CHAPTER 10

Setting Up Inline Posture

This chapter describes how to set up and configure Inline Posture nodes in standalone mode, or as a high availability pair, and contains the following topics:

Inline Posture Known Limitations

This section describes known limitations for Inline Posture in Cisco ISE:

- Inline Posture is not supported in a virtual environment, such as VMware.
- Backup and restore is not available for Inline Posture nodes.
- The Simple Network Management Protocol (SNMP) Agent is not supported by Inline Posture.
- The Cisco Discovery Protocol (CDP) is not supported by Inline Posture.

For more information on these and other known issues, see the "Known Issues" section of the *Release Notes for the Cisco Identity Services Engine, Release 1.1.x*.

Understanding the Role of Inline Posture

An Inline Posture node is a gatekeeper that enforces access policies and handles change of authorization (CoA) requests. An Inline Posture node is positioned behind the network access devices on your network that are unable to accommodate CoA, such as wireless LAN controllers (WLC) and Virtual Private Network (VPN) devices.

After the initial authentication of a client (using EAP/802.1x and RADIUS), the client must still go through posture assessment. The posture assessment process determines whether the client should be restricted, denied, or allowed full access to the network. When a client accesses the network through a WLC or VPN device, Inline Posture is responsible for the policy enforcement and CoA that these devices are unable to accommodate.

Inline Posture Policy Enforcement

Inline Posture uses RADIUS proxy and URL redirect capabilities in the control plane to manage data plane traffic for endpoints. As a RADIUS proxy, Inline Posture is able to tap into RADIUS sessions between network access devices (NADs) and RADIUS servers. NADs can open full gate to client traffic. However, Inline Posture opens only enough to allow limited traffic from clients. The restricted bandwidth allows clients the ability to have an agent provisioned, have posture assessed, and have remediation done. This restriction is accomplished by downloading and installing DACLs that are tailored for specific client flow.

Upon full compliance, a CoA is sent to the Inline Posture node by the Policy Service ISE node, and full gate is opened by the Inline Posture node for the compliant client endpoint. The RADIUS proxy downloads the full-access DACL, installs it, and associates the client IP address to it. The installed DACL can be common for a number of user groups, so that duplicate downloads are not necessary as long as the DACL content does not change at the Cisco ISE servers.

Figure 10-1 illustrates the Inline Posture policy enforcement process. This example shows the flow for WLC enforcement for traffic to the Policy Service ISE node. However, the access steps are similar for an inline deployment with VPN gateways.

Figure 10-1 Inline Posture Policy Enforcement Flow

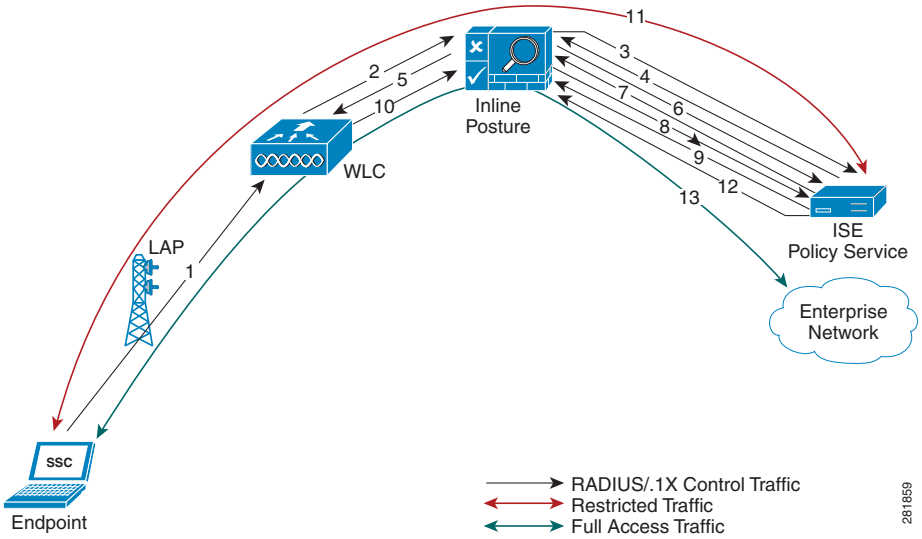

The Inline Posture policy enforcement flow illustrated in Figure 10-1 follows these steps:

1. The endpoint initiates a .1X connection to the wireless network.
2. The WLC, which is a NAD, sends a RADIUS Access-Request message to the RADIUS server (usually the Policy Service ISE node).
3. Inline Posture node, acting as a RADIUS proxy, relays the Access-Request message to the RADIUS server.

4. After authenticating the user, the RADIUS server sends a RADIUS Access-Accept message back to the Inline Posture node.

   There can be a number of RADIUS transactions between the Endpoint, WLC, Inline Posture node, and the Cisco ISE RADIUS server before the Access-Accept message is sent. The process described in this example has been simplified for the sake of brevity.

5. The Inline Posture node passes the Access-Accept message to the WLC, which in turn authorizes the endpoint access, in accordance with the profile that accompanied the message.

6. The proxied Access-Accept message triggers Inline Posture to send an Authorization-Only request to the Policy Service ISE node, to retrieve the profile for the session.

7. The Policy Service ISE node returns an Access-Accept message, along with the necessary Inline Posture profile.

8. If the access control list (ACL) that is defined in the profile is not already available on the Inline Posture node, Inline Posture downloads it from the Policy Service ISE node using a RADIUS request (to the Cisco ISE RADIUS server).

9. The Cisco ISE RADIUS server sends the complete ACL in response. It is then installed in the Inline Posture data plane so that endpoint traffic passes through it.

   There may be a number of transactions before the complete ACL is downloaded, especially if the ACL is too large for one transaction.

10. As the endpoint traffic arrives at the WLC, the WLC sends out a RADIUS Accounting-Start message for the session to the Inline Posture node.

    The actual data traffic from the endpoint may arrive at the Inline Posture untrusted side before the Accounting-Start message is received by the Inline Posture node. Upon receiving the RADIUS Accounting-Start message, the Inline Posture node learns the IP address of the endpoint involved in the session and associates the endpoint with the ACL (downloaded and installed earlier in the session). The initial profile for this client endpoint could be restrictive, to posture the client before being given full access.

11. Assuming the restrictive ACL allows only access to Cisco ISE servers, the endpoint is only allowed actions such as agent downloading and posture assessment over the data plane.

12. If the client endpoint is posture compliant (as part of the restricted communication with Cisco ISE services earlier), the Policy Service ISE node initiates a RADIUS Change of Authorization (CoA) with the new profile. Therefore a new ACL is applied at the Inline Posture node for the session. The new ACL is installed immediately and applied to the endpoint traffic.

13. The endpoint is then capable of full access to the enterprise network, as a result of the new profile that was applied to Inline Posture.

    A RADIUS stop message for a given session that is issued from the WLC, resets the corresponding endpoint access at the Inline Posture node.

In a deployment, such as outlined in the example, when more endpoints connect to the wireless network they are likely to fall into one of the identity groups that already have authenticated and authorized users connected to the network.

For example, there may be an employee, executive, and guest that have been granted access through the outlined steps. This situation means that the respective restrictive or full-access profiles for those ID groups have already been installed on the Inline Posture node. The subsequent endpoint authentication and authorization uses the existing installed profiles on the Inline Posture node, unless the original profiles have been modified at the Cisco ISE policy configuration. In the latter case, the modified profile with ACL is downloaded and installed on the Inline Posture node, replacing the previous version.

Trusted and Untrusted Interfaces

The following terminology plays a significant role in an Inline Posture deployment. For this reason, it is important that you understand the definitions as they relate to Inline Posture:

- Trusted—The interface that talks to the Policy Service ISE node and other trusted devices *inside* the Cisco ISE network. The trusted interface is always designated to Eth0.
- Untrusted—The interface that talks to the WLC, VPN, and other devices *outside* the Cisco ISE network. The untrusted interface is always designated to Eth1.

Inline Posture Dedicated Nodes

Unlike other persona services, Inline Posture is unable to share a node with other services. This inability to share a node means that Inline Posture must be a dedicated node that is registered to the primary Administration ISE node on your network.

Cisco ISE allows you to have up to two Inline Posture nodes configured as an active-standby pair for high availability.

For information on Cisco ISE distributed deployments, see Chapter 9, "Setting Up Cisco ISE in a Distributed Environment."

Planning an Inline Posture Deployment

Before you begin configuring Inline Posture for your network, you should understand the Inline Posture operating modes, deployment options, as well as the basics of filters and managed subnets as they apply to Inline Posture.

This section provides information on the following topics:

Note For information on how to install a Cisco ISE node, see the *Cisco Identity Services Engine Hardware Installation Guide, Release 1.1.x*.

About Inline Posture Configuration

Inline Posture is a dedicated node registered to the Administration ISE node. You configure Inline Posture from the administration console, and that configuration is then pushed to the Inline Posture node. A copy of the configuration is stored locally in the administration database. Registration results in the Inline Posture node being rebooted.

If you have an Inline Posture high availability (HA) pair, the configuration automatically pushes to both Inline Posture nodes. If the secondary node is down during a configuration change, you can click a database sync button on the primary node that automatically applies the latest configuration to the secondary node when it comes up. A local database maintains the configurations.

Note Registering an Inline Posture node results in system restart. Changes to infrastructure configurations, such as eth1 IP address, Inline Posture mode, and high availability changes also require a system restart.

After you register an Inline Posture node to the Administration ISE node, you are not allowed to change the eth0 (Trusted) IP address through the Admin user interface. The reason for this is that, if you change the eth0 IP address of a registered Inline Posture node, it no longer can communicate with the Administration ISE node. Any attempted communication between the Inline Posture node and Administration ISE node then fails, leading to a potential exception.

Warning **It is highly recommended that you not change the IP address of an Inline Posture node from the CLI after it has been registered on the Cisco ISE network.**

Caution The Inline Posture node's untrusted interface should be disconnected at the time the Inline Posture node is being configured. If the Inline Posture node's trusted and untrusted interfaces are connected to the same VLAN during initial configuration, and the Inline Posture node initially boots up after changing its persona, multicast packet traffic gets flooded out of the untrusted interface. This multicast storm can potentially bring down devices that are connected to the same subnet or VLAN. The Inline Posture node at at this time is in the maintenance mode.

Choosing an Inline Posture Operating Mode

The Inline Posture operating mode you choose depends largely on the architecture of your existing network. However, this choice sets a precedent for many of the other configuration options you have to specify for the deployment. For this reason, it is important that you understand the functions of each of the following Inline Posture operating modes:

- Routed mode—This mode acts as a Layer 3 "hop" in the wire, selectively forwarding packets to specified addresses. This mode provides the ability to segregate network traffic, allowing you to specify users who have access to selected destination addresses.
- Bridged mode—This mode acts as a Layer 2 "bump" in the wire, forwarding packets without regard to the destination address.
- Maintenance mode—This mode takes the node offline so that you can perform administrative procedures. This mode is also the default mode of a node when it first comes onto the network, before you perform other configurations.

