

CHAPTER 3

Enterprise Java Security Fundamentals

THE J2EE platform has achieved remarkable success in meeting enterprise needs, resulting in its widespread adoption. The security infrastructure plays a key role in the e-business strategy of a company. J2EE provides a standard approach to allow enterprise applications to be developed without hard-coded security policies. Instead, declarative policies are bundled with an assembled set of application components. Security policies specified using this security model are enforced in any operational environments and deployed in any of the application servers that host them.

The J2EE security model addresses authentication, authorization, delegation, and data integrity for the components that make up a J2EE environment. This environment includes J2EE applications—Web components, such as servlets and JSP files, EJB components, Java 2 connectors, and JavaMail—and secure interoperability requirements. The J2EE security model also considers the organizational roles that define and enforce these security policies:

- Application Component Provider
- Application Assembler
- Deployer
- System Administrator
- J2EE Product Provider

This chapter provides an overview of J2EE, exploring the J2EE security model. The chapter explains how various J2EE components are tied into enterprise security, describes how the J2EE security model addresses the security of J2EE components, and identifies the responsibility of each of the organizational roles in enforcing security. Declarative security policies and programmatic

security APIs are explained, in addition to the security requirements on JavaMail, Java connectors, client applications, and containers. This chapter also outlines the secure interoperability requirements that exist between various application servers.

Since its inception, one of the top requirements of the J2EE security model has been to support secure application deployments that do not rely on private networks or other application runtime isolation techniques. This allows application portability between containers. Another requirement has been to reduce the application developer's burden by delegating the security responsibilities to the J2EE roles. Finally, the policy-driven security model enables much of security enforcement to be handled without custom code.

3.1 Enterprise Systems

An *enterprise Java environment*, or WAS environment, is nominally viewed as a three-tier architecture (see Section 2.1.2 on page 25). Clients access the information made available through middle-tier systems, which connect to the back-end enterprise systems, as shown in Figure 3.1.

In an enterprise Java environment, the clients can be both Java based and non-Java based. Clients access the servers over a variety of protocols, including HTTP,

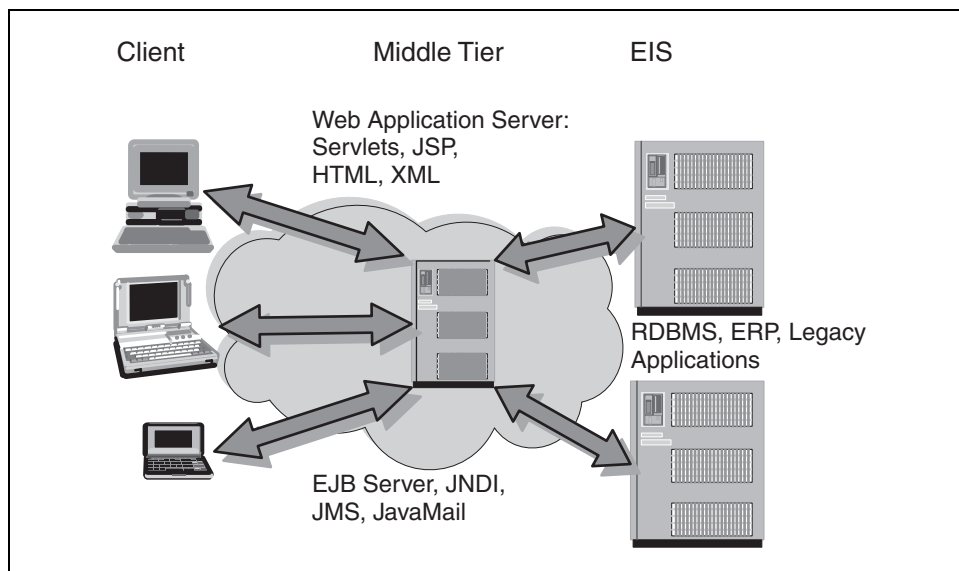


Figure 3.1. WAS Environment

IIOP, SSL, and other messaging protocols accessible through JMS. These clients connect to and access a J2EE-based server environment providing a hosting system for the enterprise components. These components constitute a presentation layer in the form of servlets, JSP files, HTML files, or XML documents. Alternatively, the components can abstract out the business logic in the form of enterprise beans. Clients may also submit their requests by using e-mail protocols through the JavaMail framework or connect to naming and directory services by using the Java Naming and Directory Interface (JNDI). In an enterprise environment, middle-tier applications are likely to connect to back-end enterprise information systems (EISs). Examples of back-end EISs include relational database management systems (RDBMSs) and ERP applications.

Before delving into the security implications of this architecturally rich environment, it is important to understand the technologies that comprise a J2EE environment.

3.2 J2EE Applications

A *J2EE application*, an enterprise application that conforms to the J2EE specification, is structured as shown in Figure 3.2 and consists of the following:

- Zero or more EJB modules
- Zero or more Web modules
- Zero or more application client modules

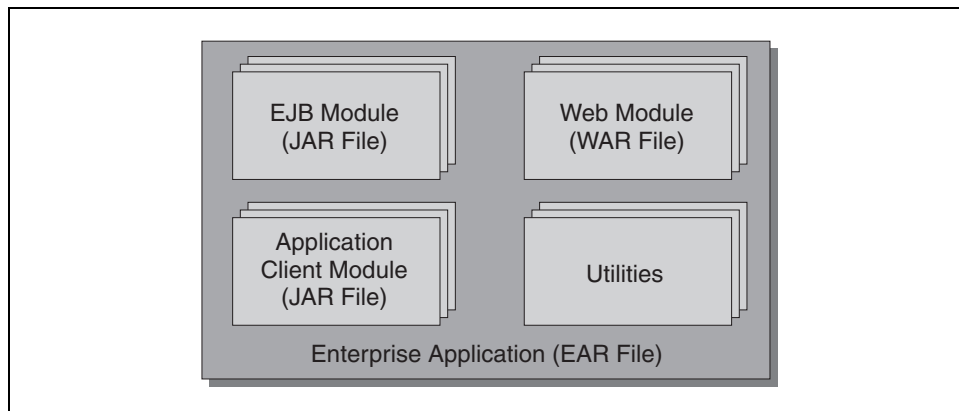


Figure 3.2. Contents of a J2EE Application

- Optionally, JAR files containing dependent classes or components required by the application
- Any combination of the preceding, as long as it contains at least one module

A J2EE application is represented by, and packaged in, an Enterprise Archive (EAR) file. The modules that comprise the EAR file are themselves packaged in archive files specific to their types. For example, a Web module is packaged in a Web Archive (WAR) file, and an EJB module, containing one or more enterprise beans, is packaged in a JAR file. WAR files can exist as independent deployment units from EAR files.

EAR files also contain a *deployment descriptor* file—an XML document describing the contents of the application and containing instructions for the deployment of the application. In particular, the deployment descriptor specifies the security settings to be enforced by the runtime environment. Each WAR file packaging a Web module, JAR file packaging enterprise beans, or JAR file packaging an application client module contains its own deployment descriptor as well.

3.2.1 EJB Modules

An *enterprise bean* is a Java component that can be combined with other resources to create distributed client/server applications. Instantiated enterprise beans reside in *enterprise bean containers*, or *EJB containers*. An EJB container provides an interface between the enterprise beans and the application server on which the enterprise beans reside. An enterprise bean is typically accessed using Java RMI-IIOP. An ORB manages the interaction between clients and enterprise beans, using IIOP. ORBs enable clients to make requests and receive responses from servers in distributed computing environments. Alternatively, enterprise beans are accessible through JMS. It is also possible to invoke an enterprise bean as a Web service via SOAP, as explained in Chapter 14 on page 497.

There are three types of enterprise beans: entity beans, session beans, and message-driven beans. *Entity beans* store persistent data and typically use database connections. Entity beans are of two types: CMP entity beans and BMP entity beans.

- Entity beans with *container-managed persistence* (CMP) let the EJB container transparently and implicitly manage the persistent state. The enterprise bean developer does not need to code any database access functions within the enterprise bean class methods.
- Entity beans with *bean-managed persistence* (BMP) manage persistent data in a manner defined by the application developer in the bean code. This usually includes writing to databases.

Session beans do not require database access, although they can obtain it indirectly, as needed, by accessing entity beans. Session beans can also obtain direct access to databases and other resources through the use of *resource references*, which include the use of JDBC. Session beans can be either stateless or stateful.

- A session bean is said to be *stateless* if it provides a stateless service to the client. A business method on a stateless session bean is similar to a procedural application or static method; there is no instance state. Therefore, all the data needed to execute a stateless session bean's method is provided by the method arguments.
- A session bean is said to be *stateful* if it acts as a server-side extension of the client that uses it. A stateful session bean is created by a client and will work for only that client until the client connection is dropped or the bean is explicitly removed. Unlike a stateless session bean, a stateful session bean has state or instance fields that can be initialized and changed by the client with each method invocation.

Message-driven beans are enterprise beans accessible asynchronously via JMS rather than synchronously through such protocols as RMI-IIOP. The EJB V2.1 specification expands the scope of message-driven beans beyond JMS to support any messaging system.

An *EJB module* is one or more enterprise beans assembled into a single deployable unit. As we have observed, an EJB module is stored in a standard JAR file, commonly referred to as *ejb-jar*. This file contains

- One or more deployable enterprise beans
- A deployment descriptor, stored in an XML file

Specifically, an EJB module's deployment descriptor file declares the contents of the module, specifies the structure and external dependencies of the enterprise beans in the module, explains how the enterprise beans are to be used at runtime, and defines the security policies applicable to the enterprise beans within the module. The format of the security policy is defined by the EJB specification (see Chapter 5 on page 157).

3.2.2 Web Modules

A *Web module* represents a *Web application*—an application that can be accessed over the Web using HTTP. A Web module is used to assemble servlets and JSP files, as well as static content, such as HTML pages, into a single deployable unit.

As we said earlier, Web modules are stored in WAR files, which are enhanced JAR files with a `.war` file extension, and contain

- One or more servlets, JSP files, and other supporting files
- A deployment descriptor, stored in an XML file

The deployment descriptor file, `web.xml`, declares the contents of the Web module. This file contains information about the structure and external dependencies of the components in the Web module and describes the components' runtime use. In addition, the deployment description file is used to declare the security policies applicable to the universal resource identifiers (URIs) that are mapped to the resources within the Web module. These security policies include both the authorization policy and the login configuration information. The format of the security policy is defined by the Java Servlet specification.

Servlets are Java programs running on a WAS and extend the Web server's capabilities. For example, servlets support generation of dynamic Web page content, provide database access, concurrently serve multiple clients, and filter data by MIME type. Servlets use the Java Servlet API. By analogy, servlets are the server-side equivalent of client-side browser applets.

JSP files enable the separation of the HTML coding from the business logic in Web pages, allowing HTML programmers and Java programmers to more easily collaborate in creating and maintaining pages. This process is described in greater detail in Section 4.1.2 on page 104.

3.2.3 Application Client Modules

Application clients are first-tier Java-based client programs. Even though it is a regular Java application, an application client depends on an *application client container* to provide system services. An *application client module* packages application client code in a JAR file. This JAR file includes a deployment descriptor XML file, which specifies the enterprise beans and external resources referenced by the application.

