



## CHAPTER 2

# Deploy the Access Control Servers

---

This chapter discusses topics that you should consider before deploying Cisco Secure Access Control Server, hereafter referred to as ACS.

This document does not describe the software installation procedure for ACS or the hardware installation procedure for the ACS SE. For detailed installation information, refer to:

- *Installation Guide for Cisco Secure ACS for Windows Release 4.2*, available on Cisco.com at:  
[http://www.cisco.com/en/US/docs/net\\_mgmt/cisco\\_secure\\_access\\_control\\_server\\_for\\_windows/4.2/installation/guide/windows/IGwn42.html](http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.2/installation/guide/windows/IGwn42.html)
- *Installation Guide for Cisco Secure ACS Solution Engine Release 4.2*, available on Cisco.com at:  
[http://www.cisco.com/en/US/docs/net\\_mgmt/cisco\\_secure\\_access\\_control\\_server\\_for\\_solution\\_engine/4.2/installation/guide/solution\\_engine/SE42.html](http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_solution_engine/4.2/installation/guide/solution_engine/SE42.html)



### Note

For more detailed information on deploying ACS, see the *Cisco Secure Access Control Server Deployment Guide* at:

[http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5712/ps2086/prod\\_white\\_paper0900aecd80737943.pdf](http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5712/ps2086/prod_white_paper0900aecd80737943.pdf)

---

This chapter contains:

- [Determining the Deployment Architecture](#), page 2-1
- [Determining How Many ACSs to Deploy \(Scalability\)](#), page 2-11
- [Deploying ACS Servers to Support Server Failover](#), page 2-13
- [Deploying ACS in a NAC/NAP Environment](#), page 2-15
- [Additional Topics](#), page 2-16

## Determining the Deployment Architecture

How your enterprise network is configured and the network topology are likely to be the most important factors in deploying ACS.

This section discusses:

- **Access types**—How users will access the network (through wireless access, LAN access through switches, and so on) and the security protocols used to control user access; for example, RADIUS, EAP-TLS, Microsoft Active Directory, and so on.
- **Network architecture**—How the network is organized (centrally through campus LANs, regional LANs, WLANs, and so on).

This section contains:

- [Access Types, page 2-2](#)
- [Placement of the RADIUS Server, page 2-11](#)

## Access Types

This section contains:

- [Wired LAN Access, page 2-2](#)
- [Wireless Access Topology, page 2-5](#)
- [Dial-up Access Topology, page 2-9](#)

## Wired LAN Access

You can use wired LAN access in a small LAN environment, a campus LAN environment, or a regionally or globally dispersed network. The number of users determines the size of the LAN or WLAN:

Size	Users
small LAN	1 to 3,000
medium-sized LAN	3,000 to 25,000
large LAN	25,000 to 50,000
very large LAN or WLAN	over 50,000

The wired LAN environment uses the following security protocols:

- **RADIUS**—RADIUS is used to control user access to wired LANs. In broadcast or switch-based Ethernet networks, you can use RADIUS to provide virtual LAN identification information for each authorized user.
- **EAP**—Extensible Authentication Protocol (EAP), provides the ability to deploy RADIUS into Ethernet network environments. EAP is defined by Internet Engineering Task Force (IETF) RFC 2284 and the IEEE 802.1x standards.

The 802.1x standard, also known as EAP over LAN (EAPoL), concerns the part of the wider EAP standard that relates to broadcast media networks. Upon connection, EAPoL provides a communications channel between an end user on a client LAN device to the AAA server through the LAN switch. The functionality is similar to what Point-to-Point Protocol (PPP) servers on point-to-point links provide.

By supporting complex challenge-response dialogues, EAP facilitates the user-based authentication demands of both conventional one-way hashed password authentication schemes such as Challenge Handshake Authentication Protocol (CHAP) and of more advanced authentication schemes such as Transport Layer Security (TLS), or digital certificates.

- **EAP-TLS**—Extensible Authentication Protocol-Transport Layer Security (EAP-TLS). EAP-TLS uses the TLS protocol (RFC 2246), which is the latest version of the Secure Socket Layer (SSL) protocol from the IETF. TLS provides a way to use certificates for user and server authentication and for dynamic session key generation.
- **PEAP**— Protected Extensible Authentication Protocol (PEAP) is an 802.1x authentication type for wireless LANs (WLANs). PEAP provides strong security, user database extensibility, and support for one-time token authentication and password change or aging. PEAP is based on an Internet Draft that Cisco Systems, Microsoft, and RSA Security submitted to the IETF.


### Small LAN Environment

In a small LAN environment (a LAN containing up to 3,000 users; see [Figure 2-1](#)), a single ACS is usually located close to the switch and behind a firewall. In this environment, the user database is usually small because few switches require access to ACS for AAA, and the workload is small enough to require only a single ACS.

However, you should still deploy a second ACS server for redundancy, and set up the second ACS server as a replication partner to the primary server; because, losing the ACS would prevent users from gaining access to the network. In [Figure 2-1](#), an Internet connection via firewall and router are included because these are likely to be features of such a network; but, they are not strictly related to the Cisco Catalyst AAA setup or required as part of it.

You should also limit access to the system hosting the ACS to as small a number of users and devices as necessary. As shown in [Figure 2-1](#), you set access by connecting the ACS host to a private LAN segment on the firewall. Access to this segment is limited only to the Cisco Catalyst Switch client and those user machines that require HTTP access to the ACS for administrative purposes. Users should not be aware that the ACS is part of the network.

**Figure 2-1 ACS Server in a Small LAN Environment**



### Campus LAN

You can use ACS for wired access in a campus LAN. A campus LAN is typically divided into subnets. [Figure 2-2](#) shows an ACS deployment in a wired campus LAN.

