster, reads the session messages, and stores the session operations in the database.

If an Access Manager instance fails due to a single hardware or software problem, a user's session associated with that instance fails over to a secondary Access Manager instance, as follows:

1. The secondary Access Manager instance publishes a query request to the Message Queue broker cluster for the user's session information.
2. The Berkeley DB clients (`amsessiondb`) subscribing to the same session request topic on the Message Queue broker cluster receive the query request retrieve the corresponding entry from the session database, and then publish the user's session information to the Message Queue broker cluster with the session response topic.
3. The secondary Access Manager instance subscribing to the session response topic receives the response with the user's session and continues without losing any session information or the user having to login again.

If a Message Queue broker fails, Access Manager continues to operate in non-session failover mode. When the Message Queue broker is later restarted, Access Manager returns to session failover mode.

For more information about the Message Queue components and the publish/subscribe delivery model, see the *Sun Java System Message Queue Technical Overview*.

Session Failover Deployment Scenario

The following figure shows a basic scenario with two host servers, each running an Access Manager instance on a web container, the Message Queue broker cluster, and the Berkeley DB client (`amsessiondb`). The load balancer distributes client requests to the Access Manager instances. Both Access Manager instances access the same Directory Server (not shown in the figure).

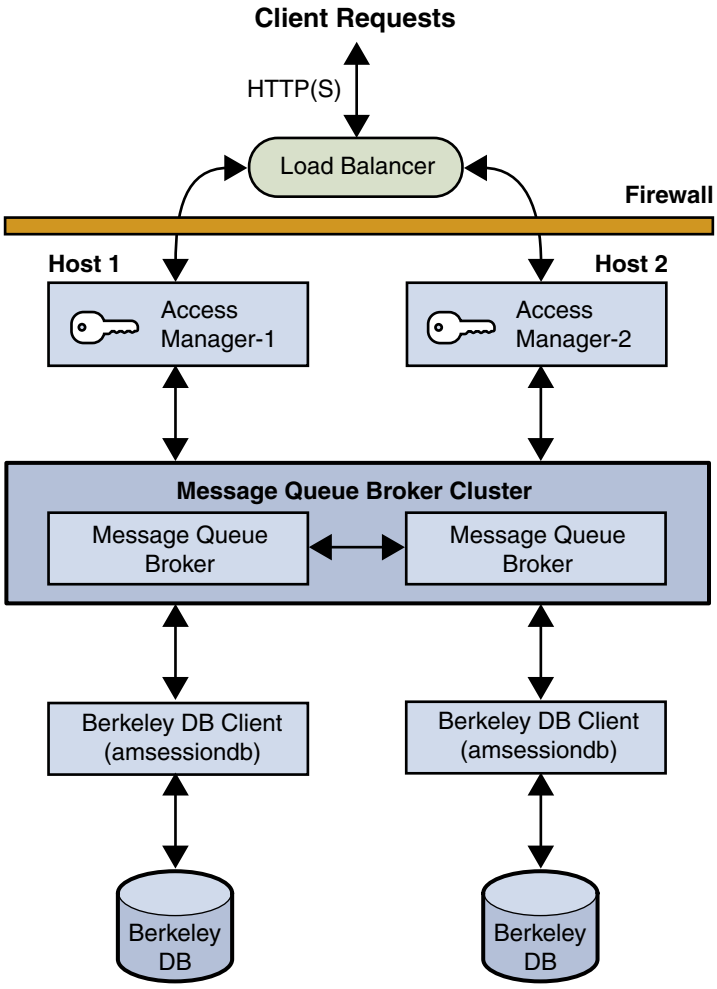

FIGURE 4–4 Access Manager Session Failover Basic Deployment Scenario

You can add additional sites similar to the one shown in the figure, with each site accessing the same Directory Server. Session failover, however, occurs only for the Access Manager instances within a site; cross-site session failover is not supported in the current release.

For more information, see "Implementing Access Manager Session Failover" on page 89.

Access Manager and Portal Server Deployment

For the Java Enterprise System 2005Q4 release, you can deploy Access Manager with Portal Server either on the same physical server or on multiple servers.

Installation on a Single Server

In this scenario, Access Manager and Portal Server are installed on the same physical server. You must also install or have access to an installed version of Directory Server, which can be either on the same server or a remote server.

To install these components, run the Java Enterprise System installer in a single session and make these selections:

- On the Component Selection panel, select these products and subcomponents:
  - Under Communication & Collaboration Services, select Portal Server.
  - Under Directory & Identity Services, select Access Manager 7 2005Q4 and its subcomponents:
    - Identity Management and Policy Services Core
    - Access Manager Administration Console
    - Common Domain Services for Federation Management
    - Access Manager SDK

      By default, when you select Portal Server, the installer installs only the Access Manager SDK, so you must specifically check the other subcomponents.

  Install and configure one of the following web containers:
  - Sun Java System Application Server
  - Sun Java System Web Server

Installation on Multiple Servers

In this scenario, Portal Server will access Access Manager on a local server from a remote server. You must also install or have access to an installed version of Directory Server, which can be either on a local or remote server:

- On the local server, install Access Manager and a web container. You must install and configure the components on this server before you install and configure the components on the remote server.
- On the remote server, install Portal Server and the Access Manager SDK. You do not need to select the other Access Manager subcomponents on the remote server.

For more information about deploying Access Manager and Portal Server, see the *Sun Java System Portal Server Deployment Planning Guide*.

Federation Management

In 2001, Sun Microsystems joined with other companies to form the Liberty Alliance Project. This project defined standards for developing identity-based infrastructures, software, and web services.

Initially, Access Manager implemented the Identity Federation Framework (Liberty ID-FF) specification, including account federation, authentication domains, and single sign-on (SSO). Subsequent releases of Access Manager added new features, as defined in version 1.2 of the Liberty ID-FF specifications and the version 1.0 specifications of the Liberty Identity Web Services Framework (Liberty ID-WSF). Web services include a framework for retrieving and updating identity data that consists of attributes stored in identity-based service providers across the Internet. A client application programming interface (API) for communication between identity providers and service providers is also provided.

Access Manager 7 2005Q4 provides additional functionality. For example, Access Manager provides the ability to bulk-federate user accounts to applications that are out-sourced to business partners and to map configured roles between the identity provider and the service provider.

For more information, see the *Sun Java System Access Manager 7 2005Q4 Federation and SAML Administration Guide*. This guide includes an introduction to the open-standard specifications used to develop these features and information about how Access Manager has implemented them. It also includes information about integrated web services and summaries of the application programming interface (API).

CHAPTER 5

Deployment Design with Access Manager

During the deployment design phase of the solution life cycle, you design a high-level deployment architecture and a low-level implementation specification, and prepare a series of plans and specifications necessary to implement the solution. Project approval occurs in the deployment design phase. This chapter includes the following sections about deployment design with Sun Java™ System Access Manager:

Using a Load Balancer

In most deployments, Access Manager is configured with a load balancer to distribute user requests between two or more Access Manager instances. The load balancer can be implemented with hardware, software, or a combination of both. The following figure shows an Access Manager deployment with a load balancer.

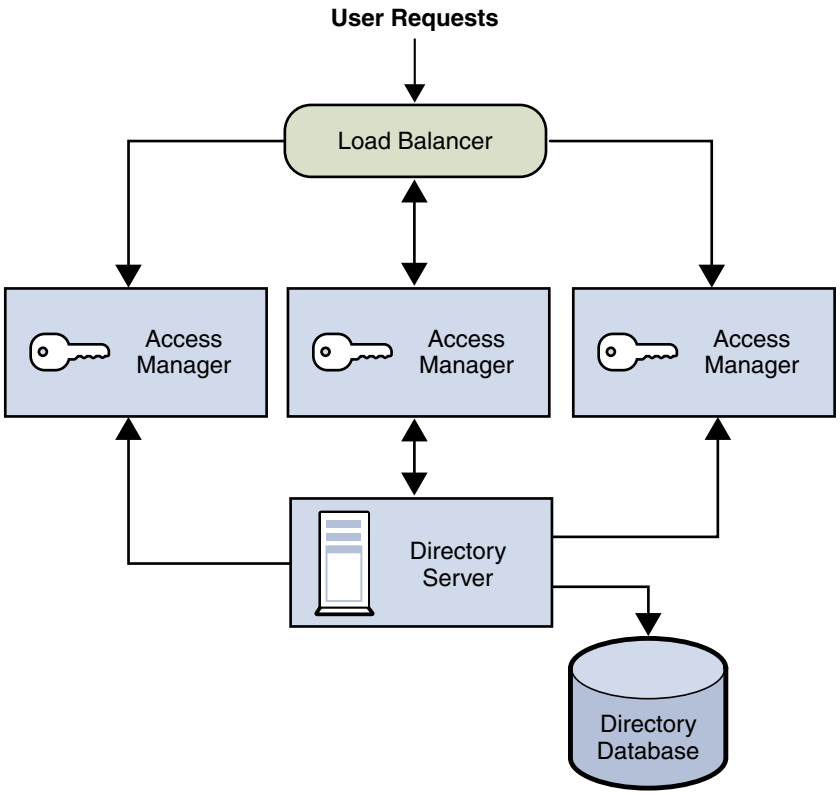

FIGURE 5–1 Access Manager Configuration With a Load Balancer

Configuring the Load Balancer for Sticky Sessions

A load balancer deployed with Access Manager must support sticky sessions. A sticky session specifies that once a session is created by a specific Access Manager instance, subsequent requests from the user will continue to be routed to that same instance, in order to preserve session information. Because Access Manager uses cookies to relay session information, the load balancer must redirect the request to the Access Manager instance that created the session. Without sticky sessions, all Access Manager instances would have to be trusted and performance could be impaired. You can implement sticky sessions using either the `setcookie` function or load balancer cookies.

For more information, see "Using a Load Balancer With Access Manager" on page 83.

Multiple JVM Environment

Access Manager services are supported in multiple Java Virtual Machine (JVM) environments. That is, an instance of Sun Java System Application Server can be configured to have multiple JVMs with Access Manager services running in all of them. The Access Manager architecture imposes no restrictions on the deployment with regards to the number of Sun Java System Application Server instances within a machine, the number of Access Manager services across multiple machines, or the number of JVMs that a single Application Server can have.

For more information about the multiple JVM environment, see the Sun Java System Application Server documentation: `http://docs.sun.com/coll/1310.1`.

Directory Server Replication Considerations

Two methods to improve Access Manager performance and response time are using load balancing across replicated Directory Servers and locating replicated servers closer to users. Directory Server can be set up in single-supplier or multi-supplier configurations. Load-balancing applications such as Sun Java System Directory Proxy Server can also be used. Directory Proxy Server dynamically performs proportional load balancing of LDAP operations across a set of configured Directory Servers. If one or more Directory Server instances become unavailable, the load is proportionally redistributed among the remaining servers. When the server comes back on line, the load is proportionally and dynamically reallocated.

Directory Server replication must be configured before installing Access Manager. This configuration ensures that the supplier and consumer databases are synchronized correctly, allowing time to verify that referrals and updates are synchronized properly.

When Access Manager is installed for replication purposes, each instance of Directory Server and each instance of Access Manager, must be configured with the same values for the following:

- Directory Manager
- Directory Manager Password
- Directory Server Administrator ID
- Server Administrator Password
- Base suffix
- Default organization

Configuring For Replication

Access Manager can be configured to work with single-supplier or multiple-supplier replication. The following figure shows a single-supplier configuration where the consumer is a read-only database. Requests for write operations are referred to the supplier database. This configuration provides some measure of enhanced server performance by distributing the workload to more than one directory.

