

Distribution and Deployment Scenarios

The SDX software comprises a set of interacting software modules that you can install on different hosts and connect to other internetworking devices and applications through a range of standard interfaces. This distributed architecture offers high scalability and extensive flexibility as you customize the SDX software for specific services or create your own applications based on the underlying SDX capabilities. SDX releases 3.1.1 and lower, in which the SDX software was an integrated application rather than a set of interacting modules, are fully compatible with this distributed architecture.

This chapter describes various distribution scenarios for SDX components. It also covers the architecture, component interactions, and deployment for key Enterprise Access Service Portal (EASP) cases. Finally, it includes deployment strategies for the SDX Workflow application according to the expected load of the system, integration with external OSSs, and distribution requirements.



Note: *This chapter describes some typical scenarios, but they are by no means the only ones; many other variations are possible.*

Topic	Page
Component Distribution Scenarios	2-2
EASP Deployment	2-10
SDX Accounting Deployment	2-15
Workflow Application Deployment Strategies	2-16

Component Distribution Scenarios

Figure 2-1 gives an overview installation of all SDX components.

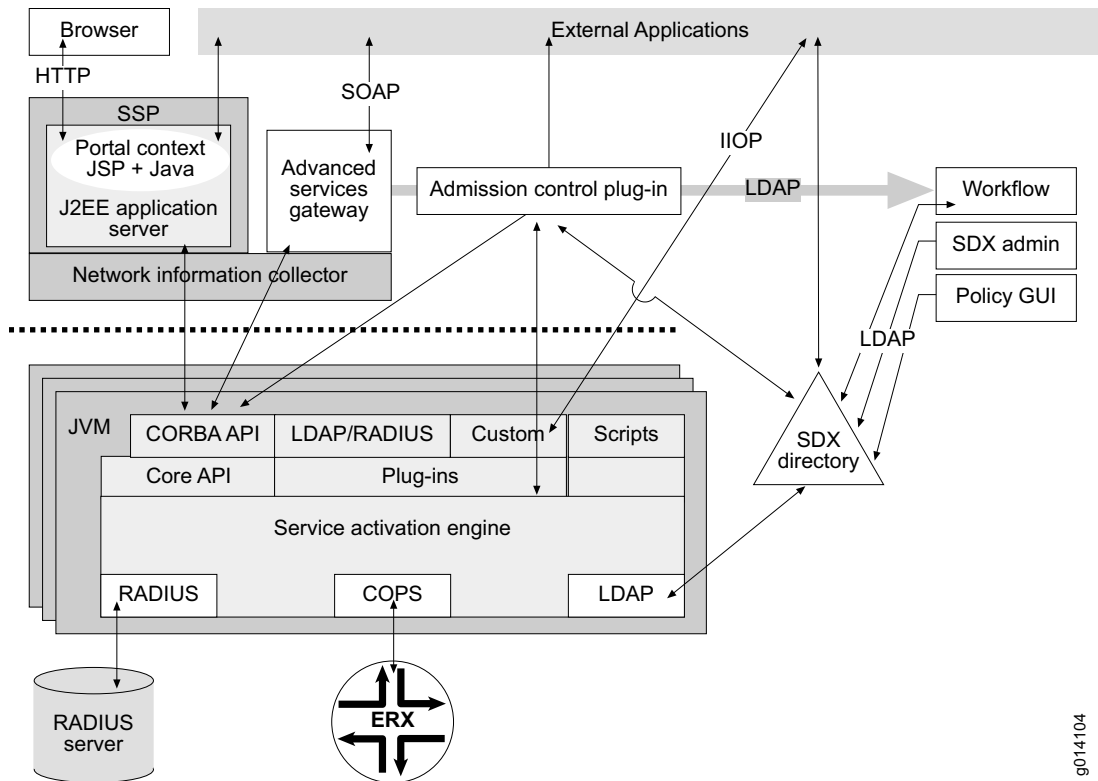


Figure 2-1 Installation of all SDX components

This section covers the following scenarios for distributing SDX components:

- Single box – a minimal installation that is suitable for small operations, demonstrations, and trials
- Distributed – an installation that distributes the SDX components among several machines in several locations and provides reliability and scalability
- Regionalized – distribution scenario for a large service provider that allows regional autonomy
- Consolidated – distribution scenario that consolidates network services into regional data centers

Single-Box Installation

Figure 2-2 shows a single-box installation that is suitable for demonstrations, trials, and small operations. The directory, RADIUS server, SAE, and a Java 2 Enterprise Edition (J2EE) Web application server that contains the portal application are all installed on the same host. The SDX Admin tool and Policy Editor GUI also run locally. The E-series router uses the single RADIUS and SAE servers, and all SDX components act as LDAP clients to the single directory.

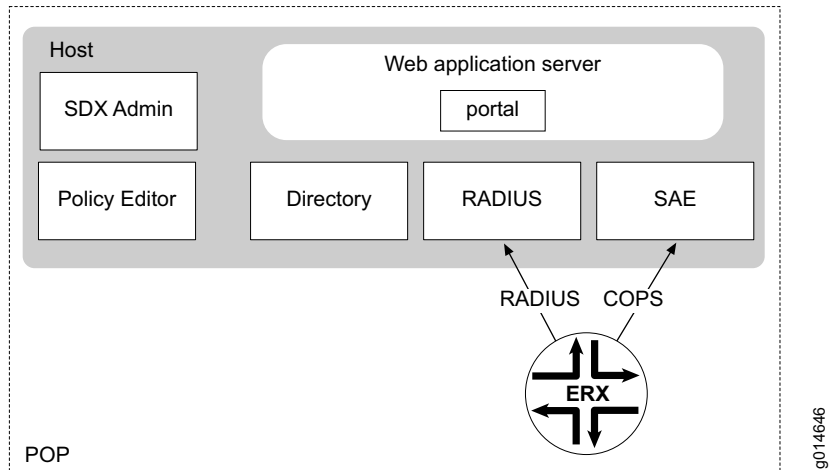


Figure 2-2 Single-box installation for small operations

Distributed Installation

Figure 2-3 shows a more complicated setup that distributes the SDX components among several machines in several locations, while still providing reliability and scalability.