Bridged mode and routed mode are discussed in greater detail throughout the rest of this section.

Inline Posture Routed Mode

In routed mode, the Inline Posture node operates as a Layer 3 router, and becomes the default gateway for the untrusted network with its managed clients. All traffic between the untrusted and trusted networks passes through the Inline Posture node, which applies the IP filtering rules, access policies, and other traffic-handling mechanisms that you decide to configure.

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

- Trusted (Eth0)
- Untrusted (Eth1)

The trusted and untrusted addresses should be on different subnets. Inline Posture can manage one or more subnets, with the untrusted interface acting as a gateway for the managed subnets.

Figure 10-2 illustrates an Inline Posture routed mode configuration. In the following routed mode example, Inline Posture is a hop for the client traffic from the VPN gateway (GW) en route to the Policy Service ISE node. Inline Posture requires that static routes be configured for subnets 10.20.80.0/24 and 10.20.90.0/24 toward the VPN gateway, just like any other router. The enterprise router on the trusted side of the network also requires that the static routes are configured for the same subnets toward the Inline Posture node.

Figure 10-2 Inline Posture Routed Mode Configuration

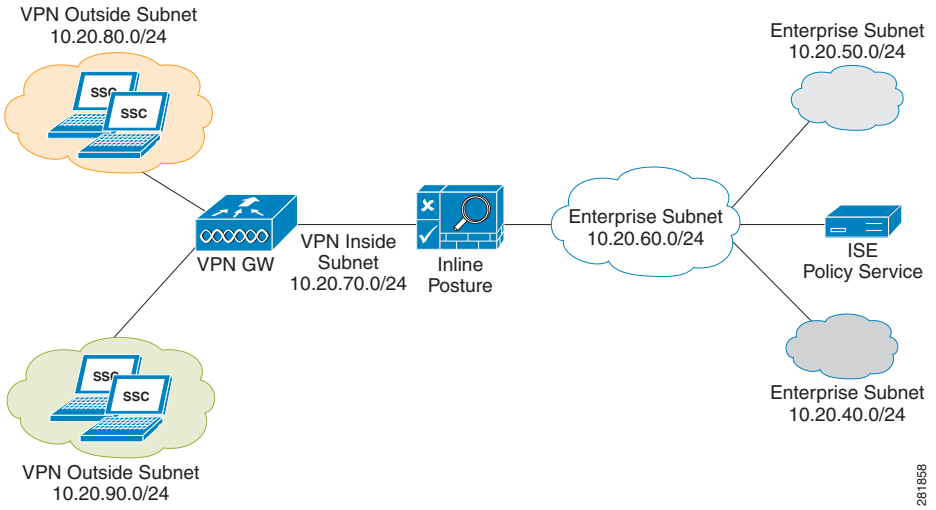

Inline Posture Bridged Mode

In bridged mode, the Inline Posture node operates as a standard Ethernet bridge. This configuration is typically used when the untrusted network already has a gateway, and you do not want to change the existing configuration.

Figure 10-3 shows the Inline Posture node acting as a bridge for the Layer 2 client traffic from the WLC into the Cisco ISE network, managed by the Policy Service ISE node. In this configuration, Inline Posture requires subnet entries for the 10.20.80.0/24 and 10.20.90.0/24 subnets to be able to respond to and send Address Resolution Protocol (ARP) broadcasts to the correct VLANs.

Figure 10-3 Inline Posture Bridged Mode Configuration

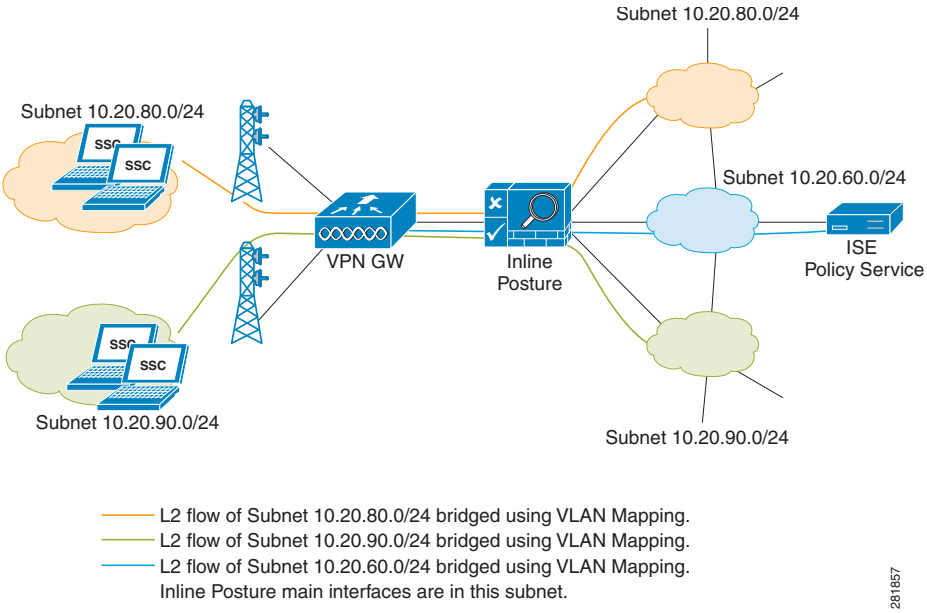

When the Inline Posture node is in bridged mode, the following conditions apply:

- Inline Posture eth0 and eth1 can have the same IP address.
- All end devices in the bridged subnet must be on the untrusted network.

Best Practices for Inline Posture

This section introduces best practice concepts for deploying Inline Posture in a distributed environment.

Using Filters to Define Access Privileges

Consider the following when configuring filters:

- As typically implemented, Inline Posture enforces authentication requirements on endpoints that attempt to access the network. Device and subnet filters are used to validate or deny WLC and VPN devices.
- For certain devices, you may want to bypass authentication, posture assessment, role assignment, or any combination thereof. Common examples of bypassed device types include printers, IP phones, servers, nonclient machines, and network devices.

  Inline Posture matches the MAC address of a device, or a MAC and IP address combination, or a subnet address to determine whether the bypass function is enabled for a device. You can choose to bypass policy enforcement, or to forcibly block access.

Warning **Do not configure the MAC address in a MAC Filter for a directly connected adaptive security appliance (ASA) VPN device without also entering the IP address. Without the addition of the (optional) IP address, VPN clients are allowed to bypass policy enforcement. This bypass happens because the VPN is a Layer 3 hop for clients, and the device uses its own MAC address (as the source address) to send packets along the network toward the Inline Posture node.**

Configuring Managed Subnets and Static Routes

Consider the following when configuring managed subnets for Inline Posture:

- Configure managed subnets for endpoints in Layer 2 proximity of the Inline Posture node. For example, a WLC that delivers packets directly to the untrusted interface of the Inline Posture node.
- When configuring subnets for endpoints in Layer 2 proximity to an Inline Posture node, you must also configure a managed subnet for Inline Posture. This configuration ensures that the Inline Posture node can send Address Resolution Protocol (ARP) queries with the appropriate VLAN IDs for the client devices on the untrusted interface. Configure the untrusted (authentication) VLAN in the VLAN ID field for the managed subnet.
- When configuring a managed subnet for Inline Posture, configure an IP address and not a subnet address. This configuration ensures that the ARP requests that Inline Posture sends have a valid source IP address.
- Subnets on the trusted side of the Inline Posture node should be dissimilar to subnets on the untrusted side.
- An Administration ISE node and Inline Posture node should not be on the same subnet, unless you have defined a static route.

Consider the following when configuring static routes for Inline Posture:

- Configure static routes for endpoints that are more than one hop away (Layer 3) from the Inline Posture node.
- Static routes should be configured for all downstream host networks that are typical of VPN address pools.

High Availability

Consider the following when configuring Inline Posture for high availability:

- Assign a service IP (also known as a virtual IP) for each side of the Inline Posture interfaces, trusted (eth0) and untrusted (eth1).
- Specify link-detect IP addresses for the trusted (eth0) and untrusted (eth1) interfaces. Link-detect appears as an optional setting in the user interface, but is highly recommended.

Standalone Mode or High Availability

One of the most important decisions you will make with regard to your Inline Posture deployment, is whether to deploy a single, standalone node, or an active-standby pair to ensure high availability.

A standalone Inline Posture node is simply a single Inline Posture node that provides services and works independently of all other nodes. You might choose to deploy a single standalone Inline Posture node for a network that serves a small facility, where redundancy is not a major concern.

An Inline Posture high availability deployment consists of two Inline Posture nodes that are configured as an active-standby pair. The active node acts as the RADIUS proxy, forwarding all the network packets until such time that it fails, then the standby node takes over. As long as the active node is functioning properly, the standby node remains passive. However, should the active node falter, the standby node takes over to perform Inline Posture functionality.

Figure 10-2 illustrates a simple Inline Posture standalone configuration, with client access through WLC and VPN devices. Figure 10-4 illustrates a routed mode high availability Inline Posture configuration.

Inline Posture High Availability

Inline Posture stateless high availability deployment has an active-standby pair node configuration, where the standby 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 an active node are automatically authorized again after a failover occurs.

The standby node monitors the active node using the heartbeat protocol (using eth2 and eth3 interfaces), which 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.

Note The heartbeat protocol that is active in an Inline Posture high availability configuration requires a direct Ethernet cable connection between the eth2 interfaces of both nodes of a high availability pair. Likewise, there must be a direct Ethernet cable connection between the eth3 interfaces of the two nodes. Figure 10-4 illustrates this principle.

In addition to the heartbeat monitor, an optional (but highly recommended) link-detect mechanism is available. With the use of link-detect, Inline Posture trusted and untrusted interfaces ping an external IP address from their respective interfaces. If both nodes are unable to ping the external IP address, then failover does not occur. However, if either of the nodes becomes unreachable, the node that is functional automatically becomes the active node.

Upon failover, the following occurs:

1. The standby Inline Posture node takes over the service IP address (SIP).
2. Once the failover happens, the administrator corrects the failed node and reverts to an earlier configuration, as needed.

   When a failed node is brought back online, a manual sync operation to update the node with the most current information is required. For information on how to perform an Inline Posture node sync operation, see Syncing an Inline Posture Node, page 10-29.
3. Active sessions are automatically reauthenticated and authorized.