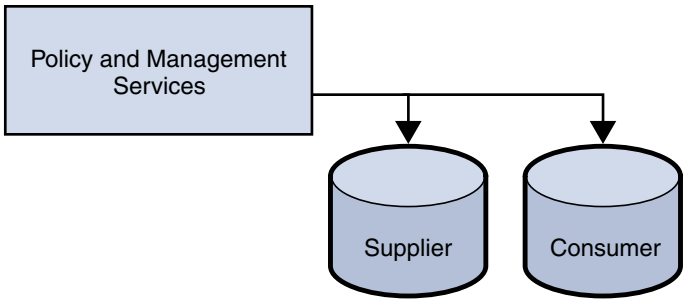

FIGURE 5–2 Single-Supplier Directory Server Replication

The following figure shows a multiple-supplier configuration, or multi-master replication (MMR), using multiple instances of Access Manager. This configuration provides failover protection as well as high availability, resulting in further enhanced server performance.

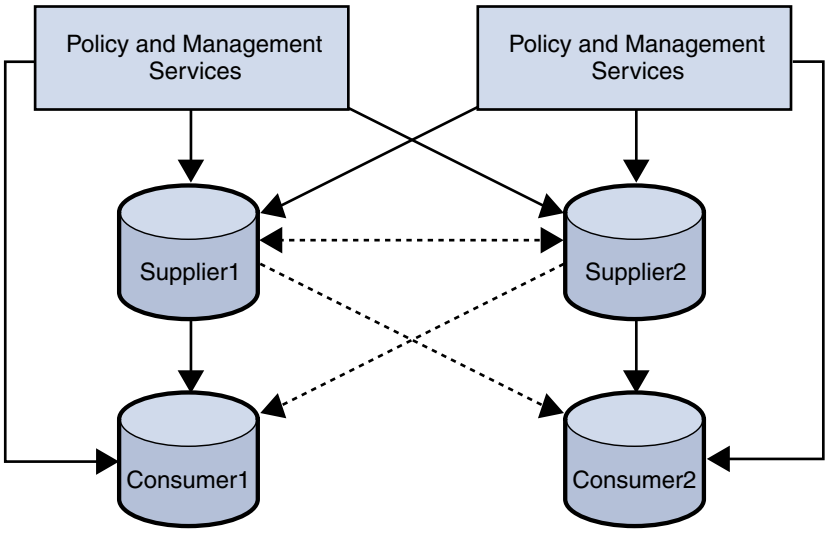

FIGURE 5–3 Multiple-Supplier Directory Server Configuration

Follow these steps to configure replication at the root or top level of the Access Manager directory tree when Access Manager has not yet been installed or to configure replication at the default organization level:

1. Install the supplier and consumer Directory Server instances.

   See the *Sun Java Enterprise System 2005Q4 Installation Guide for UNIX* for detailed instructions.

2. Set up replication agreements between the supplier and consumer and verify that the directory referrals and updates are working properly.

   You might need to migrate existing Directory Server data to work with this version of Access Manager. For information, see the *Sun Java System Access Manager 6 2005Q1 Migration Guide*.

3. If you are deploying Access Manager and Directory Server for the first time, or if there is no plan to use existing user data, run the Java ES installation program to install Access Manager.

   During installation, answer yes when asked if there is an existing Directory Server, and specify the host name and port number for a supplier Directory Server you installed in "Configuring For Replication" on page 70.

4. On the host server where Access Manager is installed, modify the `AMConfig.properties` file in the following directory, depending on your platform:

   - Solaris systems: `/etc/opt/SUNWam/config`
   - Linux systems: `/etc/opt/sun/identity/config`

5. Modify the following properties to reflect the host and port number of a consumer Directory Server installed in "Configuring For Replication" on page 70 .

   - `com.iplanet.am.directory.host`
   - `com.iplanet.am.directory.port`

6. Modify the following property to reflect the number of times Access Manager should continue to make the same request when the requested entry is not found.

   `com.iplanet.am.replica.retries`

7. Modify the following property to reflect the number of milliseconds Access Manager should allow to elapse between retries.

   `com.iplanet.am.replica.delay.between.retries`

8. In each Access Manager Authentication module enabled, use the Access Manager Console to specify the consumer directory installed in "Configuring For Replication" on page 70:

   - For the first LDAP server and port, specify the host name and port number for the primary (consumer) Directory Server. For example: `consumer1.example.com:389`.
   - For the second LDAP server and port, specify the host name and port number for the secondary (or supplier) Directory Server. For example, `supplier1.example.com:389`.

9. In the `serverconfig.xml` file, specify the host name and port number of the consumer directory installed in "Configuring For Replication" on page 70, as shown in the following example for the `serverconfig.xml` file.
10. Restart Access Manager by restarting the web container.

Example of the `serverconfig.xml` File

The following example shows the `serverconfig.xml` replication modification.

```
<iPlanetDataAccessLayer>
<ServerGroup name="default" minConnPool="1"
maxConnPool="10">
<Server name="Server1"
host="consumer1.example.com" port="389"
type="SIMPLE" />
```

Configuring With a Load Balancer

The following figure shows a multiple-supplier configuration that includes Directory Proxy Server or a hardware load balancer. This configuration takes advantage of Access Manager support for failover, high availability, and managed load-balancing.

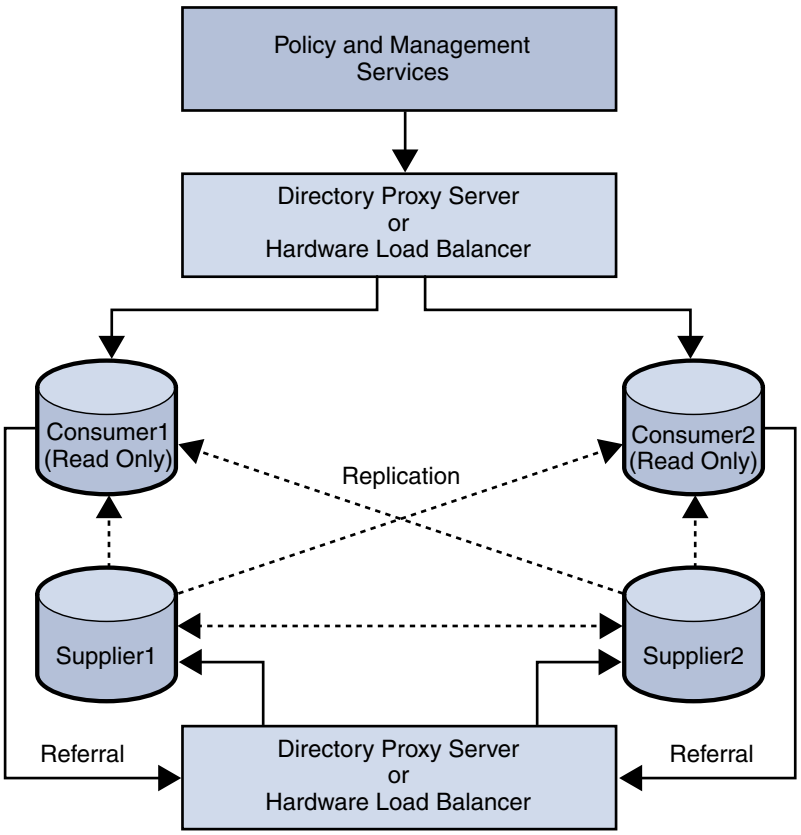

FIGURE 5–4 Multiple-Supplier Configuration With a Load Balancer

Using LDAP load balancers adds a layer of high availability and directory failover protection beyond the level that is available with Access Manager. For example, Directory Proxy Server can specify the percentage of the load that gets redistributed to each server. And, if all back-end LDAP servers become unavailable, Directory Proxy Server continues to manage requests, rejecting client queries. If you install a load balancer, Access Manager must be configured to recognize the application.

1. Before configuring Access Manager, Set up the Directory Servers for replication. For information about directory replication and for detailed setup instructions, see the Sun Java System Directory Server documentation: `http://docs.sun.com/coll/1316.1`.
2. Install and configure the LDAP load balancer. Follow the instructions in the documentation that comes with the load balancer you are using.
3. In the `AMConfig.properties` file, modify the `com.iplanet.am.directory.host` and `com.iplanet.am.directory.port` properties to point to the load balancer host and port number of a consumer Directory Server.

4. For each Access Manager Authentication module enabled, use the Access Manager Console to specify the consumer Directory Server. In the following steps, the LDAP Authentication module is used as an example:
   - For the first LDAP server and port, type the host name and port number for the primary (consumer) Directory Server using the form `proxyhostname:port`.
   - Do not enter anything for the second LDAP Server and Port.
5. In the `serverconfig.xml` file, specify the host name and port number of the consumer Directory Server, as shown in the following example for the `serverconfig.xml` file.
6. Restart Access Manager by restarting the web container.

Load Balancer Modification to the `serverconfig.xml` File

The following example shows the load balancer modification to the `serverconfig.xml` file.

```
<iPlanetDataAccessLayer>
<ServerGroup name="default" minConnPool="1"
maxConnPool="10">
<Server name="Server1"
host="idar.example.com" port="389"
type="SIMPLE"
```

Directory Server With a Firewall

If your deployment is configured with a firewall between Access Manager and Directory Server, Access Manager connections can time out if the firewall idle connection timeout value is less than the Directory Server idle connection timeout value (`nsslapd-idletimeout` attribute). This problem usually occurs during non-peak usage hours when the load on Access Manager is low.

When Directory Server connections are dropped by the firewall, Access Manager does not recognize that the connections have been dropped and then goes through the pool of LDAP connections until all connections are exhausted. Access Manager must be restarted to create a fresh pool of LDAP connections. To prevent this problem, consider the following solutions:

- "Setting the Global Timeout Attribute" on page 75
- "Setting the Timeout Value for Individual Client Connections" on page 75

Setting the Global Timeout Attribute

You might be able to set the Directory Server global `nsslapd-idletimeout` attribute to a value less than the firewall idle connection timeout value. However, this solution might not be acceptable because `nsslapd-idletimeout` is a global configuration attribute that affects applications other than Access Manager.

Setting the Timeout Value for Individual Client Connections

Directory Server allows you to set specific attributes for individual client connections. The `nsIdleTimeout` attribute specifies the idle connection timeout value for individual clients. This value takes precedence over the `nsslapd-idletimeout` value set for the global Directory Server configuration.

Set the `nsIdleTimeout` attribute for the Access Manager user that binds to the LDAP directory, which by default is `amldapuser`. This attribute also applies to the `dsameuser` and `puser` users.

To add the `nsIdleTimeout` attribute for `amldapuser`, use either the Directory Server Console or the `ldapmodify` tool. For example:

```
ldapmodify -h host-name -p port
-D "cn=Directory Manager" -w password
dn: cn=amldapuser,ou=DSAME Users, dc=example,dc=com
changetype: modify
add: nsIdleTimeout
nsIdleTimeout: timeout-value
```

For *timeout-value*, specify a value less than the connection idle timeout value set for the firewall. Thus Directory Server will close the Access Manager connections for `amldapuser` before they are closed by the firewall.

To add the timeout for `dsameuser` or `puser`, use the above syntax, except set the `dn` option to the `dsameuser` or `puser` user.

The `com.sun.am.event.connection.idle.timeout` property in the `AMConfig.properties` file specifies the timeout value in minutes after which persistent searches will be restarted. This property ensures that persistent searches are restarted when the connections are dropped. Ideally, this value should be lower than the load balancer or firewall TCP timeout value, to make sure that persistent searches are restarted before the connections are dropped. A default value of zero (0) specifies that these searches will not be restarted.

For information about the Directory Server attributes and the `ldapmodify` tool, see the Sun Java System Directory Server documentation: `http://docs.sun.com/coll/1316.1`.