In the back office, there are:

- A master directory server running on dedicated hardware
- SDX Admin and Policy Editor running on as many other machines as desired
- A pair of NIC hosts running NIC resolvers
- A Web application server with a portal (a residential portal, an enterprise portal, or an SDX gateway application)
- Non-SDX components of the service provider's OSS, which are integrated with the SDX components through the master directory as LDAP clients

In the POPs, there are primary and backup hosts that contain identical SAE, RADIUS, directory servers, and NIC hosts. The NIC hosts contain a resolver, directory agent, and SAE agent, and they communicate with the NIC hosts in the back office using CORBA. SAE, RADIUS, and directory server components within the hosts communicate via LDAP.

Clients of the NIC host need to determine which remote SAE is managing the subscriber sessions that they need to operate on. The NIC system collects and stores this information. At startup, the SAE stores its CORBA object reference in the directory. The NIC system collects this SAE reference, along with the keys to subscriber sessions (IP addresses and LDAP DNs of the subscriber profiles in the directory) managed by the SAE. Web applications can locate the SAE for a particular subscriber by querying the NIC system.

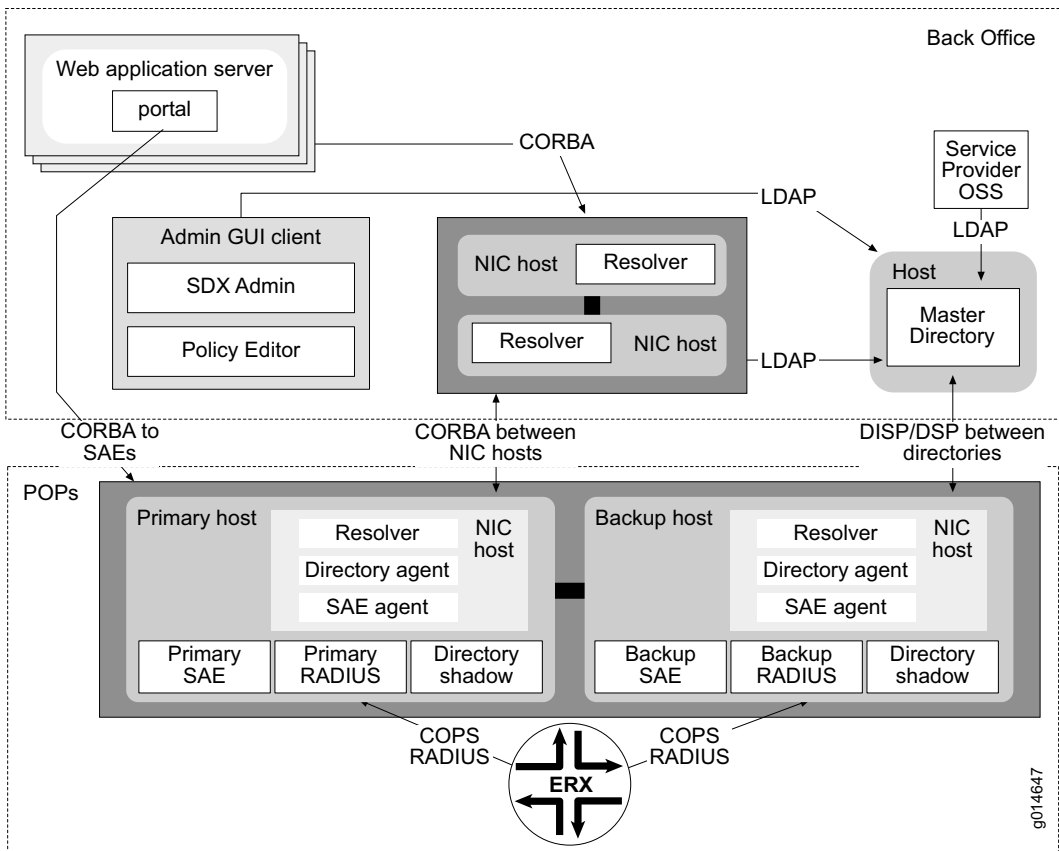


Figure 2-3 Distributed installation for reliability and scalability

Master Directory and Directory Shadows

The master directory contains all the directory data and handles all update requests, either locally via LDAP or remotely via the Directory Service Protocol (DSP) for X.500 directories, such as DirX, or via equivalent protocols for other directory types.

The information in the master directory is copied to shadow directories in the service provider's point of presence (POP). The system uses Directory Information Shadowing Protocol (DISP) for data transfer for X.500 directories, such as DirX, and equivalent protocols for other directory transfers. This type of distribution puts the directory information for SAEs and RADIUS servers physically close to the servers. A highly reliable LAN connects the hosts and provides good performance.

It is not necessary to include all information in the directory shadows. For instance, only information relevant to a particular POP, such as the information for the subscribers who can actually connect there, may be included. Also, updates generated from an SAE in a particular POP, such as cached logins, may be mastered locally and not propagated to the directory master in the back office. Finally, attributes not relevant to SAE and RADIUS operation—for instance, the subscriber's address—may be filtered from replication to the directory shadows in the POPs.

Scalability

This setup can be scaled incrementally by replicating the pattern found in the POP as the subscriber base grows.

Reliability

To avoid a single point of failure in the POPs, the RADIUS, SAE, directory servers, and NIC hosts are installed on identical primary and backup hardware. If the primary host fails, the E-series router switches over to the backup host. Also, the SAE and RADIUS servers (as LDAP clients) and NIC hosts can be configured to switch over to the directory server in the backup host in the POP or to the master directory in the back office. It is also possible to configure N to 1 and N to M redundancy schemes; such redundancy distributes the load of the routers across several hosts and reduces failover time by limiting the number of subscribers handled by any one host.