Figure 2-2 ACS in a Campus LAN

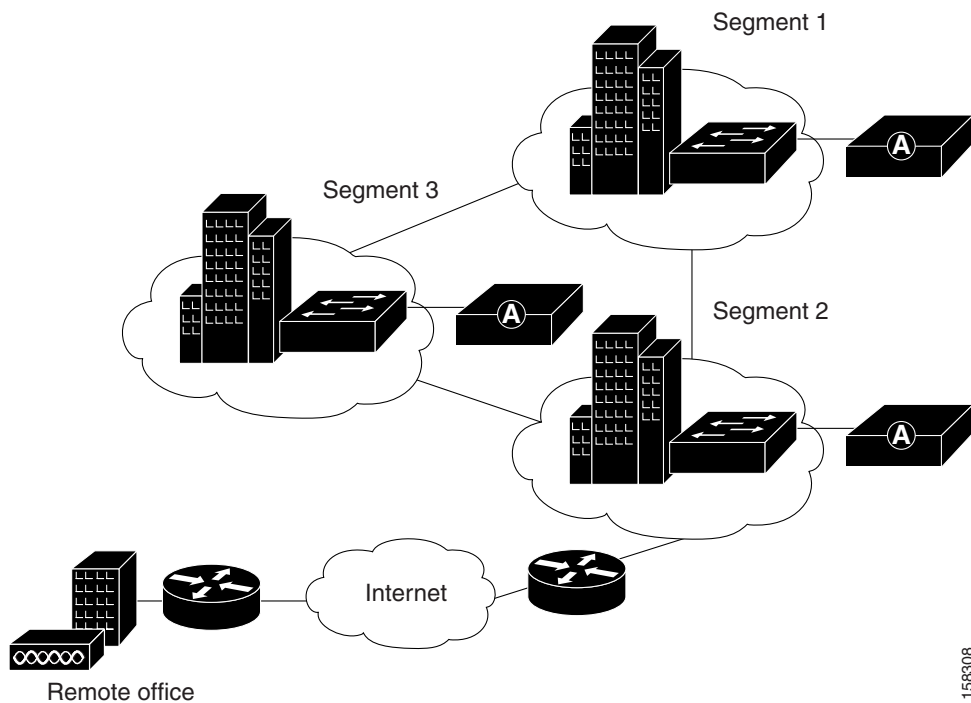


Figure 2-2 shows a possible distribution of ACS in a wired campus LAN. In this campus LAN, buildings are grouped into three segments. Each segment consists of 1 to 3 buildings and all the buildings in the segment are on a common LAN. All interbuilding and intersegment network connections use one-gigabyte fiber-optic technology. Primary network access is through switch ports over wired Ethernet.

You use ACS to provide RADIUS authentication for the network access servers, and you configure it to use an external database. One ACS is deployed for each segment of 5 to 10 buildings. A Cisco LocalDirector content switch is placed before each ACS for load balancing and failover.

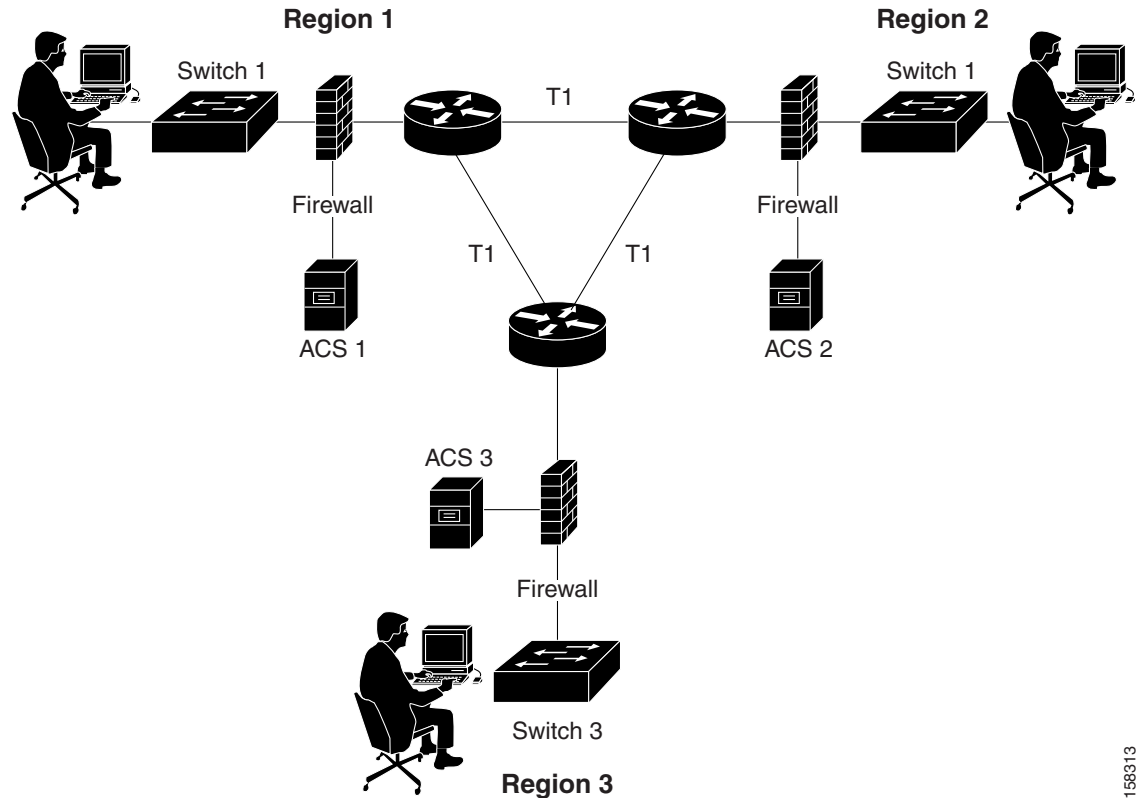
### Geographically Dispersed Wired LAN

In a larger network that is geographically dispersed, speed, redundancy, and reliability are important in determining whether to use a centralized ACS service or a number of geographically dispersed ACS units. As with many applications, AAA clients rely on timely and accurate responses to their queries. Network speed is an important factor in deciding how to deploy ACS; because delays in authentication that the network causes can result in timeouts at the client side or the switch.