CHAPTER 6

Implementation of an Access Manager Design

During the implementation phase of the solution life cycle, you implement various solutions for your deployment. This chapter describes the following implementation scenarios for Sun Java™ System Access Manager:

Installing Access Manager on Multiple Host Servers

Installing Access Manager instances on multiple host servers, with each instance accessing the same Directory Server, includes these steps:

Deploying Access Manager Instances

To install Access Manager instances on multiple host servers, with each instance accessing the same Directory Server, follow these steps.

1. Install Access Manager on a host server by running the Java Enterprise System (Java ES) installer. When you run the installer, specify either the Configure Now or Configure Later option. For information about running the installer, see the *Sun Java Enterprise System 2005Q4 Installation Guide for UNIX*.

   When you run the installer, you can also install Web Server or Application Server as the Access Manager web container. To use BEA WebLogic Server or IBM WebSphere Application Server as the web container, you must first install the product before you run the run the `amconfig` script in the following steps. For installation instructions, see the respective BEA or IBM product documentation.

2. If you specified the Configure Later option during installation or if you need to reconfigure the Access Manager instance (for example, to use BEA WebLogic Server or IBM WebSphere Application Server as the web container), you must run the `amconfig` script. The `amconfig` script and the `amsamplesilent` configuration file are located in the *AccessManager-base*`/bin` directory, where *AccessManager-base* represents the default installation directory: `/opt/SUNWam` on Solaris systems and `/opt/sun/identity` on Linux systems.

   Run the `amconfig` script as follows:

   a. Copy the `amsamplesilent` file to a writable directory and make that directory your current directory. For example, you might create a directory named `/newinstances`.

   b. Rename the copy of the `amsamplesilent` file to describe the new instance you want to configure. For example, if you plan to create a new Access Manager instance for Web Server 6.1, you might rename the file to `amwebsvr6`.

   c. Set the variables in the `amwebsvr6` file to configure the new instance. For example, to configure Access Manager in Realm mode:

      ```
      AM_REALM=true
      DEPLOY_LEVEL=1
      NEW_INSTANCE=true
      WEB_CONTAINER=WS6 # Web Server is the web container
      DIRECTORY_MODE=1
      ...
      ```

      In case you might need to reconfigure or uninstall this instance later, save the new `amwebsvr6` file.

   d. Run the `amconfig` script, specifying the new `amwebsvr6` file as the silent configuration input file. For example, on Solaris systems with Access Manger installed in the default directory:

      ```
      # cd /opt/SUNWam/bin/
      # ./amconfig -s ./newinstances/amwebsvr6
      ```

Run `amconfig` with full path to the `amsamplesilent` file (or copy of the file). The script reads the variables in the `amwebsvr6` file and then runs in silent mode (`-s` option) to configure Access manager for the web container. For more information about the `amsamplesilent` file and running the `amconfig` script, see the *Sun Java System Access Manager 7 2005Q4 Administration Guide*.

3. Repeat these steps on the other host servers to deploy additional Access Manager instances.

Several considerations for deploying additional Access Manager instances are:

- If you are running the Java ES installer and you want to use the same Directory Server as the first instance, check "Yes" for "Is Directory Server provisioned with user data?".
- If you are running the `amconfig` script, set variables in the copy of the `amsamplesilent` file. For example, to deploy Access Manager in Realm mode:

```
AM_REALM=true
DEPLOY_LEVEL=1
NEW_INSTANCE=true
WEB_CONTAINER=WS6 # Web Server is the web container
DIRECTORY_MODE=4  # Directory Server is provisioned with user data
AM_ENC_PW=password-encryption-key-value-from-the-first=instance
...
```

- If you are using non-default naming attributes and object classes, specify the custom values as appropriate for the user naming and organization naming attributes and object classes. Also, all deploy URIs (SERVER_DEPLOY_URI, CONSOLE_DEPLOY_URI, PASSWORD_DEPLOY_URI, and COMMON_DEPLOY_URI) for the web applications must match the previous installation.
- Use the same password encryption key as the first instance, as described in following Caution.

Caution – In a multiple server deployment that shares the same Directory Server, all Access Manager instances must use the same value for the password encryption key.

If you run the Java ES installer to install Access Manager on subsequent (second, third, and so on) servers in a multiple server deployment, the installer generates a new random password encryption key for each server. Therefore, when you run the installer on a subsequent server, use the encryption key value from the first Access Manager instance, which you can copy from the `am.encryption.pwd` attribute in the `AMConfig.properties` file and set as follows:

- Configure Now option. Replace the new random encryption key generated by the installer with the encryption key value from the first instance.
- Configure Later option. Set the `AM_ENC_PWD` variable in the copy of the `amsamplesilent` file with the encryption key value from the first instance before you run the `amconfig` script.

However, if you need to change the password encryption key for an Access Manager instance, see Appendix B.

Adding Additional Instances to the Platform Server List and Realm/DNS Aliases

When you install multiple instances of Access Manager on different host servers, the additional instances are not added to the platform server list or the realm/DNS aliases. You must explicitly add the values for the additional Access Manager instances, as follows:

1. Log in to the Access Manager 7 2005Q4 Console as `amadmin` on the first Access Manager host server.
2. In the Access Manager Console, click `Configuration`, `System Properties`, and then `Platform`.
3. Add each additional Access Manager instance to the Platform Server List under `Instance Name`:
   a. In the Platform Server List under `Instance Name` Name, click `New`.
   b. In `New Server Instance`, add the Server and Instance Name. For example:
      - Server: `http://amserver2.example.com:80`
      - Instance Name: `02`
   c. Click `OK` to add the instance.
   d. After you have added all instances, click `Save`.
4. Add the Realm/DNS alias for each additional Access Manager instance:
   a. In the Access Manager Console, click Access Control and then the root (top-level) realm under Realm Name.

b. Under Realm Attributes, add the Access Manager instance to Realm/DNS Aliases and then click Add. For example: `amserver2.example.com`

c. After you have added all instances, click `Save`.

Configuring an Access Manager Deployment as a Site

Access Manager 7 2005Q4 introduces the "site concept," which provides centralized configuration management for an Access Manager deployment. When Access Manager is configured as a site, client requests always go through the load balancer, which simplifies the deployment as well as resolves issues such as a firewall between the client and the back-end Access Manager servers. A site includes the following components:

- Multiple (two or more) Access Manager instances are deployed on at least two different host servers. For example, you might deploy two instances on one server and a third instance on another server. Or you might deploy all instances on different servers. You can also configure the Access Manager instances in session failover mode, if required for your deployment.
- One or more load balancers route client requests to the various Access Manager instances. You configure each load balancer according to your deployment requirements (for example, to use round-robin or load average) to distribute the load between the Access Manager instances.
- All Access Manager instances access the same Directory Server.

The following procedures refer to the Access Manager 7 2005Q4 Console in Realm Mode.

Site Configuration

If you have an Access Manager multiple server deployment, use either of these methods to configure your deployment as a site:

- If you plan to implement Access Manager session failover, the `amsfoconfig` script configures a deployment as a site. See "Implementing Access Manager Session Failover" on page 89.
- If you don't plan to implement session failover, follow the steps in this section.

When you configure a deployment as a site, you perform these functions in the Access Manager Console:

- Add the load balancer URL to the Site Name (site ID).

- Map the load balancer Site Name (site ID) to each Access Manager instance in the Platform Server List.
- Add the load balancer to the Realm/DNS Aliases.

In addition, Access Manager automatically sets the `fqdnMap` property (in memory) to include the load balancer, so you do not need to explicitly set this property in the `AMConfig.properties` file.

To configure an Access Manager deployment as a site, follow this procedure:

1. Log in to the Access Manager Console as `amAdmin`.
2. Add the load balancer URL to the `Site Name`:
   a. In the Access Manager Console, click `Configuration`, `System Properties`, and then `Platform`.
   b. Under `Site Name`, click `New` and enter the following values for the load balancer:
      - `Server`: Load balancer protocol, host name, and port. For example: `http://lb.example.com:80`
      - `Site Name`: Unique two-digit site identifier (site ID). For example: 10

      When you are finished, click `OK`.
   c. After adding the load balancer to the `Site Name`, click `Save`. The entry for the load balancer now includes the site ID. For example: `http://lb.example.com:80|10`

      The site ID must be unique with respect to server IDs and other site IDs. For example, you cannot use 01 for both a site ID and a server ID.
3. On the same Console panel, map the load balancer to each Access Manager instance:
   a. In the `Server` list under `Instance Name`, click each instance name to display the `Edit Server Instance` panel for the instance.
   b. Map the `Site Name` (site ID) for the load balancer to the Access Manager instance. For example, using a load balancer with a `Site Name` of 10, for the first server, the `Instance Name` would `01(|10)`.
   c. Click `OK` and repeat the steps for the other Access Manager instances.

      When you are finished, all Access Manager instances should be mapped to the load balancer. For example:

      ```
      http://amserver1.example.com:8080|01|10
      http://amserver2.example.com:8080|02|10
      http://amserver3.example.com:8080|03|10
      ```
   d. Click `Save` to save the configuration.
4. Add the Realm/DNS alias for the load balancer:
   a. In the Access Manager Console, click `Access Control` and then the root or top-level realm under `Realm Name`.

b. Under `Realm Attributes`, add the load balancer to `Realm/DNS Aliases` and then click `Add`. For example: `lb.example.com`.

c. Click `Save` to save your changes.

5. For clients such as a policy agent, the load balancer (as opposed to the individual Access Manager instances) should be the sole entry point. For example, if you are using a policy agent, modify the appropriate entries in the `AMAgent.properties` file to point to the load balancer.

Using a Load Balancer With Access Manager

The load balancer distributes the client requests between the Access Manager instances in multiple server deployment. Before you use this information in the section, configure your Access Manager deployment as a site, as described in "Configuring an Access Manager Deployment as a Site" on page 81. A site includes multiple (two or more) instances of Access Manager installed on different host servers. All Access Managers instances must access the same Directory Server and use the same password encryption key. For information about installing Access Manager, see "Installing Access Manager on Multiple Host Servers" on page 77.

This section include the following information about using a load balancer:

Configuring SSL Termination for a Load Balancer

Before you configure a load balancer to handle SSL requests, first configure SSL for the Access Manger web container. For instructions, see Chapter 3, "Configuring Access Manager in SSL Mode," in *Sun Java System Access Manager 7 2005Q4 Administration Guide*.

To configure SSL for a load balancer and Access Manager servers, consider the following cases:

- SSL configuration for only the load balancer: SSL termination.

  The load balancer terminates the SSL connection from the client and makes a separate SSL connection to the Access Manager servers.

- SSL configuration for only the Access Manager servers: SSL pass-through.

  The load balancer bypasses all the requests from the client to the Access Manager servers.
- SSL configuration for both the load balancer and Access Manager servers.

For all cases, except for the SSL pass-through configuration, you can use a normal server certificate to enable SSL termination for the load balancer. However, when you configure SSL pass-through for the load balancer and the Access Manager servers and the load balancer bypasses all the requests from the client to the Access Manager server, the following SSL problems exist for a normal server certificate:

- When a client accesses the Access Manager servers through the load balancer, the client gets the server certificate from the Access Manager server. The load balancer doesn't have an SSL server certificate and just bypasses the client requests to the back-end Access Manager servers. The client then receives a warning message saying that the host name and subject name in server certificate are different.
- To avoid the above problem, install a server certificate with the `SubjectDN` of the load balancer name; however, a problem occurs in the session validation between two Access Manager servers.