The security configuration of an application client determines how the application will access enterprise beans and Web resources. If the J2EE components that the client application accesses are secured, the client will be authenticated accordingly. In order for an application client to retrieve authentication data from an end user, configuration information must be specified in a deployment descriptor XML file, `application-client.xml`, associated with the client application. Application clients typically run in an environment that has a Java 2 security manager installed and the security policies enforced based on the J2SE security policy framework (see Chapter 8 on page 253).

3.3 Secure Interoperability between ORBs

J2EE applications are required to use RMI-IIOP when accessing EJB components. This allows enterprise beans to be portable between container implementations.

A *J2EE container* provides the runtime support for the J2EE components. A J2EE container vendor enables access to the enterprise beans via IIOP. This facilitates interoperability between containers by using the Common Secure Interoperability (CSI) protocol. Security is enabled and enforced by the ORBs, ensuring authenticity, confidentiality, and integrity. Version 2 of this protocol specification (CSIv2) is the accepted industry standard and is mandated by the J2EE specification.

3.4 Connectors

A *resource adapter* is defined in the J2EE Connector Architecture specification as a system-level software driver that a Java application uses to connect to an EIS. The resource adapter plugs into an application server and provides connectivity between the EIS, the J2EE application server, and the enterprise application.

The Java Connector Architecture (JCA) specification allows resource adapters that support access to non-J2EE systems to be plugged into any J2EE environment. Resource adapter components implementing the JCA API are called *connectors*.

The JCA specification describes standard ways to extend J2EE services with connectors to other non-J2EE application systems, such as mainframe systems and ERP systems. The JCA architecture enables an EIS vendor to provide a standard resource adapter for a J2EE application to connect to the EIS. A resource adapter is used by a Java application to connect to an EIS. For example, Web enablement of business applications, such as IBM's Customer Information Control System (CICS), would imply that the J2EE-based presentation layer would connect to a CICS application using a CICS connector. With this approach, protocol details of connecting to a CICS system are transparent to the Web application and are handled by the CICS connector implementation.

JCA defines a standard set of *system-level contracts* between a J2EE server and a resource adapter. In particular, these standard contracts include a *security contract* and enable secure access to non-J2EE EISs. The security contract helps to reduce security threats to the information system and protects valuable information resources managed by such a system. Given that most of these EIS systems have facilities to accept some form of authentication data representing an identity connecting to the system, the JCA security contract deals with the authentication aspects of the EIS. Essentially, it is about a J2EE application *signing on* to an EIS system. This means that the J2EE application accesses a connection to the EIS system by providing authentication information. As discussed in Section 3.9.4 on

page 87 and Section 3.10.3 on page 94, two organizational roles are involved in addressing this issue: the Application Component Provider and the Deployer. Specifically, the Application Component Provider can use either of two choices related to EIS sign-on: the declarative approach or the programmatic approach.

- The *declarative approach* allows the Deployer to set up the resource principal and EIS sign-on information. For example, the Deployer sets the user ID and password—or another set of credentials—necessary to establish a connection to an EIS instance.
- With the *programmatic approach*, the Application Component Provider can choose to perform sign-on to an EIS from the component code by providing explicit security information. For example, the user ID and password—or another set of credentials—necessary to establish a connection to an EIS instance are coded into the application code.

The Application Component Provider uses a deployment descriptor element, such as `res-auth` for EJB components, to indicate the requirement for one of the two approaches. If the `res-auth` element is set to `Container`, the application server sets up and manages EIS sign-on. If the `res-auth` element is set to `Application`, the component code performs a programmatic sign-on to the EIS.

Further details of the security aspects of a JCA-based connection to an EIS from a J2EE application are discussed in Section 3.9.4 on page 87 and Section 3.10.3 on page 94.

3.5 JMS

JMS is a standard Java messaging API that provides a common mechanism for Java-language programs to access messaging systems. Java clients and middle-tier components must be capable of using messaging systems to access the J2EE components that are enabled via a messaging layer.

- Application clients, EJB components, and Web components can send or synchronously receive a JMS message. Application clients can also receive JMS messages asynchronously.
- A new kind of enterprise bean introduced in EJB V2.0, the message-driven bean, enables the asynchronous consumption of messages. A message-driven bean can be accessed by sending a method invocation request over a messaging infrastructure. A JMS provider may optionally implement concurrent processing of messages by message-driven beans.

The J2EE specification requires JMS providers to implement both the reliable point-to-point messaging model and the publish/subscribe model. The *reliable*

point-to-point messaging model allows one entity to send messages directly to another entity that understands the format of the messages and the requests. The *publish/subscribe model* is event driven; a message is published, and the message is delivered to all subscribers of the event. One example is a StockQuote application, with multiple traders wanting to get the latest stock quote. In this scenario, the traders' applications subscribe to the stock-quote messaging service. When the stock values are published, the information is made available to all the subscribers. In a Java environment, both the StockQuote server and the traders' applications can use the JMS mechanism with a JMS provider to achieve the required messaging functionality.

However, JMS does not specify a security contract or an API for controlling message confidentiality and integrity. Security is considered to be a JMS-provider-specific feature. It is controlled by a System Administrator rather than implemented programmatically or by the J2EE server runtime.

3.6 Simple E-Business Request Flow

It would be helpful to understand a simple e-business request flow in an enterprise Java environment. Figure 3.3 presents a simple request flow that does not involve security.

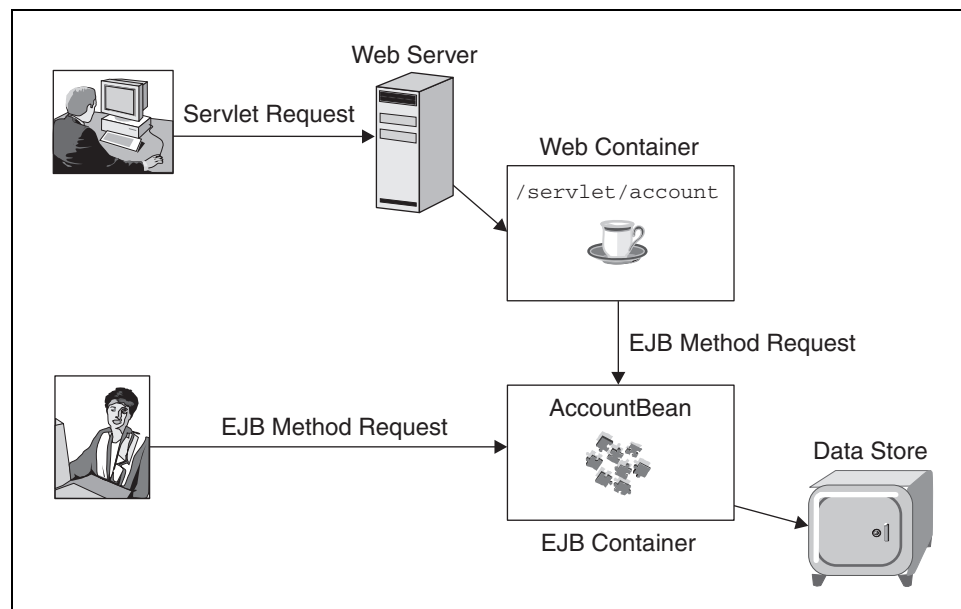


Figure 3.3. Simple E-Business Request Flow

Let us consider two types of clients: *HTTP clients*, such as Web browsers, and *IIOp clients*, regular applications capable of using IIOp to send requests and receive responses over the network. An HTTP client invokes a URL: for instance, `/servlet/account/`. The request from the user's browser gets handled by the Web server, which routes the request to a Web, or servlet, container serving the URL resource. The logic behind the URL is implemented as a Java servlet. This servlet, packaged in a Web module, is hosted in a J2EE Web container, which in turn invokes an enterprise bean, `AccountBean`, via IIOp. `AccountBean` is an entity bean, packaged in an EJB module, with its business data stored in a data store. The same enterprise bean is accessed directly from an IIOp client, packaged in an application client module. In this case, the request is not routed by the servlet but is directly accessed as a remote object from the Java client.

The request flow just described does not involve security considerations. The next sections in this chapter provide an overview of the J2EE specification as it pertains the security of an enterprise. The platform roles reflect the organizational responsibilities, from application development, application assembly, and application deployment, to administration.

3.7 J2EE Platform Roles

J2EE defines roles that reflect the responsibilities within an organization. Any person or software involved in the process of making an application available within an enterprise can usually be categorized into organization roles, called *J2EE platform roles*. The J2EE platform roles having security responsibilities are the Application Component Provider, Application Assembler, Deployer, System Administrator, J2EE Product Provider, and Tool Provider. The J2EE security model is defined with respect to these J2EE roles.

Figure 3.4 shows the interactions among the Application Component Provider, Application Assembler, Deployer, and System Administrator. These are the roles involved, from a security perspective, in the stages between development and deployment.

Figure 3.4 depicts the software process cycle from the perspective of J2EE platform roles.¹ In a typical J2EE software process cycle, application component developers build enterprise application components, such as servlets or enterprise

1. The J2EE V1.3 platform role called *Tool Provider* is responsible for supplying tools used for the development and packaging of application components. Because it depicts the software life cycle in development and deployment of a J2EE application and the users involved in the process, Figure 3.4 does not include the roles of Tool Provider and J2EE Product Provider.

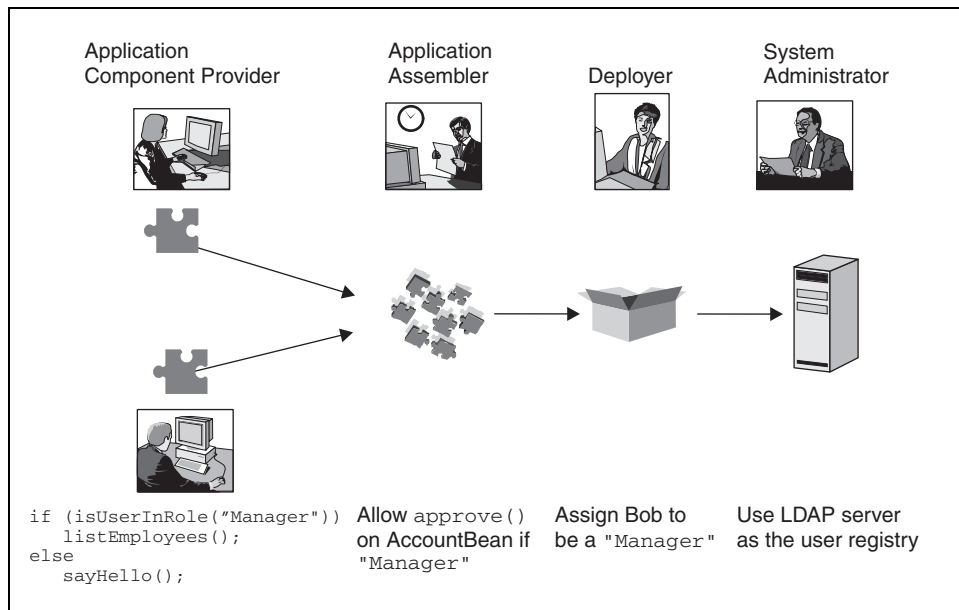


Figure 3.4. J2EE Platform Roles with Security Responsibilities in the Development and Deployment of a J2EE Application

beans. The greatest opportunity for component reuse and flexibility in reconfiguring security policy is when the components are written to be *security unaware*, meaning that they do not contain embedded security policy code. Conversely, components containing embedded security policy code are said to be *security aware*. Security-aware components are difficult to reuse, and flexibility is limited because it often requires changing the source code to reflect various security policies. For some applications, this may be unavoidable.

