CHAPTER 1

Understanding the Cisco ISE Network Deployment

This chapter provides information on how to deploy the Cisco Identity Services Engine (ISE) 3300 Series appliance and its related components, several network deployment scenarios, and describes the switch configurations that are needed to support Cisco ISE. This chapter contains the following topics:

- Before Deploying Cisco ISE, page 1-1
- Deployment Scenarios, page 1-8
- Configuration of a Cisco ISE Node, page 1-13
- Switch Configurations Required to Support Cisco ISE Functions, page 1-14
- Planning an Inline Posture Deployment, page 1-14

Before Deploying Cisco ISE

This section provides the following reference information that aids you in better understanding what is needed before you deploy the Cisco ISE appliances in your network environment:

- Understanding Node Types, Personas, Roles, and Services, page 1-1
- Types of Nodes, page 1-2
- Understanding Distributed Deployment, page 1-3
- Guidelines for Setting Up a Distributed Deployment, page 1-6
- Cisco ISE Architecture Overview, page 1-7

Understanding Node Types, Personas, Roles, and Services

Cisco ISE provides a highly available and scalable architecture that supports both standalone and distributed deployments. In a distributed environment, you configure one primary Administration ISE node and the rest are secondary nodes. The topics in this section provide information about Cisco ISE terminology, supported node types, distributed deployment, and the basic architecture.

Cisco ISE Deployment Terminology

Table 1-1 describes some of the common terms used in Cisco ISE deployment scenarios.

Table 1-1 Cisco ISE Deployment Terminology

Term	Description
Service	A service is a specific feature that a persona provides such as network access, profiler, posture, security group access, and monitoring.
Node	A node is an individual instance that runs the Cisco ISE software. Cisco ISE is available as an appliance and also as a software that can be run on a VMware server. Each instance (either running on a Cisco ISE appliance or on a VMware server) that runs the Cisco ISE software is called a node.
Node type	A node can be of two types: ISE node and Inline Posture node. The node type and persona determine the type of functionality provided by that node.
Persona	The persona or personas of a node determine the services provided by a node. An ISE node can assume any or all of the following personas: Administration, Policy Service, and Monitoring.
Role	Determines if a node is a standalone, primary, or secondary node. Applies only to Administration ISE and Monitoring ISE nodes.

Types of Nodes

A Cisco ISE network has only two types of nodes:

- ISE node—An ISE node could assume any of the following three personas:
  - Administration—Allows you to perform all administrative operations on ISE. It handles all system-related configuration and configurations related to functionality such as authentication, authorization, auditing, and so on. In a distributed environment, you can have only one or a maximum of two nodes running the Administration persona. The Administration persona can take on any one of the following roles: standalone, primary, or secondary. If the primary Administration ISE node goes down, then you must manually promote the secondary Administration ISE node. There is no automatic failover for the Administration persona.

Note At least one node in your distributed setup should assume the Administration persona.

  - Policy Service—Provides network access, posture, guest access, client provisioning, and profiling services. This persona evaluates the policies and makes all the decisions. You can have more than one node assuming this persona. Typically, there would be more than one Policy Service persona in a distributed deployment. All Policy Service ISE nodes that reside behind a load balancer share a common multicast address and can be grouped together to form a node group. If one of the nodes in a node group fails, the other nodes detect the failure and reset any pending sessions.

Note At least one node in your distributed setup should assume the Policy Service persona.

– Monitoring—Enables ISE to function as the log collector and store log messages from all the Administration and Policy Service personas on the ISE nodes in your network. This persona provides advanced monitoring and troubleshooting tools that you can use to effectively manage your network and resources.

A node with this persona aggregates and correlates the data that it collects to provide you with meaningful information in the form of reports. Cisco ISE allows you to have a maximum of two nodes with this persona that can take on primary or secondary roles for high availability. Both the primary and secondary Monitoring personas collect log messages. In case the primary Monitoring persona goes down, the secondary Monitoring persona automatically assumes the role of the primary Monitoring persona.

Note At least one node in your distributed setup should assume the Monitoring persona.

- Inline Posture node—A gatekeeping node that is positioned behind network access devices such as wireless LAN controllers (WLCs) and virtual private network (VPN) concentrators on the network. Inline Posture enforces access policies after a user has been authenticated and granted access, and handles Change of Authorization (CoA) requests that a WLC or VPN are unable to accommodate. Cisco ISE allows you to have two Inline Posture nodes that can take on primary or secondary roles for high availability.

Note An Inline Posture node is dedicated solely to that service, and cannot operate concurrently with other ISE services. Likewise, due to the specialized nature of its service, an Inline Posture node cannot assume any persona. Inline Posture nodes are not supported on VMware server systems.

Note Each ISE node in a deployment can assume more than one of the three personas (Administration, Policy Service, or Monitoring) at a time. By contrast, each Inline Posture node operates only in a dedicated gatekeeping role.

In a distributed deployment, you can have the following combination of nodes on your network:

- Primary and secondary Administration ISE nodes
- Primary and secondary Monitoring ISE nodes
- One or more Policy Service ISE nodes
- One or more Inline Posture nodes

Understanding Distributed Deployment

An ISE distributed deployment consists of one primary Administration ISE node and multiple secondary nodes. Each ISE node in a deployment can assume any of the following personas: Administration, Policy Service, and Monitoring.

Note The Inline Posture node cannot assume any other persona, due to its specialized nature. The Inline Posture node must be a dedicated node. Inline Posture nodes are not supported on VMware server systems. For more information, see the *Cisco Identity Services Engine User Guide, Release 1.0*.

After you install ISE on all your nodes as described in this guide, the nodes come up in a standalone state. You must then define one node to be your primary Administration ISE node. After defining a primary Administration ISE node, you can choose to configure other personas on that node, such as Policy Service or Monitoring. After you define personas on the primary Administration ISE node, you can register other secondary nodes with the primary Administration ISE node and then define personas for the secondary nodes.

When you register an ISE node as a secondary node, ISE immediately creates a database link from the primary to the secondary node and begins the process of replicating or sharing ISE configuration data from the primary to the secondary nodes. This process ensures consistency between the configuration data that is present in all the ISE nodes that are part of your deployment.