  For example, if a user gets a session from `amserver1` and a second request for the same user is directed to `amserver2`, then `amserver2` has to validate the users session to `amserver1`. When `amserver2` sends a session validation request to `amserver1`, it makes an SSL connection to `amserver1` and then gets the server certificate with the `SubjectDN` of the load balancer from `amserver1`. Because those two names (host name of `amserver1` and `subjectDN` in certificate) differ, `amserver2` stops the SSL handshaking, and the session validation fails.

To solve these problems, Access Manager provides these properties:

- `com.iplanet.am.jssproxy.trustAllServerCerts`

  If enabled (true), Access Manager ignores all certificate related issues (such as a name conflict) and continues the SSL handshaking.

Caution – To prevent a possible security risk, enable this property only for testing or when the enterprise network is tightly controlled. Avoid enabling this property if a security risk might occur (for example, if a server connects to a server in a different network).

- `com.iplanet.am.jssproxy.SSLTrustHostList`

  If enabled (true), Access Manager checks the platform server list in the `AMConfig.properties` file. If the server FQDNs of the two servers in the platform server list match, Access Manager continues the SSL handshaking.
- `com.iplanet.am.jssproxy.checkSubjectAltName`

If enabled (by specifying a comma separated list of trusted FQDNs) and a server certificate includes the Subject Alternative Name (SubjectAltName) extension, Access Manager checks all name entries in the extension. If one of names in the SubjectAltName extension is the same as the server FQDN, Access Manager continues the SSL handshaking. Using this property is more secure than enabling the com.iplanet.am.jssproxy.trustAllServerCerts property. With a Public-Key Infrastructure (PKIX) definition, a certificate can have multiple subject names with SubjectAltName extension.

To enable this property, set it to a comma separated list of trusted FQDNs. For example:

```
com.iplanet.am.jssproxy.checkSubjectAltName=
amserv1.example.com,amserv2.example.com
```

To get a certificate with SubjectAltName extension, see the next section.

Generating a CSR with the SubjectAltName Extension

To generate a certificate signing request (CSR) with the SubjectAltName extension, use the Certificate Database Tool (certutil). If certutil is not available in the /usr/sfw/bin directory, first install the SUNWtlsu package on Solaris systems or the sun-nss-sun-nss-devel RPM on Linux systems. If necessary, set the LD_LIBRARY_PATH environment variable to the appropriate certutil path.

For information about certutil, see: http://www.mozilla.org/

This section describes how to use the certutil if you are using Web Server or Application Server as the web container. If you are using BEA WebLogic Server or IBM WebSphere Application Server as the web container, refer to the respective BEA or IBM product documentation.

To generate a CSR with the SubjectAltName extension, follow these steps:

1. Log in as or become superuser (root}.
2. Create a new certificate database (cert8.db) using the certutil -N option. If necessary, first create a directory for your database. For example:

   ```
   # mkdir certdbdir
   # cd certdbdir
   # certutil -N -d .
   ```

   When prompted by certutil, enter the password to encrypt your keys:

   ```
   Enter a password which will be used to encrypt your keys.
   The password should be at least 8 characters long,
   and should contain at least one non-alphabetic character.

   Enter new password: your-password
   Re-enter password:  your-password
   ```

3. Generate the CSR with the SubjectAltName extension. For example:

```
# certutil -R -s "cn=lb.example.com,o=example.com,c=us"
-o server.req -d . -a -8 amserv1.example.com,amserv2.example.com
```

When prompted by `certutil`, enter the password (or pin) and then type keys to generate the random seed to create your key:

```
Enter Password or Pin for "NSS Certificate DB": your-password

A random seed must be generated that will be used in the
creation of your key.  One of the easiest ways to create a
random seed is to use the timing of keystrokes on a keyboard.

To begin, type keys on the keyboard until this progress meter
is full.  DO NOT USE THE AUTOREPEAT FUNCTION ON YOUR KEYBOARD!

Continue typing until the progress meter is full:

|************************************************************|

Finished.  Press enter to continue:

Generating key.  This may take a few moments...
```

4. Send the CSR (`server.req` file in the example) to the Certificate Authority (CA). Get the server certificate and add it to the certificate database using the `certutil -A` option.
5. Copy the certificate database (`cert8.db`) to the web container directory.
   - Web Server. Copy the `cert8.db` and `key3.db` databases to the `/opt/SUNWwbsrv/alias` directory and rename them using the Web Server instance name. For example:

     ```
     https-webserver.example.com-webserver-cert8.db
     https-webserver.example.com-webserver-key3.db
     ```

   - Application Server. Copy the `cert8.db` and `key3.db` databases to the instance `/config` directory. For example:

     ```
     /var/opt/SUNWappserver/domains/domain1/config/cert8.db
     /var/opt/SUNWappserver/domains/domain1/config/key3.db
     ```

Configuring Access Manager For Load Balancer Cookies

To configure Access Manager for load balancer cookies, update the configuration for all Access Manager instances in the deployment so that the instances can recognize the load balancer. In this scenario, multiple (two or more) Access Manager instances are deployed on different host servers. A load balancer routes client requests to the various Access Manager instances. All Access Manager instances use the same Directory Server.

1. In the Access Manager Console, configure the Access Manager deployment as a site, as described in "Configuring an Access Manager Deployment as a Site" on page 81. When you configure a deployment as a site, Access Manager

automatically sets the `fqdnMap` property (in memory) to include the load balancer.

2. In the `AMConfig.properties` file for each Access Manager instance, add the following properties:

   ```
   com.iplanet.am.lbcookie.name=amlbcookie
   com.iplanet.am.lbcookie.value=amserver
   ```

   where *amlbcookie* is the load balancer cookie, and *amserver* is the name of the Access Manager host server for the instance.

3. Restart all Access Manager instances by restarting the respective web container.

Configuring a Load Balancer with SAML

In this scenario, an Access Manager site is using a load balancer to distribute client requests to various Access Manager instances, and the site has implemented the Security Assertions Markup Language (SAML) service. When a request is sent to an Access Manager instance through a load balancer, the instance must know which other Access Manager server in the deployment issued the original assertion or artifact in order to retrieve the SAML assertion.

The deployment must first be configured as a site. Multiple Access Manager instances are installed on host servers, and a load balancer routes client requests to the various instances. All Access Manager instances access the same Directory Server. Access Manager session failover is optional.

To configure a site to use a load balancer with SAML, follow these steps:

1. The Access Manager deployment must be configured as a site in order for SAML load balancing to work. If you haven't configured the Access Manager deployment as a site, follow the instructions in "Configuring an Access Manager Deployment as a Site" on page 81.
2. Log in to the Access Manager Console as `amadmin`.
3. In the Access Manager Console, click `Federation` and then `SAML`.
4. Under the `Properties` section in `SAML Profile`, add or modify the following entries:
   - Site Identifiers. Add each Access Manager instance in the deployment. All Access Manager instances must share the same Site ID and Site Issuer Name.
   - Trusted Partners. Add your partner's deployment site's Source ID (site ID), Issuer Name, and Host List. The unique Source ID (site ID) and Issuer Name for the Access Manager servers and the URL or IP address or host name of the load balancer will identify the deployment and will be given out to your partner's site for configuration.

     For information about these fields, see the *Sun Java System Access Manager 7 2005Q4 Federation and SAML Administration Guide*.
5. Click `Save` to save your changes.

Setting the `fqdnMap` Property

If you have configured an Access Manager deployment as a site, Access Manager automatically sets the `fqdnMap` property (in memory) to include the load balancer , and you do not need to set this property in the `AMConfig.properties` file. However, for the following Access Manager deployments, you must explicitly set the property:

- The deployment is not configured as a site.
- The deployment has virtual hosts that are mapped to a physical host.

If you need to set the `fqdnMap` property, set the property to the load balancer in the `AMConfig.properties` file for each Access Manager instance in the deployment. If necessary, first remove the comment character (#) from the property. For example:

```
com.sun.identity.server.fqdnMap[lb.example.com]=lb.example.com
```

Accessing an Access Manager Instance Through a Load Balancer

Accessing an Access Manager instance through a load balancer depends on the mode (realm or legacy) and the console you want to access. Use the following syntax to access an Access Manager instance through a load balancer:

`http://`*loadbalancer*`.`*domain*`:`*port*`/amserver/console|/amconsole`

In legacy mode, you can access both consoles:

- New Access Manager 7 2005Q4 Console. For example:

  ```
  http://loadbalancer.example.com:80/amserver/console
  ```

- Access Manager 6 2005Q1 Console. For example:

  ```
  http://loadbalancer.example.com:80/amconsole
  ```

In realm mode, you can access only the new Access Manager 7 2005Q4 Console. For example:

```
http://loadbalancer.example.com:80/amserver/console
```

Implementing Access Manager Session Failover

Access Manager provides a web container independent session failover implementation using Sun Java System Message Queue (Message Queue) as the communications broker and the Berkeley DB by Sleepycat Software, Inc. as the session store database. Access Manager 7 2005Q4 enhancements includes the `amsfoconfig` script to configure the session failover environment and the `amsfo` script to start and stop the Message Queue broker and Berkeley DB client.

This section covers these topics:

Access Manager Session Failover Scenario

The following figure shows an Access Manager session failover deployment scenario that includes these components:

- Three Access Manager instances, running on different host servers on supported web containers. All Access Manager instances access the same Directory Server (not shown in the figure).
- Message Queue brokers, running in cluster mode on different servers.
- Berkeley DB client (`amsessiondb`), running on the same servers as the Message Queue brokers.
- Load balancer to improve performance and security.
- Client requests can originate from a Web browser, C or Java application using the Access Manager SDK, or a J2EE/Web agent.

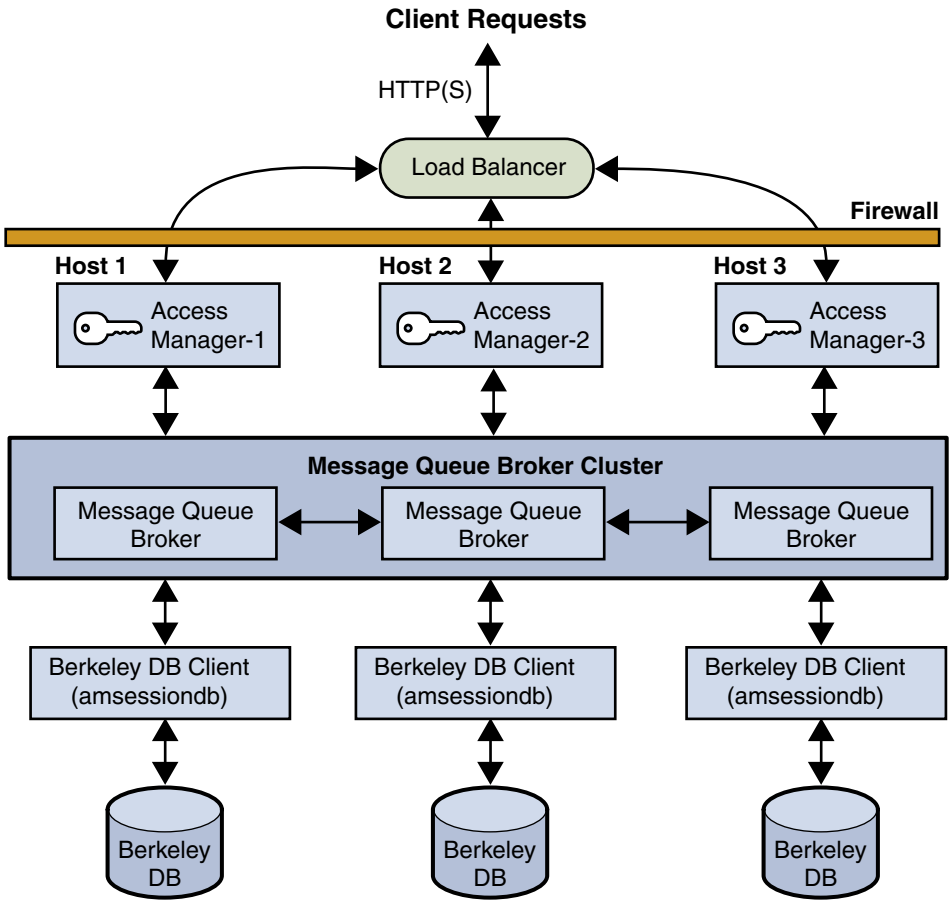

FIGURE 6–1 Access Manager Session Failover Scenario

Installing the Session Failover Components

The following table describes how to install the components required for Access Manager session failover.