This setup avoids service outages in the case of any single network, server, or software failure. Existing subscribers are even unaffected by long periods of disconnection between their POP and the back office. The directory server protocols ensure that all information is properly distributed regardless of the pattern of intermittent connectivity between

the sites. Since relatively static directory information is cached locally in the directory shadow in the POP, very high transaction rates for SAE and the RADIUS server are achieved.

Simplified Management and Security

Additional benefits of this setup at the POP are simplified management because of the use of identical hardware and software, and an added level of security because SAE, RADIUS, directory, and NIC hosts are all on the same machine.



Note: *In this and subsequent scenarios, protection of the data in the back office, such as subscriber names and passwords, is a critical issue. Consequently, the back-office site is typically heavily protected by firewalls. One key advantage of this setup is that only directory protocols need be passed through firewalls, and these protocols have rich and flexible security properties.*

Regionalized Installation

Figure 2-4 extends the scheme shown in the last section with an additional layer of directory replication for very large service providers who partition their organization into regions with regional data centers.

A single back office still houses the master directory, some centralized management servers and clients, and a pair of NIC hosts. There are also still primary and backup hosts at the POP, with SAE and RADIUS servers and NIC hosts with a resolver, directory agent, and SAE agent.

In this case, there is also a middle layer of regional data centers that house the first level of replication from the master directory in the back office. The regional data centers may also contain a complete set of SDX components and other OSS management components integrated with the local directory. If the local directory fails, these regional components can switch over to the master directory in the back office and switch back once the local failure is corrected. Also, directory administrative controls can be defined to limit the access of regional management operators to an appropriate scope according to the service provider's policies.

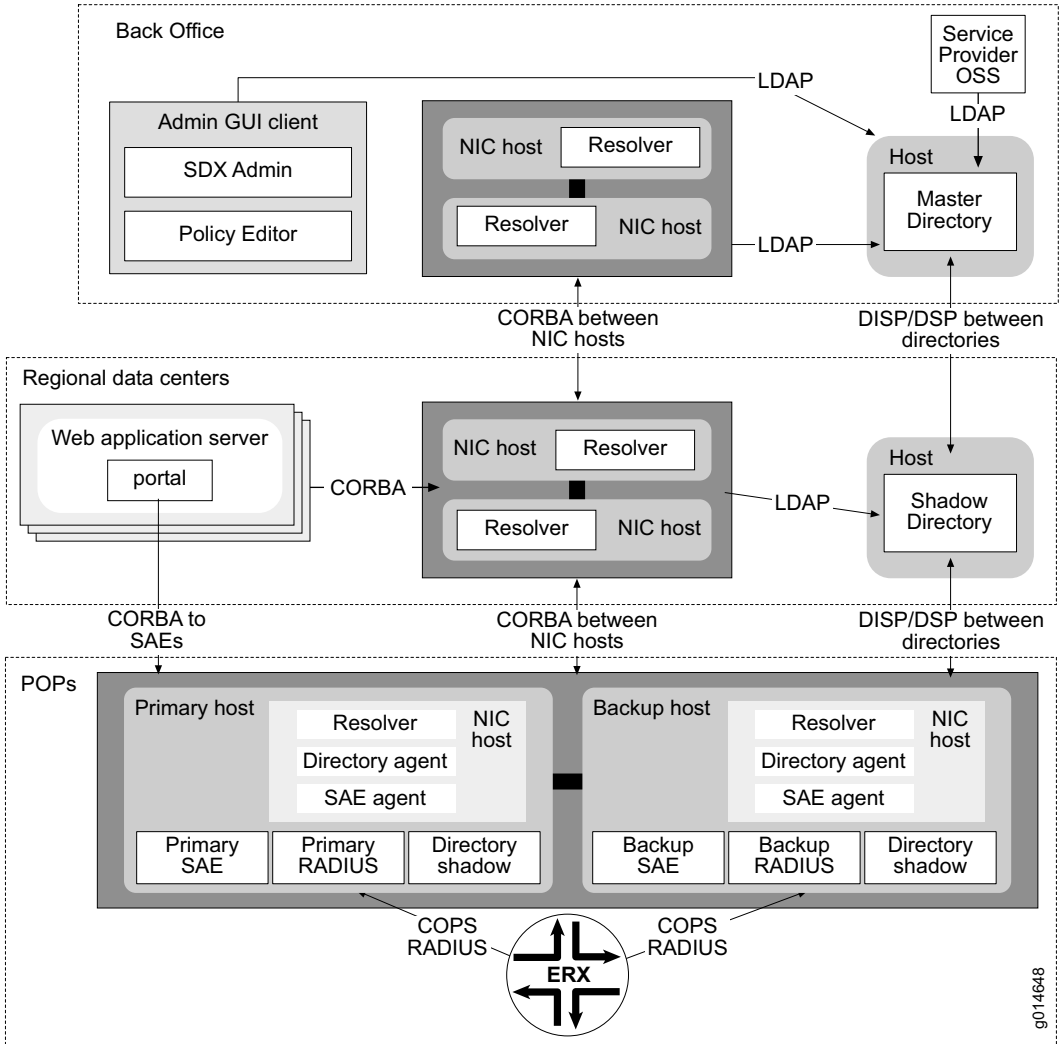


Figure 2-4 Regionalized directory installation for regional autonomy

Consolidated Installation

All of the previous scenarios provide top reliability because all of the network services—that is, the SAE and RADIUS servers—as well as DNS and DHCP servers and NIC hosts, are at the same site as the E-series router and are connected by a reliable LAN. However, to maintain this reliability, hardware must be dedicated to this function in every POP, no matter how small, and economies of scale cannot be achieved through consolidation in large hosts.