An Application Assembler integrates a set of components supplied by one or more Application Component Providers. The Application Assembler has the in-depth knowledge of the application. The Application Assembler specifies security policies as hints to the Deployer. For example, the Application Assembler can provide hints such that the `approve()` method of an enterprise bean should be accessed only by those principals granted the role of a Manager.

A Deployer deploys enterprise applications, assembled by Application Assemblers, into an operational environment. When tailoring the security policies to the operational environment, the Deployer consults the security policy hints provided by the Application Assembler. For example, the Deployer can assign the role of a Manager to a user named Bob.

A System Administrator is responsible for administering the system, including security. This may include configuring the J2EE product to use an LDAP server for managing security information, including user and group membership.

The following subsections provide a more detailed description of the four J2EE roles. The J2EE platform roles that we have listed are considered to be the typical roles, although in practice, the roles are adapted to better match the organization's application development and deployment work flow. The rest of this section summarizes the major responsibilities of the individual J2EE platform roles with respect to security management.

3.7.1 Application Component Provider

The *Application Component Provider* is the J2EE platform role responsible for implementing the business logic as a set of J2EE application components—enterprise beans, servlets, and/or JSP files. These components are packaged in an *ejb-jar* file containing one or more enterprise beans and/or a *aWAR* file containing one or more servlets and/or JSP files, and/or a *JAR* file containing an application client module.

The Application Component Provider has numerous responsibilities in code development. These responsibilities range from resource access to programmatic access of the caller's security context. Following are the Application Component Provider's key security responsibilities:

3.7.1.1 Access of Resources in the Underlying Operating System

The J2EE architecture does not define the operating system principal—for example, the operating system user—under which EJB methods or servlets execute. Therefore, the Application Component Provider cannot rely on a specific principal for accessing the underlying operating system resources. The Application Component Provider should design the applications so that special privileges are not required to access system resources.

3.7.1.2 Security Recommendations

The Application Component Provider should avoid implementing security mechanisms or hard-coded security policies in the component but instead should rely on the security mechanisms provided by the J2EE container. The Application Component Provider should let the Application Assembler and the Deployer define the appropriate security policies for the application. The Application Component Provider can use the deployment descriptors to convey security-related information to the Application Assembler.

3.7.1.3 Programmatic Access to the Caller's Security Context

Programmatic-security APIs should be avoided when possible. However, they should be used when the J2EE component methods need access to security-context information because the J2EE declarative security model is insufficient to implement application security requirements.

3.7.1.4 Conveying the Use of Role References

A *security role* is a set of J2EE authorizations. The Application Component Provider may build a security-aware application and use *role references*—security role names within the application components. For example, a component may call `isUserInRole("Manager")` on a `javax.servlet.http.HttpServletRequest` object. When security role names are hard-coded in an application component, the Application Component Provider must identify these role names for the Application Assembler so that it can map component-defined security role references in each of the components in the deployment to a single application-level security role name. For example, two components in an application may use the security role references `Manager` and `Boss` within the component, whereas both of these roles imply the application security role of a Supervisor.

3.7.2 Application Assembler

The *Application Assembler* is the J2EE platform role responsible for combining J2EE components into deployable application units. The Application Assembler also simplifies the Deployer's job by providing a security view of the enterprise beans, servlets, and JSP files in the relevant deployment descriptors. A *security view* consists of a set of J2EE security roles. A security role is a semantic grouping of J2EE authorizations, or permissions—implemented as `java.security.Permission` objects—that a given type of application users must have in order to successfully use the application. The Application Assembler defines one or more security roles in the deployment descriptor and specifies and associates J2EE permissions with these roles. For example, the security role `Manager` could be granted the J2EE permissions to invoke an enterprise bean to grant loans and view the loan status of all the customers using Web applications. In contrast, the security role `HelpDesk` could be granted only a subset of these J2EE permissions—for example, only the J2EE permission to view the loan status of the customers—by having been granted access to the relevant URIs.

Following are some of the Application Assembler's security responsibilities.

3.7.2.1 Defining EJB Method Permissions

The home, local home, remote, and local interfaces of an enterprise bean are defined as part of the EJB specification.

- The *home interface* of an enterprise bean is a Java interface used to create, find, or delete an instance of the enterprise bean. The methods defined in the remote interface can be accessed from within the same container or remotely via RMI-IIOP.
- The *local home interface* of an enterprise bean is functionally similar to the home interface, but the methods defined in the local home interface are accessible only from within the same container.
- The *remote interface* of an enterprise bean is a Java interface that defines the operations that can be performed on the enterprise bean to access the business logic associated with the enterprise bean itself. The methods defined in the remote interface can be accessed from within the same container or remotely via RMI-IIOP.
- The *local interface* of an enterprise bean is functionally similar to the remote interface, but the methods defined in the local interface are accessible only from within the same container.

The home, local home, remote, and local interfaces of an enterprise bean define which methods the enterprise bean exposes to a client. An EJB method permission is defined by an XML `method-permission` element in an EJB module's deployment descriptor and is used to assign groups of methods of the home, local home, remote, and local interfaces of an enterprise bean packaged in that EJB module to the security roles. This way, the Application Assembler can define the security view of the enterprise bean.

An EJB method permission specifies the methods of the home, local home, remote, and local interfaces that each of the listed security roles is allowed to invoke. This implies that an EJB method permission may include a list of one or more security roles and a list of one or more methods. In addition, a security role or a method may appear in multiple XML `method-permission` elements. Users of particular security roles are granted access to all the methods listed in all the EJB method permission elements where those security roles appear. EJB method permissions and the deployment descriptor are discussed further in Chapter 5 on page 157.

3.7.2.2 Defining Web Resources Security Constraints

An Application Assembler uses a Web module's deployment descriptor to define security constraints for a Web application packaged in that module. The Web

module's deployment descriptor's `auth-constraint` element is used for this purpose. This element consists of one or more security roles and a list of URL patterns that users with any of those security roles are authorized to invoke. Specifically, deployment descriptors are used to assign groups of URL patterns to the security roles, thus defining security views of Web applications. Login configuration information, such as requiring a user to be authenticated using a form-based login mechanism, and transport guarantee constraints, such as requiring access to a URL pattern to be submitted only using an HTTPS connection, can also be specified in the deployment descriptor.

3.7.2.3 Declaring Security Roles within a J2EE Application

An Application Assembler defines each security role by using the `security-role` XML element in the relevant deployment descriptor.

- If the deployment descriptor belongs to an `ejb-jar` file, the `security-role` element is scoped to that `ejb-jar` file and applies to all the enterprise beans in that EJB module.
- If the deployment descriptor belongs to a `WAR` file, the security role element is scoped to that `WAR` file and applies to all the servlets and/or JSP files in that Web module.
- If the deployment descriptor belongs to an `EAR` file, the `security-role` element applies to all the `JAR` and `WAR` files that are packaged within that `EAR` file. Effectively, the set of security roles declared in the `EAR` file's deployment descriptor is the union of the security roles defined in the deployment descriptors of the `JAR` and `WAR` files packaged within that `EAR` file. Technically, however, the security roles described in the constituent modules of an `EAR` file are the ones that are used to enforce authorization, because those roles are associated with authorization policies. The roles declared in the `EAR` file's deployment descriptor are typically used for administration and management purposes only. For example, they can be used to assign security roles to principals with regard to the whole application packaged in the `EAR` file.

The deployment descriptor of an application client module does not contain security role elements, because security roles are specific to the server side of a J2EE application, not for the client.

Within each `security-role` element, the Application Assembler will use the `role-name` subelement to define the name of the security role and, optionally, will use the `description` subelement to provide a description of the security role.

3.7.3 Deployer

For each J2EE application, the Deployer takes the modules comprising that application and deploys the module components into a specific operational, or runtime, environment. The modules were produced by an Application Assembler. The operational environment in which the application is deployed includes a specific J2EE container. The Deployer is also responsible for ensuring the security of an assembled application when it is deployed in the target operational environment.

The Deployer has the following responsibilities with respect to security management.

3.7.3.1 Reading the Security View of the J2EE Application

The Deployer uses the deployment tools supplied by the J2EE Product Provider to read the security view of the application. The Deployer should treat the security policies specified in a deployment descriptor as hints and modify those policies as appropriate to the operational environment in which the application is being deployed.

3.7.3.2 Configuring the Security Domain

A *security domain* within an enterprise represents an instance of an authentication authority and relevant security infrastructure. For example, a security domain may point to a particular Kerberos domain for authentication and an LDAP user repository to deduce user and group membership to be used for authorization. In the case of multiple security domains within the enterprise, the Deployer is responsible for configuring the J2EE product to use the appropriate security domains.

3.7.3.3 Assigning of Principals to Security Roles

The Deployer is responsible for assigning principals and/or groups of principals used for managing security in the runtime to the security roles defined in the XML `security-role` elements of the deployment descriptors. The process of assigning the logical security roles defined in the J2EE application's deployment descriptor to the operational environment's security concepts is specific to the configuration capabilities of a particular J2EE product.

3.7.3.4 Configuring Principal Delegation

Delegation allows an intermediary to perform a task, initiated by a client, under an identity based on a *delegation policy*. The Deployer is responsible for configuring principal delegation for intercomponent calls by using the appropriate deployment descriptor elements, as follows.

- If the deployment descriptor belongs to an `ejb-jar` file, the Deployer uses the `security-identity` deployment descriptor element for this purpose.

When the value of the `security-identity` element is `use-caller-identity`, the identity of the caller of the enterprise bean will be used when calling other components from the enterprise bean. When the value specified is `run-as`, the identity of the caller to the enterprise bean will be propagated in terms of the security role name defined in the `run-as` element of the descriptor. For example, if the caller to an enterprise bean is user Bob, Bob's identity will be used if the `security-identity` element is set to `use-caller-identity`. If the `security-identity` element is set to `run-as` and the role name is `Teller`, the downstream calls from the enterprise bean will be performed in terms of the `Teller` role.

- If the deployment descriptor belongs to a WAR file, the Deployer uses the `run-as` deployment descriptor element for this purpose. If the `run-as` element is not declared in the WAR file, the identity of the servlet's caller will be used for components called from the servlet. If the `run-as` element is declared in the deployment descriptor, the identity passed when the servlet makes calls to other components will be that of the security role name defined in the `run-as` element of the descriptor. For example, if the caller to a servlet is user Bob, Bob's identity will be used if no `run-as` element is declared in the WAR file's deployment descriptor. If the `run-as` element is declared in the deployment descriptor and the role name is `Teller`, the downstream calls from the servlet will be performed in terms of the `Teller` role.