A useful approach in large extended networks, such as for a globally dispersed corporation, is to have at least one ACS deployed in each major geographical region. Depending on the quality of the WAN links, these servers may act as backup partners to servers in other regions to protect against failure of the ACS in any particular region.

Figure 2-3 shows ACS deployed in a geographically dispersed wired LAN. In the illustration, Switch 1 is configured with ACS 1 as its primary AAA server but with ACS 2 of Region 2 as its secondary. Switch 2 is configured with ACS 2 as its primary but with ACS 3 as its secondary. Likewise, Switch 3 uses ACS 3 as its primary but ACS 1 as its secondary. Using a local ACS as the primary AAA server minimizes AAA WAN traffic. When necessary, using the primary ACS from another region as the secondary further minimizes the number of ACS units.

Figure 2-3 ACS in a Geographically Dispersed LAN



158313

## Wireless Access Topology

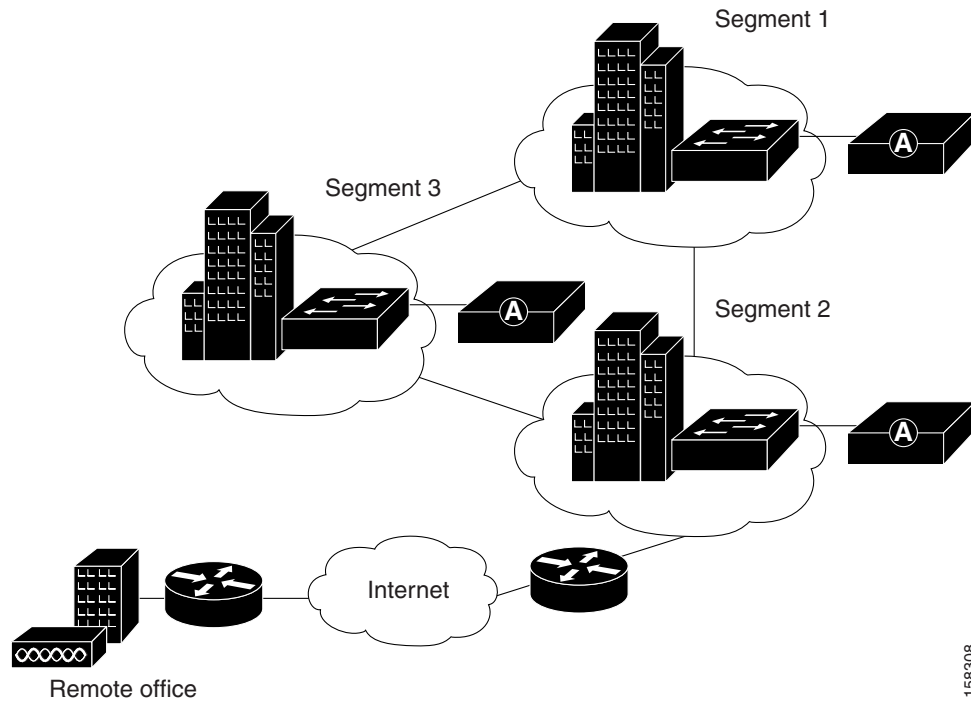
A wireless access point (AP), such as the Cisco Aironet series, provides a bridged connection for mobile end-user clients into the LAN. Authentication is absolutely necessary, due to the ease of access to the AP. Encryption is also necessary because of the ease of eavesdropping on communications.

Scaling can be a serious issue in the wireless network. The mobility factor of the WLAN requires considerations similar to those given to the dial-up network. Unlike the wired LAN, however, you can more readily expand the WLAN. Though WLAN technology does have physical limits as to the number of users who can connect via an AP, the number of APs can grow quickly. As with the dial-up network, you can structure your WLAN to allow full access for all users, or provide restricted access to different subnets among sites, buildings, floors, or rooms. This capability raises a unique issue with the WLAN: the ability of a user to roam among APs.

### Simple WLAN

A single AP might be installed in a simple WLAN (Figure 2-4). Because only one AP is present, the primary issue is security. An environment such as this generally contains a small user base and few network devices. Providing AAA services to the other devices on the network does not cause any significant additional load on the ACS.


Figure 2-4 Simple WLAN



158308

### Campus WLAN

In a WLAN where a number of APs are deployed, as in a large building or a campus environment, your decisions on how to deploy ACS become more complex. Depending on the processing needs of the installation, all of the APs might be on the same LAN. [Figure 2-5](#) shows all APs on the same LAN; however, the APs might also be distributed throughout the LAN, and connected via routers, switches, and so on.