Key Points for High Availability

- The terms primary and secondary have different meanings with regard to Inline Posture high availability than they do in relation to Cisco ISE nodes. For Inline Posture high availability, primary and secondary denote the device that takes over the active state and the device that takes the standby role in case there is a contention, such as when both nodes boot up at the same time.
- The terms active and standby are representative of high availability states. A primary or secondary Inline Posture node can be in either an active or standby state.
- If the heartbeats simultaneously go down for both Inline Posture high availability nodes, a partitioning state may ensue. A partitioning state is a condition where both nodes assume that the other has totally failed, and both try to take over active control.
- The secondary Inline Posture node is read-only, and cannot be used for configuration of any kind, even high availability.
- The eth2 and eth3 interfaces of both nodes in an Inline Posture high availability pair (primary and secondary) communicate with heartbeat protocol exchanges to determine the health of the nodes. For the heartbeat to work, you must connect the eth2 interface of the primary Inline Posture node to

the eth2 interface of the secondary node using an Ethernet cable. Likewise, the eth3 interface of the primary Inline Posture node must be connected to the eth3 interface of the secondary node with an Ethernet cable. Figure 10-4 illustrates this principle.

Note A heartbeat is a message that is sent from one node in an Inline Posture high availability pair to the other member of the pair at regular intervals. If a heartbeat is not received for an extended period of time, usually several heartbeat intervals, the node that should have sent the heartbeat is assumed to have failed. If it is the primary Inline Posture node that fails, the secondary node takes over so there is no disruption in service.

- When a node in a high availability pair is down and configuration changes are made to the single active node, there is no mechanism that automatically populates the failed node with the new configuration when it comes back up. The Sync-up Peer Node button that appears in the Inline Posture high availability user interface on the active node, allows you to manually sync the standby node with the latest Inline Posture database from the active node.
- For high availability, you register two Inline Posture nodes, then choose one node to be primary and enable high availability. For more information, see Configuring Inline Posture for High Availability, page 10-25.

Configuring Inline Posture High Availability in Routed Mode

An Inline Posture high availability (HA) pair consists of two physical Inline Posture nodes configured as a cluster that have heartbeat links on the eth2 and eth3 interfaces, connected by dedicated cables. Each Inline Posture node has its own physical IP addresses on the trusted and untrusted Ethernet interfaces, but a separate service IP address must be assigned to the cluster as a whole.

Note The service IP address, also called a virtual IP address, is required for RADIUS authentication purposes. You assign the SIP to both the trusted and untrusted interfaces for both nodes of the active-standby pair, thus making the SIP the address of the cluster, representing it as a single entity to the rest of the network.

For example, the untrusted IP address for IPEP1 can be10.20.70.101, and the untrusted IP address for IPEP2 can be 10.20.70.102. However, the service IP address for both nodes on the untrusted side of the network would be10.20.70.100. The active Inline Posture node in the pair, at any point of time, assumes the service IP address on the untrusted side of the network. The same holds true for the trusted side of the network.

Figure 10-4 shows an example of an Inline Posture high availability routed mode configuration. Note the dedicated cables that connect the eth2 and eth3 interfaces between the two nodes to facilitate the heartbeat communication that checks for failure in the active node.

Figure 10-4 Inline Posture Routed Mode High Availability Example

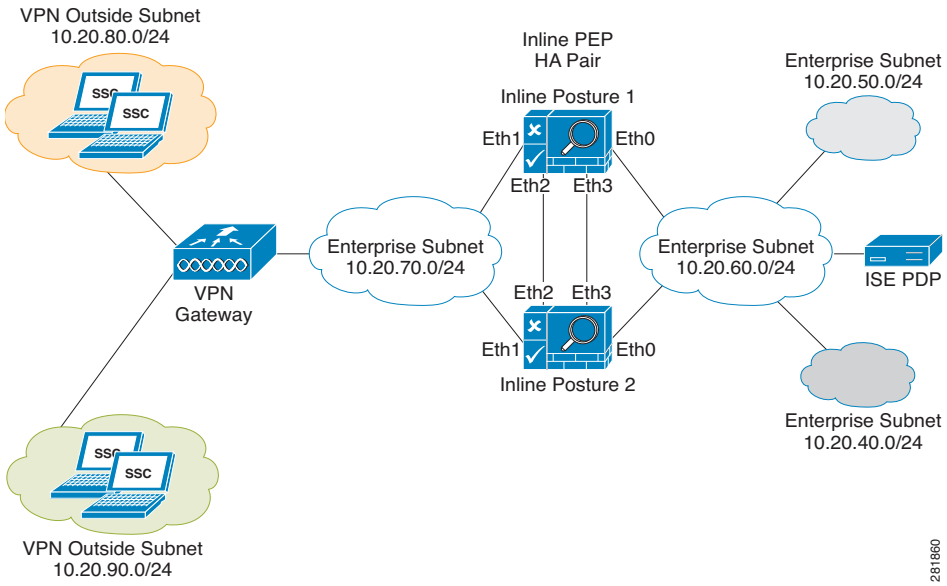

Configuring Inline Posture High Availability in Bridged Mode

The following guidelines apply to an Inline Posture bridged mode high availability configuration:

- Inline Posture eth0 and eth1 should have IP addresses in the same subnet. Having the same IP address is recommended.
- Any devices on the trusted side of the network that have IP addresses in the subnets that are managed by an Inline Posture in bridged mode, must have an explicit static route configured at the Inline Posture node. This configuration is necessary because by default, Inline Posture assumes that the subnet that it manages (as configured on the Managed Subnets user interface page) lies entirely on the untrusted side of the network.

Inline Posture Guidelines for Distributed Deployment

Before you begin configuring an Inline Posture node in a distributed deployment, be sure you understand the following statements:

1. Inline Posture is unable to run concurrently with Administration, Policy Service, or Monitoring personas, and therefore is a dedicated node.
2. An Inline Posture node must be registered as a secondary node to the primary Administration ISE node on your network.
3. You can deploy a standalone Inline Posture node, or an active-standby pair.
4. You can have up to two Inline Posture nodes configured on your network at any one time. For an Inline Posture high availability active-standby pair, two nodes are configured. One node is designated as the primary node and the other as the secondary node. The primary node has the preference for being the active node when both nodes come up at the same time.
5. For an Inline Posture active-standby pair configuration, all configuration related to functionality must be done from the active node of the pair. The user interface for the standby node, in the Cisco ISE user interface, shows only basic configuration tables.

6. You can sync an Inline Posture active node configuration to its peer standby node from the Failover tab of the active node. For more information, see Syncing an Inline Posture Node, page 10-29.

Note If you have a WLC authentication, authorization, and accounting (AAA) server (Cisco 2100 or 4400 Series Wireless LAN controllers) on your network, the RADIUS authentication server timeout value needs to be set to a minimum of 30 seconds. This minimum value ensures that RADIUS failover will work in conjunction with Inline Posture. See the WLC server hardware documentation for more information.

Deploying an Inline Posture Node

The initial process for configuring an Inline Posture node is the same, whether it is intended to be a standalone node or part of an active-standby pair. This section contains the series of tasks you must complete to configure an Inline Posture node on your Cisco ISE network.

To configure an Inline Posture node, complete the following tasks:

1. Configuring Inline Posture in Bridged or Routed Mode, page 10-12
2. Creating Inline Posture Downloadable Access Control Lists, page 10-20
3. Creating Inline Posture Node Profiles, page 10-22
4. Creating an Inline Posture Authorization Policy, page 10-23

Configuring Inline Posture in Bridged or Routed Mode

To introduce an Inline Posture node in your Cisco ISE network you must first register the Inline Posture node with the primary Policy Service ISE node, configure the Inline Posture settings, and then create authorization profiles and policies that establish the Inline Posture gatekeeping policies.

The Inline Posture node is a RADIUS proxy that interfaces with NADs as their RADIUS server, making the NADs (VPN gateway, WLC) RADIUS clients. As a proxy, Inline Posture interfaces with the Policy Service ISE node as a client, making the Policy Service ISE node its RADIUS server.

Note Upon completing the following procedure, a NAD entry is automatically created for the Inline Posture node. For a standalone node, the IP address for that node is used. For an HA pair, the service IP address for the active node is used.

Guidelines for Configuring Certificates for Inline Posture

Secure communication between Administration and Inline Posture nodes requires mutual authentication. This means that not only must the Inline Posture node prove its identity to the Administration node, but the reverse is also true.

For a proper communication between Administration and Inline Posture nodes, the primary Administration node local certificate should have both Client Authentication and Server Authentication EKU attributes.

Observe the following guidelines when configuring certificates on these nodes:

- The presence of certain combinations of attributes in the local certificates of the Administration and Inline Posture nodes can prevent mutual authentication from working.

  The attributes are:

  - Extended Key Usage (EKU)—Server Authentication
  - Extended Key Usage (EKU)—Client Authentication
  - Netscape Cert Type—SSL Server Authentication
  - Netscape Cert Type—SSL Client Authentication

  Either of the following combinations is required for the Administration certificate:

  - Both EKU attributes should be disabled, if both EKU attributes are disabled in the Inline Posture certificate, or both EKU attributes should be enabled, if the server attribute is enabled in the Inline Posture certificate.
  - Both Netscape Cert Type attributes should be disabled, or both should be enabled.

  Either of the following combinations is required for the Inline Posture certificate:

  - Both EKU attributes should be disabled, or both should be enabled, or the server attribute alone should be enabled.
  - Both Netscape Cert Type attributes should be disabled, or both should be enabled, or the server attribute alone should be enabled.

- Where self-signed local certificates are used on the Administration and Inline Posture nodes, you must install the self-signed certificate of the Administration node in the trust list of the Inline Posture node. In addition, if you have both primary and secondary Administration nodes in your deployment, you must install the self-signed certificate of both Administration nodes in the trust list of the Inline Posture node.
- Where CA-signed local certificates are used on the Administration and Inline Posture nodes, mutual authentication should work correctly. In this case, the certificate of the signing CA is installed on the Administration node prior to registration, and this certificate is replicated to the Inline Posture node.
- If CA-issued keys are used for securing communication between the Administration and Inline Posture nodes, before you register the Inline Posture node, you must add the public key (CA certificate) from the Administration node to the CA certificate list of the Inline Posture node.

Prerequisites

- You should have administrative permissions on the primary Administration ISE node.
- Follow and apply the Guidelines for Configuring Certificates for Inline Posture, page 10-12.
- Register the Inline Posture node with the primary Administration ISE node, as described in Registering and Configuring a Secondary Node, page 9-13. All nodes must be registered with the primary Administration ISE node to function as a member of the Cisco ISE distributed system. Be sure to check the Inline Posture check box. The Administration, Monitoring, and Policy Service check boxes are automatically unchecked.

Note Registering an Inline Posture node results in a system restart. Likewise, changes to infrastructure configurations, such as the eth1 IP address, Inline Posture mode, and high availability changes also require a system restart. The restart is automatic. However to manually restart the node from the CLI, use the **application stop ise** and **application start ise** commands.

- RADIUS configuration is mandatory. At least one client and one server configuration is necessary. You need the corresponding shared secret information for both sides to complete this procedure.