A full replication typically occurs when you first register an ISE node as a secondary node. An incremental replication occurs after a full replication, and ensures that any new changes such as additions, modifications, or deletions to the configuration data in the primary Administration ISE node are reflected in the secondary nodes. The process of replication ensures that all ISE nodes in a deployment are in sync. You can view the status of replication from the deployment pages of the ISE administrative user interface.

The Policy Service ISE nodes that reside in a single location behind a load balancer and share a common multicast address can be grouped together. In such scenarios, you can define node groups and assign the nodes to the particular group.

To remove a node from a deployment, you must deregister it. When you deregister a secondary node from the primary Administration ISE node, the status of the deregistered node changes to standalone and the connection between the primary and the secondary node will be lost. Replication updates are no longer sent to the deregistered secondary node.

Note You cannot deregister a primary Administration ISE node.

The application server in an ISE node restarts when you make any of the following changes:

- Register a node (standalone to secondary)
- Deregister a node (secondary to standalone)
- Primary node is changed to standalone (if no other nodes are registered with it; primary to standalone)
- Administration ISE node is promoted (secondary to primary)
- Change the personas (when you assign or remove the Policy Service or Monitoring persona from a node)
- Modify the services in the Policy Service ISE node (enable or disable the session and profiler services)
- Restore a backup on the primary and a sync up operation is triggered to replicate data from the primary to secondary nodes

Note For example, if your deployment has two nodes and you deregister the secondary node, both nodes in this primary-secondary pair are restarted. (The former primary and secondary nodes become standalone.)

Note When you make any of these changes, the application services are restarted. You must expect a delay while these services restart.

Note You can have only one primary node in your deployment. The other Cisco ISE nodes are secondary nodes that can be configured for one or more of the roles previously described. When the primary node is lost, you must promote one of the secondary nodes to become the primary. Cisco ISE supports the promotion of any secondary appliance to serve as the primary node.

When the Cisco ISE installation has been completed, you must configure one of your Cisco ISE instances as the primary node. You can edit the primary node and enable any service that you want to run on the primary.

Before Registering Secondary Nodes

Prerequisites:

- The fully qualified domain name (FQDN) of the standalone node that you are going to register, for example, *ise1.cisco.com* must be DNS-resolvable from the primary Administration ISE node. Otherwise, node registration will fail. You must enter the IP addresses and FQDNs of the ISE nodes that are part of your distributed deployment in the DNS server.
- The primary Administration ISE node and the standalone node that you are about to register as a secondary node should be running the same version of Cisco ISE.
- Node registration fails if you provide the default credentials (username: admin, password: cisco) while registering a secondary node. Before you register a standalone node, you must log into its administrative user interface and change the default password (cisco).
- You can alternatively create an administrator account on the node that is to be registered and use those credentials for registering that node. Every ISE administrator account is assigned one or more administrative roles. To register and configure a secondary node, you must have one of the following roles assigned: Super Admin, System Admin, or RBAC Admin. See “Cisco ISE Admin Group Roles and Responsibilities” in Chapter 4 of the *Cisco Identity Services Engine User Guide, Release 1.0*, for more information on the various administrative roles and the privileges associated with each of them.
- If you plan to register a secondary Administration ISE node for high availability, we recommend that you register the secondary Administration ISE node with the primary first before you register other Cisco ISE nodes. If Cisco ISE nodes are registered in this sequence, you do not have to restart the secondary ISE nodes after you promote the secondary Administration ISE node as your primary.
- If you plan to register multiple Policy Service ISE nodes running Session services and you require mutual failover among those nodes, you must place the Policy Service ISE nodes in a node group. You must create the node group first before you register the nodes because you need to select the node group to be used on the registration page. See “Creating, Editing, and Deleting Node Groups” in Chapter 9 of the *Cisco Identity Services Engine User Guide, Release 1.0*, for more information.
- Ensure that the Certificate Trust List (CTL) of the primary node is populated with the appropriate Certificate Authority (CA) certificates that can be used to validate the HTTPS certificate of the standalone node (that you are going to register as the secondary node). See “Creating Certificate Trust Lists in the Primary Cisco ISE Node” in Chapter 12 of the *Cisco Identity Services Engine User Guide, Release 1.0*, for more information.
- After registering your secondary node to the primary node, if you change the HTTPS certificate on the registered secondary node, you must obtain appropriate CA certificates that can be used to validate the secondary node’s HTTPS certificate and import it to the CTL of the primary node. See “Creating Certificate Trust Lists in the Primary Cisco ISE Node” in Chapter 12 of the *Cisco Identity Services Engine User Guide, Release 1.0*, for more information.

Note We recommend that you set all Cisco ISE nodes to the UTC time zone. This procedure ensures that the reports and logs from the various nodes in your deployment are always in sync with regard to the timestamps.

You can register the secondary nodes and edit their configuration profiles by using the user interface of the primary node. After you install a secondary node, Cisco ISE immediately creates a database link between the primary and the secondary node for replicating and synchronizing all changes. In addition, you can remove a node from the deployment by deregistering it. This action deletes it from the deployment.

When you deregister a node from the primary, the status of the deregistered node changes to standalone. Any connection between the primary and the secondary nodes is lost, no replication updates are sent to the secondary node.

Next Steps:

For more information on configuring Cisco ISE nodes, see:

- *Cisco Identity Services Engine User Guide, Release 1.0*
  - Chapter 10, "Setting Up ISE in a Distributed Environment" and "Registering and Configuring a Secondary Node"

Guidelines for Setting Up a Distributed Deployment

Observe the following guidelines before you attempt to set up Cisco ISE appliances in a distributed deployment:

- You must have a properly configured, working DNS for a distributed deployment to work correctly.
- A Cisco ISE node can run any of the ISE node personas at the same time.
- A Cisco ISE node can be designated to perform as a standalone node, or as either a primary or a secondary node in a primary-secondary pair, depending upon configuration and settings.
- You can have only one primary Cisco ISE node in your deployment.