3.7.4 System Administrator

The System Administrator is responsible for the configuration and administration of the enterprise's computing and networking infrastructure, including the J2EE container. The System Administrator is also responsible for the overall management and operational aspects of the J2EE applications at runtime. The following list describes the security-related responsibilities of the System Administrator. Some of these responsibilities may be carried out by the Deployer or may require the cooperation of both the Deployer and the System Administrator.

3.7.4.1 Administering the Security Domain

The System Administrator is responsible for administering the security domain. This includes the principal administration, user account management, group membership assignment, deployment of J2EE products within an enterprise environment, including configuration of DMZs, firewalls, user registries, and so on. These are typically performed using the tools provided by the relevant product vendor; for example, user registry management is performed using an LDAP server product, firewall configuration using the firewall product, and so on.

3.7.4.2 *Assigning Application Roles to Users and Groups*

The System Administrator is responsible for assigning principals and/or groups of principals used for managing security in the runtime to the security roles defined in the XML `security-role` elements of the deployment descriptors. The process of assigning the logical security roles defined in the J2EE application's deployment descriptor to the operational environment's security concepts is specific to the configuration capabilities of a particular J2EE product. For example, using the tools provided by a J2EE product, the System Administrator can assign a Teller role scoped to a J2EE application FinanceApp to user Bob and group TellerGroup.

3.7.5 J2EE Product Provider

The J2EE Product Provider has the following areas of responsibility.

3.7.5.1 *Supplying Deployment Tools*

The J2EE Product Provider is responsible for supplying the deployment tools that the Deployer uses to perform all the deployment tasks, including the security-related tasks. For example, the J2EE Product Provider will supply tools to perform security-role-to-principal and/or -user and/or -group assignment.

3.7.5.2 *Configuring Security Domains*

The J2EE Product Provider is responsible for configuring the J2EE product to use appropriate security domains. For example, the J2EE Product Provider needs to supply facilities to configure the J2EE product to use a particular authentication mechanism—for example, to use a Kerberos domain.

3.7.5.3 *Supplying Mechanisms to Enforce Security Policies*

The J2EE Product Provider is responsible for supplying the security mechanisms necessary to enforce the security policies set by the Deployer. This includes authentication of principals, authorization to perform EJB/servlet calls, configuration of resource adapters defined in the JCA, and secure communication with remote clients, integrity, and confidentiality.

3.7.5.4 *Providing Tools for Principal Delegation*

The J2EE Product Provider is responsible for passing principals on EJB/servlet calls. In particular, the J2EE Product Provider is responsible for providing the deployment tools that allow the Deployer to configure principal delegation for calls from one J2EE component to another.

3.7.5.5 *Providing Access to the Caller's Security Context*

The J2EE Product Provider is responsible for providing access to the caller's security context information when programmatically queried from enterprise beans

and servlets using the J2EE-defined security APIs. For example, when a servlet calls `getUserPrincipal()` on a `javax.servlet.http.HttpServletRequest` object, the J2EE Product Provider must return the `java.security.Principal` object representing the caller of the servlet.

3.7.5.6 *Supplying Runtime Security Enforcement*

One of the most significant responsibilities of the J2EE Product Provider is to supply *runtime security enforcement*, as follows.

- Provide enforcement of the client access control as specified by the current security policy.
- Isolate an enterprise bean instance from other instances and other application components running on the server, thus preventing unauthorized access to privileged information.
- Provide runtime facilities to implement the principal-delegation policies set in the deployment descriptor.
- Allow a J2EE application to be deployed independently multiple times, each time with a different security policy.

3.7.5.7 *Providing a Security Audit Trail*

Optionally, the J2EE Product Provider may provide a *security audit trail* mechanism whereby secure access to enterprise beans and Web resources is logged. Such audit logs can be used to determine the information about the activity on the J2EE components. For example, these logs can be used to discover unauthorized attempts to access enterprise beans and Web resources.

3.8 J2EE Security Roles

The J2EE authorization model is based on the concept of security roles. Security roles are different from J2EE platform roles. As noted in Section 3.7.2 on page 67, a security role is a semantic grouping of permissions that a given type of application users must be granted to be authorized to use the application. In contrast, a J2EE platform role represents the organizational responsibility in making a J2EE application available to an enterprise, as described in Section 3.7 on page 64. Both declarative and programmatic security are based on the security roles.

Security roles are defined by the Application Component Provider and the Application Assembler. The Deployer then maps each security role to one or more security identities, such as users and groups, in the runtime environment. Listing 3.1 is an example of an XML description of two security roles declared within the deployment descriptor of an application's EAR file.

Listing 3.1. Description of Security Roles Teller and Supervisor

```
<assembly-descriptor>
  <security-role>
    <description>
      This role is intended for employees who provide
      services to customers (tellers).
    </description>
    <role-name>Teller</role-name>
  </security-role>
  <security-role>
    <description>
      This role is intended for supervisors.
    </description>
    <role-name>Supervisor</role-name>
  </security-role>
</assembly-descriptor>
```

Declarative authorization can be used to control access to an enterprise bean method. This contract is specified in the deployment descriptor. Enterprise bean methods can be associated with `method-permission` elements in EJB modules' deployment descriptors. As described in Section 3.7.2 on page 67, a `method-permission` element specifies one or more EJB methods that are authorized for access by one or more security roles. In Section 3.7.3 on page 70, we observed that the mapping of principals to security roles is performed by the Deployer. If the calling principal is in one of the security roles authorized access to a method, the caller is allowed to execute the method. Conversely, if the calling principal is not a member of any of the roles, the caller is not authorized to execute the method. Listing 3.2 is an example of an XML `method-permission` element in an EJB module's deployment descriptor.

Listing 3.2. Example of an XML `method-permission` Element in an EJB Module's Deployment Descriptor

```
<method-permission>
  <role-name>Teller</role-name>
  <method>
    <ejb-name>AccountBean</ejb-name>
    <method-name>getBalance</method-name>
  </method>
  <method>
    <ejb-name>AccountBean</ejb-name>
    <method-name>getDetails</method-name>
  </method>
</method-permission>
```

Access to Web resources can be similarly protected. An action on a Web resource URI can be associated with an XML `security-constraint` element in a

Web module's deployment descriptor. The `security-constraint` element contains one or more URI patterns that can be authorized for access by one or more security roles. If the calling principal is a member of one or more of the security roles authorized to access an HTTP method on a URI, the principal is authorized to access the URI. Conversely, if the calling principal is not a member of any of the roles, the caller is not allowed to access the URI. Listing 3.3 shows a Web module's deployment descriptor fragment that defines access authorization requirements for a Web application.

Listing 3.3. Example of an XML `security-constraint` Element in a Web Module's Deployment Descriptor

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>
      Account servlet protected area
    </web-resource-name>
    <url-pattern>/finance/account/</url-pattern>
    <http-method>GET</http-method>
  </web-resource-collection>
  <auth-constraint>
    <description>Teller can access the URIs</description>
    <role-name>Teller</role-name>
  </auth-constraint>
</security-constraint>
```

In a J2EE environment, the security roles form the basis for the security provided by the containers that host the components. A container can provide two types of security: declarative and programmatic.

In the *declarative security model*, an application expresses its security policies through security constraints in a form external to the application. In J2EE, security constraints are specified in the deployment descriptors. This allows the application to be *security-mechanism agnostic*, or *security unaware*. Application code is not required to enable security or enforce the application security policy.

With declarative security, the application's logical security requirements are defined in the deployment descriptors and then mapped by the Deployer and the System Administrator to a deployment environment. The Deployer uses container-deployment tools to process the deployment descriptors. At runtime, the container uses the security policy configured by the Deployer and the System Administrator to enforce authorization.

Declarative security allows greater opportunities for application portability because all the security issues related to the underlying container and operating system are defined in configuration files external to the application. In addition, an application that makes use of declarative security is easier to develop because security and policy configuration issues are managed outside the application.

Application developers need not to be security experts, and an application based on declarative security can be more easily extended. Therefore, declarative security should always be used instead of programmatic security unless declarative security alone is insufficient to describe the security requirements of an application.

In the *programmatic security model*, the application programmer is responsible for explicitly writing the code that defines and enables security. The application security policy is an integral part of the application. An application conforming to this model is said to be *security aware*.

Programmatic security makes application development more difficult and severely limits the portability and extensibility of an application, because security issues related to the specific application, container, and operating system on which the application is running must be hard-coded. For these reasons, programmatic security should be used only when declarative security alone is insufficient to express the security model of an application. For example, the declarative security capabilities of J2EE V1.3 do not allow expressing a policy whereby a user cannot withdraw more than \$1,000 from an automatic teller machine (ATM). Similarly, instance-level authorization to impose that only Bob can access Bob's account cannot be defined declaratively. In these cases, the application needs to enforce these rules programmatically.

3.9 Declarative Security Policies

Security policies associated with URIs and enterprise beans include the following:

- Login configurations associated with URIs: for example, use of form-based login
- Authorization policies associated with URIs and enterprise beans based on J2EE security roles
- Principal-delegation policies that apply to Web applications and enterprise beans
- Connection policies associated with JCA connectors that dictate how applications access EIS in a secure manner

Such authorization and delegation policies can be specified declaratively within the relevant deployment descriptors.

3.9.1 Login-Configuration Policy

Authentication is the process of proving the identity of an entity. Authentication generally is performed in two steps: (1) acquiring the authentication data of a principal and (2) verifying the authentication data against a user (principal) registry.

J2EE security authenticates a principal on the basis of the authentication policy associated with the resource the principal has requested. When a user requests a protected resource from a Web application server, the server authenticates the user. J2EE servers use authentication mechanisms based on validating credentials, such as digital certificates (see Section 10.3.4 on page 372), and user ID and password pairs. Credentials are verified against a user registry that supports the requested authentication scheme. For example, authentication based on user ID and password can be performed against an LDAP user registry, where authentication is performed using an LDAP bind request.

A Web server is responsible for servicing HTTP requests. In a typical J2EE environment, a Web server is a component of a J2EE WAS. In this case, the WAS hosts servlets, JSP files, and enterprise beans. The login—authentication—configuration is managed by the WAS, which drives the authentication challenges and performs the authentication. Similarly, if the Web server is independent of the WAS and the Web server is the front end for the WAS, the Web server acts as a proxy for J2EE requests. Again, the authentication is typically performed by the WAS.

The authentication policy for performing authentication among a user, a Web server, and a WAS can be specified in terms of the J2EE login configuration elements of a Web application's deployment descriptor. The authentication policy can specify the requirement for a secure channel and the authentication method. The requirement to use a secure channel when accessing a URI is specified through the `user-data-constraint` descriptor.

The authentication method is specified through the `auth-method` element in the `login-config` descriptor. There are three types of authentication methods:

- 1. HTTP authentication method.** The credentials that the client must submit to authenticate are user ID and password, sent to the server as part of an HTTP header and typically retrieved through a browser's dialog window. The two modes of HTTP authentication are *basic* and *digest*. In both cases, the user ID is sent as cleartext.² In basic authentication, the password is transmitted in cleartext as well; in digest authentication, only a hash value of the password is transmitted to the server (see Section 10.2.2.4 on page 356).

2. More precisely, the cleartext is encoded in *base64 format*, a commonly used Internet standard. Binary data can be encoded in base64 format by rearranging the bits of the data stream in such a way that only the six least significant bits are used in every byte. Encoding a string in base64 format does not add security; the algorithm to encode and decode is fairly simple, and tools to perform encoding and decoding are publicly available on the Internet. Therefore, a string encoded in base64 format is still considered to be in cleartext.

2. **Form-based authentication method.** The credentials that the client must submit to authenticate are user ID and password, which are retrieved through an HTML form.
3. **Certificate-based authentication method.** The credential that the client must submit is the client's digital certificate, transmitted over an HTTPS connection.

3.9.1.1 Authentication Method in Login Configuration

The `auth-method` element in the `login-config` element specifies how a server challenges and retrieves authentication data from a user. As noted previously, there are three possible authentication methods: HTTP (user ID and password), form based (user ID and password), and certificate based (X.509 certificate).

With the *HTTP authentication method*, the credentials provided by the user consist of a user ID and password pair, transmitted as part of an HTTP header. When HTTP authentication is specified, a user at a Web client machine is challenged for a user ID and password pair. The challenge usually occurs in the following way:

1. A WAS issues an HTTP unauthorized client error code (401) and a `WWW_Authenticate` HTTP header.
2. The Web browser pops up a dialog window.
3. The user enters a user ID and password pair in this dialog window.
4. The information is sent to the Web server.
5. The WAS extracts the information and authenticates the user, using the authentication mechanism with which it has been configured.

With HTTP authentication, a realm name also needs to be specified. *Realms* are used to determine the scope of security data and to provide a mechanism for protecting Web application resources. For example, a user defined as `bob` in one realm is treated as different from `bob` in a second realm, even if these two IDs represent the same human user, Bob Smith.

Once specified, the realm name is used in the HTTP 401 challenge to help the Web server inform the end user of the name of the application domain. For example, if the realm is `SampleAppRealm`, the dialog window prompting the user for a user ID and password pair during authentication will include that the user ID and password are to be supplied for the `SampleAppRealm` realm.

HTTP authentication can be either basic or digest. In *basic authentication*, the credentials requested of the user are user ID and password, and both are transmitted as cleartext. In order for the authentication method to be basic, the `auth-method` element in the `login-config` descriptor must be set to `BASIC`. Listing 3.4

is a deployment descriptor fragment showing an example of login configuration requiring basic authentication.

Listing 3.4. Login Configuration for Basic Authentication