TABLE 6–1 Installation of Access Manager Session Failover Components

Component	How to install ...
Access Manager	Install the first instance of Access Manager on each host server using the Java ES Installer. The installer adds the required session failover Solaris packages or Linux RPMs. Reference: *Sun Java Enterprise System 2005Q4 Installation Guide for UNIX* When you install Access Manager using the Java ES installer, you can select either Realm Mode (version 7.x) or Legacy Mode (version 6.x). Access Manager session failover is supported in both modes. After you run the Java ES installer, run the `amconfig` script to: ■ Configure the first Access Manager instance, if you specified the Configure Later option during installation. ■ Redeploy or reconfigure an installed Access Manager instance. For information, see "Installing Access Manager on Multiple Host Servers" on page 77.
Message Queue	Install Message Queue using the Java ES installer. Reference: *Sun Java Enterprise System 2005Q4 Installation Guide for UNIX*
Berkeley DB Client (Access Manager subcomponent)	The Java ES installer and `amconfig` script adds the Access Manager packages or RPMs required for the Berkeley DB client. However, if you want to install the Berkeley DB client on a server where you have not installed Access Manager, you must manually add the following packages or RPMs, depending on your operating system. For the Solaris OS, add the following packages using the `pkgadd` command: `SUNWamsfodb`, `SUNWbdb`, and `SUNWbdbj`. Reference: Solaris documentation For the Linux OS, add the following RPMs using the `rpm` command: `sun-identity-sfodb`, `sun-berkeleydatabase-core`, and `sun-berkeleydatabase-java`. Reference: Linux online man pages.

Caution – In a multiple server deployment that shares the same Directory Server, all instances of Access Manager must use the same password encryption key value. When you install the first Access Manager instance, save the password encryption key value from the `am.encryption.pwd property` in the `AMConfig.properties` file. Then, when you run the Java ES installer or `amconfig` script to deploy Access Manager instances on other host servers, use this same value for the password encryption key.

Configuring Access Manager for Session Failover

To configure Access Manager for session failover, follow these steps:

Each step is described in detail in the following sections.

To determine if session failover is enabled for a deployment, change the `com.iplanet.services.debug.level` property from `error` to `message` in the `AMConfig.properties` file. Then, check the `amSession` logs in the `/var/opt/SUNWam/debug` directory on Solaris systems or the `/var/opt/sun/identity/debug` directory on Linux systems.

1–Disable Cookie Encoding

On each host server that is running an Access Manager instance, disable cookie encoding:

- If Web Server is the web container, make sure the following property in the `AMConfig.properties` file is set to false (which is the default value set by the Java ES installer):

  ```
  com.iplanet.am.cookie.encode=false
  ```

  In the `sun-web.xml` file, set the `encodeCookies` property to false. For example:

  ```
  <sun-web-app>
  <property name="encodeCookies" value="false"/>
  ...
  </sun-web-app>
  ```

- If Application Server, BEA WebLogic, or IBM WebSphere Application Server is the web container, set the following property in the `AMConfig.properties` file to false:

  ```
  com.iplanet.am.cookie.encode=false
  ```

The Access Manager client should not do any cookie encoding or decoding. A remote SDK client must be in sync with the Access Manager server side settings, either in the `AMConfig.properties` file or the web container's `sun-web.xml` file.

2–Edit the Web Container `server.xml` File

On each host server that is running an Access Manager instance, add the installed locations of `imq.jar` and `jms.jar` in the `server.xml` (or equivalent) configuration file for the Access Manager web container. For example, on Solaris systems:

```
<JAVA javahome="/usr/jdk/entsys-j2se" serverclasspath=
"/usr/share/lib/imq.jar:/usr/share/lib/jms.jar:
/opt/SUNWwbsvr/bin/https/jar/webserv-rt.jar:
${java.home}/lib/tools.jar:
/opt/SUNWwbsvr/bin/https/jar/webserv-ext.jar:
/opt/SUNWwbsvr/bin/https/jar/webserv-jstl.jar:
/usr/share/lib/ktsearch.jar"
```

3–Add a New User in the Message Queue Server

If you don't want to use the `guest` user as the Message Queue user name and password, add a new user and password to connect to the Message Queue broker on servers where Message Queue is installed. For example, on Solaris systems, to add a new user named `amsvrusr`:

```
# /usr/bin/imqusermgr add -u amsvrusr -p password
```

Then, make the `guest` user inactive by issuing the following command:

```
# /usr/bin/imqusermgr update -u guest -a false
```

4–Edit the `amsessiondb` Script (if Needed)

The `amsessiondb` script is called by the `amsfo` script to start the Berkeley DB client (`amsessiondb`), create the database, and set specific database values. The script contains variables that specify various default paths and directories:

```
JAVA_HOME=/usr/jdk/entsys-j2se/
IMQ_JAR_PATH=/usr/share/lib
JMS_JAR_PATH=/usr/share/lib
BDB_JAR_PATH=/usr/share/db.jar
BDB_SO_PATH=/usr/lib
AM_HOME=/opt/SUNWam
```

If any of these components are not installed in their default directories, edit the `amsessiondb` script and set the variables, as needed, to the correct locations.

5–Run the `amsfoconfig` Script

Access Manager 7 2005Q4 provides the `amsfoconfig` script to configure an Access Manager deployment for session failover.

Requirements to Run the `amsfoconfig` Script

To run the `amsfoconfig` script, an Access Manager deployment must meet the following requirements:

- Two or more Access Manager instances must be installed and configured in the deployment, but the deployment cannot be configured as a site. If the `amsfoconfig` script determines that the deployment is configured as a site or that any of the server entries in the platform server list are site enabled, the script displays a message and exits. To configure session failover manually, see "Configuring Session Failover Manually" on page 102
- The Java Message Queue (MQ) broker must be installed and configured on at least two servers in the deployment.
- The Berkeley DB client and database must be installed and configured in the deployment.
- Directory Server must be running, accessible to the script, and configured with Access Manager data.

Functions of the `amsfoconfig` *Script*

The `amsfoconfig` script reads the `amsfo.conf` configuration file and then configures an Access Manager deployment for session failover by performing these functions:

- Configures a new site. The script uses the Access Manager instances in the platform server list and the load balancer information from the `amsfo.conf` file to create a new site for the Access Manager session failover deployment. The script modifies the existing platform server list, so that after the site is configured, all server entries under the platform server list then belong to the site.

  For example, `http://server1.example.com:80|01` changes to `http://server1.example.com:80|01|10`, if the default value of 10 is used as the `SiteID`.
- Modifies the existing Realm/DNS alias list. The script appends the host name of the load balancer to the list. This host name is obtained from the `lbServerHost` variable of the `amsfo.conf` file.
- Loads session failover configuration XML into Directory Server. The script dynamically generates the session configuration XML file based on the configuration information and loads the generated XML into Directory Server. This information corresponds to the Secondary Configuration Instance under Session in the Access Manager Console.

The following table lists the Access Manager session failover scripts and configuration files.

TABLE 6–2 Access Manager Session Failover Scripts and Configuration Files

Name	Description and Location
`amsofconfig`	Script to configure Access Manager for session failover.
	Solaris systems: *AccessManager-base*`/SUNWam/bin`
	Linux systems: *AccessManager-base*`/identity/bin`
`amsfo`	Script to start and stop the Message Queue broker and `amsessiondb` client.
	Solaris systems: *AccessManager-base*`/SUNWam/bin`
	Linux systems: *AccessManager-base*`/identity/bin`
`amsfopasswd`	Script to generate the encrypted Message Queue broker user password.
	Solaris systems: *AccessManager-base*`/SUNWam/bin`
	Linux systems: *AccessManager-base*`/identity/bin`

TABLE 6–2 Access Manager Session Failover Scripts and Configuration Files *(Continued)*

Name	Description and Location
`amsfo.conf`	Session failover configuration file.
	Solaris systems: *AccessManager-base*`/SUNWam/lib`
	Linux systems: *AccessManager-base*`/sun/identity/lib`
`amProfile.conf`	Session failover environment file.
	Solaris systems: `etc/opt/SUNWam/config`
	Linux systems: `etc/opt/sun/identity/config`

AccessManager-base represents the base installation directory for Access Manager. The default values are:

Solaris systems: `/opt`

Linux systems: `/opt/sun`

Running the `amsfoconfig` Script

To run the `amsfoconfig` script to configure Access Manager for session failover, follow these steps.

1. Log in as or become superuser (`root`).
2. Set the variables in the `amsfo.conf` file, as described in Table 6–3.
3. Run the script. For example, on a Solaris system with Access Manager installed in the default directory:

   ```
   # cd /opt/SUNWam/bin
   # ./amsfoconfig
   ```

   The script displays status information as it runs.
4. When the `amsfoconfig` script prompts you, enter the following passwords:
   - Access Manager administrator (`amAdmin`) password
   - Message Queue broker user password
5. To check the results, see the `/var/tmp/amsfoconfig.log` file.

The following table describes the variables in the `amsfo.conf` file that are used by the `amsfoconfig` script. Set these variables as needed for your deployment before you run the `amsfoconfig` script.

TABLE 6–3 Variables in the `amsfo.conf` File Used by the `amsfoconfig` Script

Variable	Description
`CLUSTER_LIST`	Message Queue broker list participating in the cluster. The format is: *host1*`:`*port*`,`*host2*`:`*port*`,`*host3*`:`*port* For example: `jmq1.example.com:7777,jmq2.example.com:7777, jmq3.example.com:7777` There is no default.
`lbServerPort`	Port for the load balancer. The default is 80.
`lbServerProtocol`	Protocol (`http` or `https`) used to access the load balancer. The default is `http`.
`lbServerHost`	Name of the load balancer. For example: `lbhost.example.com`
`SiteID`	Identifier for the new site (and the load balancer) that the `amsfoconfig` script will create. `SiteID` can be any value greater than the Server IDs that already exist in the platform server list. The default is 10.

`amsfoconfig` Script Sample Run

The following example shows a sample run of the `amsfoconfig` script.

```
Welcome to Sun Java System Access Manager 7 2005Q4

Session Failover Configuration Setup script.
============================================================
============================================================
Checking if the required files are present...
============================================================

Running with the following Settings.
------------------------------------------------
Environment file: /etc/opt/SUNWam/config/amProfile.conf
Resource file: /opt/SUNWam/lib/amsfo.conf
         ------------------------------------------------
Using /opt/SUNWam/bin/amadmin

Validating configuration information.
Done...

Please enter the LDAP Admin password:
(nothing will be echoed): password1
Verify: password1
```

```
Please enter the JMQ Broker User password:
(nothing will be echoed): password2
Verify: password2

Retrieving Platform Server list...
Validating server entries.
Done...

Retrieving Site list...
Validating site entries.
Done...

Validating host: http://amhost1.example.com:7001|02
Validating host: http://amhost2.example.com:7001|01
Done...

Creating Platform Server XML File...
Platform Server XML File created successfully.

Creating Session Configuration XML File...
Session Configuration XML File created successfully.

Creating Organization Alias XML File...
Organization Alias XML File created successfully.

Loading Session Configuration schema File...
Session Configuration schema loaded successfully.

Loading Platform Server List File...
Platform Server List server entries loaded successfully.

Loading Organization Alias List File...
Organization Alias List loaded successfully.

Please refer to the log file /var/tmp/amsfoconfig.log for additional
information.
################################################################
Session Failover Setup Script. Execution end time 10/05/05 13:34:44
################################################################
```

Starting the Session Failover Components

Access Manager 7 2005Q4 provides the `amsfo` script to perform these functions:

- Start and stop the Java Message Queue (MQ) broker specified for the session failover deployment.
- Start and stop the `amsessiondb` client specified for the session failover deployment.
- Read the `amsfo.conf` configuration file and take specific actions based on variables in the file. For example, you can have the script first delete and then recreate the Berkeley DB database.