Note Other Cisco ISE nodes are considered to be secondary nodes that can be configured for one or more other roles depending upon licenses and settings. When the primary node is lost, you need to promote a valid secondary node to become the primary. Cisco ISE only supports the promotion of a secondary node appliance with the Administration persona to serve as the "new" primary node. In addition, it must possess a valid license as a secondary node with an Administration persona.

- The primary Cisco ISE node must run the Administration persona.
- All Cisco ISE system-related configuration and configuration that is related to functionality should be made only on the primary Cisco ISE node.
- The configuration changes that you perform on the primary node are replicated to all the secondary nodes in your deployment.
- The Inline Posture node requires a dedicated Cisco ISE node. No other service can run on a node that is designated as an Inline Posture node.

Note The Inline Posture node is not supported on VMware server systems.

When the Cisco ISE installation is complete, you must configure one of your Cisco ISE nodes as the primary node. You can edit the primary node and enable any service that you want to run on the primary. You can register secondary nodes and edit their configuration by using the user interface of the primary node. After you install a secondary node, Cisco ISE immediately creates a database link between the primary and secondary nodes for replicating and synchronizing all changes.

When you deregister a node from the primary, the status of the deregistered node changes to standalone. To register a deregistered node back with the primary, you must first reset the database configuration on the node and bring it back to a freshly installed node state and then register it again.

For more information:

See the *Cisco Identity Services Engine User Guide, Release 1.0* for more information about:

- Cisco ISE Admin group roles and responsibilities
- Cisco ISE node services
- Resetting the configuration of a node

Cisco ISE Architecture Overview

Figure 1-1 illustrates a basic overview of the Cisco ISE architecture that includes the following components:

- Nodes and persona types
  - ISE node—Administration, Policy Service, Monitoring
  - Inline Posture node—Gatekeeping and access policy enforcer
- Network resources
- Endpoints

Note Figure 1-1 shows ISE nodes and persona types (Administration, Policy Service, and Monitoring), an Inline Posture node, and a policy information point.

The policy information point represents the point at which external information is communicated to the Policy Service persona. For example, external information could be a Lightweight Directory Access Protocol (LDAP) attribute.

Figure 1-1 Cisco ISE Architecture

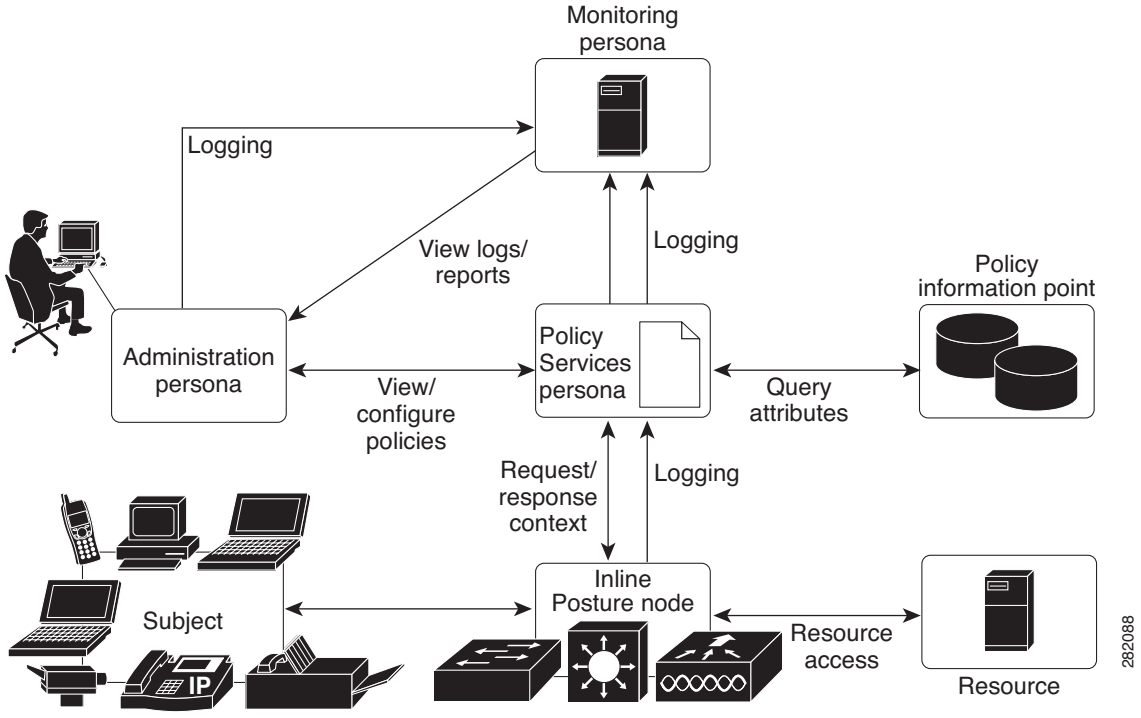

Deployment Scenarios

This section describes three scenarios in which Cisco ISE can be deployed in a distributed deployment:

- Small Cisco ISE Network Deployments, page 1-8
- Medium Cisco ISE Network Deployments, page 1-10
- Large Cisco ISE Network Deployments, page 1-11

Small Cisco ISE Network Deployments

The smallest Cisco ISE deployment consists of two Cisco ISE nodes as shown in Figure 1-2, with one Cisco ISE node functioning as the primary appliance in a small network that supports up to 3,000 concurrent endpoints.

Note Concurrent endpoints represent the total number of supported users and devices. This can be any combination of users, personal computers, laptops, IP phones, smart phones, gaming consoles, printers, fax machines, or other types of network devices.

The primary node provides all the configuration, authentication, and policy capabilities that are required for this network model, while the secondary Cisco ISE node functions in a backup role. The secondary node supports the primary node and maintains a functioning network whenever connectivity is lost between secondary network appliances, network resources, or RADIUS.

RADIUS is where the centralized AAA operations are performed between clients and the primary Cisco ISE node. As a result, the key requirement is to ensure that you can synchronize or replicate all of the content that resides on the primary Cisco ISE node with the secondary Cisco ISE node(s).