```
<login-config>
  <auth-method>BASIC</auth-method>
  <realm-name>SampleAppRealm</realm-name>
</login-config>
```

This scheme is not considered to be a secure method of user authentication, unless used in conjunction with some external secure systems, such as SSL.

In *digest authentication*, the user ID and a hash value of the password are transmitted to the server as part of an HTTP header. Therefore, the password does not appear in cleartext, which is the biggest weakness of basic authentication.

When digest authentication is specified, the Web server responds to the client's request by requiring digest authentication. A one-way hash of the password (see Section 10.2.2.4 on page 356), as specified by the Request for Comments (RFC) 2617,³ is computed by the client, based on a random number, called *nonce*, uniquely generated by the server each time a 401 response is made. The hash value of the password is sent to the server, which computes the digest of the password for the user ID and compares the resulting hash value with the one submitted by the client. The requesting user is considered to be authenticated if the hash values are identical.

This mode of authentication assumes that the server has access to the user's password in cleartext—a necessary requirement in order for the server to compute the hash of the password. However, this is rarely the case in most enterprise environments, as the password in cleartext is not retrievable from a user repository containing the user ID and password information. Rather, the server typically delegates responsibility to the user repository to validate a user's password. Therefore, digest authentication is not widely adopted in enterprise environments and hence is not required to be supported by a J2EE container.

J2EE servers that do support digest authentication can be configured to issue a digest authentication challenge by setting the value of the `auth-method` element in the `login-config` descriptor to `DIGEST`. Listing 3.5 is a deployment descriptor fragment illustrating how a J2EE server can be configured to require digest authentication.

3. See <http://www.ietf.org/rfc/rfc2617.txt>.

Listing 3.5. Login Configuration for Digest Authentication

```
<login-config>
  <auth-method>DIGEST</auth-method>
  <realm-name>SampleAppRealm</realm-name>
</login-config>
```

The second authentication method is form based. With this method, the `auth-method` element in the `login-config` element must be set to `FORM`. The form-based authentication method assumes that the server is configured to send the client an HTML form to retrieve the user ID and password from the Web user, as opposed to sending a 401 HTTP unauthorized client error code as in the basic challenge type.

The configuration information for a form-based authentication method is specified through the `form-login-config` element in the `login-config` element. This element contains two subelements: `form-login-page` and `form-error-page`.

- The Web address to which a user requesting the resource is redirected is specified by the `form-login-page` subelement in the Web module's deployment descriptor. When the form-based authentication mode is specified, the user will be redirected to the specified `form-login-page` URL. An HTML form on this page will request a user ID and password.
- If the authentication fails, the user is redirected to the page specified by the `form-error-page` subelement.

Listing 3.6 is a sample HTML page for the login form.

Listing 3.6. Login Page Contents

```
<HTML>
  <HEAD>
    <TITLE>Sample Login page.</TITLE>
  </HEAD>
  <BODY>
    <TR><TD>
      <HR><B>Please log in!</B><BR><BR>
    </TD></TR>
    <CENTER>
      Please enter the following information:<BR>
      <FORM METHOD=POST ACTION="j_security_check">
        Account <INPUT TYPE=text NAME="j_username"
          SIZE=20><BR>
        Password <INPUT TYPE=password
          NAME="j_password" SIZE=20><BR>
```

```
        <INPUT TYPE=submit NAME=action
            VALUE="Submit Login">
    </FORM><HR>
</CENTER>
</BODY>
</HTML>
```

Listing 3.7 is a deployment descriptor fragment showing an example of login configuration that requires form-based authentication.

Listing 3.7. Login Configuration for Form-Based Authentication

```
<login-config>
  <auth-method>FORM</auth-method>
  <form-login-config>
    <form-login-page>/login.html</form-login-page>
    <form-error-page>
      /login-failed.html
    </form-error-page>
  </form-login-config>
</login-config>
```

The third type of authentication method is *certificate based* (X.509 certificate). In order for the authentication method to be certificate based, the `auth-method` element in the `login-config` descriptor must be set to `CLIENT-CERT`. The certificate-based authentication method implies that the Web server is configured to perform mutual authentication over SSL. The client is required to present a certificate to establish the connection. When the `CLIENT-CERT` mode is specified, the client will be required to submit the request over an HTTPS connection. If the request is not already over HTTPS, the J2EE product will redirect the client over an HTTPS connection. Successful establishment of an SSL connection implies that the client has presented its own certificate and not anyone else's. The details of how the server ensures that the client certificate really belongs to the client are explained in Section 10.3.4 on page 372 and Section 13.1.2 on page 452. The certificate used by the client is then mapped to an identity in the user registry the J2EE product is configured to use.

Listing 3.8 is a deployment descriptor fragment showing an example of login configuration that requires certificate-based authentication.

Listing 3.8. Login Configuration for Certificate-Based Authentication

```
<login-config>
  <auth-method>CLIENT-CERT</auth-method>
</login-config>
```

Note that the user registry is not specified in this XML deployment descriptor fragment because it is not part of the J2EE specification.

3.9.1.2 *Secure-Channel Constraint*

Establishing an HTTPS session between the client and the Web server is often a necessary requirement to provide data confidentiality and integrity for the information flowing between the HTTP client and the server. In a J2EE environment, the security policy can require the use of a secure channel, specified through the `user-data-constraint` deployment descriptor element. When the requirement for a secure channel is specified, the request to the URI resource should be initiated over an HTTPS connection. If access is not already via a HTTPS session, the request is redirected over an HTTPS connection.

Specifying `INTEGRAL` or `CONFIDENTIAL` as the value for the `transport-guarantee` element in the `user-data-constraint` descriptor will be treated as a requirement for the HTTP request to be over SSL. This requirement can be specified as part of the `user-data-constraint` element in a Web application's login configuration. In theory, `INTEGRAL` should enforce communication integrity, whereas `CONFIDENTIAL` should enforce communication confidentiality, and it could be possible to select different cipher suites to satisfy these requirements. However, a J2EE server typically does not differentiate `INTEGRAL` from `CONFIDENTIAL` but instead treats both of these values to indicate the need to require an SSL connection with a particular cipher suite, not based on whether `INTEGRAL` or `CONFIDENTIAL` was specified.

Listing 3.9 is a deployment descriptor fragment showing an example of login configuration that contains the `user-data-constraint` element. More details are provided in Section 4.6.6 on page 132.

Listing 3.9. Specifying the Requirement for a Secure Channel