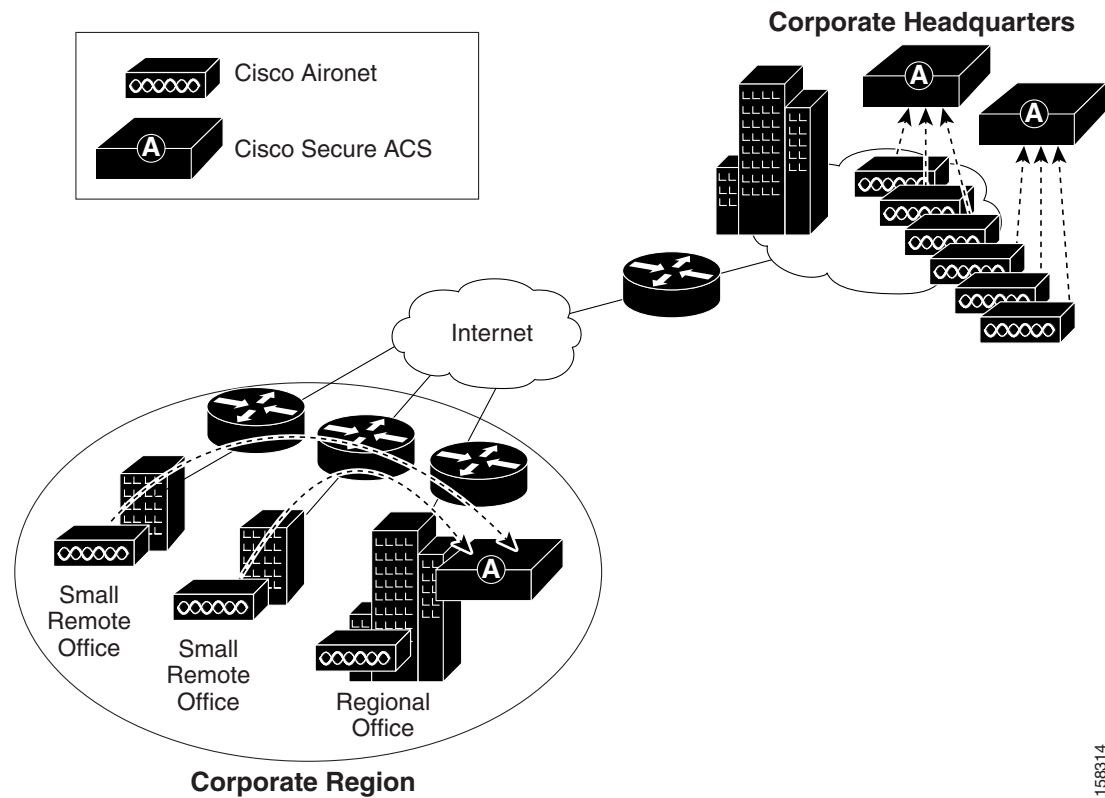
**Figure 2-5** Campus WLAN

### Regional WLAN Setting

In a given geographical or organizational region, the total number of users might or might not reach a critical level for a single ACS. Small offices would not qualify for separate installations of ACSs and a regional office might have sufficient reserve capacity. In this case, the small offices can authenticate users across the WAN to the larger regional office. Once again, you should determine that this does not pose a risk to the users in the remote offices. Assess critical connectivity needs against the reliability and throughput to the central ACS.

Figure 2-6 shows a regional WLAN.

Figure 2-6 ACS in a Regional WLAN




### Large Enterprise WLAN Setting

In a very large geographically dispersed network (over 50,000 users), access servers might be located in different parts of a city, in different cities, or on different continents. If network latency is not an issue, a central ACS might work; but, connection reliability over long distances might cause problems. In this case, local ACSs may be preferable to a central ACS.

If the need for a globally coherent user database is most important, database replication or synchronization from a central ACS may be necessary. For information on database replication considerations, see [Database Replication Considerations, page 2-13](#) and [Database Synchronization Considerations, page 2-14](#). Authentication by using external databases, such as a Windows user database or the Lightweight Directory Access Protocol (LDAP), can further complicate the deployment of distributed, localized ACSs.

Figure 2-7 shows ACS installations in a geographically dispersed network that contains many WLANs.

**Figure 2-7 ACS in a Geographically Dispersed WLAN**



For the model in Figure 2-7, the location of ACS depends on whether all users need access on any AP, or require only regional or local network access. Along with database type, these factors control whether local or regional ACSs are required, and how database continuity is maintained. In this very large deployment model (over 50,000 users), security becomes a more complicated issue, too.

#### Additional Considerations for Deploying ACS in a WLAN Environment

You should also consider the following when deploying ACS in a WLAN environment, consider if:

- Wireless is secondary to wired access, using a remote ACS as a secondary system is acceptable.
- Wireless is the primary means of access, put a primary ACS in each LAN.
- The customer uses ACS for user configuration, data replication is critical.

## Dial-up Access Topology


Until recently, dial-up access was the most prevalent method for providing remote access to network resources. However, DSL access and access through VPNs have largely replaced dial-up access through modems.

ACS is still used in some LAN environments to provide security for dial-up access. You can provide dial-up access for a small LAN or for a large dial-in LAN.

#### Small Dial-Up Network Access

In the small LAN environment, see Figure 2-8, network architects typically place a single ACS internal to the AAA client, which a firewall and the AAA client protect from outside access. In this environment, the user database is usually small; because, few devices require access to the ACS for authentication, authorization and accounting (AAA), and any database replication is limited to a secondary ACS as a backup.

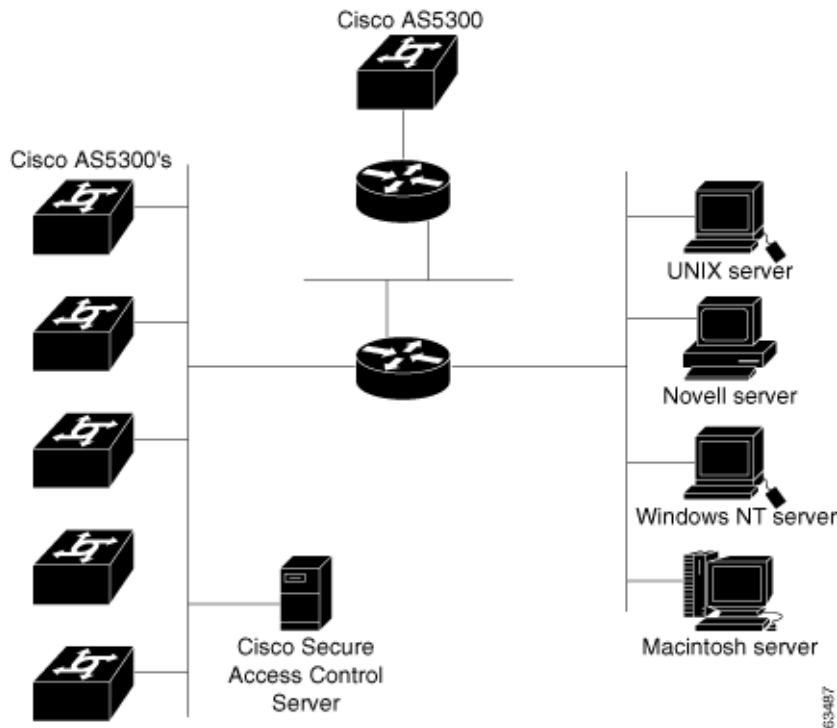
Figure 2-8 Small Dial-up Network



### Large Dial-Up Network Access

In a larger dial-in environment, a single ACS with a backup may be suitable, too. The suitability of this configuration depends on network and server access latency. Figure 2-9 shows an example of a large dial-in network. In this scenario, the addition of a backup ACS is recommended.