Being able to synchronize between the primary and secondary node makes it possible to keep the secondary node current with the state of your primary node. In a small network deployment, this type of configuration model allows you to configure both your primary and secondary node on all RADIUS clients by using this type of deployment or a similar approach.

Figure 1-2 Small Cisco ISE Network Deployment

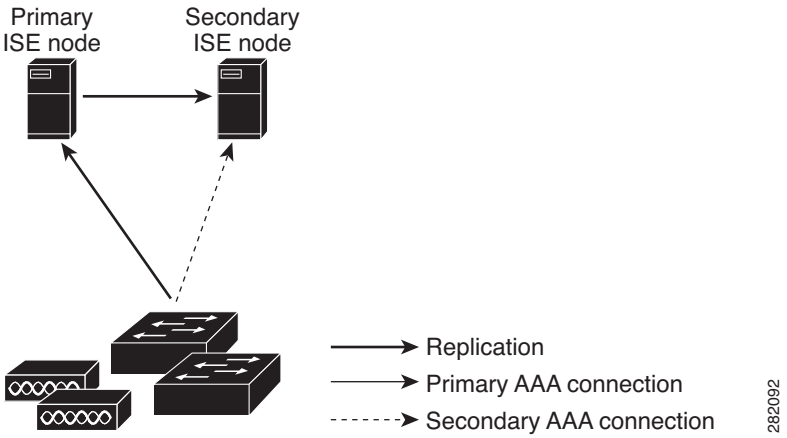

As the number of devices, network resources, users, and AAA clients increases in your network environment, we recommend that you change your deployment configuration from the basic small model and use more of a split or distributed deployment model, as shown in Figure 1-3.

Note Figure 1-2 shows the secondary Cisco ISE node acting as a Policy Service persona performing AAA functions. The secondary Cisco ISE node could also be acting as a Monitoring or Administration persona.

Split Cisco ISE Deployments

In the case of split Cisco ISE deployments, you will continue to maintain primary and secondary nodes as described in the small Cisco ISE deployment. However, the AAA load is split between these two Cisco ISE nodes to optimize the AAA workflow. Each Cisco ISE appliance (primary or secondary) needs to be able to handle the full workload if there are any problems with AAA connectivity. When running under normal network operations, neither the primary or secondary node carries the full load of handling AAA requests because this workload is distributed between the two nodes.

The ability to split the load in this way directly reduces the stress on each Cisco ISE node in the system. In addition, splitting the load also provides better loading while still maintaining the functional status of the secondary node during the course of normal network operations.

Another advantage is that each node can perform its own specific operations, such as network admission or device administration, and still perform all the AAA functions in the event of a failure. If you have two Cisco ISE nodes that process authentication requests and collect accounting data from AAA clients, we recommend that you set up one of the Cisco ISE nodes to act as a log collector. Figure 1-3 shows the secondary Cisco ISE node in this role.

Figure 1-3 Split Cisco ISE Network Deployment

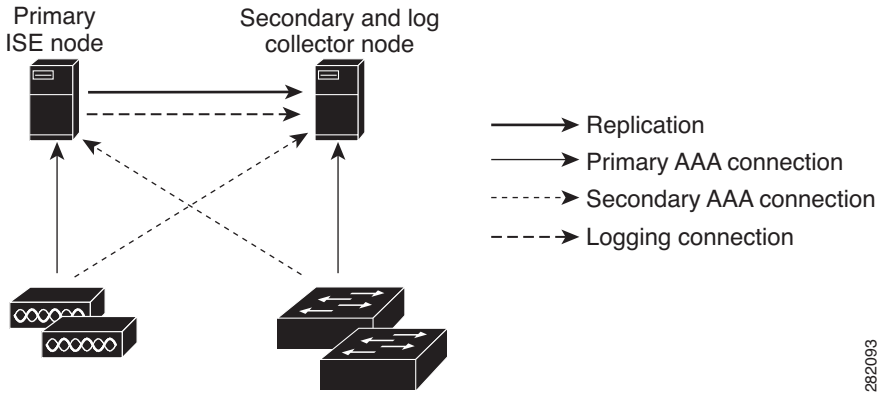

In addition, the split Cisco ISE node deployment design provides an advantage because it also allows for growth, as shown in Figure 1-4.

Medium Cisco ISE Network Deployments

As small, local networks grow, you can keep pace and manage network growth by adding additional Cisco ISE nodes to create a medium network that supports up to 6,000 concurrent endpoints. In medium network deployments, consider promoting one Cisco ISE node to perform as the primary to handle all the configuration services, and secondary Cisco ISE nodes to manage all your AAA functions.

As the amount of log traffic increases in the network, you can choose to either use the primary Cisco ISE node as your centralized log collector or dedicate one of the secondary Cisco ISE nodes to serve in this capacity for your network.