- Have all necessary configuration information for your installation on hand. For example, you might need the trusted and untrusted IP addresses, service IP address, the IP addresses for other Cisco ISE nodes, shared secret for RADIUS configuration, management VLAN ID, WLC, or VPN IP address, and so on. Check with your system architect for a complete list of the information you will need.

Warning **Do not configure the MAC address in a MAC Filter for a directly connected ASA VPN device without also entering the IP address. Without the addition of the (optional) IP address, VPN clients are allowed to bypass policy enforcement. This access happens because the VPN is a Layer 3 hop for clients, and the device uses its own MAC address (as the source address) to send packets along the network toward the Inline Posture node.**

To configure Inline Posture in bridged or routed mode, complete the following steps:

Step 1 From the primary Administration ISE node, choose **Administration > System > Deployment**.

Step 2 Click **Deployment** in the Deployment navigation pane, and then in the Deployment Nodes page, check the **Inline Posture node** check box and click **Edit**.

Step 3 On the General Settings tab, check the **Inline PEP** check box. The Administration, Monitoring, and Policy Service check boxes are automatically unchecked.

Figure 10-5 Edit Inline Posture Node

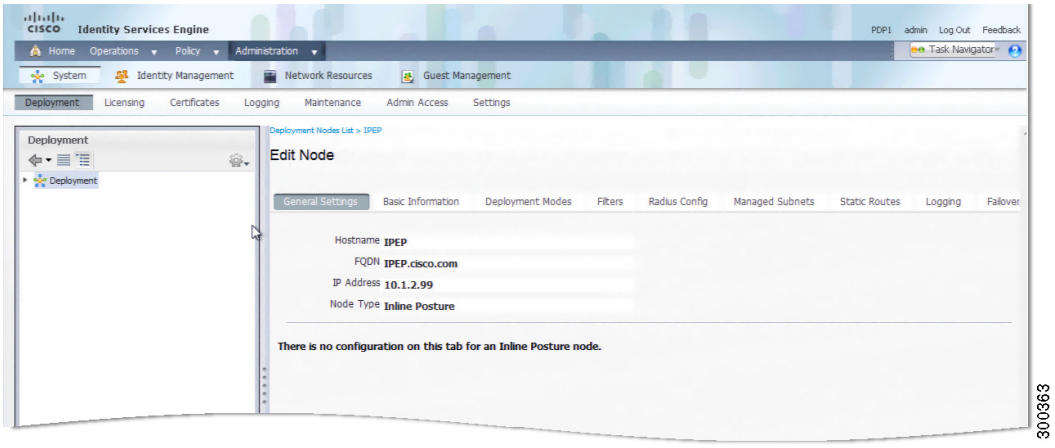

The tabs change to General Settings, Basic Information, Deployment Modes, Filters, Radius Config, Managed Subnets, Static Routes, Logging, and Failover.

Note A newly registered Inline Posture node comes up with a default IP address of 192.168.1.100, a subnet mask of 255.255.255.0, and a default gateway of 192.168.1.1. Change these values to fit your deployment in Step 3.

Step 4 Click the **Basic Information** tab and enter the appropriate information for the following options:

- Time Sync Server: Primary, Secondary, Tertiary
- DNS Server: Primary, Secondary, Tertiary

- Trusted Interface (to protected network): Set Management VLAN ID (all the other information is automatically populated for these options)
- Untrusted Interface (to management network): IP Address, Subnet Mask, Default Gateway, Set Management VLAN ID

Figure 10-6 is an example of a bridged mode configuration. Figure 10-7 is an example of a routed mode configuration.

Figure 10-6 Basic Information (Bridged)

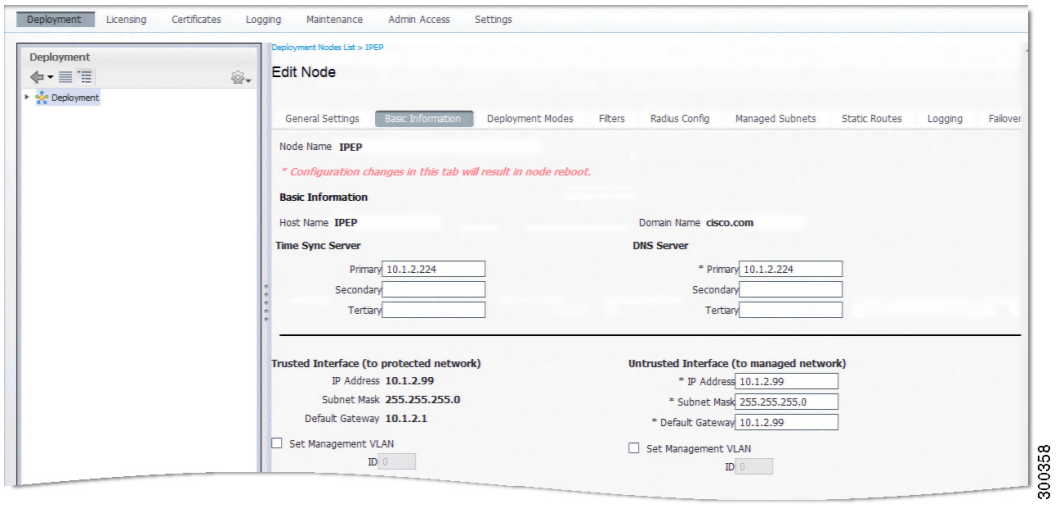

Figure 10-7 Basic Information (Routed)

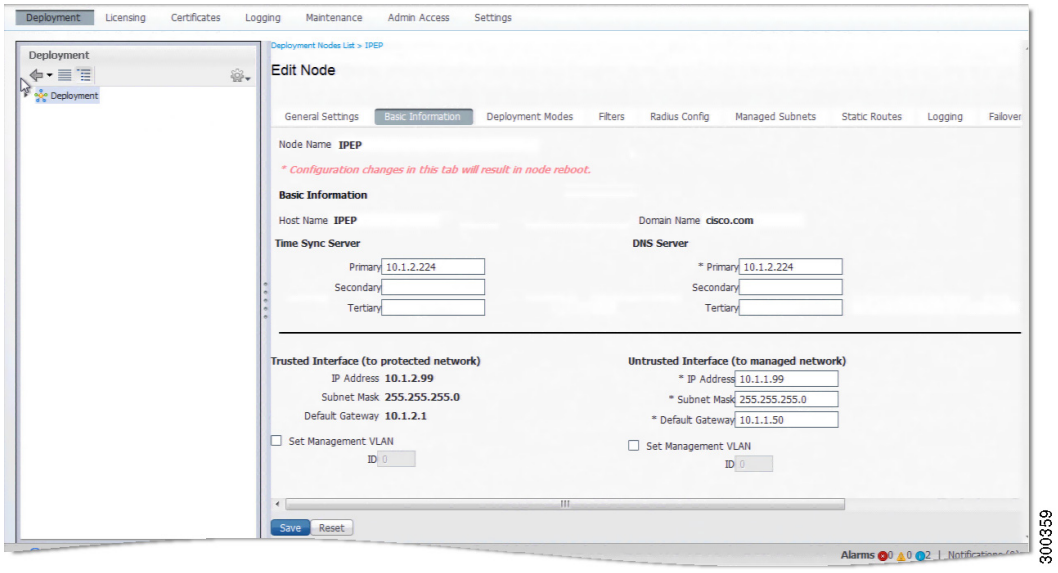

Step 5 Click the **Deployment Modes** tab. A newly registered Inline Posture node comes up in maintenance mode. For production purposes, choose one of the following:

- Routed Mode—Provides router (hop in the wire) functionality for Inline Posture. Figure 10-8 provides an example for routed mode.

- Bridged Mode—Provides VLAN mapping functionality for the subnets to be managed by Inline Posture. After checking the Bridged Mode check box, enter the Untrusted Network and Trusted Network VLAN ID information. Figure 10-9 provides an example for bridged mode.

  For VLAN mapping, you should also do the following:

  – Add a mapping for management traffic by entering the appropriate VLAN ID for the trusted and untrusted networks.

  – Add a mapping for client traffic by entering the appropriate VLAN ID for the trusted and untrusted networks.

Figure 10-8 Deployment Modes (Routed)

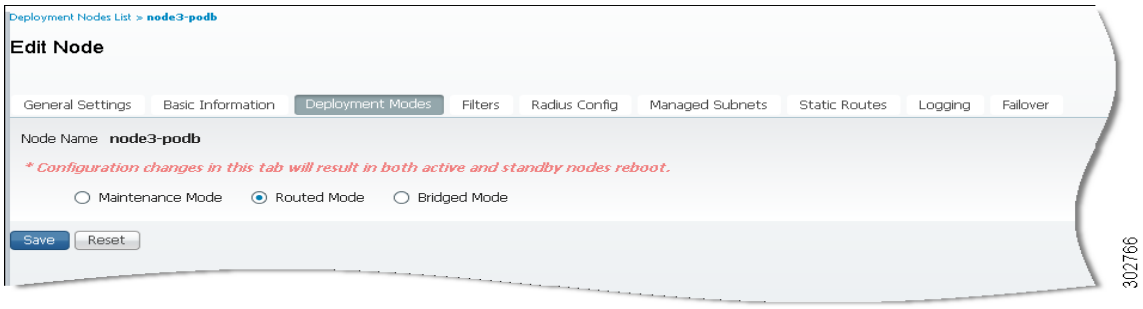

Figure 10-9 Deployment Modes (Bridged)

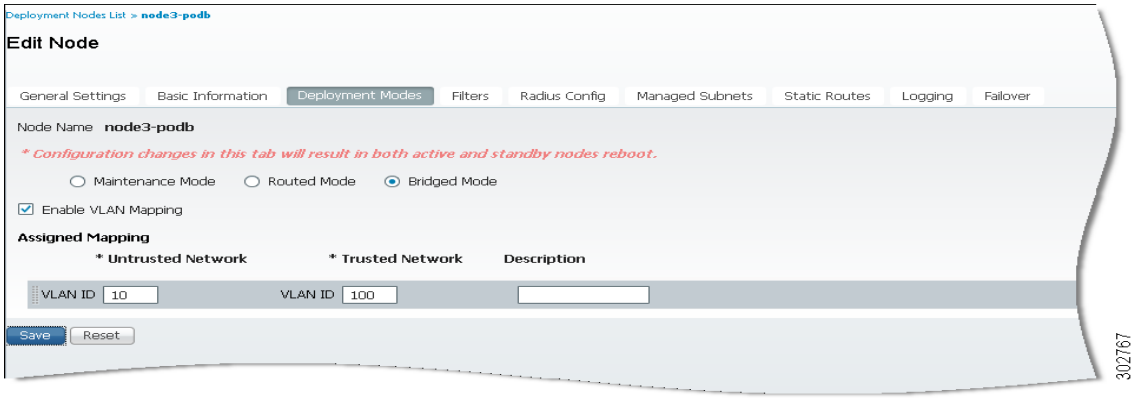

Step 6 Click the **Filters** tab and enter the subnet address and subnet mask for the client device, or the MAC address and IP address of the device on which to filter.