```
<user-data-constraint>
  <transport-guarantee>CONFIDENTIAL</transport-guarantee>
</user-data-constraint>
```

3.9.2 Authorization Policy

The *role-permission interpretation* of the J2EE security model treats a security role to be a set of permissions. The security role uses the `role-name` label defined in the `method-permission` element of an EJB module's deployment descriptor and in the `security-constraint` element of a Web module's deployment descriptor as the name of the set of permissions. The set of permissions defines a number of *resources*—the enterprise beans and the Web resources to which the

method-permission and security-constraint elements refer, respectively—and a set of *actions*—the methods listed by the method-permission and the security-constraint descriptors. For example, in Listing 3.2 on page 74, the security role Teller is associated with the permissions to invoke the `getBalance()` and `getDetails()` methods on the AccountBean enterprise bean. Similarly, in Listing 3.3 on page 75, the security role Teller is associated with the permission to perform a GET invocation over HTTP to the `/finance/account/` URI. If multiple method-permission and security-constraint descriptors refer to the same security role, they are all taken to contribute to the same role permission. In other words, the sets of permissions associated with that security role are merged to form a single set.

This model has the advantage of dramatically reducing the number of objects in a *security object space*—a set of pairs (*subject*, \langle *target*, *operation* \rangle), where the *subject* is an entity requesting to perform a security-sensitive *operation* on a given *target*. The Deployer and the System Administrator can define authorization policies, associated with EJB or URI targets and the operations of enterprise bean methods and HTTP methods, respectively, for the security roles in their applications. Then, they associate subjects to security roles; by extension, those subjects are granted the permissions to perform the operations permitted by the security roles.

Based on the J2EE security model, a protected action can be performed by a subject who has been granted at least one of the security roles associated with the action. The security roles associated with a protected action are the *required security roles*—the permissions necessary to perform the action itself. The roles associated with a subject are the *granted security roles*—the permissions that have been given to that subject. This means that the subject will be allowed to perform an action if the subject's granted security roles contain at least one of the required security roles to perform that action. For example, if the action consisting of accessing the EJB method `getDetails()` on the AccountBean enterprise bean can be performed only by the security roles Teller and Supervisor and if subject Bob has been granted the security role of Teller, Bob will be allowed to perform that action, even if Bob has not been granted the security role of Supervisor.

The table that represents the association of security roles to sets of permissions is called the *method-permission table*. A method-permission table (see Table 3.1) can be used to deduce the set of required security roles. The rows in the table represent security roles; the columns represent protected actions.

It can be inferred from Table 3.1 that in order to access the `getBalance()` method on AccountBean, the required security roles are Teller and Supervisor. In order to access any URI that matches the pattern `/public/*`, a PublicRole is required.

Table 3.1. Example of Method-Permission Table

	<code>/finance/ accountGET</code>	<code>/finance/ accountPUT</code>	<code>/public/*</code>	<code>AccountBean. getBalance()</code>	<code>AccountBean. getDetails()</code>
Teller	Yes	No	No	Yes	Yes
Supervisor	Yes	Yes	No	Yes	Yes
PublicRole	No	No	Yes	No	No

The table that represents the association of roles to subjects is called the *authorization table*, or *protection matrix*. In such a table, the security role is defined as the *security object*, and users and groups are defined as *security subjects*. An authorization table (see Table 3.2) can be used to deduce the set of granted security roles. The rows in the table refer to the users and user groups that are security subjects in the protection matrix; the columns represent the J2EE security roles that are security objects in the protection matrix.

The method-permission table and the protection matrix reflect the configuration specified in the deployment descriptors. For example, the first row in Table 3.1 reflects the deployment descriptor obtained from the deployment descriptor fragments of Listing 3.2 on page 74 and Listing 3.3 on page 75. It can be inferred from Table 3.2 that user Bob and group TellerGroup are granted the security role of Teller, everyone is granted the PublicRole, and only users in the ManagerGroup are granted the security role of Supervisor.

Combining Table 3.1 and Table 3.2, it follows that Bob can access the `getBalance()` and `getDetails()` methods on the `AccountBean` enterprise bean and can issue an HTTP `GET` request on the `/finance/account/` URI. Bob cannot, however, issue an HTTP `PUT` request on the `/finance/account/` URI. Note that Bob will be able to access any URI that matches `/public/*`, as everyone has been granted the role `PublicRole`, which is the role necessary to get access to `/public/*`.

In the J2EE security model, the Application Assembler defines the initial mapping of actions on the protected resources to the set of the required security

Table 3.2. Example of Authorization Table

	Teller	Supervisor	PublicRole
TellerGroup	Yes	No	No
ManagerGroup	No	Yes	No
<i>Everyone</i>	No	No	Yes
Bob	Yes	No	No

roles (see Section 3.7.2 on page 67). This can be done using the application assembly tool. Subsequently, the Deployer will refine the policies specified by the Application Assembler when installing the application into a J2EE environment (see Section 3.7.3 on page 70). The Deployer also can use the application assembly tool to redefine the security policies, when necessary, and then install the application into the J2EE container. The method-permission table is formed as a result of the required security roles getting specified through the process of application assembly and refinement during deployment.

Authorization policies can be broadly categorized into *application policies*, which are specified in deployment descriptors and map J2EE resources to roles, and *authorization bindings*, which reflect role to user or group mapping. As discussed in Section 3.7.2 on page 67, a set of security roles is associated with actions on J2EE protected resources. These associations are defined in the J2EE deployment descriptors when an application is assembled and deployed. The security roles specified in this way are the required security roles—the sets of permissions that users must be granted in order to be able to perform actions on protected resources. Pragmatically, before a user is allowed to perform an action on a protected resource, either that same user or one of the groups that user is a member of should be granted at least one of the required security roles associated with that protected resource. The authorization table that relates the application-scoped required security roles to users and user groups is managed within the J2EE Product Provider using the J2EE Product Provider configuration tools.

3.9.3 Delegation Policy

Earlier in this chapter, we defined *delegation* as the process of forwarding a principal's credentials with the cascaded downstream requests. Enforcement of delegation policies affects the identity under which the intermediary will perform the downstream invocations on other components. By default, the intermediary will impersonate the requesting client when making the downstream calls. The downstream resources do not know about the real identity, prior to impersonation, of the intermediary. Alternatively, the intermediary may perform the downstream invocations using a different identity. In either case, the access decisions on the downstream objects are based on the identity at the outbound call from the intermediary. To summarize, in a J2EE environment, the identity under which the intermediary will perform a task can be either

- The client's identity—the identity under which the client is making the request to the intermediary
- A specified identity—an identity in terms of a role indicated via deployment descriptor configuration

The application deployment environment determines whether the client or a specified identity is appropriate.

The Application Assembler can use the `security-identity` element to define a delegation identity for an enterprise bean's method in the deployment descriptor. Consider an example in which a user, Bob, invokes methods on a `SavingsAccountBean` enterprise bean. `SavingsAccountBean` exposes three methods—`getBalance()`, `setBalance()`, and `transferToOtherBank()`—and its delegation policy is defined as in Table 3.3. Figure 3.5 shows a possible scenario based on the delegation policy specified in Table 3.3.

The method `setBalance()` will execute under the client's identity because the delegation mode is set to `use-caller-identity`. The method `getBalance()` will execute under the client's identity as well because no delegation mode is specified, and the default is `use-caller-identity`. Therefore, if Bob invokes the method `getBalance()` on `AccountBean`, the method will execute under Bob's identity, `bob`. Suppose that the `getBalance()` method invokes a `lookup()` method on

Table 3.3. `SavingsAccountBean` Enterprise Bean's Delegation Policy

Method	Delegation Mode	Specified Role
<code>getBalance()</code>		
<code>setBalance()</code>	<code>use-caller-identity</code>	
<code>transferToOtherBank()</code>	<code>run-as</code>	Supervisor

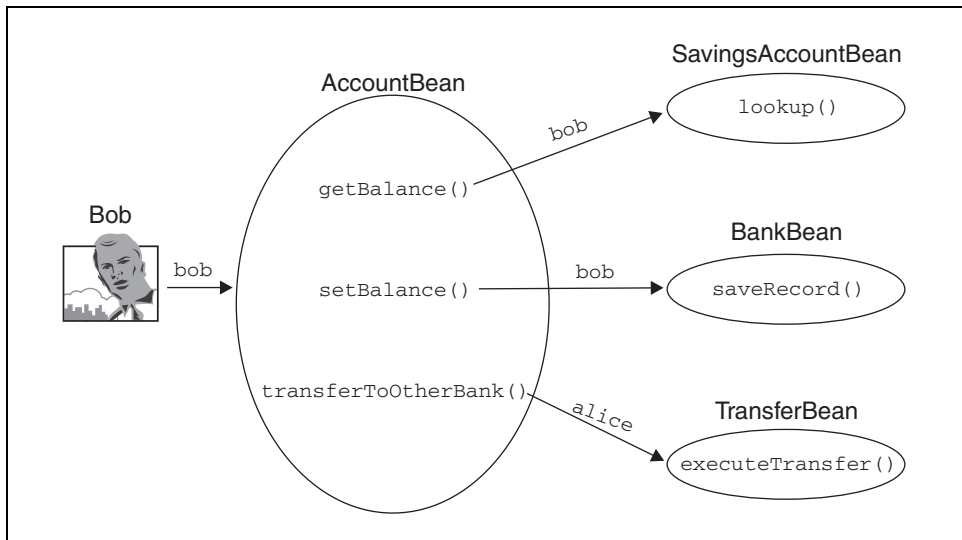


Figure 3.5. Delegation Policy Scenario

SavingsAccountBean. This invocation will still be executed under Bob's identity and will succeed only if Bob has been granted the permission to invoke `lookup()` on SavingsAccountBean.

Any downstream call from `transferToOtherBank()` will perform method calls on a TransferBean enterprise bean. These invocations will need to execute under a principal that has been granted the Supervisor role. The Deployer or the System Administrator needs to map the Supervisor role to a principal that has been granted the Supervisor role. This can be done by specifying a valid user ID and password pair corresponding to a user who has been granted that role. For example, if user Alice has been granted the Supervisor role and if the user ID and password pair for Alice is associated with the Supervisor role, the calls to `transferToOtherBank()` will occur under Alice's identity.

3.9.4 Connection Policy

Information in any EIS must be protected from unauthorized access. An EIS system is likely to have its own authorization model. At a minimum, most of these systems have facilities to accept some form of authentication data representing an identity connecting to the EIS. The JCA is designed to extend the end-to-end security model for J2EE-based applications to include integration with EISs. A WAS and an EIS collaborate to ensure the proper authentication of a resource principal when establishing a connection to a target EIS. As discussed in Section 3.4 on page 61, the JCA allows for two ways to sign on to an EIS: container-managed sign-on and component-managed sign-on.

With container-managed sign-on, the connection to an EIS is obtained through declarative security. In order for a connection to be container managed, the deployment descriptor will indicate that the `res-auth` element associated with a resource definition is declared as `Container`. If the connection is obtained by passing the identity information programmatically, the value for `res-auth` should be set to `Application`. Details of component-managed sign-on are discussed in Section 3.10.3 on page 94.

A deployment descriptor fragment that declares that the authentication facilitated by the resource adapter should be set to be `Container` is shown in Listing 3.10.

Listing 3.10. An XML `res-auth` Element in a Deployment Descriptor

```
<resource-ref>
  <description>Connection to myConnection</description>
  <res-ref-name>eis/myConnection</res-ref-name>
  <res-type>javax.resource.cci.ConnectionFactory</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

The container is responsible for obtaining appropriate user authentication information needed to access the EIS. The connection to the EIS is facilitated by the specified resource adapter. The JCA allows specifying the authentication mechanism. The `authentication-mechanism-type` element in the deployment descriptor is used to specify whether a resource adapter supports a specific authentication mechanism. This XML element is a subelement of the `authentication-mechanism` element. The JCA specification supports the following authentication mechanisms:

- **Basic authentication.** The authentication mechanism is based on user ID and password. In this case, the `authentication-mechanism-type` XML element in the deployment descriptor is set to `BasicPassword`.
- **Kerberos V5.** The authentication mechanism is based on Kerberos V5. In this case, the `authentication-mechanism-type` element in the deployment descriptor is set to `Kerbv5`.

Other authentication mechanisms are outside the scope of the JCA specification.

In a secure environment, it is likely that a J2EE application component, such as an enterprise bean, and the EIS system that is accessed through the component are secured under different security domains, where a *security domain* is a scope within which certain common security mechanisms and policies are established. In such cases, the identity under which the J2EE component is accessed should be mapped to an identity under which the EIS is to be accessed. Figure 3.6 depicts a possible scenario.

In this scenario, an enterprise bean in a J2EE container is accessed by a user, Bob Smith. The enterprise bean is protected in a way that it allows only users from a specified LDAP directory to access it. Therefore, the identity under which Bob Smith will access the enterprise bean must be registered in that LDAP directory. Bob Smith uses the identity of `bsmith` when he accesses the enterprise bean.

In a simplistic case, where the `run-as` policy of the enterprise bean is set to be the caller identity, the connections to the EIS will be obtained on behalf of Bob Smith. If the connections are obtained through user ID and password, when the enterprise bean obtains a connection to a back-end system, such as a CICS system, the J2EE container will retrieve a user ID and password to act on behalf of user `bsmith`. The application invokes the `getConnection()` method on the `javax.resource.cci.ConnectionFactory` instance (see Listing 3.10 on page 87) with no security-related parameters, as shown in Listing 3.11, a fragment of Java code.