Figure 1-4 Medium Cisco ISE Network Deployment

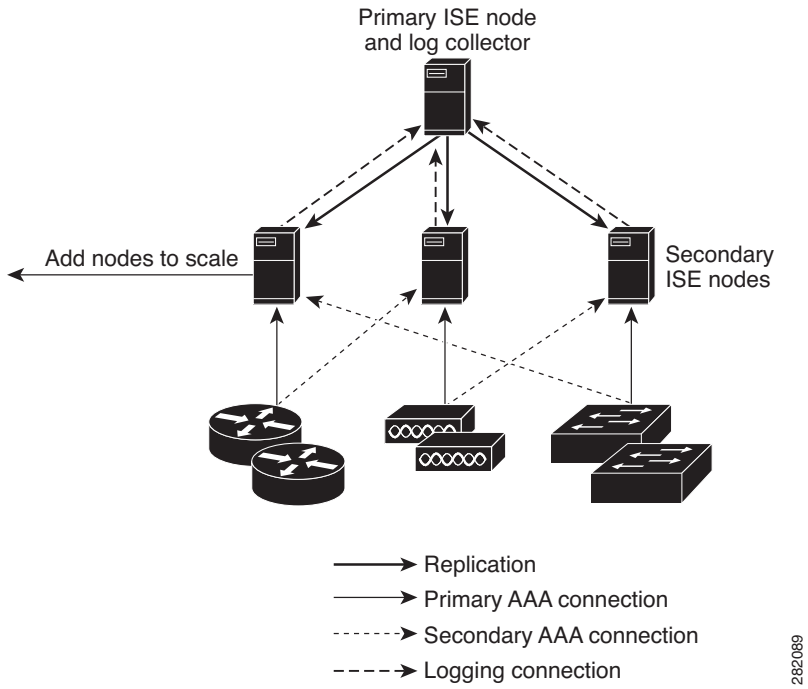

Large Cisco ISE Network Deployments

We recommend that you use centralized logging (as shown in Figure 1-5) for larger Cisco ISE networks that support up to 10,000 concurrent endpoints. To use centralized logging, you must set up a dedicated logging server that serves as a Monitoring persona (for monitoring and logging) to handle the potentially high syslog traffic that a large, busy network can generate.

Because syslog messages are generated for outbound log traffic, any RFC-3164-compliant syslog appliance can serve as the collector for outbound logging traffic. A dedicated logging server enables you to use the reports and alert features that are available in Cisco ISE to support all the Cisco ISE nodes. See Understanding the Setup Program Parameters, page 3-3 when configuring the Cisco ISE software to support a dedicated logging server.

You can also consider having the appliances send logs to both a Monitoring persona on the Cisco ISE node and a generic syslog server. Adding a generic syslog server provides a redundant backup if the Monitoring persona on the Cisco ISE node goes down.

In large centralized networks, you should use a load balancer (as shown in Figure 1-5), which simplifies the deployment of AAA clients. Using a load balancer requires only a single entry for the AAA servers, and the load balancer optimizes the routing of AAA requests to the available servers.

However, having only a single load balancer introduces the potential for having a single point of failure. To avoid this potential issue, deploy two load balancers to ensure a measure of redundancy and failover. This configuration requires you to set up two AAA server entries in each AAA client, and this configuration remains consistent throughout the network.

Figure 1-5 Large Cisco ISE Network Deployment

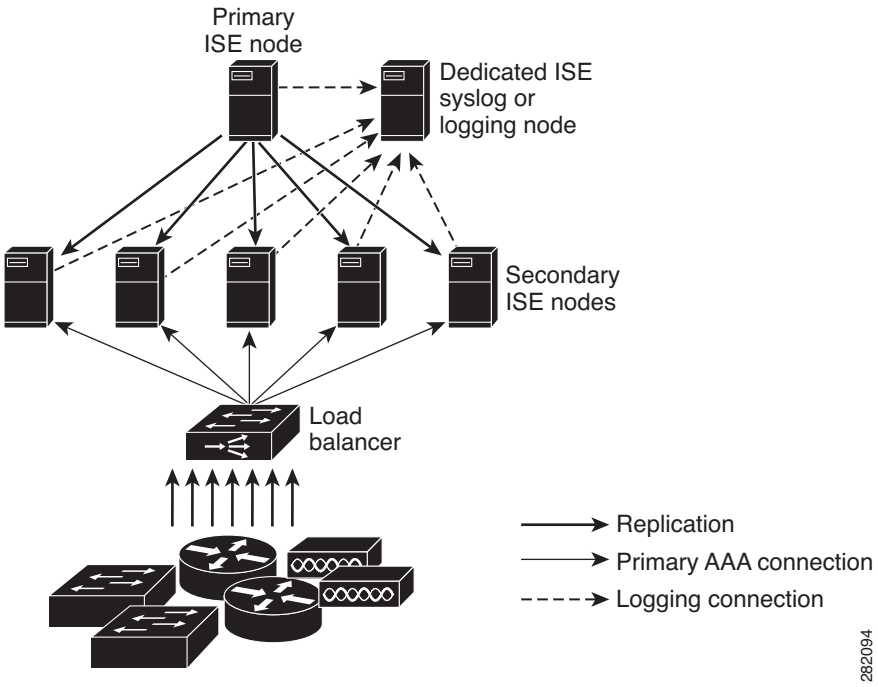

Dispersed Cisco ISE Network Deployments

Dispersed Cisco ISE network deployments are most useful for organizations that have a main campus with regional, national, or satellite locations elsewhere. The main campus is where the primary network resides, is connected to additional LANs, ranges in size from small to large, and supports appliances and users in different geographical regions or distant locations.

To optimize AAA performance, each remote site should have its own AAA infrastructure (as shown in Figure 1-6). A centralized management model helps maintain a consistent, synchronized AAA policy. A centralized configuration model uses a primary Cisco ISE node with secondary Cisco ISE nodes. We still recommend that you use a separate Monitoring persona on the Cisco ISE node, but each remote location should retain its own unique network requirements.