You can use MAC and subnet filters to bypass Inline Posture enforcement to certain endpoints or devices on the untrusted side of the network. For example, if VPN or WLC management traffic is required to pass through Inline Posture, you would not want to subject those particular NADs to Cisco ISE policy enforcement. By providing the MAC address and IP address for these NADs on a filter, you can then access the user interface or configuration terminal by way of Inline Posture without restrictions.

- MAC filters—MAC address and/or IP address on which to avoid policies
- Subnet Filters—Subnet address and subnet mask on which to avoid policies

Note For security reasons, we recommend that you always include the IP address along with the MAC address in a MAC filter entry. For more information, see the Warning in Prerequisites, page 10-13.

Figure 10-10 Filters

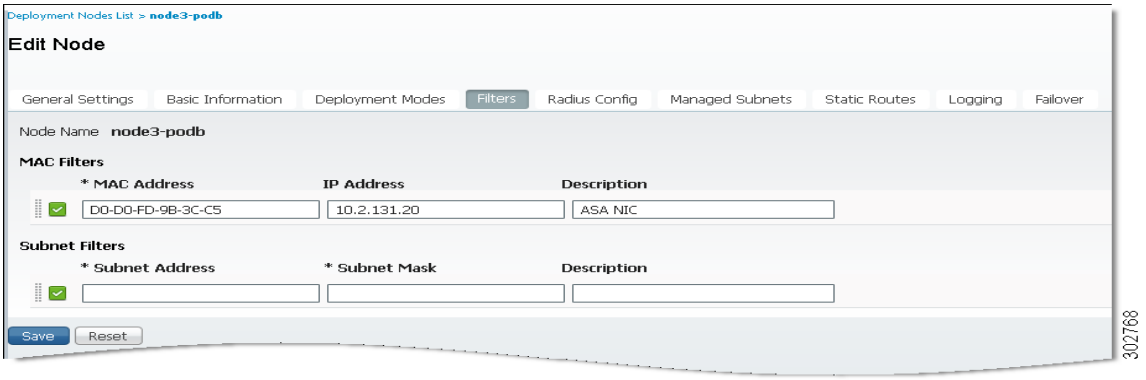

Step 7 Click the **RADIUS Config** tab and enter the IP address and shared secret for the following:

- Primary Server—Primary RADIUS server, usually the Policy Service ISE node
- Secondary Server—Optional
- Client—Device that requests access on behalf of clients, WLC or VPN

Note WLC roaming is not supported in Cisco ISE Release 1.1.x.

RADIUS configuration is mandatory. At least one client and one server configuration is necessary for Inline Posture. For more information on RADIUS proxy services, see Proxy Service, page 16-21.

Figure 10-11 RADIUS Configuration

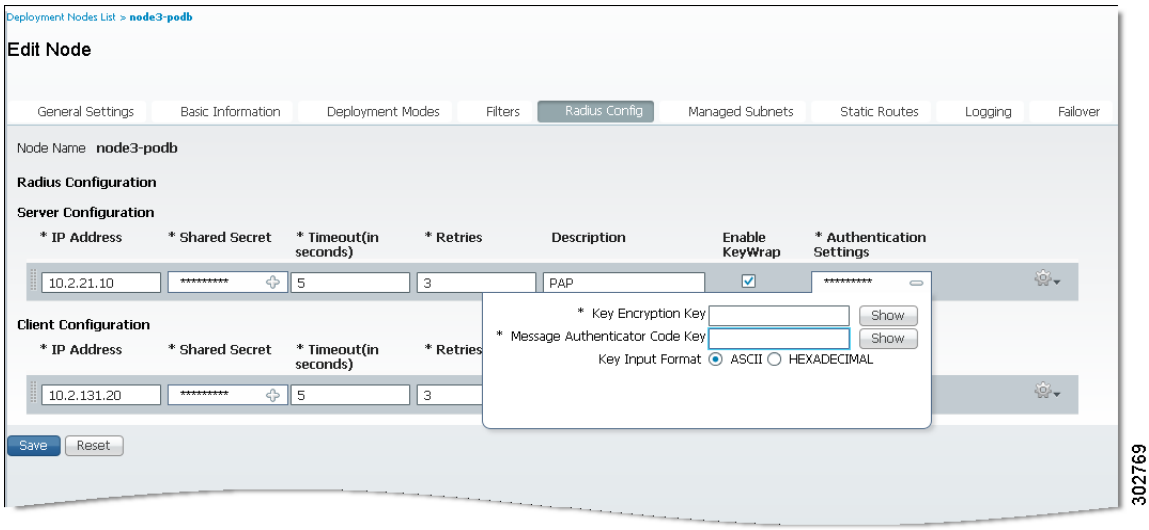

Step 8 (Optional) Check the Enable KeyWrap check box and specify the following Authentication Settings:

- Key Encryption Key
- Message Authenticator Code Key
- Key Input Format: ASCI or Hexidecimal

Deployments that utilize wireless LAN technology require secure transmission from a RADIUS server to a network access point. KeyWrap attributes provide stronger protection and more flexibility.

Step 9 Click the **Managed Subnets** tab, and enter the following information for each Managed Subnet:

- IP Address
- Subnet Mask
- VLAN ID
- Description

For subnets of endpoints that are in Layer 2 proximity to the Inline Posture node (such as a WLC), you must configure managed subnets. This configuration requires an unused IP address in the same subnet as the managed subnet, along with the VLAN (if any) of the subnet. You can have multiple managed subnet entries.

Figure 10-12 Managed Subnets

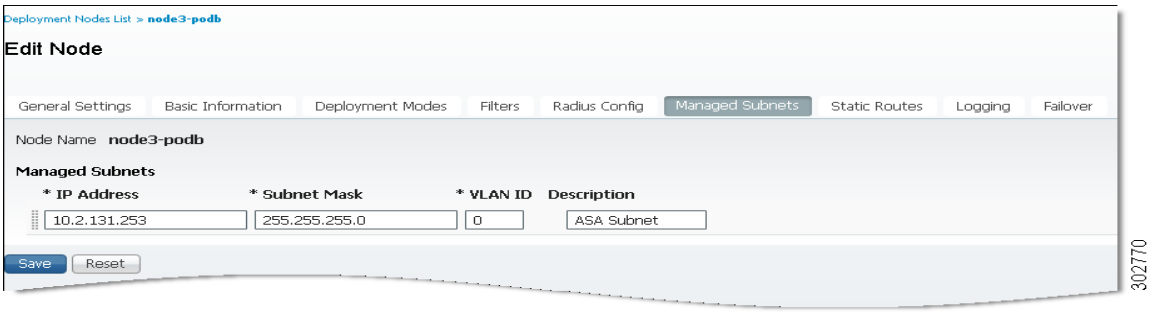

Step 10 Click the **Static Routes** tab, then enter the subnet address, subnet mask, and choose **Trusted** or **Untrusted** from the Interface Type drop-down list. Repeat this step as needed for your configuration.

When the subnets of the endpoints under Cisco ISE control are Layer 3 away from the Inline Posture node, a static route entry is needed. For example, if a VPN gateway device (that sends managed subnet traffic to the Inline Posture untrusted interface) is two hops away, its client subnet needs to have a static route defined for Inline Posture. The network on the trusted side should know to send traffic to the Inline Posture trusted interface.

Figure 10-13 Static Routes

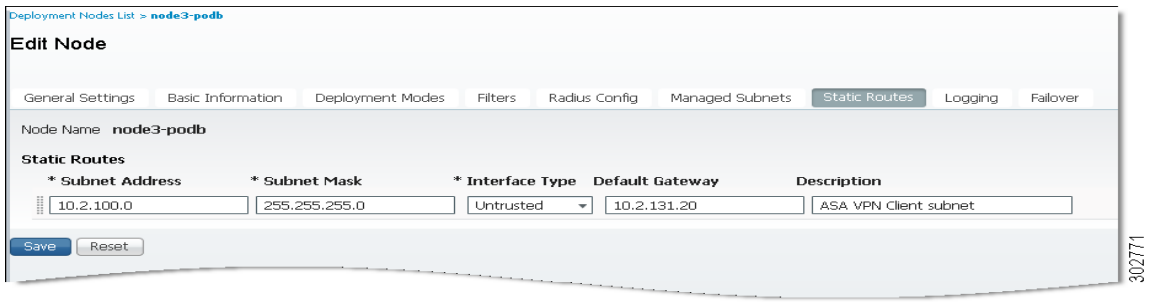

Step 11 Click the **Logging** tab and enter the IP address and port number for the logging server, which is typically the Monitoring ISE node.

An IP address and port (default 20514) for logging Inline Posture events are mandatory. This requirement ensures that the viable status of the Inline Posture node is displayed in the Cisco ISE dashboard in the System Summary dashlet, and that other log information regarding the nodes is available.

Figure 10-14 Logging

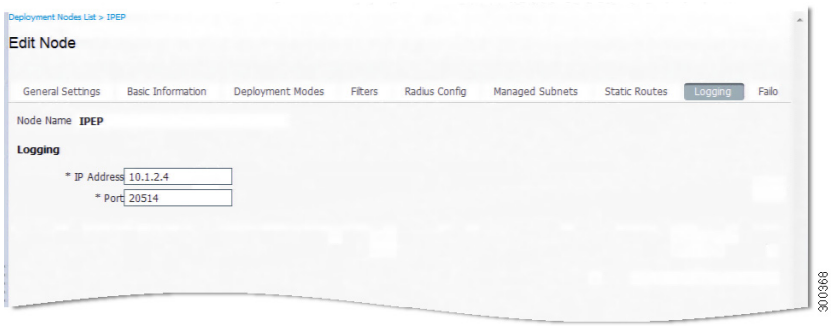

Step 12 Click **Save**. The node restarts.

Step 13 To verify the automatically generated Inline Posture NAD listing, go to **Administration > Network Resources > Default Device**.

For a standalone node, the IP address for that node is used. For an HA pair, the service IP address for the active node is used.

Next Steps

To complete the configuration setup of the Inline Posture node, complete the following tasks, creating three DACLs, authorization profiles, and authorization policy rules: unknown, compliant, and noncompliant.

1. Creating Inline Posture Downloadable Access Control Lists, page 10-20
2. Creating Inline Posture Node Profiles, page 10-22

Note It is important to associate the appropriate downloadable access control list (DACL) with the corresponding profile. For example, the unknown DACL should be associated with the unknown authorization profile.