- Write the `amsessiondb.log`, `jmq.pid`, and `amdb.pid` files in the `/tmp/amsession/logs/` directory. The default log directory is determined by the `LOG_DIR` variable in the `amsfo.conf` file.

To start the Access Manager session failover components, follow this sequence:

1. Set the variables in the in the `amsfo.conf` configuration file, as required by your deployment. For a description of these variables, see Table 6–4
2. Run the `amsfo` script to start the Java Message Queue (MQ) broker and the `amsessiondb` client. For detailed information, see "Running the `amsfo` Script" on page 99.
3. Start each Access Manager instance by starting the respective web container. For information, see the *Sun Java System Access Manager 7 2005Q4 Administration Guide*.

Running the `amsfo` Script

The `amsfo` script includes the start and stop options:

Usage: `amsfo { start | stop }`

To run the `amsfo` script, follow these steps:

1. Log in as or become superuser (`root`).
2. Set the variables in the `amsfo.conf` file, as required for your deployment. For a description of these variables, see Table 6–4.
3. Run the script. For example, to start the session failover components on a Solaris system with Access Manager installed in the default directory:

```
# cd  /opt/SUNWam/bin
# ./amsfo start
```

4. To check the results of the script, see the `/tmp/amsession/logs/amsessiondb.log` file.

The following table describes the variables in the `amsfo.conf` configuration file. Set these variables as needed for your deployment before you run the `amsfo` script.

TABLE 6–4 `amsfo.conf` Configuration File

Variable	Description
`AM_HOME_DIR`	Access Manager default installation directory. The default directory depends on the platform: Solaris systems: *AccessManager-base*`/SUNWam` Linux systems: *AccessManager-base*`/identity` *AccessManager-base* represents the base installation directory for Access Manager. The default values are `/opt` on Solaris systems and `/opt/sun` on Linux systems.
`AM_SFO_RESTART`	Specifies (true or false) whether the script should automatically restart the `amsessiondb` client. The default is true (restart the `amsessiondb` client).
`CLUSTER_LIST`	Message Queue broker list participating in the cluster. The format is: *host1*`:`*port*`,`*host2*`:`*port*`,`*host3*`:`*port* For example: `jmq1.example.com:7777,jmq2.example.com:7777, jmq3.example.com:7777` There is no default.
`DATABASE_DIR`	Directory where the session database files will be created. The default is `"/tmp/amsession/sessiondb"`.
`DELETE_DATABASE`	Specifies (true or false) whether the script should delete and then create a new database when the `amsessiondb` process is restarted. The default is true.
`LOG_DIR`	Location of the log directory. The default is `"/tmp/amsession/logs"`.
`START_BROKER`	Specifies (true or false) whether the Message Queue broker should be started with the `amsessiondb` process. Set this variable as follows: true - The Message Queue broker will run on the same machine as the `amsessiondb` process. false - The Message Queue broker and the `amsessiondb` process will run on different machines. The default is true.
`BROKER_INSTANCE_NAME`	Name of the Message Queue broker instance to start. The default is `aminstance`.

TABLE 6–4 `amsfo.conf` Configuration File *(Continued)*

Variable	Description
`BROKER_PORT`	Port for the local Message Queue broker instance. The default is 7777.
`BROKER_VM_ARGS`	Java VM arguments. The default is `"-Xms256m -Xmx512m"`, which sets the maximum value based on the system resources.
`USER_NAME`	User name used to connect to the Message Queue broker. The default is guest. If you specified a different user name under step "3–Add a New User in the Message Queue Server" on page 93, set `USER_NAME` to that name.
`PASSWORDFILE`	Location of the password file that contains the encrypted password used to connect to the Message Queue broker. To generate the encrypted password, use the `amsfopasswd` script, as described in "`amsfopasswd` Script" on page 101 The default is *$AM_HOME_DIR*`/.password`, where *$AM_HOME_DIR* specifies the Access Manager default installation directory.

`amsfopasswd` Script

The `amsfopasswd` script accepts the Message Queue broker password in clear text and returns the encrypted password in a file. You can then use this file as input to the `amsfo` script (`PASSWORDFILE` variable).

The `amsfopasswd` script is located in the following directory:

- Solaris systems: *AccessManager-base*`/SUNWam/bin`
- Linux systems: *AccessManager-base*`/identity/bin`

The default *AccessManager-base* installation directory is `/opt` on Solaris systems and `/opt/sun` on Linux systems.

Use the following syntax to run the `amsfopasswd` script.

```
amsfopasswd -f filename | --passwordfile filename
            -e password | --encrypt password
amsfopasswd -h | --help
```

The following table describes the `amsfopasswd` script arguments.

TABLE 6–5 `amsfopasswd` Script Arguments

Argument	Description
`-f` *filename* | `--passwordfile` *filename*	Path to the destination file where `amsfopasswd` stores the encrypted password.
`-e` *password* | `--encrypt` *password*	Clear text password that `amsfopasswd` encrypts.
`-h` | `--help`	Display the `amsfopasswd` command usage and then exit.

The following example shows the `amsfopasswd` script. The encrypted password is stored in the `/opt/SUNWam/.password` file.

```
# ./amsfopasswd -f /opt/SUNWam/.password -e mypassword
```

Configuring Session Failover Manually

In some situations, you might need to manually configure Access Manager for session failover. For example, you do not plan to run the `amsfoconfig` script. Or, the `amsfoconfig` script exited with one of the following messages before finishing the configuration: “Site is already configured” or “Server entry is already site configured”.

These steps describe how to manually configure Access Manager for session failover:

- “1–Install the Required Components in the Deployment” on page 102
- “2–Configure the Access Manager Deployment as a Site” on page 102
- “3–Create a New Secondary Configuration Instance for the Load Balancer” on page 103
- “4–Perform Session Failover Miscellaneous Configuration Tasks” on page 103
- “5–Start the Session Failover Components” on page 103

These steps are equivalent to the previous steps that described how to install the required components, configure session failover using the `amsfoconfig` script and then start the various components.

1–Install the Required Components in the Deployment

Install all components in the deployment, including Access Manager instances, load balancer, Message Queue, and the Berkeley DB client. For more information, see “Installing the Session Failover Components” on page 90.

2–Configure the Access Manager Deployment as a Site

If you do not plan to run the `amsfoconfig` script, which configures multiple Access Manager instances and a load balancer as a site, you must configure the deployment, as described in “Configuring an Access Manager Deployment as a Site” on page 81.

3–Create a New Secondary Configuration Instance for the Load Balancer

To create a new secondary configuration instance for your load balancer, follow these steps:

1. Log in to the Access Manager 7 2005Q4 Console as `amAdmin`.
2. Click `Configuration`, `Global Properties`, `Session`, and then `Secondary Configuration Instance`.
3. c. Click `New`, and add the following values:
   - `Name`. Load balancer URL. For example: `http://lb.example.com:80`
   - `Session Store User`. Name you are using to connect to the Message Queue Server (if other than guest).
   - `Session Store Password`. Password for the Session Store User.
   - `Maximum Wait Time`. 5000 (Use the default unless you require another value).
   - `Database Url`: Message Queue broker address list. For example:

     ```
     mqsvr1.example.com:7777,mqsvr2.example.com:7777,
     mqsvr3.example.com:7777
     ```

     The default Message Queue port is 7676. If you are using Application Server as the web container, however, consider using another port, because port 7676 might already be in use by Application Server. For the range of the valid port numbers, refer to the Message Queue documentation.
4. Click `Add` to save your changes.

4–Perform Session Failover Miscellaneous Configuration Tasks

Perform the following tasks (which are the same as if you are running the `amsfoconfig` script):

- "1–Disable Cookie Encoding" on page 93.
- "2–Edit the Web Container `server.xml` File" on page 93.
- "3–Add a New User in the Message Queue Server" on page 93.
- "4–Edit the `amsessiondb` Script (if Needed)" on page 94.

5–Start the Session Failover Components

Run the `amsfo` script to start the Message Queue broker and Berkeley DB client (`amsessiondb`). Then, start each Access Manager instance by starting the respective web container. See "Starting the Session Failover Components" on page 98.

amsessiondb Script

The amsessiondb script is called by the amsfo script to start the Berkeley DB client (amsessiondb), create the database, and set specific database values.

Note – The recommended method to start and stop the Access Manager session failover components is to run the amsfo script and let it call the amsessiondb script. The following information is included only in case you might need to run the amsessiondb script independently.

Before you run the amsessiondb script, make sure you have the paths set correctly, as described under "4–Edit the amsessiondb Script (if Needed)" on page 94.

When you run the amsessiondb script, you can enter the Message Queue broker password on the command line as clear text (-w or --password option). However, if you prefer to use an encrypted password in a file (-f or --passwordfile option), first run the amsfopasswd script to encrypt the Message Queue broker clear text password to a file. Then run the amsessiondb script, using this file for the -f or --passwordfile option.

Use the following syntax to run the amsessiondb script.

```
amsessiondb [ -u username | --username username ]
[ -w password | --password password |
-f filename | --passwordfile filename ]
[ -c cachesize | --cachesize cachesize ]
[ -b dbdirectory | --dbdirectory dbdirectory ]
-a MQServerAddressList | --clusteraddress MQServerAddressList
[ -s numcleanexpiredsessions | --numcleansessions numcleanexpiredsessions ]
[ -v | --verbose ]
[ -i statsinterval | --statsInterval statsinterval ]
amsessiondb -h | --help
amsessiondb -n | --version
```

The following table describes the amsessiondb script arguments.

TABLE 6–6 amsessiondb Script Arguments

Argument	Description
-u *username* | --username *username*	User name to connect to the Message Queue broker. Specify the user you specified under "3–Add a New User in the Message Queue Server" on page 93. Default is "guest".

TABLE 6–6 `amsessiondb` Script Arguments *(Continued)*

Argument	Description
-w *password* | --password *password*	Clear text password for the user name used to connect to the Message Queue broker. Specify the password you specified under "3–Add a New User in the Message Queue Server" on page 93. Default is "guest".
-f *filename* | --passwordfile *filename*	File that contains the encrypted password for accessing the Message Queue broker. Note If you specify this option, do not specify the -w or --password option.
-c *cachesize* | --cachesize *cachesize*	Cache size in MB. Default is 8 MB.
-b *dbdirectory* | --dbdirectory *dbdirectory*	Base directory where the Berkeley DB database (`amsessions.db`) is created. Default is "sessiondb", created in the directory where you are running the `amsessiondb` script. **Note** To ensure that you have sufficient disk space where you are creating the database, allow 1 GB for each 100,000 sessions.
-a *MQServerAddressList* | --clusteraddress *MQServerAddressList*	Message Queue broker address list, in the format: *host*`1:`*port*`[,`*host*`2:`*port*`,`*host*`3:`*port*`,...]` For example: `mqsvr1:7777,mqsvr2:7777`
-s numcleanexpiredsessions | --numcleansessions numcleanexpiredsessions	Number of expired sessions to be deleted for each cleanup interval. Default is 1000.
`-v | --verbose`	Run in verbose mode. Results are sent to the standard output. Default is non-verbose mode.
-i statsinterval | --statsInterval statsinterval	Interval in seconds to print the statistics for total requests, reads, writes, and deletes to the standard output. Default is 60 seconds.
`-h | --help`	Display `amsessiondb` command usage and then exit.
`-n | --version`	Return the version of Access Manager currently installed and then exit.