The SDX software also supports a deployment scenario that allows a trade-off between consolidation of components in large hosts and the risk of less reliable MAN/WAN connections between sites. This scenario, shown in Figure 2-5, consolidates the network services in regional data centers. Here, the regional data center has:

- Two directory servers for reliability.
- A pair of very large SAE hosts that can be used as the primary or backup for different E-series routers in remote POPs.
- A set of RADIUS hosts that can be load balanced across the various E-series routers and the SAEs for the region.
- A pair of NIC hosts, each running a resolver, directory agent, and SAE agent.
- A Web application server with a portal (a residential portal, an enterprise portal, or an SDX gateway application)

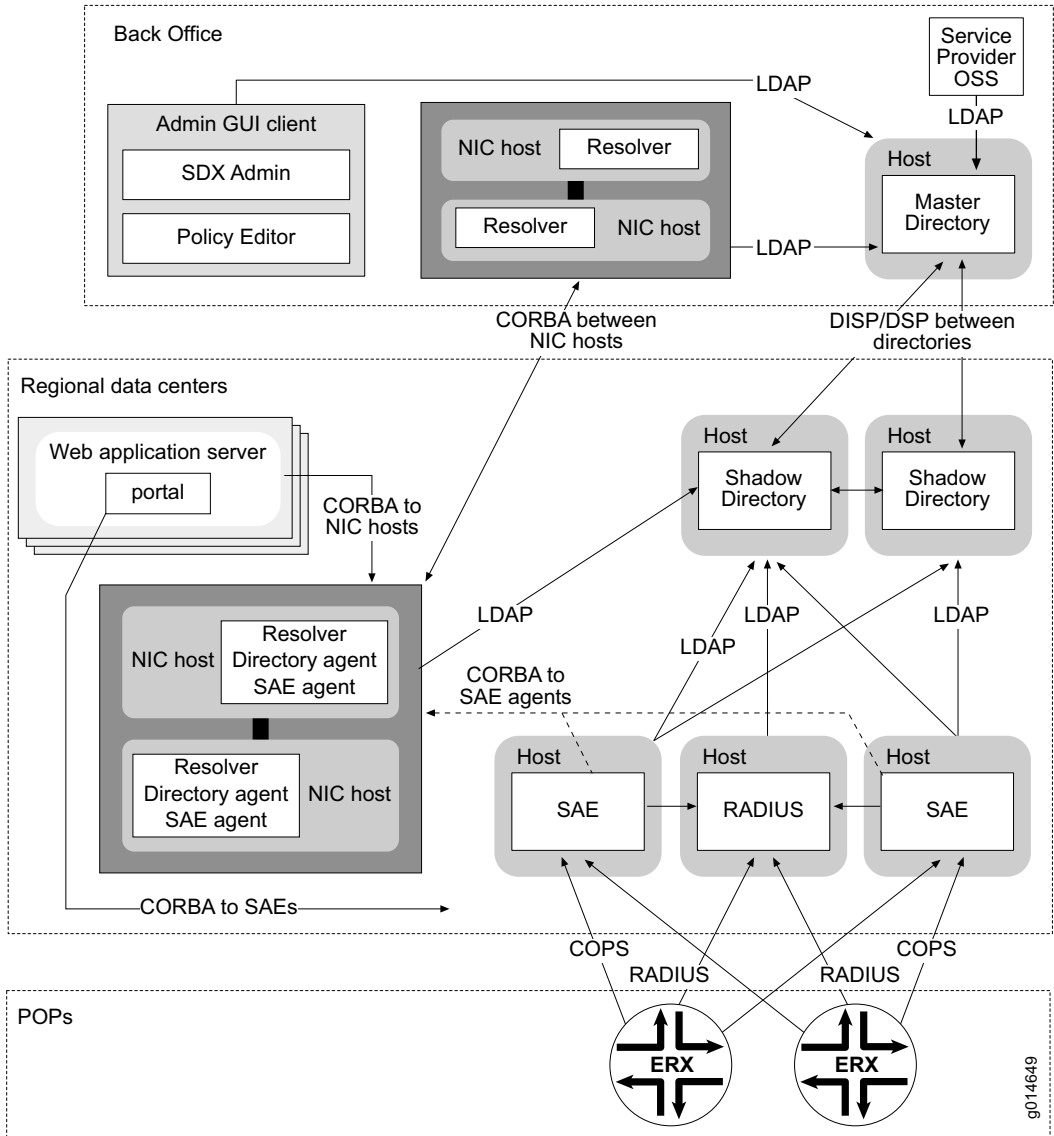


Figure 2-5 Consolidated network services

Redundancy Schemes

The N to 1 and N to M redundancy schemes are even more important in regional data centers because a server could be serving a very large number of subscribers.

RADIUS Because RADIUS is stateless, it is enough to configure a sufficient number of RADIUS servers for the load and configure both the E-series routers and SAE to load balance across them.

NIC Hosts NIC hosts do not need to be redundant in every regional data center. Regional data centers may not even have a NIC host. You can also install NIC servers and NIC directory agents on separate hosts. Each NIC server needs to connect to a NIC directory agent. It is up to service providers to add enough NIC hosts to achieve the desired level of availability and performance.

COPS Connection For the COPS connection between the SAE and E-series routers, special care must be taken. During a failover, existing activated services are not affected; but subscribers cannot log in, activate, or deactivate services until failover synchronization is complete. Thus, it may be desirable to configure multiple SAE machines (for example, tens) in the regional data center to limit the number of subscribers served by any one machine. The E-series routers can be configured with primary, secondary, and tertiary COPS servers, so it is possible to configure many failover schemes. It is also possible for SAE servers to redirect existing and new COPS connections to other, more lightly loaded SAE servers. This COPS connection redirection can be triggered manually during a scheduled maintenance window or automatically based on SAE load monitoring.

Adding or Replacing Hardware Start-up is simplified because there is always a pool of SAE hosts to manage any new E-series routers as they are brought online. In the case of a disastrous server failure, the offending hardware can simply be removed and replaced as time and resources allow. Also, in regularly scheduled maintenance windows, incremental software upgrades can be achieved in the same fashion.