Figure 2-9 Large Dial-up Network



## Placement of the RADIUS Server

From a practical standpoint, the RADIUS server should be inside the general network, preferably within a secure subnet designated for servers, such as DHCP, Domain Name System (DNS), and so on. You should avoid requiring RADIUS requests to travel over WAN connections because of possible network delays and loss of connectivity. Due to various reasons, this type of configuration is not always possible; for example, with small remote subnets that require authentication support from the enterprise.

You must also consider backup authentication. You may use a system that is dedicated as the RADIUS secondary. Or, you may have two synchronized systems that each support a different network segment but provide mutual backup if one fails. Refer to the documentation for your RADIUS server for information on database replication and the use of external databases.

## Determining How Many ACSs to Deploy (Scalability)

A number of factors affect the scalability of an ACS installation (that is, how effectively each ACS can process user access requests) and how many ACS servers you should deploy in the network.

For detailed information on scalability considerations, see the following white papers on ACS deployment, which are available on Cisco.com at:

[http://www.cisco.com/en/US/products/sw/secursw/ps2086/prod\\_white\\_papers\\_list.html](http://www.cisco.com/en/US/products/sw/secursw/ps2086/prod_white_papers_list.html)

- *Building a Scalable TACACS+ Device Management Framework*
- *Catalyst Switching and ACS Deployment Guide*
- *Deploying Cisco Secure ACS for Windows in Cisco Aironet Environment*
- *EAP-TLS Deployment Guide for Wireless LAN Networks*
- *Guidelines for Placing ACS in the Network*

This section contains:

- [Number of Users, page 2-11](#)
- [Number of Network Access Servers, page 2-12](#)
- [LAN Versus WAN Deployment \(Number of LANs in the Network\), page 2-12](#)
- [WAN Latency and Dependability, page 2-12](#)
- [Determining How Many ACS Servers to Deploy in Wireless Networks, page 2-13](#)

## Number of Users

In all topologies, the number of users is an important consideration. For example, assuming that an ACS can support 21,000 users, if a wireless access point can support 10 users, then a given ACS could support 2,100 wireless access points in a WLAN environment.

The size of the LAN or WLAN is determined by the number of users who use the LAN or WLAN:

Size	Users
Small LAN	1 to 3,000
Medium-sized LAN	3,000 to 25,000
Large LAN	25,000 to 50,000
Very large LAN or WLAN	Over 50,000

For a detailed formula, see the white paper *Deploying Cisco Secure ACS for Windows in Cisco Aironet Environment*, which is available on Cisco.com at this location:

[http://www.cisco.com/en/US/products/sw/secursw/ps2086/prod\\_white\\_papers\\_list.html](http://www.cisco.com/en/US/products/sw/secursw/ps2086/prod_white_papers_list.html)

## Number of Network Access Servers

An ACS can support up 5,000 discrete network access servers (NASs). You can use the multi-NAS capability of ACS to increase this number.

## LAN Versus WAN Deployment (Number of LANs in the Network)

In general, you should provide one ACS server per LAN. If a backup ACS is required, the backup ACS may reside on the same LAN or can be an ACS on another LAN.

## WAN Latency and Dependability

The distance between LANs in a large network (25,000 to 50,000 users) is also a consideration.

If the network is centralized, one primary ACS and one secondary ACS might be sufficient.

If the network is geographically dispersed, the number of ACS servers required varies with the needs of the regions. For example:

- Some regions may not need a dedicated ACS.
- Larger regions (regions with over 10,000 users), such as corporate headquarters, might need several ACSs.

The distance between subnets is also a consideration. If subnets are close together, the connections will be more reliable, and fewer ACS servers will be needed. Adjacent subnets could serve other buildings with reliable connections. If the subnets are farther apart, more ACS servers might be needed.

The number of subnets and the number of users on each subnet is also a factor. For example, in a WLAN, a building may have 400 potential users and the same subnet might comprise four buildings. One ACS assigned to this subnet will service 1,600 users (about one tenth of the number of current users). Other buildings could be on adjacent subnets with reliable WAN connections. ACSs on adjacent subnets could then be used as secondary systems for backup.

If the WAN connections between buildings in this subnet are short, reliable, and pose no issue of network latency, two ACSs can service all of these buildings and all the users. At 40-percent load, one ACS would take half of the access points as the primary server, and the other ACS would take the remaining APs. Each ACS would provide backup for the other. Again, at 40-percent load, a failure of one ACS would

only create an 80-percent load on the other ACS for the duration of the outage. If the WAN is not suitable for authentication connections, we recommend using two or more ACSs on the LAN in a primary or secondary mode or load balanced.

## Determining How Many ACS Servers to Deploy in Wireless Networks

In planning how many ACS servers to deploy in a wireless network, consider:

- The location and number of access points. For example, with 4,200 APs:
  - One ACS could handle half of the APs as primary server.
  - Other ACSs could handle the remaining APs.
- The number of EAP-TLS clients together with EAP-TLS authentications per second
- The number of clients
- Scalability with different protocols

For example, if you use EAP-TLS, you will need more ACS servers; but, if you use PEAP, you will need fewer. EAP-TLS is slower than PEAP due to public-key infrastructure (PKI) processing time.