Figure 1-6 Dispersed Cisco ISE Deployment

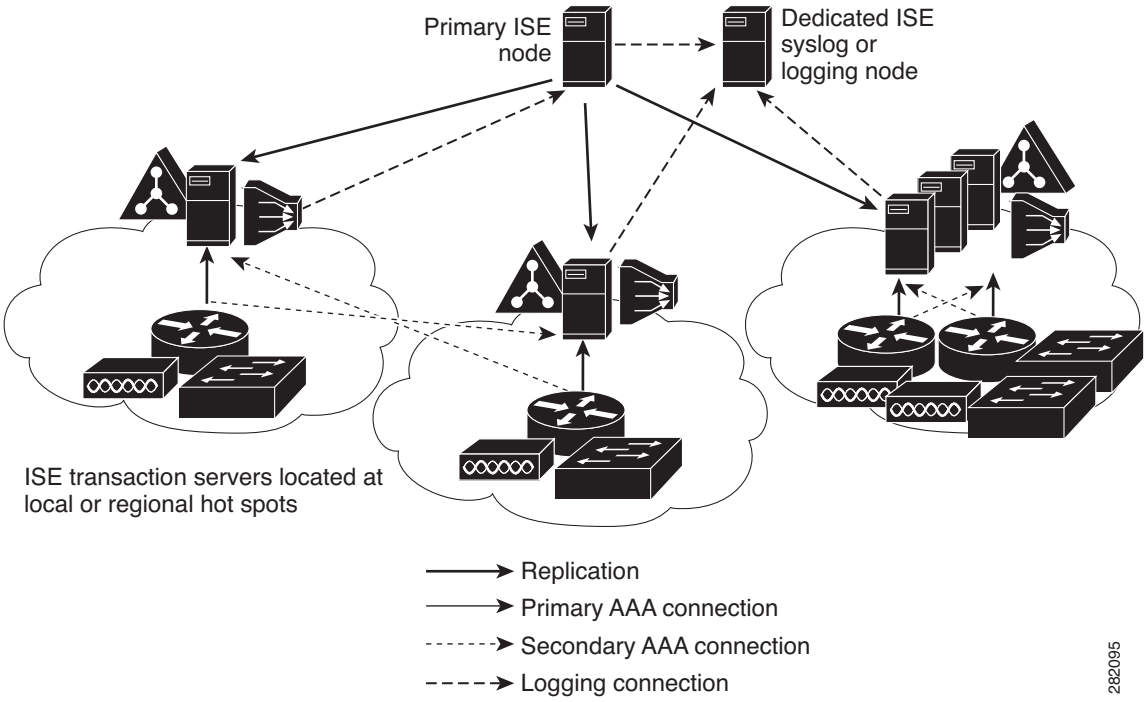

Some factors to consider when planning a network that has several remote sites include the following:

- Verify if a central or external database is used, such as Microsoft Active Directory or LDAP. For optimizing the process, each remote site should have a synchronized instance of the external database that is available for Cisco ISE to access.
- Locating the AAA clients is important. You should locate your Cisco ISE nodes as close as possible to the AAA clients to reduce network latency effects and the potential for loss of access that is caused by WAN failures.
- Cisco ISE has console access for some functions such as backup. Consider using a terminal at each site, which allows for direct, secure console access that bypasses network access to each node.
- If small, remote sites are in close proximity and have reliable WAN connectivity to other sites, consider using a Cisco ISE node as a backup for the local site to provide redundancy.
- DNS should be properly configured on all Cisco ISE nodes to ensure access to the external databases.

Configuration of a Cisco ISE Node

This section briefly describes the roles that various Cisco ISE appliances play in a network deployment and how to configure them. For more information on assigning a role to a node and configuring it, see the *Cisco Identity Services Engine User Guide, Release 1.0*. This section contains:

- Primary Node, page 1-13
- Secondary Node, page 1-13
- Logging Server, page 1-14

All Cisco ISE appliances have a similar installation procedure. For specific details, see the following sections:

- Chapter 3, "Configuring the Cisco ISE 3300 Series Appliance," for installing Cisco ISE software on the Cisco ISE 3300 Series appliance.
- Chapter 4, "Installing Cisco ISE 3300 Series Software in a VMware Virtual Machine," for installing Cisco ISE software on a VMware ESX server.

Note For any Cisco ISE network deployment, your first hardware installation must be performed on the node that is designated as the primary node in your network.

Primary Node

In a Cisco ISE deployment, only one appliance can serve as a Cisco ISE primary node. This primary node provides configuration capabilities and is the source for all replication operations.

When in a primary-secondary pair, only the primary and secondary nodes that operate as the Administration persona need to be configured in the license file. When you install the license file on the primary, the license requirements for the secondary node are met.

Secondary Node

Because the network can only have a single primary Cisco ISE node, all other Cisco ISE nodes function as secondary nodes. Although the Cisco ISE secondary nodes receive all the system configurations from the primary node, you must configure the following on each secondary node:

- License—When the base license is installed on the primary, replication copies the license onto each of the Cisco ISE secondary nodes in the deployment.
- New local certificates—You can either configure the local certificates on the secondary nodes or import the local certificates from the primary node onto each secondary node.
- Logging server—You can configure either the primary or the secondary node to serve as the dedicated logging server for your Cisco ISE network. We strongly recommend that you configure a secondary Cisco ISE node as the dedicated logging server.

In a primary-secondary node pair, the secondary node is registered and it begins to receive the full synchronization of the configuration and replication updates from the primary node in the network.

Logging Server

You can configure to use either a primary node or one of the secondary nodes as the dedicated logging server for your network. In this role, the logging server receives logs from the primary node and all the secondary nodes deployed in the Cisco ISE network. We recommend that you designate one of the Cisco ISE secondary nodes as the Monitoring persona and exclude this particular secondary node from any of the AAA activities. Three main logging categories are captured:

- Audit
- Accounting
- Diagnostics

For a complete description that provides more details on logging categories and best practices for configuring the logging server, see Chapter 13, "Logging" in the *Cisco Identity Services Engine User Guide, Release 1.0*.

Switch Configurations Required to Support Cisco ISE Functions

To ensure that Cisco ISE is able to interoperate with network switches, and functions from Cisco ISE are successful across the network segment, you must configure your network switches with certain required NTP, RADIUS/AAA, 802.1X, MAB, and other settings.

For more information:

- For more switch configuration requirements, see Appendix C, "Switch Configuration Required to Support Cisco ISE Functions" in the *Cisco Identity Services Engine User Guide, Release 1.0*.

Planning an Inline Posture Deployment

This section is only intended to provide a brief overview of what is needed to plan and deploy Inline Posture in a Cisco ISE network. It is the responsibility of your network or system architect to research the issues involved in Inline Posture deployment to determine what best suits your network needs and requirements.

Before you start any planning for deploying or configuring Inline Posture for your network, you must first understand what types of Inline Posture operating modes and deployment options are supported.

Note For more details about Inline Posture operating modes, filters, managed subnets, and Inline Posture high availability as these topics correspond to the Cisco ISE network, see Chapter 10, "Setting Up an Inline Posture Node," in the *Cisco Identity Services Engine User Guide, Release 1.0*.