EASP Deployment

The SDX software includes an API and framework for building Enterprise Access Service Portals (EASPs). Service providers use this API and framework to develop Web applications that their enterprise customers can use to manage their services.

For more information on EASP, see *Chapter 4, Enterprise Services*.

The following sections explain the architecture, component interactions, and deployment for key EASP cases.

Enterprise Portal Architecture

Figure 2-6 shows the basic elements and communication protocols of an EASP.

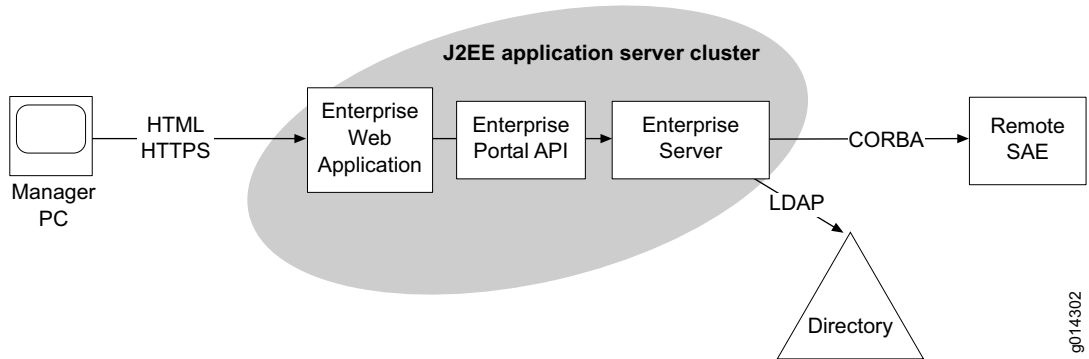


Figure 2-6 EASP elements and communication protocols

EASP Elements

The EASP consists of a server cluster that communicates with the following network elements:

- Directory system – a distributed set of directories with information shadowing and chaining agreements between master and slave servers
- Remote SAE
- Manager PC – a client PC where the enterprise manager runs a Web browser to communicate with the EASP

Internally, the EASP consists of a J2EE application server cluster that implements an Enterprise Portal API, an enterprise Web application that uses this API, and an enterprise server. The enterprise server requires persistent sessions in the cluster. That is, the cluster member that receives the first manager session request must receive all subsequent requests for the same session. The enterprise server works with the same NIC deployments as described in the previous sections.

Communication Protocols

Table 2-1 describes the communication protocols that are used between elements in the EASP network.

Table 2-1 EASP communication protocols

Protocol	Used for Communication Between
HTML/HTTPS (HyperText Markup Language over Secure HyperText Transmission Protocol)	Enterprise manager's Web browser and the enterprise portal Web application running in the EASP
Enterprise Portal API	Enterprise Web application and the enterprise server
CORBA (common object request broker architecture)	Enterprise server and remote SAEs running in a different Web application container than the enterprise server
LDAP	Enterprise server and SDX directories

EASP Deployment Scenario

This section covers deployment and component interactions for EASP deployment as shown in Figure 2-7.

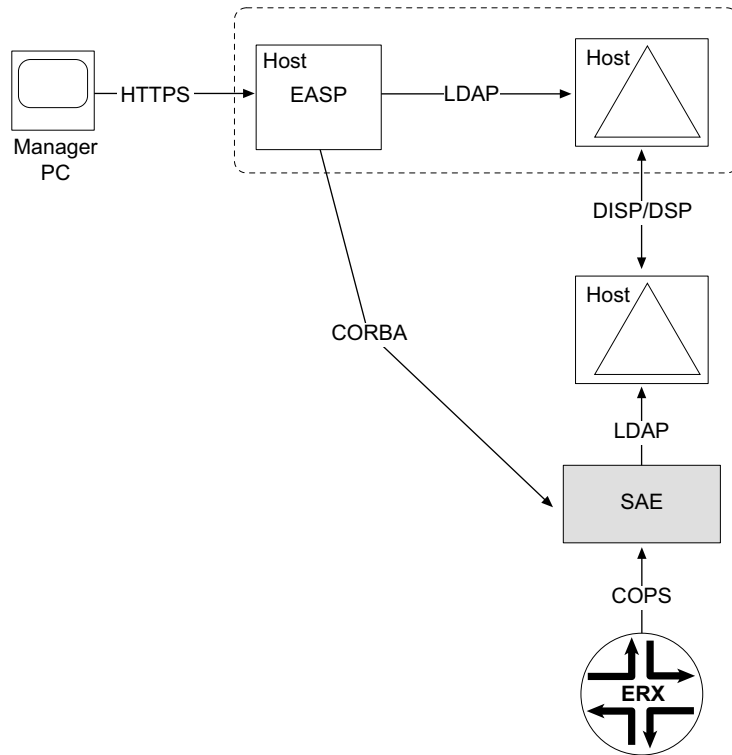


Figure 2-7 EASP deployment

The directory servers are synchronized by means of server-to-server protocols, such as DISP and DSP in the case of X.500 directories, and DirX and equivalent protocols in the case of native LDAP directories, such as Sun ONE Directory Server.

In this configuration, bulk service session requests and implicit subscription reactivation caused by substitution changes are made through replication of directory information. The EASP writes new information to its local directory, and the server-to-server protocols transfer the information to the SAE's local directory. Then the SDX directory eventing system notifies the SAE of the new information, and the SAE reacts by activating and deactivating subscriptions.

The EASP gets feedback on the session state and parameter values of a session using remote procedure calls via the CORBA connection directly to the SAE managing the session.

SDX Gateway Architecture

Figure 2-8 shows the architecture for the SDX gateway. The SDX gateway allows a gateway client—an application that is not part of the SDX network—to interact with SDX components. The SDX gateway supports several Web applications that allow gateway clients to interact with the SDX network. These Web applications communicate with gateway clients through SOAP interfaces.