3. Creating an Inline Posture Authorization Policy, page 10-23

Troubleshooting Topics

- Primary and Secondary Inline Posture Nodes Heartbeat Link Not Working, page D-7

Creating Inline Posture Downloadable Access Control Lists

Downloadable access control lists (DACLs) are building blocks for authorization profiles, and they provide the rules for the profiles to follow. Access control lists (ACLs) prevent unwanted traffic from entering the network by filtering source and destination IP addresses, transport protocols, and other variables, using the RADIUS protocol.

After you create DACLs as named permission objects, add them to authorization profiles, which you then specify as the result of an authorization policy. For more information on DACLs, see Understanding Authorization Policies, page 17-1.

Figure 10-15 Inline Posture DACLs

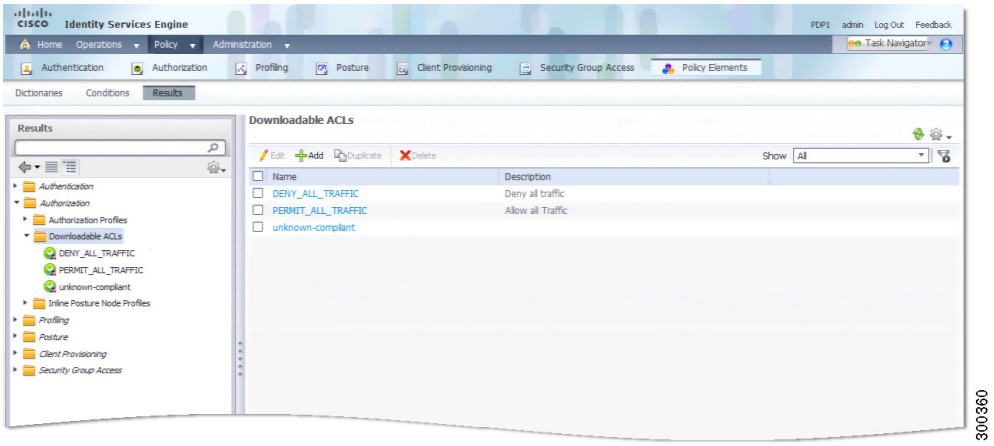

Note Every administrator account is assigned one or more administrative roles. Depending upon the roles assigned to your account, you may or may not be able to perform the operations or see the options described in the following procedure.

To create a DACL for Inline Posture, complete the following steps:

Step 1 Following the instructions as described in Configuring Permissions for Downloadable ACLs, page 17-34, create the following DACLs:

- ipep-unknown (Pre-Posture): Use at least one ACL to allow supplicants and the Policy Service to have access to each other for posture evaluation. This DACL can be used to block or quarantine users until they pass authentication. See Figure 10-16 for an example.
- ipep-compliant (Permit All): Use the following: permit ip any any
- ipep-noncompliant (Deny All): Use the following: deny ip any any

Figure 10-16 Inline Posture DACL Compliance Unknown

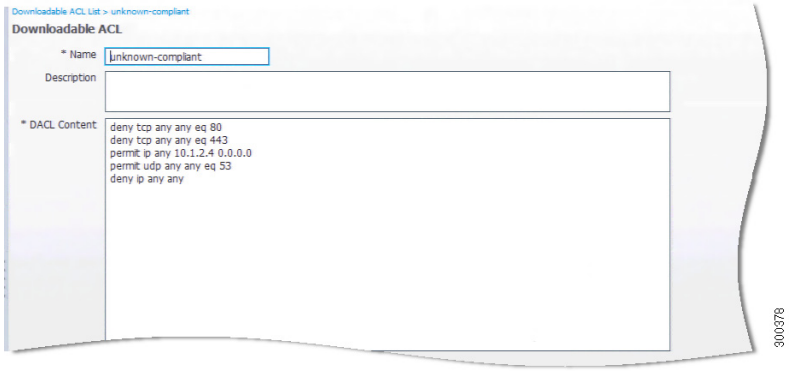

Figure 10-17 Inline Posture DACL Compliant

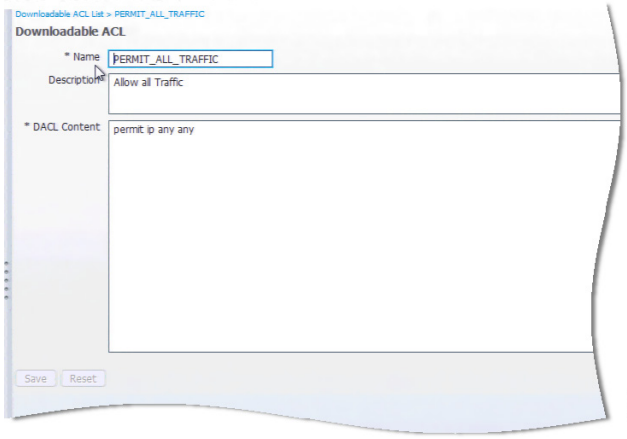

Step 2 Save the DACLs, and then go to Creating Inline Posture Node Profiles, page 10-22.

Troubleshooting Topics

- Primary and Secondary Inline Posture Nodes Heartbeat Link Not Working, page D-7

Creating Inline Posture Node Profiles

This section describes how to create authorization profiles for Inline Posture. You create three Inline Posture authorization profiles, as well an authorization profile for a NAD. For more information, see Cisco ISE Authorization Policies and Profiles, page 17-5.

All Inline Posture inbound profiles are automatically set to cisco-av-pair=ipep-authz=true so that the Inline Posture node is sure to apply these rules, instead of proxying them on to the NADs. The URL redirect is essential for client provisioning, as well as agent discovery redirection.

To create authorization profiles for NAD and Inline Posture, complete the following steps:

Step 1 Create a NAD authorization profile as described in Creating and Configuring Permissions for a New Standard Authorization Profile, page 17-29.

Note You can configure a RADIUS Reply Message = NAD Profile, to see *NAD Profile* in the RADIUS log messages for Inline Posture. This configuration can be helpful for troubleshooting at a later time.

Step 2 Create authorization profiles to Inline Posture that correspond to the DACLs you created in Creating Inline Posture Downloadable Access Control Lists, page 10-20.

Figure 10-18 Inline Posture Profiles

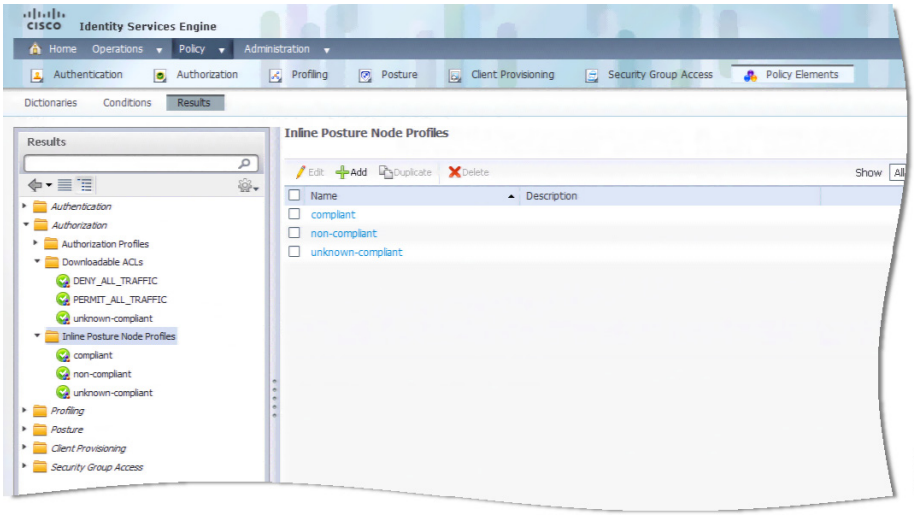

Specify the appropriate DACL for each of the following authorization profiles:

- Unknown-Compliant (Pre-Posture): This profile requires that you enter a URL redirect.

  From the Inline Posture Authorization Profiles page, select the Unknown-Compliant DACL name from the drop-down list, enter the following URL redirect in the text field, and click **Submit**:

  url-redirect=https://ip:8443/guestportal/gateway?sessionId=SessionValue&Action=cpp

  The URL redirect appears in the Attributes Details field.

Figure 10-19 Unknown-Compliant Authorization Profile

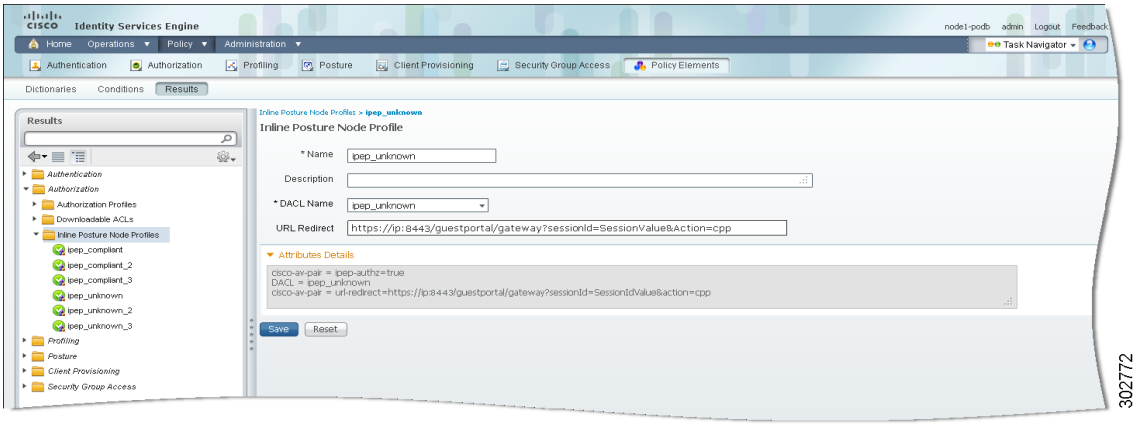

You are redirected to a web page where you download and install an agent. The agent then scans your system. If your system passes, you are automatically granted full access. If your system does not pass, you are denied access.