The following example shows the `amsessiondb` script.

```
amsessiondb -u amsvrusr -f pwfile -c 128 -b sessiondb
-a host1:7777,host2:7777
```

Performance Tests With the amsessiondb Client

Performance tests with the amsessiondb client include this criteria:

- The approximate number of operations on the Berkeley DB was two times the number of authentications per second.
- Tests were conducted with the following configuration:
  - Write data – 3 Kbytes
  - Duration – 1 minute
  - Berkeley DB cache size – 28 Mbytes (default cache size in Access Manager 7 2005Q4 is 32 Mbytes)

The following table shows the results of the tests.

TABLE 6–7 Performance Tests With the amsessiondb Client

Disk	Notes
Normal IDE disk: 666 writes per second	Each site can support up to 300 authentications per second. Therefore, IDE disks are not recommended.
Normal 10K RPM SCSI disk on Sun Blade server: 1520 writes per second	Each site can support up to 750 authentications per second.
Seagate Cheetah 15K RPM SCSI disk: 1860 writes per second	Each site can support up to 900 authentications per second.
Sun T-300 disk array: 2700 writes per second	Each site can support up to 1300 authentications per second.
Disk using swap space in /tmp: 3300 writes per second	Each site can support up to 1600 authentications per second.

Setting Session Quota Constraints

Access Manager 7 2005Q4 includes the new session quota constraints feature, which allows Access Manager to limit users to a specific number of active, concurrent sessions based on configurable attributes. An Access Manager administrator can set session quota constraints at the following levels:

- Globally. Constraints apply to all users.
- To an entity (organization or realm, role, or user). Constraints apply only to the specific users that belong to the entity.

Deployment Scenarios for Session Quota Constraints

The following Access Manager deployments support session quota constraints:

- Access Manager Single Server Deployment

  In this scenario, Access Manager is deployed on a single host server. Access Manager maintains the active session counts in memory for all logged in users. When a user attempts to log in to the server, Access Manager checks whether the number of the valid sessions for the user exceeds the session quota and then takes action based on the configured session quota constraints options.

- Access Manager Session Failover Deployment

  In this scenario, multiple instances of Access Manager are deployed on different host servers in a session failover configuration. The Access Manager instances are configured for session failover using Sun Java System Message Queue (Message Queue) as the communications broker and the Berkeley DB by Sleepycat Software, Inc. as the session store database. For more information about Access Manager session failover, see "Implementing Access Manager Session Failover" on page 89.

In a session failover deployment, when a user attempts to log in, the Access Manager server receiving the session creation request first retrieves the session quota for the user from the Access Manager identity repository. Then, the Access Manager server fetches the session count for the user directly from the centralized session repository (accumulating all the sessions from all the Access Manager servers within the same site) and checks whether the session quota has been exhausted. If the session quota has been exhausted for the user, the Access Manager server takes action based on the configured session quota constraints options.

If session constraints are enabled in a session failover deployment and the session repository is not available, users (except superuser) are not allowed to log in.

In a session failover deployment, if an Access Manager instance is down, all the *valid* sessions previously hosted by that instance are still considered to be valid and are counted when the server determines the actual active session count for a given user. An Access Manager multiple server deployment that is not configured for session failover does not support session quota constraints.

Configuration of Session Quota Constraints

To configure session quota constraints, the top-level Access Manager administrator (such as `amAdmin`) must set the following attributes in the Access Manager Console for one of the Access Manager instances. If you reset any of these attributes, you must restart the server for the new value to take effect.

- **Enable Quota Constraints** is a global attribute that enables or disables the session quota constraints feature. If this attribute is enabled, Access Manager enforces session quota constraints whenever a user attempts to logs in via a new client (and thus create a new session).

The default is disabled (OFF).

- **Read Timeout for Quota Constraint** defines the time in milliseconds that an inquiry to the session repository for the active user session counts continues before timing out. If the maximum wait time is reached due to the unavailability of the session repository, the session creation request is rejected.

  The default is 6000 milliseconds.

- **Resulting Behavior If Session Quota Exhausted** determines the behavior if a user exhausts the session constraint quota. This attribute takes effect only if the "Enable Quota Constraints" attribute is enabled. Values can be:

  - DENY_ACCESS. Access Manager rejects the login request for a new session.
  - DESTROY_OLD_SESSION. Access Manager destroys the next expiring existing session for the same user and allows the new login request to succeed.

  The default is DESTROY_OLD_SESSION.

- **Exempt Top-Level Admins From Constraint Checking** specifies whether session constraint quotas apply to the administrators who have the Top-level Admin Role. This attribute takes effect only if the "Enable Quota Constraints" attribute is enabled.

  The default is NO.

  The super user defined for Access Manager in the `AMConfig.properties` file (`com.sun.identity.authentication.super.user`) is always exempt from session quota constraint checking.

- Active User Sessions defines the maximum number of concurrent sessions for a user. Access Manager includes both a dynamic attribute and a user attribute, with same attribute name.

  The default is 5.

Multiple Settings For Session Quotas

If a user has multiple settings for session quotas at different levels, Access Manager follows this precedence to determine the actual quota for the user:

- user (highest)
- role/organization/realm (based on the conflict resolution levels)
- global (lowest)

For example, Ken is a member of both the marketing and management roles. Session quotas are defined as follows (all have the same conflict resolution level):

- organization - 1
- marketing role - 2
- management role - 4
- user Ken - 3

Ken's quota is 3.

For more information about the session quota constraints attributes, see the Access Manager Console online help.

Enabling Session Property Change Notifications

The session property change notification feature causes Access Manager to send a notification to all registered listeners when a change occurs on a specific session property. This feature takes effect when the "Enable Property Change Notifications" attribute is enabled (ON) in the Access Manager Console.

For example, in a single sign-on (SSO) environment, one Access Manager session can be shared by multiple applications. When a change occurs on a specific session property defined in the "Notification Properties" list, Access Manager sends a notification to all registered listeners.

All client applications participating in the SSO automatically get the session notification if they are configured in the notification mode. The client cached sessions are automatically updated based on the new session state (including the change of any session property, if there is any). An application that wants to take a specific action based on a session notification can write an implementation of the `SSOTokenListener` interface and then register the implementation through the `SSOToken.addSSOTokenListener` method. For more information, see the *Sun Java System Access Manager 7 2005Q4 Developer's Guide*.

To configure session property change notifications, follow these steps:

1. Log in to Access Manager Console as `amAdmin`.
2. Click the Configuration tab.
3. Under Global Properties, click Session.
4. Set "Enable Property Change Notifications" to ON.
5. In the "Notification Properties" list, add each property for which you want a notification sent when the property is changed.
6. When you have finished adding properties to the list, click Save.

Tuning Your Deployment

After you install Access Manager, you can tune your deployment for optimum performance using the `amtune` and related scripts. These scripts allow you to tune Access Manager, the Solaris™ Operating System (OS), the web container, and Directory Server.

The Java Enterprise System installer installs the tuning scripts and related files in the `bin/amtune` directory.

The `amtune` script is not interactive. Before you run `amtune`, you must edit the parameters in the `amtune-env` configuration file to specify the tuning you want `amtune` to perform for your specific environment. The `amtune-env` configuration file includes two major sections:

- Performance related parameters that you set to control the tuning.
- An internal section that is maintained by Access Manager engineering and should not be modified.

You can run the `amtune` script in two modes:

- Review mode: `amtune` reports tuning recommendations but does not make any actual changes to your environment.
- Change mode: `amtune` makes actual changes, except for Directory Server, depending on parameters in the `amtune-env` configuration file.

The `amtune` script does not automatically tune Directory Server. Most deployments have applications other than Access Manager that also access Directory Server, so you don't want to make tuning changes without considering how they would affect your other applications.

Before you tune Directory Server, first back up your Directory Server data using `db2bak`.

When you run `amtune`, the script creates a tar file that contains the Directory Server tuning script, `amtune-directory`. Untar this file in a temporary directory and then run the script in review mode. When you are certain that your changes are acceptable for all applications at your deployment, run `amtune-directory` in change mode.

For detailed information about running the tuning scripts and setting tuning parameters in the `amtune-env` configuration file, see the *Sun Java System Access Manager 7 2005Q4 Performance Tuning Guide*.

APPENDIX **A**

Installed Product Layout

This appendix describes the directory layout after you install Sun Java™ System Access Manager 7 2005Q4 using the Sun Java™ System Enterprise System (Java ES) installer.

The following table shows a summary of the Access Manager default directories after installation.

Summary of Access Manager Directories

TABLE A–1 Summary of Access Manager Directories

Description	Default Directory
Base Installation Directory See "Base Installation Directory" on page 112.	Solaris systems: `/opt/SUNWam` Linux systems: `/opt/sun/identity` During installation, you can specify a different base installation directory for `/opt` or `/opt/sun`, if you prefer. However, do not change the `/SUNWam` or `/identity` product directory name.
Configuration Directory See "Configuration (`/config`) Directory" on page 117.	Solaris systems: `/etc/opt/SUNWam/config` Linux systems: `/etc/opt/sun/identity/config`
Temporary Files Directory	Solaris systems: `/var/opt/SUNWam/tmp` Linux systems: `/var/opt/sun/identity/tmp`

TABLE A–1 Summary of Access Manager Directories *(Continued)*

Description	Default Directory
Debug Files Directory	Solaris systems: `/var/opt/SUNWam/debug`
	Linux systems: `/var/opt/sun/identity/debug`
Log Files Directory	Solaris systems: `/var/opt/SUNWam/logs`
	Linux systems: `/var/opt/sun/identity/logs`

Base Installation Directory

The default base installation directory depends on the platform where you are installing Access Manager:

- Solaris systems: `/opt`
- Linux systems: `/opt/sun`

In the Access Manager documentation, the *AccessManager-base* variable represents the base installation directory.

Within the base installation directory, Access Manager packages, shared binary files, command-line tools, and other files are installed in the `/SUNWam` directory on Solaris systems and the `/identity` directory on Linux systems. Therefore, the default base and product directory also depend on the platform:

- Solaris systems: `/opt/SUNWam`
- Linux systems: `/opt/sun/identity`

Note – During installation, you can specify a different base installation directory if you wish. However, do not change the `/SUNWam` or `/identity` product directory name.

The `/SUNWam` or `/identity` directory contains the following files and directories:

- Web application archive (WAR) files, such as `amcommon.war`, `amconsole.war`, `ampassword.war` , and `amserver.war`.

  For information about WAR files, see the *Sun Java System Access Manager 7 2005Q4 Developer's Guide*.

  Subdirectories include:

  - "`/bin` Directory" on page 113
  - "`/docs` Directory" on page 114
  - "`/dtd` Directory" on page 114
  - "`/include` Directory" on page 115

After installing Access Manager, check the package installation accuracy by using the `pkgchk (1M)` utility. For example:

```
pkgchk -l -p /opt/SUNWam
```

/bin Directory

The following table describes the command-line tools and utilities in the `/bin` directory. For information about running these tools and utilities, see the *Sun Java System Access Manager 7 2005Q4 Administration Guide*.