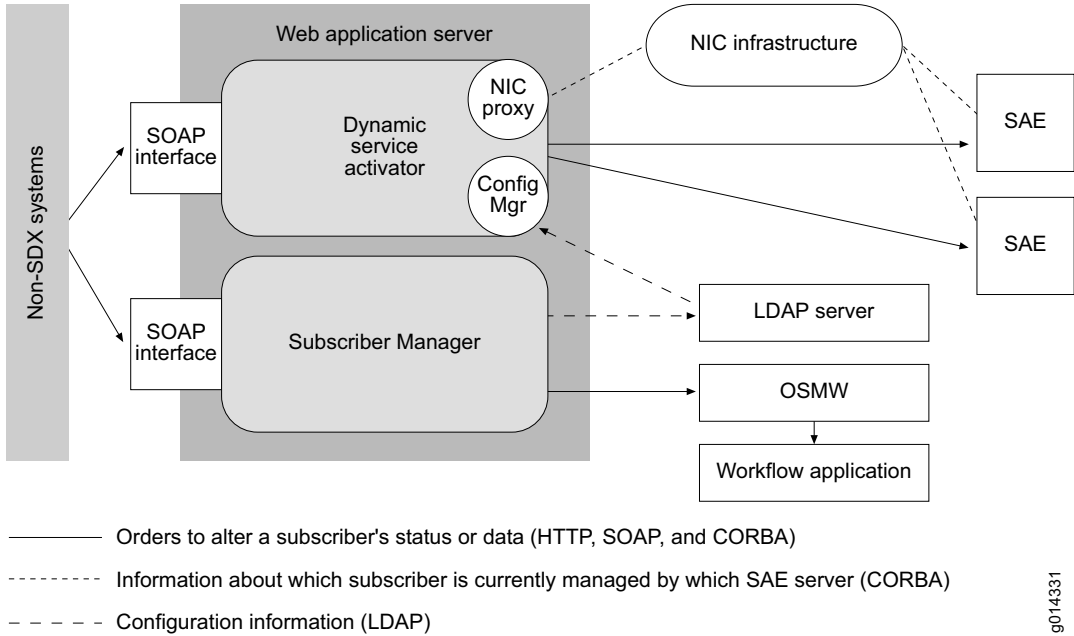



Figure 2-8 SDX gateway architecture

SDX Accounting Deployment

The SDX software allows a variety of accounting deployments. This section shows the standard deployment that we supply, a second option that does not depend on a RADIUS server, and a third option where customers develop their own deployment by choosing a CORBA plug-in.

In the standard SDX deployment, the E-series router and the SAE are clients of the RADIUS accounting server. They pass subscriber accounting information to a designated RADIUS accounting server in an accounting request. The RADIUS accounting server receives the accounting request and creates accounting log files.

Figure 2-9 shows a standard SDX deployment.




g014309

Figure 2-9 Standard SDX accounting deployment



Note: The SDX software works with other AAA RADIUS servers; however, we validate the SDX software only with Merit, Interlink RAD-Series AAA RADIUS Server, or Funk software.


A second option for SDX deployment, shown in Figure 2-10, uses an accounting flat file generated directly by the SAE, with no dependence on a RADIUS server.



g014310

Figure 2-10 An optional SDX accounting deployment

Figure 2-11 illustrates a third possibility for SDX deployment, one in which the customer uses a CORBA plug-in of his or her own choice.



g014311

Figure 2-11 Customer choice for SDX accounting deployment

Workflow Application Deployment Strategies

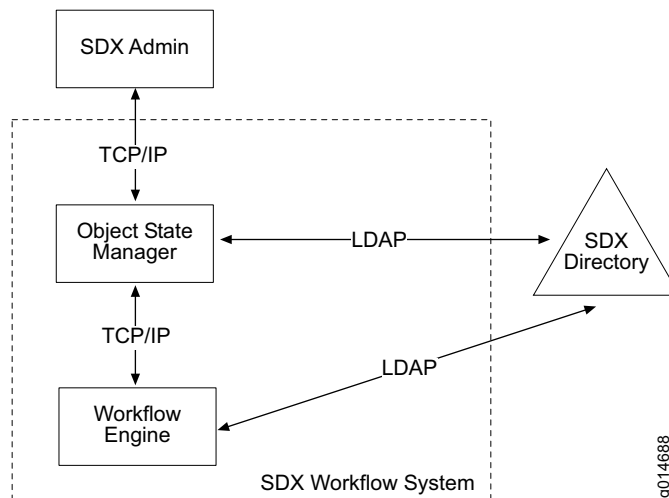
This section describes deployment strategies for the SDX Workflow application. The strategies are divided based on the expected load of the system, integration with external OSSs, and distribution requirements. The following deployment strategies are covered:

- Basic deployment
- Centralized deployment
- High-availability deployment
- Externalized deployment
- Distributed deployment
- Web-based deployment
- Mixed deployment

Basic Deployment

A basic deployment is defined as low load, no external OSS, and a single site with no redundancy.

In this situation, a single OSM manages the state of all transactional objects, and a single workflow engine executes workflows. The OSM can even be located in the same host as the workflow engine. This setup is typical for integration with the DirX directory server for offline processing.



9014688

Figure 2-12 Basic deployment architecture

Because only one site hosts the Workflow application, the transactional objects, which are the directory entries that trigger workflows and whose consistency is governed by the OSM, need to be replicated and accessible in only one of the directory servers, the one used by the OSM.

Centralized Deployment

Centralized deployment is characterized by high load, no external OSS, and a single site with no redundancy.

In this situation (see Figure 2-13), a cluster of workstations execute workflows and manage the state of the transactional objects. To handle the high load, workflow engines are placed in several hosts. Only one OSM is necessary, since no redundancy is required. The OSM is located in one of the hosts that is allocated to one of the workflow engines. Each workflow engine is configured to report to the same OSM, and the OSM is configured to be associated with all workflow engines.

In this case, the SDX component, like the SDX Admin tool, is the client of the Workflow application.