- IPEP-Compliant (Permit Any)
- IPEP-Noncompliant (Deny All)

Figure 10-20 Non-Compliant Authorization Profile

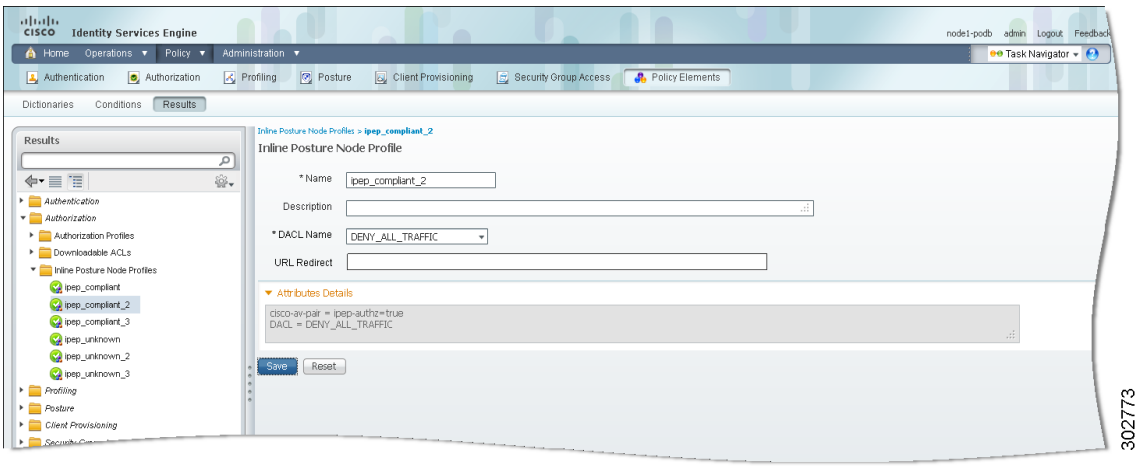

Step 3 After you have saved each of the authorization profiles, continue with Creating an Inline Posture Authorization Policy, page 10-23.

Creating an Inline Posture Authorization Policy

Authorization policies provide the means for controlling access to the network and its resources. Cisco ISE lets you define a number of different authorization policies.

The elements that define the authorization policy are referenced when you create policy rules. Your choice of conditions and attributes defines the authorization profile. Figure 10-21 shows the authorization rules that are necessary for VPN and WLC access.

Figure 10-21 Authorization Rules for VPN and WLC Access

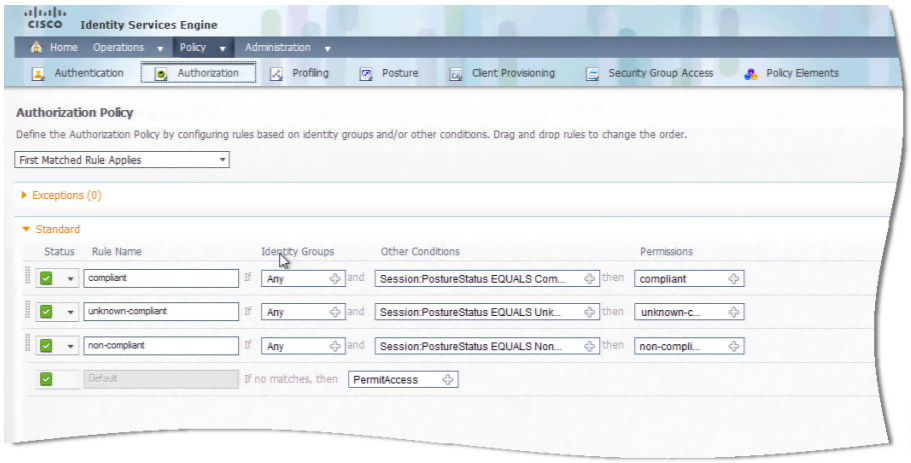

For more information on authorization policies, see Cisco ISE Authorization Policies and Profiles, page 17-5.

To create authorization policies, complete the following steps:

Step 1 Create an authorization policy as described in Creating a New Authorization Policy, page 17-15, leaving the default rule as is.

Step 2 Create the following Unknown Posture Status Rule:

- Identity Group: Any
- Condition: Session:PostureStatus EQUALS = Unknown
- Permissions: ipep-unknown-compliant + nad-authorization-profile

Step 3 Create the following Compliant Posture Rule:

- Identity Group: Any
- Condition: Session:PostureStatus EQUALS = Compliant
- Permissions: ipep-compliant + nad-authorization-profile

Step 4 Create the following Noncompliant Posture Rule:

- Identity Group: Any
- Condition: Session:PostureStatus EQUALS = Noncompliant
- Permissions: ipep-noncompliant + nad-authorization-profile

Step 5 Save the policy. The Inline Posture node configuration process is now complete.

Next Step

Complete the following task: Adding Inline Posture as a RADIUS Client, page 10-30.

Configuring Inline Posture for High Availability

This section explains how to configure two Inline Posture nodes for high availability. One node is specified as the primary unit in the pair and becomes the active node by default. The other becomes the secondary node, which is a standby unit in case of default.

A high availability node failover prompts the standby node to take over the service IP address. After this process occurs, an administrator must correct the failed Inline Posture node and revert it to the earlier configuration, as needed because high availability failover is stateless, all active sessions are automatically reauthorized after a failover occurs.

This section contains the following topics:

Configuring a High Availability Pair

This section shows you how to define a high availability relationship between two registered Inline Posture nodes.

In the example that is presented, the service IP address used for the bridged mode high availability pair is different from the physical IP addresses of the Inline Posture nodes, effectively creating a cluster. The WLC interacts with the cluster as a single unit, using the service IP address. For this reason, the service IP is defined for the trusted and untrusted networks.

Configuring Primary and Secondary Inline Posture Nodes

Warning **Both nodes in a high availability pair must use the same mode, either bridged or router. Mixed modes are not supported on Inline Posture high availability pairs.**

Prerequisites

- You should have administrative permissions on the primary Administration ISE node.
- You should have successfully configured two (2) Inline Posture nodes, and registered them on the Cisco ISE network as described in Configuring Inline Posture in Bridged or Routed Mode, page 10-12.
- The eth2 and eth3 interfaces of both nodes in an Inline Posture high availability pair (primary and secondary) communicate with heartbeat protocol exchanges to determine the health of the nodes. For the heartbeat to work, you must connect the eth2 interface of the primary Inline Posture node to the eth2 interface of the secondary node using an Ethernet cable. Likewise, the eth3 interface of the primary Inline Posture node must be connected to the eth3 interface of the secondary node with an Ethernet cable. Figure 10-4 illustrates this principle.
- For RADIUS purposes, you need a service IP address that you will assign to both the trusted and untrusted interfaces of the Inline Posture active-standby cluster during in the course of this procedure.
- Have all necessary network configuration information for your installation on hand. For example, you will need the IP addresses for both Inline Posture nodes, a service IP address for the cluster, the IP address for the Policy Service ISE node, and the shared secret for RADIUS configuration. You might also need the management VLAN ID, WLC IP address, VLAN IP address, and so on. Check with your system architect for a complete list of the information you will need.

To configure an Inline Posture high availability pair, complete the following steps:

Step 1 From the primary Administration ISE node, choose **Administration > System > Deployment**.

Step 2 Click the **Deployment** link in the Deployment navigation pane. Then, in the Deployment Nodes page, check the check box next to the Inline Posture node that you want to designate as the primary node, and click **Edit**.

Step 3 On the General Settings tab, verify the node name, that the Inline PEP check box is selected, then choose **Active** as the HA Role from the drop-down list.

Step 4 Click the **Failover** tab, and check the **HA Enabled** check box.

Step 5 Choose the **HA Peer Node** from the drop-down list. A list of eligible standalone Inline Posture nodes appears from which to choose.

Step 6 Specify the following for the active node:

a. Enter the Trusted Service IP address (eth0) and the Untrusted Service IP address (eth1) for the traffic interfaces of the primary node. In the bridged mode example that follows, the service IP address is the same for both trusted and untrusted networks.

b. Optionally (but recommended as a best practice), enter the IP address for the Link-Detect system for both the trusted and untrusted sides. This address is usually the IP address for the Policy Service ISE node, because both the active and standby nodes should always be able to reach the Policy Service ISE node.

Then, Enter a Link-Detect Timeout value. The default value of 30 seconds is recommended. However, there is no maximum value.

Link-detect ensures that the Inline Posture node maintains communication with the Policy Service ISE node. If the active node does not receive notification (ping) from the Policy Service ISE node at the specified intervals, the active node fails over to the standby node.

Step 7 Enter a Heart Beat Timeout value. The default value of 30 seconds is recommended. However, there is no maximum value.

The heartbeat is a message that is sent between the two Inline Posture nodes at specified intervals. The heartbeat happens on eth2 and eth3 interfaces. If the heartbeat stops or does not receive a response in the allotted time, failover occurs.

Figure 10-22 Failover

Step 8 Choose the **HA Peer Node** from the drop-down list. The secondary node syncs to the primary node.

- Replication Status—(Only appears for secondary nodes) Indicates whether incremental replication from the primary node to the secondary node is complete or not. You will see one of the following states:
  - Failed—Incremental database replication has failed.
  - In-Progress—Incremental database replication is currently in progress.
  - Complete—Incremental database replication is complete.
  - Not Applicable—Displayed if the ISE node is a standalone or primary node.
- Sync Status—(Only appears for secondary ISE nodes) Indicates whether replication from the primary node to the secondary node is complete or not. A replication happens when a node is registered as secondary or when you click Syncup to force a replication. You will see one of the following states:
  - Sync Completed—Full database replication is complete.
  - Sync in Progress—Database replication is currently in progress.
  - Out of Sync—Database was down when the secondary node was registered with the primary ISE node.
  - Not Applicable—Displayed if the ISE node is a standalone node.

Step 9 If the sync status for any secondary node is out of sync, check the check box next to that node, and click **Syncup** to force a full database replication.

Note You must use the Syncup option to force a full replication if the Sync Status is *Out of Sync* or the Replication Status is *Failed*.

Step 10 Click **Save**. Both Inline Posture nodes restart.

When the nodes come back up, they are configured as primary and secondary, according to the settings you specified. You can view the state of a node by selecting the node to edit, as described in Step 2, and then clicking the Failover tab.

Note that the primary node has more options available for editing. That is because you make all configuration changes on the primary node. Configuration changes made to the primary node are automatically populated onto the secondary node. For this reason, the secondary node is read-only.

The following figures compare the Failover tabs of the active primary and standby secondary Inline Posture nodes.

Figure 10-23 Inline Posture Active Options

Figure 10-24 Inline Posture Standby Options

Next Step

Complete the following task: Adding Inline Posture as a RADIUS Client, page 10-30.

Troubleshooting Topics

- Primary and Secondary Inline Posture Nodes Heartbeat Link Not Working, page D-7

Syncing an Inline Posture Node

The procedure that is covered in this section assumes that you have already configured two Inline Posture nodes in an active-standby pair. The purpose of this section is to show you how to sync one node in an active-standby pair to the other node.

Prerequisites

- You should have successfully configured two Inline Posture nodes, as described in Configuring Inline Posture in Bridged or Routed Mode, page 10-12.
- You should have successfully established the relationship between the two nodes, as described in Configuring a High Availability Pair, page 10-25.
- You should have administrative permissions on the primary Administration ISE node.

To sync one Inline Posture node to another, complete the following steps:

Step 1 From the primary Administration ISE node, choose **Administration > System > Deployment**.

Step 2 Click the **Deployment** link in the Deployment navigation pane.

Step 3 In the Deployment Nodes page, check the check box next to the Inline Posture node to which you want to sync the other node (usually the active node), and click the **Edit** icon.

Step 4 Click the **Failover** tab.