TABLE A–2 Access Manager Command-Line Tools and Utilities

Utility	Description
`am2bak`	Backs up the Access Manager components.
`amadmin`	Load XML service files into Directory Server and performs batch administrative tasks on the DIT.
`amsfo`, `amsfoconfig`, `amsfopassword`	Access Manager session failover scripts.
`ampassword`	Changes passwords for Access Manager administrator or users.
`amsamplesilent`	Sample silent install file for use with the installation and configuration scripts.
`amconfig`, `amutils`, `amdsconfig`, `amsdkconfig`, `amsvcconfig`, `amas70config`, `amwas51config`, `amwl81config`, `amws61config`	Installation and configuration scripts for installing, configuring, and uninstalling Access Manager instances. For information about these scripts, see the *Sun Java System Access Manager 7 2005Q4 Administration Guide*.
`amserver`	Start and stops the `amunixd` and `amsecuridd` daemons.
`amtune`	Sets operating system, Access Manager, web container, and Directory Server parameters to improve performance.

TABLE A–2 Access Manager Command-Line Tools and Utilities *(Continued)*

Utility	Description
`amverifyarchive`	Verifies the log archives to detect possible tampering and/or deletion of any files in the archive.
`bak2am`	Restores Access Manager components backed up by the am2back utility.
`ldapmodify`	Edits the contents of an LDAP directory, either by adding new entries or by modifying existing ones.
`ldapsearch`	Issues search requests to an LDAP directory and displays the result as LDIF text.
`amGenerateLDIF.pl` and `amGenerateNI.pl`	Access Manager bulk federation scripts.
`am2bak.template`, `amserver.template`, `amadmin.template`, `amverifyarchive.template`, `ampassword.template`, and `bak2am.template`	Access Manager template files.

/docs Directory

The `/docs` directory contains the HTML. JAR, CSS, and related files used for the Java API reference (Javadocs).

/dtd Directory

The `/dtd` directory contains the Document Type Definition (DTD) files used by Access Manager. A DTD defines the structure for XML files accessed by Access Manager. For more information, see the *Sun Java System Access Manager 7 2005Q4 Developer's Guide*.

The following table describes the Access Manager DTD files in the `/dtd` directory.

TABLE A–3 Access Manager DTD Files

File	Description
`Auth_Module_Properties.dtd`	Defines the structure for XML files used by the authentication modules to specify their properties.
`amAdmin.dtd`	Defines the structure for XML files used to perform batch LDAP operations on the directory tree using the `amAdmin` command-line tool.

TABLE A–3 Access Manager DTD Files *(Continued)*

File	Description
`amWebAgent.dtd`	Defines the structure for XML files used to handle requests from, and send responses to, web agents. This file is deprecated and remains for purposes of backward compatibility.
`policy.dtd`	Defines the structure for XML files used to store policies in Directory Server.
`remote-auth.dtd`	Defines the structure for XML files used by the Authentication Service's remote Authentication API.
`server-config.dtd`	Defines the structure for `serverconfig.xml` which details ID, host and port information for all server and user types.
`sms.dtd`	Defines the structure for XML service files.
`web-app_2_2.dtd`	Defines the structure for XML files used by the Access Manager deployment container to deploy J2EE applications.

`/include` Directory

The `/include` directory contains header (`.h`) files.

`/ldaplib` Directory

The `/ldaplib/ldapsdk` subdirectory contains the shared object (`.so`) files needed to run the LDAP utilities included with Access Manager.

`/lib` Directory

The `/lib` directory contains JAR files and additional shared object (`.so`) files. It also contains a link to the `/etc/opt/SUNWam/config/AMConfig.properties` file.

`/locale` Directory

The `/locale` directory contains the localization properties files. Each properties file includes a corresponding English localization file. For example, `amAdminCLI_en.properties` is the corresponding file for `amAdminCLI.properties`.

/migration Directory

The /migration directory contains the scripts and supporting files used to migrate data from earlier versions of Access Manager. For example, the /opt/SUNWam/migration/61to62/scripts subdirectory contains the Upgrade61DitTo62 script, which is used to migrate a DIT to Access Manager 2005Q1.

For more information about migration, the *Sun Java Enterprise System 2005Q4 Upgrade Guide* (part number 819–2331).

/public_html Directory

The /public_html directory and subdirectories contain the HTML and related files used for the online help.

/samples Directory

The /samples directory contains the following subdirectories: /admin, /appserver, /authentication, /console, /csdk, /liberty, /logging, /phase2, /policy, /saml, and /sso.

Each subdirectory contains samples for the respective functionality, which is indicated by the subdirectory name. For more specific information about these samples, the Readme.html file.

/share Directory

The /share/bin subdirectory contains the following additional utilities used internally by Access Manager:

- amtune/amtune-utils
- amsecuridd, amunixd, amwar, checkport, and wsutils.ksh

/upgrade Directory

The /upgrade directory contains the following directories:

- The /scripts contains the upgrade scripts and files.
- The /services directory contains directories for the Access Manager services.

/web-src Directory

The /web-src directory contains the subdirectories in which Access Manager J2EE web applications are deployed on a web container. It contains the following subdirectories:

- applications/ directory where the Access Manager Console is deployed. It contains the index.html file and various subdirectories. The /console directory contains various console related subdirectories.
- The /common directory (and subdirectories) is where the Access Manager Liberty Common Domain component is deployed.
- The /password directory (and subdirectories) is where the Access Manager Password Synchronization component is deployed. It contains the index.html file and the various subdirectories.
- The /services directory (and subdirectories) is where Access Manager Core Services are deployed. It contains the index.html file and the various subdirectories.

Configuration (/config) Directory

The default location of the configuration (/config) directory depends on the platform where you are installing Access Manager:

- Solaris systems: /etc/opt/SUNWam
- Linux systems: /etc/opt/sun/identity

The /config directory contains configuration, XML, and LDIF files, including:

- The .version file contains the current version of Access Manager.
- The AMConfig.properties file, SSOConfig.properties, and LogConfig.properties contain Access Manager configuration attributes.
- The serverconfig.xml file provides configuration information for the Access Manager for Directory Server.
- The /ldif subdirectory contains the LDIF files needed for populating the Directory Server data store when installing Access Manager. For example:
  - During installation, the ds_remote_schema.ldif file loads the Access Manager specific LDAP schema object classes and attributes (such as the iplanet-am-managed-people-container) needed to store Access Manager data in Directory Server. The sunone_schema2.ldif file loads the Access Manager specific LDAP schema object classes and attributes.
  - During uninstallation, The ds_remote_schema_uninstall.ldif file removes the Access Manager LDAP schema object classes and attributes from Directory Server.

- The `/ums` subdirectory contains XML files, including:
  - The `ums.xml` file provides a set of templates that contain LDAP configuration information for objects managed using Access Manager.
  - The `/xml` subdirectory contains these files:
    - The `amserveradmin` script loads the Access Manager services.
    - The XML files are generally not used for configuration. If they are modified, they must be manually reloaded into the Directory Server data store. (Any changes in the server are not synchronized with these files.) For information about the XML files in this directory, see the *Sun Java System Access Manager 7 2005Q4 Developer's Guide*.

APPENDIX **B**

Changing the Password Encryption Key

Sun Java™ System Access Manager uses a password encryption key to encrypt user passwords. All Access Manager subcomponents must use the same password encryption key value. If you plan to deploy multiple instances of Access Manager on different host servers, you must use the same password encryption key for all instances.

Installation Considerations

When you install the first Access Manager instance, the Java Enterprise System installer generates a default password encryption key string. You can either accept this default value or specify another value produced by a J2EE random number generator. The installer stores the password encryption key value in the `am.encryption.pwd` attribute in the `AMConfig.properties` file.

If you specify a value for the password encryption key, the string must be at least 12 characters long.

To deploy multiple instances of Access Manager, save the password encryption key value from the `am.encryption.pwd` attribute after you install the first instance. Then, use this key value to set the value when you deploy additional instances:

- If you run the Java ES installer, copy this value into the Password Encryption Key field on the "Access Manager: Administration" panel.
- If you run the `amconfig` script, set the `AM_ENC_PWD` variable to this value in the `amsamplesilent` configuration file (or copy of the file).

The following scenarios explain why you might need to retrieve and change the password encryption key. In these scenarios, the Access Manager instances use the same Directory Server.

- If you are doing a multiple server installation of Access Manager and you did not save the password encryption key when you installed the first Access Manager instance, you must retrieve the key to use when you deploy additional instances.
- If you have deployed an additional Access Manager instance that uses a different password encryption key from the first Access Manager instance, you must modify the encryption key value to match the first instance.

What else needs to be changed if you change the password encryption key?

Passwords and the password encryption key must be consistent throughout a deployment. If you change a password in one place or instance, you must also update the password in all other places and instances.

The `serverconfig.xml` file contains the encrypted user passwords, which are identified by the `<DirPassword>` element. For example:

```
<DirPassword>
Adfhfghghfhdghdfhdfghrteutru
</DirPassword>
```

The `puser` and `dsameuser` passwords in `serverconfig.xml` are encrypted using the password encryption key defined in `am.encryption.pwd` in the `AMConfig.properties` file. If you change the password encryption key, you must also re-encrypt these passwords in the `serverconfig.xml` file using the `ampassword` utility.

For information about the `ampassword` utility, see the *Sun Java System Access Manager 7 2005Q4 Administration Guide*.

Changing the Key Value

To change the password encryption key value, follow these steps:

1. Log in as or become superuser (`root`) on the host server where the first Access Manager instance is installed.
2. In the `AMConfig.properties` file for the first Access Manager instance, get and save the values of the following attributes:
   - Password encryption key: `am.encryption.pwd`
   - Shared secret: `com.iplanet.am.service.secret`

   The `AMConfig.properties` file is installed in the following directories:

   - Solaris systems: `/etc/opt/SUNWam/config`
   - Linux systems: `/etc/opt/sun/identity/config`
3. Log in as or become superuser (`root`) on the server where the second Access Manager instance is deployed.

4. As a precaution, back up the `AMConfig.properties` and `serverconfig.xml` files, which are in the `/config` directory.
5. Stop the web container of the second Access Manager instance. For example, on Solaris systems, with Web Server as the web container:

   ```
   # cd /opt/SUNWwbsvr/https-host2-name
   #./stop
   ```

6. Edit the `AMConfig.properties` file and replace the values for `am.encryption.pwd` and `com.iplanet.am.service.secret` with the values that you saved from the first Access Manager instance in Step 2.
7. Because the encryption key defined in `am.encryption.pwd` is changed, you must run the `ampassword` utility to re-encrypt and replace the passwords in the `serverconfig.xml` file. The passwords in `serverconfig.xml` are identified by the `<DirPassword>` element. Consider the following cases:

   Passwords are the same. If the password for `puser` and `dsameuser` is the same as the `amadmin` password in `serverconfig.xml`, run `ampassword` to re-encrypt the `amadmin` password. For example on Solaris systems:

   ```
   # cd /opt/SUNWam/bin
   # ./ampassword --encrypt password
   ```

   where *password* is the password you used for `amadmin` when you installed the first instance. Use the `ampassword` output (new encrypted password) to replace the two passwords in the `serverconfig.xml` file for the second instance.

   Passwords are different. If the passwords for `puser` and `dsameuser` are different from the `amadmin` password in `serverconfig.xml`, run `ampassword` to re-encrypt each password (`type="proxy"` and `type="admin`).

   Use the `ampassword` output (new encrypted passwords) to replace the `puser` and `dsameuser` passwords in `serverconfig.xml` for the second instance.
8. Restart the web container for Access Manager for the second instance. For example, on Solaris systems, with Web Server as the web container:

   ```
   # cd /opt/SUNWwbsvr/https-host2-name
   # ./start
   ```

9. Repeat Step 3 through Step 8 for additional instances of Access Manager in the deployment.

Index

W