Because only one site hosts the Workflow application, the transactional objects in the SDX directory need to be replicated and accessible in only one of the directory servers, the one used by the OSM.

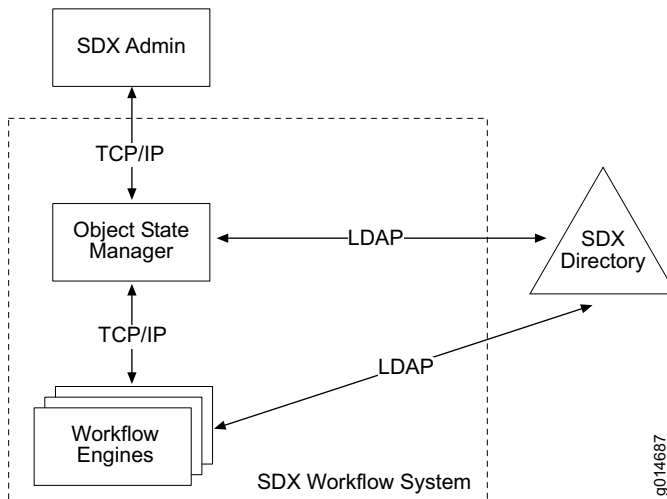


Figure 2-13 Centralized deployment architecture

High-Availability Deployment

High-availability deployment is characterized by high load, no external OSS, and a single site with redundancy.

This deployment is very similar to the centralized deployment; the only difference is the use of multiple OSMs (see Figure 2-14). The OSMs are replicated in different servers to improve the overall availability of the Workflow application. As a bonus, the load can be shared among them.

Like the previous case, an SDX component is the client of the OSMs. The client can issue requests to the various OSMs in a round-robin fashion. For this process to work, each OSM is associated with all workflow engines, each of which reports to a different OSM. With this setup, an engine can be reached via different OSMs, and the load is also shared. If an OSM fails, the reports destined to the OSM are queued in the engine until the OSM comes back up. If an engine fails, the requests are routed to the other OSMs.

Because only one site hosts the Workflow application, the transactional objects in the SDX directory need to be replicated and accessible in only one of the directory servers, the one used by the OSM. For increased availability, you can have more than one directory server. However, you must add directory servers with care so that the additions do not affect the performance of the OSM because of the shadowing agreements of the transactional objects. If you add a directory server, only one server should shadow the transactional objects, the one used by the OSMs. The workflow engines can use all directory servers.

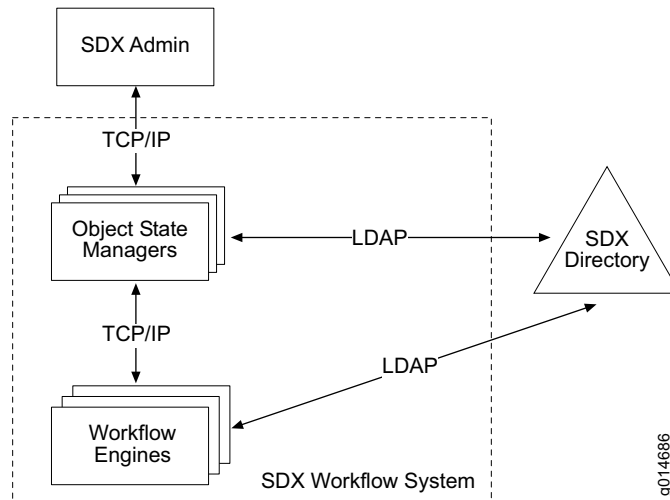


Figure 2-14 High-availability deployment architecture

Externalized Deployment

Externalized deployment is characterized by high load, external OSS, and a single site with no redundancy.

This situation is basically the same as centralized deployment, but an external OSS controls the provisioning process (see Figure 2-15). This case can be used, for example, for a customer care system that concentrates customer requests and communicates with the Workflow application to take care of the provisioning tasks.

Instead of using the regular version of the OSM, the OSMW is preferred, because it provides a friendlier interface based on XML over HTTP. The OSMW is hosted as a Web application in a Web server that is compatible with Java servlets.

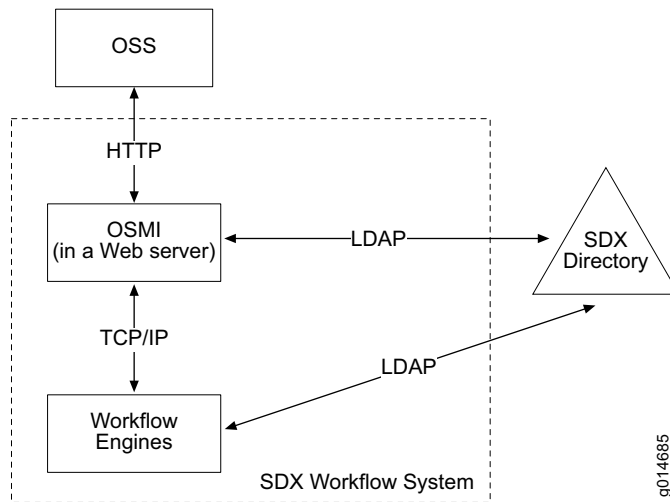


Figure 2-15 Externalized deployment architecture

If availability is an issue, the Web server takes care of it, in which case the usual techniques for Web server replication can be used.

Distributed Deployment

Distributed deployment is characterized by a high load, no external OSS, and multiple sites with redundancy.

This case is equivalent to high-availability deployment, but is extended to multiple sites. A multisite deployment is necessary only if the OSS issues provisioning requests in more than one site; for example, for geographic reasons. This deployment guarantees a consistent view of the global state of all provisioning requests from every point of the service provider network that accesses the directory.

A multisite deployment of the Workflow application uses the SDX directory as the only means to communicate between sites. This means that extra configuration on the service provider administrative network is not necessary for the system components to cooperate.

Provisioning requests cannot trigger workflows across sites; that is, the OSM must be associated with workflow engines located on the same site. The workflow definitions, on the other hand, can be deployed through the directory and be accessed everywhere.