For a detailed formula that you can use to calculate the number of ACS servers required in a wireless network, see the white paper titled *Deploying Cisco Secure ACS for Windows in an Aironet Environment*, available on Cisco.com at:

[http://www.cisco.com/en/US/products/sw/secursw/ps2086/prod\\_white\\_papers\\_list.html](http://www.cisco.com/en/US/products/sw/secursw/ps2086/prod_white_papers_list.html)

## Deploying ACS Servers to Support Server Failover

This section discusses deployment topologies for implementing server failover. This section contains:

- [Load Balancing and Failover, page 2-13](#)
- [Database Replication Considerations, page 2-13](#)
- [Database Synchronization Considerations, page 2-14](#)

### Load Balancing and Failover

To implement load balancing, you can set up user groups and then assign groups to a specific RADIUS server (usually the nearest RADIUS server).

### Database Replication Considerations

Database replication replicates selected database information, such as user and group information, from a primary ACS to one or more ACS backups or clients. The following aspects of replication are configurable with ACS:

- **Configuration components for replication**—What is replicated.
- **Replication scheduling**—When replication occurs.
- **Replication frequency**—How often systems are replicated.
- **Replication partners**—Which systems are replicated.

- **Client configuration**—How to configure the client.
- **Reports and event (error) handling**—What information to include in the logs.

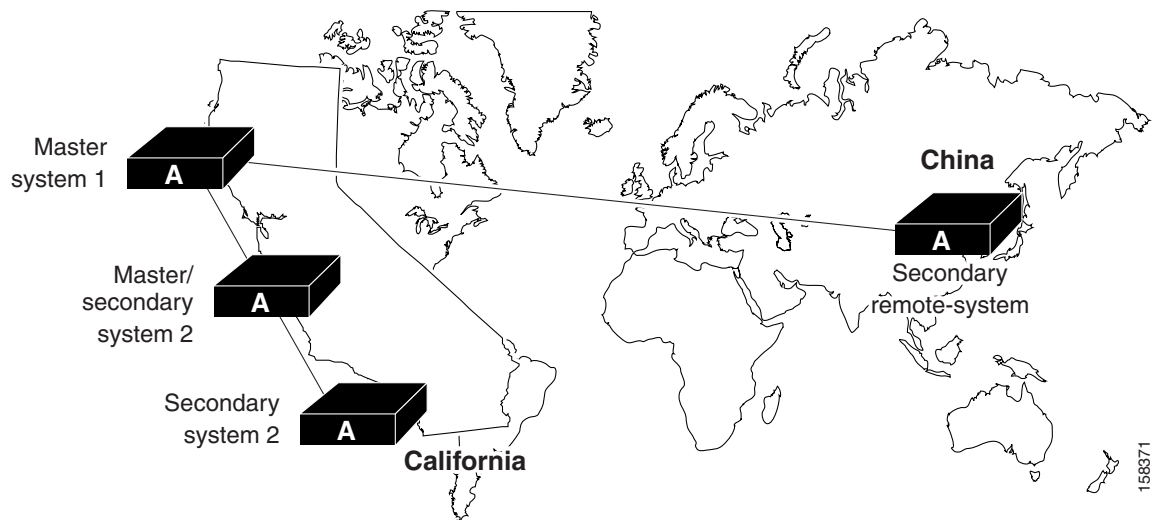
## Replication Design

Because database replication in a ACS is a top-down approach, using the cascade method minimizes replication-induced downtime on the master server. If the primary server is not used for authentication services, but for database maintenance only, the cascade method may not be as critical.

However, when traveling across time zones, particularly international time zones, it may be necessary to use the cascade method going to remote secondaries. In this case, when you configure database replication on the Database replication setup page, click *At specific times* instead of *Automatically triggered cascade*.

Use the automatically triggered cascade method so that local replication occurs during a time that will minimize the impact on user authentication. During these long-distance replications, replicating to the backup or secondary server first also helps reduce this impact. [Figure 2-10](#) shows a hypothetical deployment for replication where each region has a primary and a secondary ACS deployed. In this scenario, replication is made to the secondary servers to avoid replication downtime to the primary, but, may not be needed if the primary is used mainly for database maintenance but not for authentication.

**Figure 2-10 ACS Database Replication Scenario**



## Database Synchronization Considerations

An alternative to database replication is the use of Relational Database Management System (RDBMS) synchronization. You use the RDBMS synchronization feature to update the ACS user database with information from an Open Database Connectivity (ODBC)-compliant data source. The ODBC-compliant data source can be the RDBMS database of a third-party application. It can also be an intermediate file or database that a third-party system updates. Regardless of where the file or database resides, ACS reads the file or database via the ODBC connection. RDBMS synchronization supports addition, modification, and deletion for all data items it can access.

## Deploying ACS in a NAC/NAP Environment

You can deploy ACS in a Cisco Network Admission Control and Microsoft Network Access Protection (NAC/NAP) environment. In the NAC/NAP environment, NAP client computers authorize with ACS by using EAP over UDP (EoU) or EAP over 802.1x.

Table 2-1 describes the components of a NAC/NAP deployment.

**Table 2-1** *Components of a NAC/NAP Deployment*

Component	Description
NAP client	A computer running Windows Vista or Windows Server 2008. NAP clients send their health credentials as Statements of Health (SoHs) or as a health certificate.
NAP agent	A process running on a NAP client that sends SoHs or health certificates to ACS.
Network access devices	Cisco devices through which you can access the network, such as routers, switches, wireless access points, and VPN concentrators.
ACS	Cisco AAA server product.
Network Policy Server (NPS)	A Microsoft server that validates health certificates from NAP clients and provides remediation instructions if needed.
Health Registration Authority	A Microsoft certificate server that obtains health certificates on behalf of NAP clients from a public key infrastructure (PKI).
Policy Servers	Servers that provide current system health state for Microsoft NPSs.