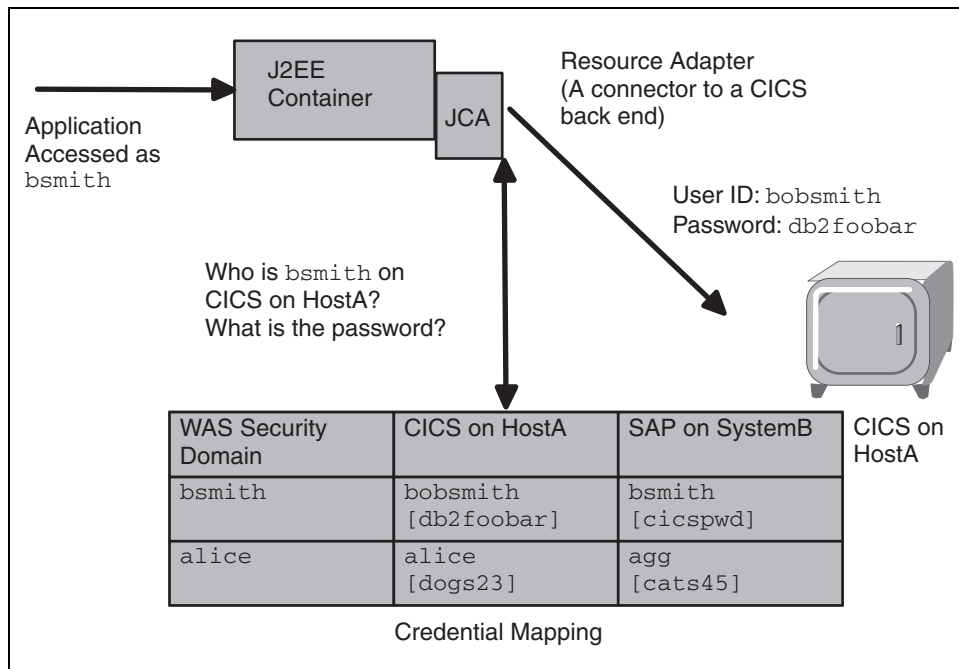


Figure 3.6. Credential Mapping when Accessing an EIS from a J2EE Container

Listing 3.11. Getting a Connection to an EIS with Container-Managed Sign-On

```
// Construct the InitialContext
Context initctx = new InitialContext();

// Perform a JNDI lookup to obtain a ConnectionFactory
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory) initctx.lookup
    ("java:comp/env/eis/MyEIS");

// Invoke the ConnectionFactory to obtain a connection.
// The security information is not passed to the
// getConnection() method
javax.resource.cci.Connection cx = cxf.getConnection();
```

The application relies on the container to manage the sign-on to the EIS instance. This is possible in simple deployment scenarios in which the identity under which the EIS system is accessed is specified by the Deployer. This effectively means that all identities accessing the application are mapped to a single identity to access the EIS system: a *many-to-one identity mapping*.

In more sophisticated deployment scenarios, a many-to-one identity mapping may not be sufficient for security policy reasons. For example, it may be necessary for the EIS system to log all the identities that accessed it. For this logging facility to be useful, the identities accessing a J2EE application must not all be mapped to the same identity on the EIS system. A *one-to-one* or *many-to-many identity mapping* is recommended in this case. In particular, the container may use a credential mapping facility whereby `bsmith` is mapped to user ID `bobsmith` and password `db2foobar`, as shown in Figure 3.6.

If connections require Kerberos credentials or other generic credentials to be passed, the mapping facility is responsible for mapping one form of the credential to another that can be used by the target security domain. The manner in which these mappings happen and the level of sophistication in mapping available in J2EE application servers are server specific and not dictated by the J2EE specification.

In enterprise environments consisting of multiple departments, organizations, and even acquired companies, it is typical for systems to be interconnected and the applications shared. In such environments in which J2EE applications are deployed, it is a good architectural approach to design the application integration in a way that applications use JCA to obtain connections to other applications and to follow the declarative approach to define connection sign-on, as explained in this section. The use of JCA will make applications unaware of cross-security domains when accessing non-J2EE systems, and the use of declarative security will enhance application flexibility and portability. JCA with declarative security will also help manage the mapping of credentials and identities outside the application as enforced and facilitated by the enterprise-level mapping infrastructure.

3.10 Programmatic Security

Declarative security should always be used instead of programmatic security whenever possible. However, when declarative security is insufficient, it may be necessary to retrieve security-sensitive information programmatically from the container. This section explains how to retrieve the user's identity and privilege information programmatically.

Applications that make use of programmatic security typically invoke the following EJB and servlet/JSP security APIs:

- **EJB method `isCallerInRole()` in interface `javax.ejb.EJBContext`.** This method is used to test whether the current caller, the client, has been assigned to a specified security role.

- **EJB method `getCallerPrincipal()` in interface `javax.ejb.EJBContext`.** This method is used to obtain a `Principal` object representing the current caller, the client.
- **Servlet/JSP method `isUserInRole()` in interface `javax.servlet.http.HttpServletRequest`.** This method, similar to the EJB method `isCallerInRole()`, returns a `boolean` indicating whether the authenticated user, the client, is a member of the specified security role.
- **Servlet/JSP method `getUserPrincipal()` in interface `javax.servlet.http.HttpServletRequest`.** This method, similar to the EJB method `getCallerPrincipal()`, returns a `Principal` object representing the current authenticated user, the client.

3.10.1 Retrieving Identity Information

The Java Servlet and EJB specifications provide mechanisms to programmatically obtain identity information about the user invoking a method on a servlet or an enterprise bean.

3.10.1.1 From a Servlet or JSP File

The `HttpServletRequest` object passed to a servlet method can be used to obtain information about the user invoking the method. Invoking the `getRemoteUser()` method on the `HttpServletRequest` object returns the name of the user if the user has been authenticated, `null` otherwise.

The `getRemoteUser()` method can be invoked as shown in Listing 3.12.

Listing 3.12. Retrieving the User Name from a Servlet

```
public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
{
    // other code...

    // obtain the user name
    String userName = req.getRemoteUser();

    // other code...
}
```

The `getUserPrincipal()` method in the `HttpServletRequest` object returns the `Principal` object corresponding to the user if the user has been authenticated, `null` otherwise. The name of the user can then be obtained by calling the `getName()` method on the `Principal` object, if this not `null`, as shown in Listing 3.13.

Listing 3.13. Retrieving the Principal Object and the User Name from a Servlet

```
public void doGet(HttpServletRequest req,
    HttpServletResponse res)
{
    // other code...

    // obtain the user Principal
    Principal userPrincipal = req.getUserPrincipal();

    // obtain the user name
    String userName;

    if (userPrincipal != null)
        userName = userPrincipal.getName();

    // other code...
}
```

3.10.1.2 From an Enterprise Bean

The `getCallerPrincipal()` method can be called on a `javax.ejb.EJBContext` object to obtain the `Principal` object corresponding to the user making the enterprise bean method invocation. The `Principal` object can then be used to obtain information about the user. A code example is shown in Listing 3.14.

Listing 3.14. Retrieving the User Name from an Enterprise Bean

```
public String getUserName(EJBContext context)
{
    // obtain and return the user name
    return context.getCallerPrincipal().getName();
}
```

3.10.2 Proactive Authorization

The Java Servlet and EJB specifications provide mechanisms to programmatically obtain information about the user's privileges by invoking a method on a servlet or an enterprise bean.

3.10.2.1 From a Servlet or JSP File

The `HttpServletRequest` object passed to a servlet method can be interrogated to obtain information about whether the user invoking the method has been granted a particular security role. Based on the result, the servlet may make decisions on how to proceed. For example, if the caller is granted the `Boss` role, the servlet redirects to a page that has managerial capabilities; otherwise, it might redirect to a different page. Note that when the servlet checks for the `Boss` role, this

role is scoped to the servlet. The Application Assembler performs the mapping of this role reference to an enterprise application role by using the `role-link` tag in the deployment descriptor, as shown in Listing 3.15.

Listing 3.15. An XML `role-link` Element in a Deployment Descriptor

```
<security-role-ref>
  <role-name>Boss</role-name>
  <role-link>Manager</role-link>
</security-role-ref>
```

A code example is shown in Listing 3.16.

Listing 3.16. Retrieving the User's Role Information from a Servlet

```
public void doGet(HttpServletRequest req,
  HttpServletResponse res)
{
  // other code...

  if (req.isUserInRole("Boss"))
  {
    // code to redirect to Manager's page...
  }
  else
  {
    // code to redirect to generic page...
  }

  // other code...
}
```

3.10.2.2 From an Enterprise Bean

The `isCallerInRole()` method on an `EJBContext` object can be used to obtain information about the roles granted to a particular user. This is similar to what we saw in Section 3.10.1.1 on page 91. A code fragment example is shown in Listing 3.17.

Listing 3.17. Retrieving the User's Role Information from an Enterprise Bean

```
public Object getOrganizationInfo(EJBContext context)
{
  // other code...

  // obtain the user name
  if (context.isCallerInRole ("Boss")
  {
```

(continues)

Listing 3.17. Retrieving the User's Role Information from an Enterprise Bean

```
        // code to access the Boss entity bean and get the
        // budget info...
    }
    else
    {
        // code to access the employee bean and get the
        // organization chart...
    }

    // other code...
}
```

3.10.3 Application-Managed Sign-On to an EIS

Section 3.9.4 on page 87 described container-managed sign-on, whereby a connection to an EIS is obtained through declarative security. An alternative approach is to use programmatic security by allowing the sign-on to an EIS to be managed directly by the application. In order for a connection to be application managed, the value of the `res-auth` XML element associated with a resource definition in the deployment descriptor must be set to `Application`, as shown in Listing 3.18.

Listing 3.18. Setting the `res-auth` Deployment Descriptor Tag to `Application`

```
<resource-ref>
  <description>Connection to myConnection</description>
  <res-ref-name>eis/myConnection</res-ref-name>
  <res-type>javax.resource.cci.ConnectionFactory</res-type>
  <res-auth>Application</res-auth>
</resource-ref>
```

In the case of application-managed sign-on, the application is responsible for retrieving appropriate user information, such as the user ID and password, necessary to connect to the EIS. The connection is facilitated by a resource adapter. The application invokes the `getConnection()` method on the `ConnectionFactory` instance with the security information: user ID and password. Specifying security information is dependent on the resource adapter type and the way in which the adapter accepts the user ID and password. For example, in order to connect to an EIS system called MyEIS, the application may be required to pass user ID and password through a `com.myeis.ConnectionSpecImpl` object. Listing 3.19 is a Java code fragment showing such a scenario, which is similar to the one discussed in Section 3.9.4 on page 87, except that the security information is coded into the application and is passed to the `getConnection()` method.