Inline Posture Planning Considerations

This section poses some basic questions and considerations that must be addressed by your network or system architect when planning to deploy Inline Posture nodes. Ensure that you have understood the following planning and deployment issues prior to starting any Inline Posture node configuration in a distributed Cisco ISE network deployment:

- How do you plan to deploy your Inline Posture node?
- How will you deploy your Inline Posture node(s)?
- Will the Inline Posture node be run as a standalone node, or as part of a primary-secondary pair of Inline Posture nodes?

Note Cisco ISE networks support up to two Inline Posture nodes configured on your network at any one time. If you plan to deploy an Inline Posture high-availability primary-secondary pair, then two Inline Posture nodes must be configured. In this mode, one node is designated as the primary and the other as the secondary node. The primary node assumes the primary role when both nodes come up at the same time.

- Will your deployment plans include an Inline Posture primary-secondary pair configuration? If so, be aware that all configuration related to functionality can only be done from the primary node of this pair (the Cisco ISE user interface only shows basic configuration tables for the secondary node in this configuration).
- Note that you can synchronize an Inline Posture primary node configuration with its peer secondary node using the Failover tab of the primary node in this Inline Posture pair. For more information, see Chapter 10, "Setting Up an Inline Posture Node," in the *Cisco Identity Services Engine User Guide, Release 1.0*.

The following topics in this section provide some basic information on Inline Posture nodes, but these topics are not intended to provide you with all the information needed to complete a comprehensive deployment plan for your network.

Choosing an Inline Posture Operating Mode

Which Inline Posture operating mode you choose largely depends on your existing network architecture. The choice you make limits many of the other configuration options you may want in your Cisco ISE deployment. Therefore, you need to fully understand each of the following primary Inline Posture operating modes:

- Routed mode—This mode acts as a Layer 3 "hop" in the network connections. The routed mode selectively forwards packets to specified addresses. The routed mode ensures it can segregate network traffic, which allows you to specify access to users who can access selected destination addresses.
- Bridged mode—This mode acts as a Layer 2 "bump in the wire" in the network connections. The bridged mode forwards packets regardless of the destination address.

Note Inline Posture nodes also support a maintenance mode, which takes the node offline so that you can perform administrative procedures. This mode is also the default when an Inline Posture node is initially brought online in the network.

Inline Posture Routed Mode

In the routed mode, an Inline Posture node operates as a Layer 3 router and functions as the default gateway for an untrusted (outside Cisco ISE) network with its managed clients. All traffic between an untrusted and trusted network passes through this Inline Posture routed mode. The routed mode applies IP filtering rules, the configured access policies, and other traffic-based policies you have set up for your network.

When you configure an Inline Posture node in its routed mode, specify the IP addresses of its two interfaces:

- Trusted (Eth0)
- Untrusted (Eth1)

The trusted and untrusted addresses should be on different subnets. An Inline Posture node can manage one or more subnets, and the untrusted interface acts as a gateway for the managed subnets. Figure 1-7 illustrates an example of an Inline Posture routed mode configuration.

Figure 1-7 Inline Posture Routed Mode Configuration

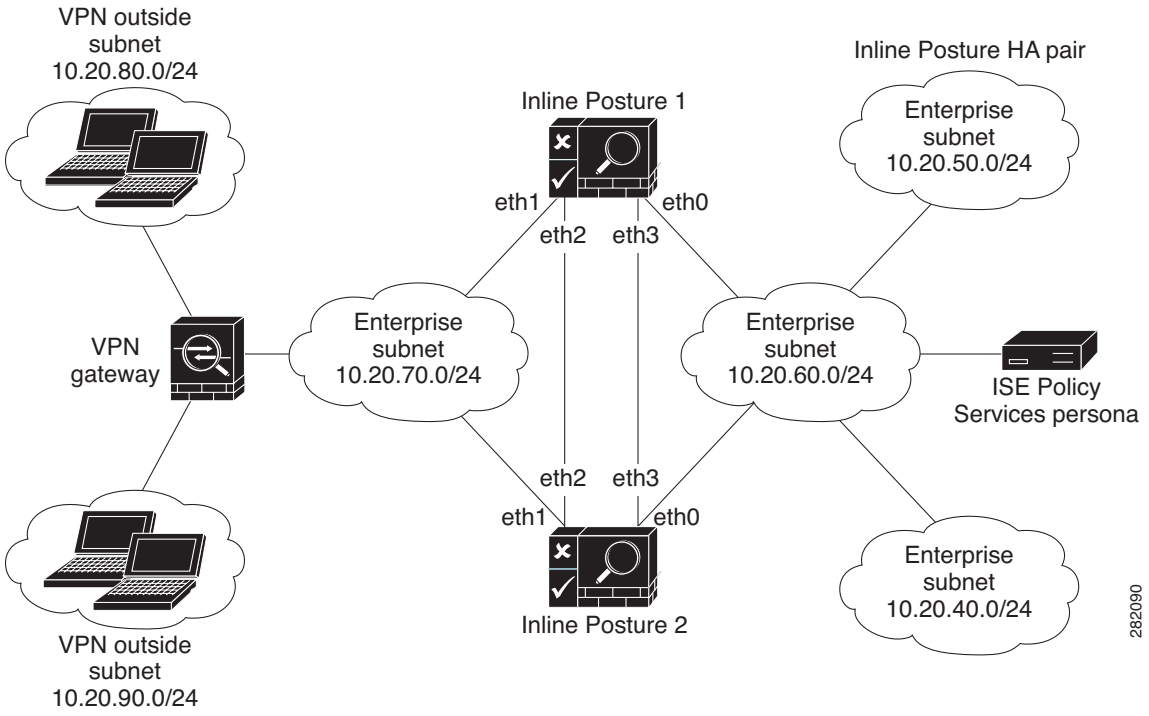

Inline Posture Bridged Mode

When operating in a bridged mode, the Inline Posture node operates like a standard Ethernet bridge. This configuration is used most often when the untrusted network already contains a gateway, and you do not want or plan to make any changes to the existing configuration.

Figure 1-8 shows the Inline Posture node acting as a bridge for the Layer 2 client traffic from the WLC into the Cisco ISE network. While in this configuration, the Inline Posture node requires subnet entries for the subnets to be able to respond to and send ARP broadcasts to the correct VLANs.