When a NAP client connects, it uses a NAP agent to send ACS one of the following:

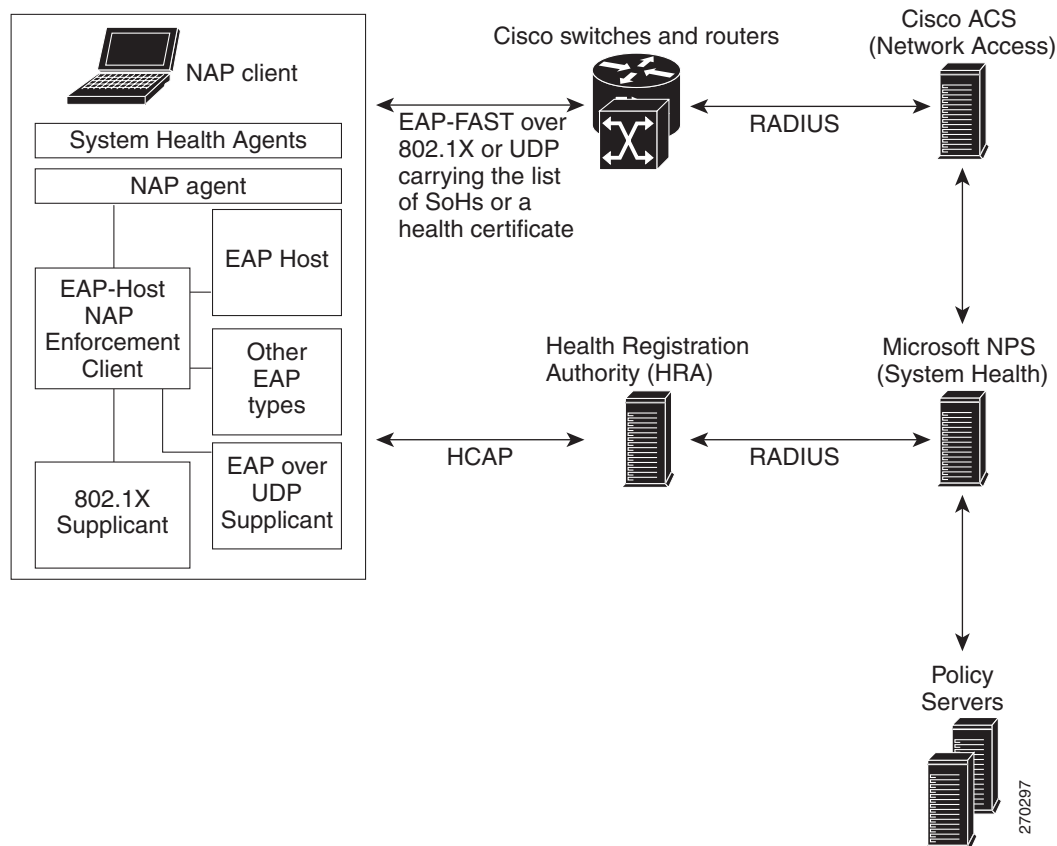
- A list of SoHs.
- A certificate that the client has received from a Microsoft Health Registration Authority (HRA).

The ACS host validates the client credentials. If the NAP agent sends a:

- List of SoHs, the ACS sends the list to a Microsoft NPS by using the Cisco Host Credentials Authorization Protocol (HCAP). The NPS evaluates the SoHs. The ACS then sends an appropriate NAP to the network access device (switch, router, VPN, and so on) to grant the authorized level of access to the client.
- Health certificate rather than a list of SoHs, then ACS validates the certificate as the EAP-FAST session is established to determine the overall health of the client. The ACS then sends the appropriate NAP to the network to grant the authorized level of access to the client.

Figure 2-11 illustrates the architecture of a NAC/NAP network.

**Figure 2-11 NAC/NAP Deployment Architecture**



## Additional Topics

This section describes additional topics to consider when deploying ACS. This section contains:

- [Remote Access Policy, page 2-16](#)
- [Security Policy, page 2-17](#)
- [Administrative Access Policy, page 2-17](#)
- [Database Considerations, page 2-19](#)
- [Network Latency and Reliability, page 2-19](#)

## Remote Access Policy

Remote access is a broad concept. In general, it defines how the user can connect to the LAN, or from the LAN to outside resources (that is, the Internet). Connectivity is possible in many ways: dial-in, ISDN, wireless bridges, and secure Internet connections. Each method incurs its own advantages and disadvantages, and provides a unique challenge to providing AAA services. In addition to the method of

access, other decisions can also affect how ACS is deployed; these include specific network routing (access lists), time-of-day access, individual restrictions on AAA client access, access control lists (ACLs), and so on.

You can implement remote-access policies for employees who telecommute, or mobile users who dial in over ISDN or a public switched telephone network (PSTN). Such policies are enforced at the corporate campus with ACS and the AAA client. Inside the enterprise network, remote-access policies can control wireless access by individual employees.

ACS remote-access policies provide control by using central authentication and authorization of remote users. The Cisco user database maintains all user IDs, passwords, and privileges. You can download ACS policies in the form of ACLs to network access servers such as the Cisco AS5300 Network Access Server, or by allowing access during specific periods, or on specific access servers.

Remote-access policies are part of the overall Cisco corporate security policy.

## Security Policy

Every organization that maintains a network should develop a security policy for the organization. The sophistication, nature, and scope of your security policy directly affect how you deploy ACS.

For more information about developing and maintaining a comprehensive security policy, refer to these documents:

- [Network Security Policy: Best Practices White Paper](#)
- [Cisco IOS Security Configuration Guide](#)

## Administrative Access Policy

Managing a network is a matter of scale. Providing a policy for administrative access to network devices depends directly on the size of the network and the number of administrators required to maintain the network. A network device can be authenticated locally; but, this ability is not scalable. The use of network management tools can help in large networks (25,000 to 50,000 users); but, if local authentication is used on each network device, the policy usually entails a single login on the network device. A single login on the network device does not provide adequate network device security.