Listing 3.19. Getting a Connection to an EIS with Container-Managed Sign-On

```
// Method in an application component
Context initctx = new InitialContext();

// Perform a JNDI lookup to obtain a ConnectionFactory
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory) initctx.lookup
        ("java:comp/env/eis/MyEIS");

// Insert here the code to get a new ConnectionSpec
com.myeis.ConnectionSpecImpl props = // ...

// Set user ID and password
props.setUserName("bobsmith");
props.setPassword("db2foobar");

// Invoke the ConnectionFactory to obtain a connection.
// The security information is passed explicitly to the
// getConnection() method.
javax.resource.cci.Connection cx = cxf.getConnection(props);
```

3.11 Secure Communication within a WAS Environment

A fundamental aspect of any distributed computing system is remote communication. In an enterprise environment, components, such as Web servers, plug-ins, and WASs; external servers, such as LDAP directory servers; and clients communicate with one another over multiple protocols.

- HTTP clients invoke URL requests to Web servers over HTTP.
- WASs communicate with one another over IIOP.
- Some WASs may communicate with external systems by using other protocols. For instance, a WAS can communicate with an LDAP directory server over LDAP.

Because these components can host and distribute security-sensitive information, it is necessary to provide secure communication channels. The quality-of-service (QoS) should include encryption, integrity, and, possibly, authentication. The SSL protocol is generally used to meet these QoS requirements. Typically, two modes of SSL connections are used in J2EE:

1. **Server-side SSL.** The client connects to the server and attempts to verify the authenticity of the server's identity. The server does not verify the client's identity. If the client can authenticate the server successfully, the client/server communication is performed over an encrypted channel.

- 2. Mutual-authentication SSL.** The client connects to the server. Both client and server attempt to authenticate to each other. If the mutual-authentication process is successful, client and server communicate over an encrypted channel.

The SSL configuration on the server side dictates whether a client connecting to the server should connect over server-side SSL or mutual-authentication SSL. In both cases, the strength of encryption depends on the configured cipher suite. In a Java environment, JSSE-compliant SSL providers are used to establish secure communication using SSL between the end points (see Chapter 13 on page 449).

The following list shows possible combinations of securely communicating parties:

- **Web client to Web server.** Any Web browser can issue a request to a Web server over a secure connection. This communication can be over either server-side SSL or mutual-authentication SSL. The Web server should be configured to accept connections from Web browsers over a secure-socket port—typically, port 443. If a WAS requires a client to present a client certificate in order to be authenticated to access a servlet, the underlying Web server should be configured to require mutual-authentication SSL connections from Web browsers. This configuration is specific to the underlying Web server that is used, whereas the client certificate information is configured on the Web browser and is specific to the browser settings.
- **Web server to WAS.** In general, a Web server needs a plug-in to communicate with a back-end WAS, unless the Web server and the WAS are integrated to form a single component. In a typical scenario, a Web server plug-in may communicate with an application server over HTTP. This can be configured to be over SSL, in which case the resulting protocol is HTTPS. In order for this to happen, the WAS transport must be configured to accept only secure connections. In the case of mutual-authentication SSL, the WAS transport can also be configured to trust only a set of selected clients to connect to the WAS. By properly configuring the digital-key storage facility—for example, a key file—both the plug-in and the WAS can be configured to accept a list of trusted Certificate Authorities (CAs), so that a trusted communication link can be established with only the clients whose certificates have been authenticated by one of the trusted CAs (see Section 10.3.4 on page 372).
- **WAS to WAS.** WASs communicate to other WASs by using IIOP. In a secure environment, all these communications are over SSL. The digital-key databases can be configured to reflect the trust policy by using the

tools provided with J2EE. In general, it is also possible to configure the strength of encryption enforced in such a secure connection.

- **Application client to WAS.** Similar to a Web browser's making a request to a Web resource, Java clients can use the EJB programming model to invoke methods on an enterprise bean by connecting to the WAS that hosts the enterprise bean. In a secure environment, the communication from the client to the protected resources should be protected by the SSL protocol. At that point, an application client can securely communicate with a WAS over SSL. This is achieved by configuring the WAS with a list of trusted application clients. Alternatively, the WAS can be configured to accept a list of trusted CAs, so that a trusted communication link can be established with only the application clients whose digital certificates have been authenticated by one of the trusted CAs.
- **WAS to LDAP directory server.** A WAS may need to connect to external systems. For example, a WAS may be configured to use an LDAP directory server as a user registry. In this case, the WAS will make calls against the LDAP directory. The protocol of communication between these two entities is LDAP. User ID and password pairs are verified by performing LDAP bind operations against the LDAP directory. In a scenario like this, these values will flow over the wire unencrypted unless the communication is protected by the SSL protocol. For this reason, an LDAP directory server should be configured to require that all connections be over SSL. The set of digital certificates trusted by the directory server can be imported into the WAS's digital-key storage facility so that the WAS can successfully establish an SSL connection with the LDAP directory server.

3.12 Secure E-Business Request Flow

We are now ready to revisit the simple e-business request flow presented in Section 3.6 on page 63 with the security semantics added to the flow. The secure e-business request flow takes the security model and security technologies available in the J2EE environment into account and depicts the use of those components.

Figure 3.7 enhances the diagram depicted in Figure 3.3 on page 63 and shows how the J2EE security technologies presented in this chapter can play a role in a secure e-business request flow.

Let us consider the request flow discussed in Section 3.6 on page 63 and understand how security technologies play a role when that request is secured. A user, say, Bob, invokes URL `http://samples.com/servlet/account`, which is handled by the Web container that hosts the account servlet. Based on the security

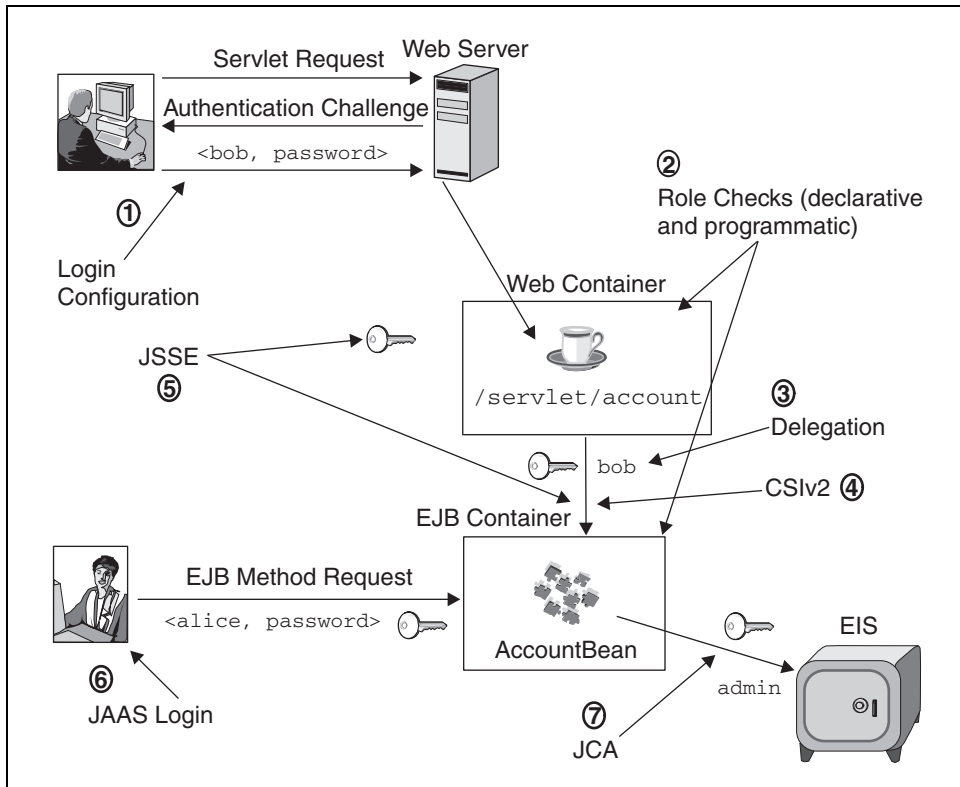


Figure 3.7. Secure E-Business Request Flow

constraints defined, the URI `/servlet/account` is a protected resource. The login configuration associated with the Web application is set to `BASIC`. Therefore, when the request arrives, the Web container issues an HTTP 401 response back to the browser. Bob resubmits his request by providing his user ID (`bob`) and his password (see step 1 in Figure 3.7). The Web container validates the user ID and password by authenticating the pair against the underlying user registry.

After the credentials are validated, the Web container performs an authorization check. The URI is protected in a way to grant access only to the Teller role. The container checks whether user Bob is granted the role (see step 2 in Figure 3.7). As Bob has been granted the Teller role, he is allowed to access the URI.

The servlet invokes the AccountBean enterprise bean, and the request is dispatched to the enterprise bean. The delegation policy on the servlet is not set to `run-as`, which means that the downstream requests will be performed by using the caller's identity. In this case, the caller is Bob, and therefore the identity under

which the enterprise bean is called consists of Bob's credentials, as depicted in step 3 in Figure 3.7.

The Web container dispatches the request to the EJB container when the account servlet makes a call to the AccountBean enterprise bean. The request is sent over IIOP. Given that this request is sent over a secure environment, the CSIV2 protocol is in effect (see Section 3.3 on page 61). The servers hosting the Web container and the EJB container establish a secure association using the CSIV2 technology as depicted in step 4 in Figure 3.7. Based on a successful establishment of the connection and the validity of Bob's credentials, the received identity at the EJB container is Bob's.


The connections between the Web server and the Web container, and the Web container and the EJB container are over SSL. This ensures confidentiality and integrity of the messages sent over the wire. It is essential for all communications in an enterprise to be over a secure connection. JSSE is used to establish the SSL connection as depicted in step 5 in Figure 3.7.

The AccountBean enterprise bean can also be invoked from a Java client. In this case, a user, say, Alice, sends the request directly over IIOP. The request is made from a J2EE client. In this case, the J2EE security technology used for authentication and authorization is JAAS, which is discussed in detail in Chapter 9 on page 289. The client is configured with a `javax.security.auth.callback.CallbackHandler` and performs a JAAS login against the server hosting the EJB container (step 6 in Figure 3.7). Alice needs to provide a valid user ID and password pair. Then, she can perform a method invocation on the AccountBean enterprise bean. J2EE supports a Java client to submit a user ID and password pair over an IIOP message. Using the CSIV2 protocol, the user ID and password pair is in a Generic Security Services Username Password (GSSUP) token within the CSIV2 `ESTABLISH_CONTEXT` message.⁴ In order for this communication to be protected, transport-level SSL is recommended.

The AccountBean enterprise bean is protected by J2EE declarative security. A method permission definition declares the enterprise bean's methods to be accessible only by those who are granted the Teller role. Both users, Bob and Alice, are granted the Teller role. Therefore, the requests that come through the servlet and the one directly submitted are allowed to be invoked. The AccountBean enterprise bean is successfully accessed after the authorization check.

The AccountBean enterprise bean needs to access the data source in order to retrieve the account information. The bean uses a JCA connection manager to obtain a connection to an EIS. When the connection is obtained, based on the connector security configuration, an identity is associated with the connection. The

4. Note that there is no standard declarative way to specify this type of authentication. The details are specific to the J2EE container.



`admin` identity is used to obtain the connection and access the back-end EIS. This access is enforced by the JCA capability and ensures that the EIS is accessed securely, based on the security configuration, as depicted in step 7 in Figure 3.7.

As illustrated in the simple request flow enhanced with security characteristics, the J2EE security model provides the capability and infrastructure to perform secure transactions in an e-business environment.