The Layer 2 flow of traffic from the three example subnets (10.20.80.0/24, 10.20.90.0/24, and 10.20.60.0/24) all reflect the use of the bridged mode on the Inline Posture node using VLAN mapping. The only difference between the three subnet examples is that for the 10.20.60.0/24 subnet, the Inline Posture main interfaces reside within this subnet.

Figure 1-8 Inline Posture Bridged Mode Configuration

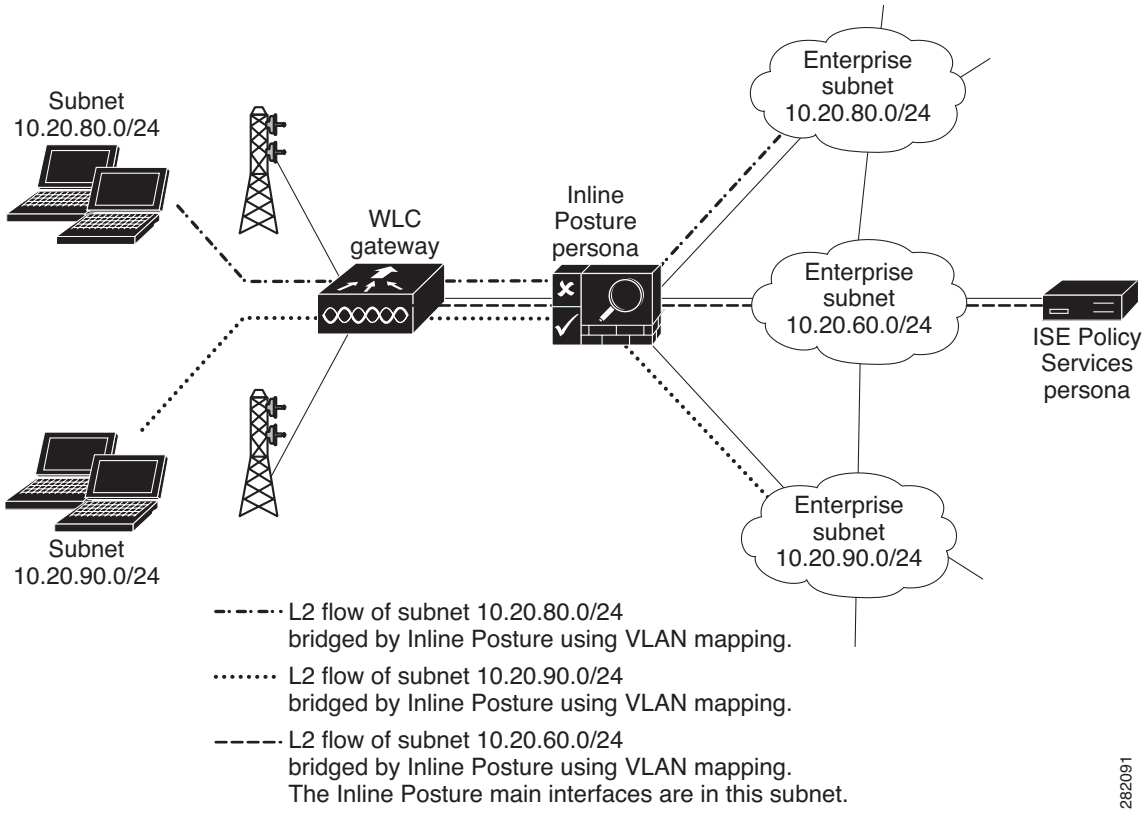

Deploying Inline Posture as Standalone or High Availability

The most important decision you may make about your Inline Posture deployment is whether to deploy it as a single, standalone Inline Posture node, or as a primary-secondary pair to ensure high availability and provide redundancy for network reliability.

A standalone Inline Posture node is a single Inline Posture node that provides Inline Posture services, while working independently of all other nodes in your Cisco ISE network. You may decide to deploy a single standalone Inline Posture node for a network that serves a smaller facility or for a small network where network redundancy is not a major concern.

When you configure a pair of Inline Posture nodes for high availability, they act as primary-secondary pair to provide additional redundancy and reliability. This primary-secondary pair ensures that your network continues functioning even if one node in the pair fails. If the primary node fails, the secondary node takes over and provides the needed Inline Posture functionality.

About Inline Posture High Availability

Inline Posture high availability consists of two Inline Posture nodes that are configured as a primary-secondary pair. In this configuration, the primary node acts as the RADIUS proxy and forwards all network packets. If the primary node fails, the secondary Inline Posture node in this pair takes over.

In an Inline Posture stateless high-availability deployment that has a primary-secondary pair configuration, the secondary node acts as a backup unit and does not forward any packets between the interfaces. Stateless means that sessions that have been authenticated and authorized by the primary node are automatically authorized again once a failover occurs.

The secondary node monitors the primary node using the heartbeat protocol (on the eth2 and eth3 interfaces). The heartbeat protocol requires that messages are sent at regular intervals between the two nodes. If the heartbeat stops or does not receive a response back in the allotted time, failover occurs and recovery action takes place.

When the heartbeat protocol is active in an Inline Posture high-availability configuration, it requires a network connection between the eth2 and eth3 interfaces of the Inline Posture primary-secondary pair. The eth2 and eth3 interfaces of each node in an Inline Posture high-availability pair (primary and secondary) are configured to use heartbeat protocol exchanges between the two nodes. For this reason, you must make a direct cable connection between the eth2 interfaces of both Inline Posture nodes, and likewise, there must also be a direct cable connection between the eth3 interfaces of both nodes to ensure redundancy.

Note The heartbeat protocol requires a direct cable connection between the eth2 interfaces of both nodes in a high-availability pair, as well as a direct cable connection between the eth3 interfaces of the two nodes. You can use any Ethernet cable to make these connections. Figure 1-9 illustrates this cable requirement.

Figure 1-9 Heartbeat Protocol: eth2 and eth3 Interface Ethernet Cable Connections