Step 5 Click **Sync Peer Node**.

Data from the selected node is automatically transferred to its peer node.

Figure 10-25 Sync Peer Node

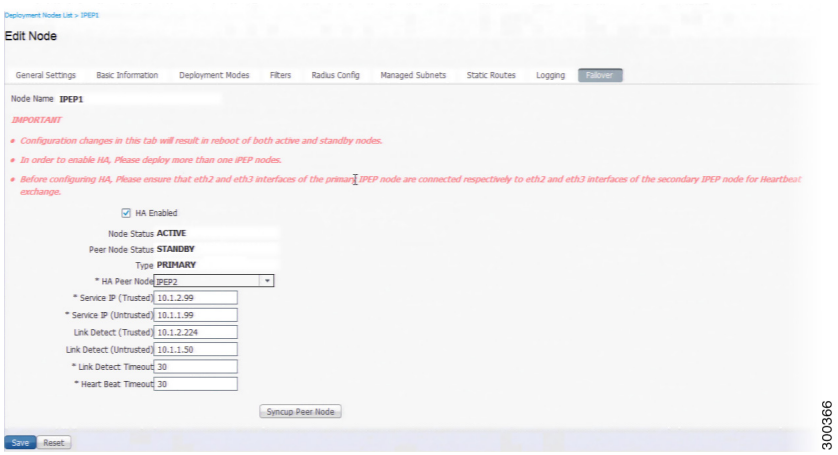

Troubleshooting Topics

- Primary and Secondary Inline Posture Nodes Heartbeat Link Not Working, page D-7

Adding Inline Posture as a RADIUS Client

For a standalone Inline Posture node, you must add the trusted IP address as a RADIUS client. For a high availability pair, add the service IP address for the trusted interface as a RADIUS client. This section contains the basic steps for this task. For more in-depth information, see Chapter 6, "Managing Network Devices."

Prerequisites

You must have completed the tasks in the appropriate section:

- Deploying an Inline Posture Node, page 10-12
- Configuring Inline Posture for High Availability, page 10-25

To add Inline Posture as a RADIUS client, complete the following steps:

Step 1 Choose **Administration > Network Resources > Network Devices**.

Step 2 In the Network Devices navigation panel, choose Network **Devices**.

Step 3 Enter a Name and Description for the device.

Step 4 Do one of the following:

- For a standalone Inline Posture node, enter the IP address for the trusted interface.
- For a high availability pair, enter the service IP address for the trusted interface.

Step 5 Enter a Model Name and Software Version, as necessary.

Step 6 For the Network Device Group, specify a Location and Device Type, as necessary.

Step 7 Check the **Authentication Settings** check box, and enter the shared secret.

Step 8 Click **Save**.

Next Step

- Monitoring an Inline Posture Node, page 10-30

Monitoring an Inline Posture Node

You can monitor the health of a deployed Inline Posture node from the Cisco ISE dashboard, that is running on the Administration ISE node. The Inline Posture node appears on the System Summary dashlet. A green icon with a check mark means that the system is healthy. A yellow icon indicates a warning, and a red icon indicates of a critical system failure. Sparklines indicate the utilization of CPU, memory, and latency over time. You can choose to display data for the past 24 hours or the last 60 minutes.

When you hover your mouse cursor over the health icon, a quick view dialog appears showing detailed information on system health.

Figure 10-26 System Summary Quick View Status

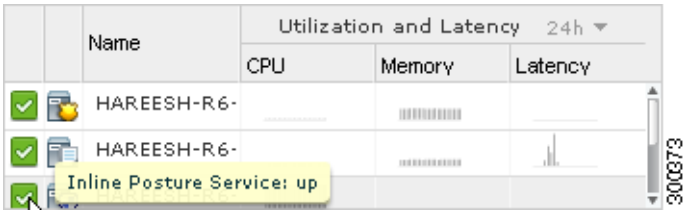

For more information, see Cisco ISE Dashboard Monitoring, page 24-3.

Removing an Inline Posture Node from Deployment

To remove an Inline Posture node from the deployment, you must first change it to maintenance mode, and then you can deregister it. Maintenance mode is a neutral state that allows the node to smoothly transition to the network or from a deployment.

Prerequisites

- You should have administrative permissions on the primary Administration ISE node.

To remove a node from deployment, complete the following steps:

Step 1 From the primary Administration ISE node, choose **Administration > System > Deployment**.

Step 2 Click **Deployment** on the left pane, and then check the check box next to the Inline Posture node that you want to remove from the deployment, and click **Edit**.

Step 3 Click the **Deployment Modes** tab.

Step 4 Click the **Maintenance Mode** radio button, and then click **Save**.

Step 5 Click **Deployment** on the left pane, and then check the check box next to the Inline Posture node that you want to remove from the deployment, and then click **Deregister**.

You are prompted with the following message: Are you sure you want to deregister the selected items?

Step 6 Click **OK** to remove the node from the deployment.

Troubleshooting Topics

- Primary and Secondary Inline Posture Nodes Heartbeat Link Not Working, page D-7

Remote Access VPN Use Case

This section describes how to use an Inline Posture node with a VPN device such as ASA in a Cisco ISE network. Figure 10-27 shows a Cisco ISE deployment that uses an Inline Posture node for remote VPN access. The term iPEP in this illustration refers to the Inline Posture node and PDP refers to the Policy Service node. All the traffic from the VPN gateway must go through the Inline Posture node to ensure that Cisco ISE can apply policies and secure a network.

Figure 10-27 Cisco ISE Deployment with Inline Posture Node

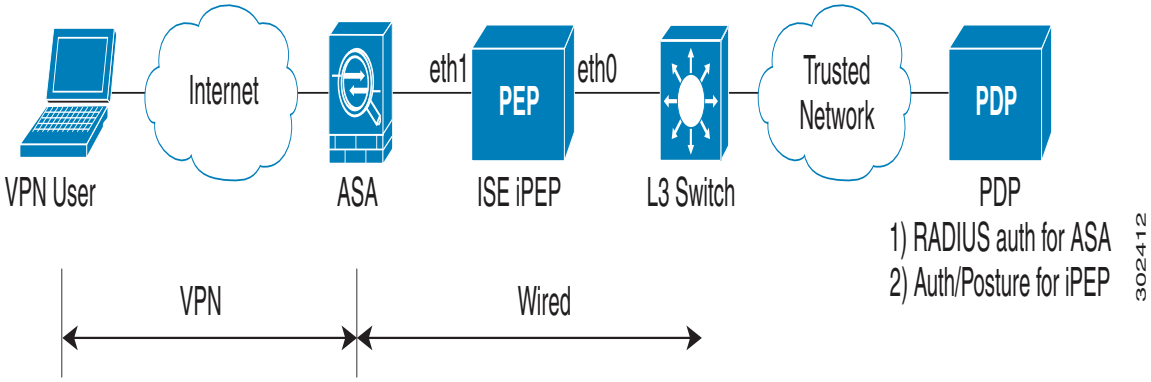

Process Flow

1. Remote user authenticates to VPN gateway (ASA) using the RADIUS protocol.
2. As a RADIUS client, the ASA sends an authentication request to the AAA server (Inline Posture node).
3. As a RADIUS proxy, the Inline Posture node relays the RADIUS authentication request to the Cisco ISE node that acts as the RADIUS Server (Policy Service node).
4. The Cisco ISE Policy Service node authenticates the remote user using the configured identity store and returns the RADIUS response to the Inline Posture node which in turn relays it to the ASA (the network access device (NAD)).
5. Based on the authorization policy that is applicable for the user, the Policy Service node returns the appropriate attributes to the Inline Posture node and optionally to the ASA.
6. Each authorization policy rule entry can reference separate authorization profiles for both the Inline Posture node profile and the NAD (standard authorization profile).
   a. Inline Posture Node Profile: Specifies RADIUS attributes to be applied to the Inline Posture node such as a URL for redirection to the Client Provisioning service and downloadable ACLs (dACLs) for policy enforcement by the Inline Posture node.
   b. Standard Authorization Profile: Specifies any RADIUS attributes intended for NAD, or ASA in this example.
7. If the authorization policy determines that the endpoint is NonCompliant with the posture policy, or if the posture status is Unknown, then the Policy Service node returns a URL redirect attribute value to the Inline Posture node along with a dACL to specify the traffic to be allowed. All HTTP traffic denied by the dACL is redirected to the specified URL.
8. When the posture becomes Compliant, a reauthorization occurs and the Policy Service node sends a new dACL to the Inline Posture node, which provides the user privileged access to the internal network.

Configuring a Cisco ISE Deployment Using an Inline Posture Node

Before You Begin

1. Ensure that your network infrastructure is configured correctly to route or switch traffic to and from the Inline Posture node and its downstream networks.
2. For third-party VPN concentrators to integrate with Cisco ISE and Inline Posture nodes, the following AAA attributes must be included in RADIUS communication:

- NAS_PORT_TYPE
- MAC_ADDRESS
- USER_NAME
- DEVICE_LOCATION

3. For VPN devices, the RADIUS accounting message must have the framed-ip-address attribute set to the VPN client's IP address pool.

To configure your Cisco ISE deployment with an Inline Posture node for remote VPN access, complete the following steps:

Step 1 Configure a standalone Cisco ISE node. For more information, refer to Configuring an ISE Node, page 9-7.

Step 2 Register the standalone Cisco ISE node as an Inline Posture node to an existing primary Administration ISE node, and configure the Inline Posture node from the primary Administration ISE node. For more information, refer to Deploying an Inline Posture Node, page 10-12.

Step 3 Optionally, you can configure a second Inline Posture node and configure an Active/Standby pair. For more information, refer to Configuring Inline Posture for High Availability, page 10-25.

Step 4 Set up a Policy Service ISE node (PDP) to be the RADIUS server for the Inline Posture node. Configure the Policy Service ISE node with the same RADIUS shared secret that is configured on the Inline Posture node.

Step 5 Configure authorization profiles (Inline Posture node profiles) for use by the Inline Posture node. You can optionally configure standard authorization profiles for the NAD's use. For more information, refer to Creating Inline Posture Node Profiles, page 10-22 and Creating Inline Posture Downloadable Access Control Lists, page 10-20.

Step 6 Configure authorization policy to apply the Inline Posture node profiles to remote VPN users based on identity and posture status. For more information, refer to Creating an Inline Posture Authorization Policy, page 10-23.

Step 7 Add the VPN gateway's inside IP address as a RADIUS client in the Inline Posture node's RADIUS configuration along with the NAD's (ASA in this example) RADIUS shared secret.

Step 8 Configure the VPN gateway (ASA) for RADIUS authentication and accounting with the Inline Posture node configured as the RADIUS server. To do this:

a. Choose **Policy > Authentication**.

b. Ensure that the Default Rule is configured to authenticate users against the identity source that contains the user records.

c. Click **Save**.