ACS provides a centralized administrator database, and you can add or delete administrators at one location. TACACS+ is the recommended AAA protocol for controlling AAA client administrative access because of its ability to provide per-command control (command authorization) of AAA client administrator access to the device. RADIUS is not well suited for this purpose because of the one-time transfer of authorization information at the time of initial authentication.

The type of access is also an important consideration. In the case of different administrative access levels to the AAA clients, or if a subset of administrators is to be limited to certain systems, you can use ACS with command authorization per network device to restrict network administrators as necessary. Using local authentication restricts the administrative access policy to no login on a device or by using privilege levels to control access.

Controlling access by means of privilege levels is cumbersome and not very scalable. Such control requires altering the privilege levels of specific commands on the AAA client device and defining specific privilege levels for the user login. You can easily create more problems by editing command privilege levels. Using command authorization on ACS does not require that you alter the privilege level of controlled commands. The AAA client sends the command to ACS to be parsed and ACS determines whether the administrator has permission to use the command. The use of AAA allows authentication on any AAA client for any user on ACS and limits access to these devices on a per-AAA-client basis.

A small network with a small number of network devices may require only one or two individuals to administer it. Local authentication on the device is usually sufficient. If you require more granular control than what authentication can provide, some means of authorization is necessary. As discussed earlier, controlling access by using privilege levels can be cumbersome. ACS reduces this problem.

In large enterprise networks, with many devices to administer, the use of ACS practically becomes a necessity. Because administration of many devices requires a larger number of network administrators, with varying levels of access, the use of local control is simply not a viable way to track network-device configuration changes that are required when changing administrators or devices.

The use of network management tools, such as CiscoWorks, helps to ease this burden; but, maintaining security is still an issue. Because ACS can comfortably handle up to 300,000 users, the number of network administrators that ACS supports is rarely an issue. If a large remote-access population is using RADIUS for AAA support, the corporate IT team should consider separate TACACS+ authentication by using ACS for the administrative team. Separate TACACS+ authentication would isolate the general user population from the administrative team and reduce the likelihood of inadvertent access to network devices. If the use of TACACS+ is not a suitable solution, using TACACS+ for administrative (shell or exec) logins, and RADIUS for remote network access, provides sufficient security for the network devices.

## Separation of Administrative and General Users

You should prevent the general network user from accessing network devices. Even though the general user may not intend to gain unauthorized access, inadvertent access could accidentally disrupt network access. AAA and ACS provide the means to separate the general user from the administrative user.

The easiest and recommended method to perform such separation is to use RADIUS for the general remote-access user and TACACS+ for the administrative user. One issue is that an administrator may also require remote network access, like the general user. If you use ACS, this issue poses no problem. The administrator can have RADIUS and TACACS+ configurations in ACS. By using authorization, RADIUS users can set PPP (or other network access protocols) as the permitted protocol. Under TACACS+, only the administrator would be configured to have shell (exec) access.

For example, if the administrator is dialing in to the network as a general user, a AAA client would use RADIUS as the authenticating and authorizing protocol, and the PPP protocol would be authorized. In turn, if the same administrator remotely connects to a AAA client to make configuration changes, the AAA client would use the TACACS+ protocol for authentication and authorization. Because this administrator is configured on ACS with permission for shell under TACACS+, the administrator would be authorized to log in to that device. This does require that the AAA client have two separate configurations on ACS, one for RADIUS and one for TACACS+.

An example of a AAA client configuration under IOS that effectively separates PPP and shell logins is:

```
aaa new-model
tacacs-server host ip-address
tacacs-server key secret-key
radius-server host ip-address
radius-server key secret-key
aaa authentication ppp default group radius
aaa authentication login default group tacacs+ local
aaa authentication login console none
aaa authorization network default group radius
aaa authorization exec default group tacacs+ none
aaa authorization command 15 default group tacacs+ none
username user password password
line con 0
login authentication console
```

Conversely, if a general user attempts to use his or her remote access to log in to a network device, ACS checks and approves the username and password; but, the authorization process would fail because that user would not have credentials that allow shell or exec access to the device.

## Database Considerations

Aside from topological considerations, the user database is one of the most influential factors in deployment decisions for ACS. The size of the user base, distribution of users throughout the network, access requirements, and type of user database are all factors to consider when you decide how to deploy ACS.

### Number of Users

ACS is designed for the enterprise environment, and can handle 300,000 users. This capacity is usually more than adequate for a corporation. In an environment that exceeds these numbers, the user base would typically be geographically dispersed, which requires the use of more than one ACS configuration. A WAN failure could render a local network inaccessible because of the loss of the authentication server. In addition, reducing the number of users that a single ACS handles improves performance by lowering the number of logins occurring at any given time and reducing the load on the database.

### Type of Database

ACS supports several database options, including the ACS internal database or by using remote authentication with any of the external databases that ACS supports. Each database option has its own advantages and limitations in scalability and performance.

## Network Latency and Reliability

Network latency and reliability are also important factors in how you deploy ACS. Delays in authentication can result in timeouts for the end-user client or the AAA client.

The general rule for large, extended networks, such as those in a globally dispersed corporation, is to have at least one ACS deployed in each region. This configuration may not be adequate without a reliable, high-speed connection between sites. Many corporations use secure VPN connections between sites so that the Internet provides the link. Although this option saves time and money, it does not provide the speed and reliability of a dedicated frame relay or T1 link. If a reliable authentication service is critical to business functionality, such as a WLAN of retail outlets with cash registers that are linked by a WLAN, the loss of WAN connection to a remote ACS could be catastrophic.

The same issue can be applied to an external database that ACS uses. You should deploy the database close enough to ACS to ensure reliable and timely access. Using a local ACS with a remote database can result in the same problems as using a remote ACS. Another possible problem in this scenario is that a user may experience timeout problems. The AAA client would be able to contact ACS; but, ACS would wait for a reply that might be delayed or never arrive from the external user database. If the ACS were remote, the AAA client would time out and try an alternate method to authenticate the user; but, in the latter case, it is likely the end-user client would time out first.