For this setup to be possible, the only requirement from the directory point of view is that the transactional objects be accessible on at least one directory server per site, because they are used to communicate between OSMs in different sites.

Figure 2-16 shows a typical distributed deployment.

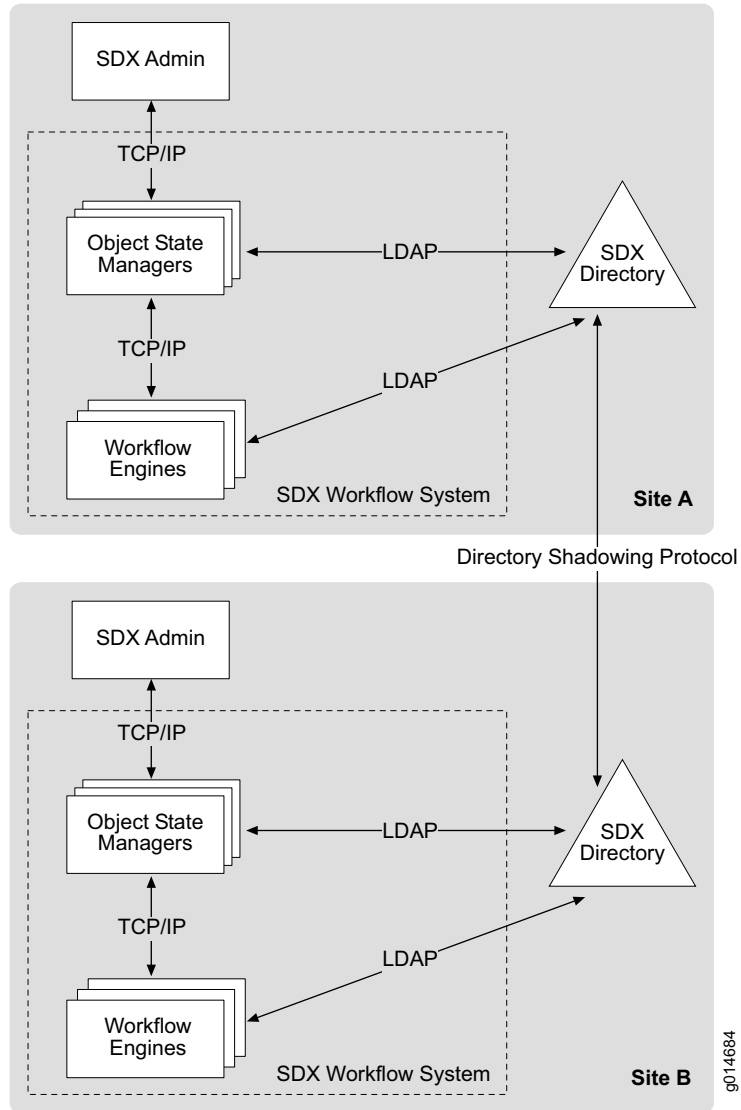


Figure 2-16 Distributed deployment architecture

Requests issued by the SDX component on site A cause workflows to be run on site A only. The same is true for site B. However, if requests are issued on different sites for the same object—for example, the same subscription—the OSMs that handle these requests coordinate via the directory to ensure that these requests are properly executed. This coordination means that the SDX component on site B can see the state of

an ongoing or completed process without needing to communicate with the other SDX component on site A.

Other sites can be easily added just by providing a directory server that replicates and accesses the transactional objects, which allows for a gradual expansion of the system.

We do not recommend that the transactional objects be included in any shadowing agreement. Instead, the transactional objects must reside in the master directory server and all accesses to it be done via referrals.

Web-Based Deployment

Web-based deployment is characterized by high load, external OSS, and multiple sites with redundancy.

Web-based deployment is the natural extension of the distributed deployment architecture in which, instead of having an SDX component as an OSS, Web-based software (for example, a portal) is used to allow customer self-care. The only difference from the distributed scenario is use of the OSMW instead of OSM. It can be deployed together with the Web servers that host the OSS application or on dedicated servers, according to security or other requirements.

Mixed Deployment

Mixed deployment is characterized by high load, external and internal OSS, and a single site with redundancy.

Mixed deployment is similar to the high-availability scenario, but it also includes the externalized case (see Figure 2-17). This architecture solves the problem of using both external (Web-based, for example) and internal OSS applications.

The important detail here is the use of private workflow engines for the OSMW, because the external OSS application does not communicate with the OSMs that communicate with the internal OSS application.

The natural extension of this case is to distribute it on multiple sites. The idea is the same, but the distributed and Web-based cases are mixed.

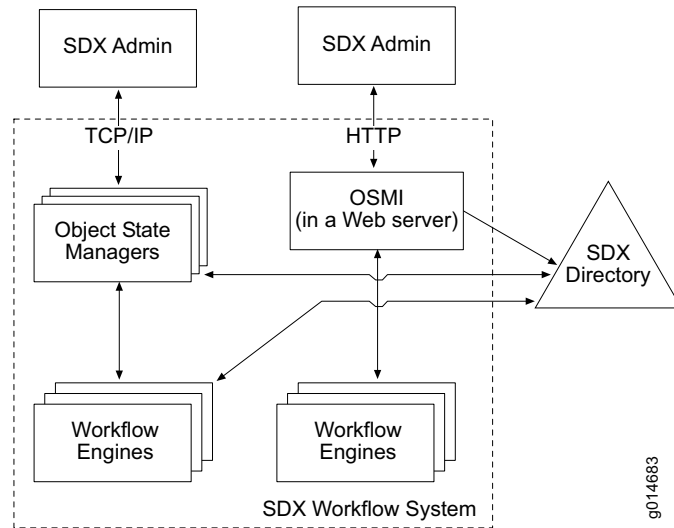


Figure 2-17 Mixed-deployment architecture (most protocol labels omitted for clarity)

