



Setting Up Inline Posture

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

- [Role of Inline Posture Node in a Cisco ISE Deployment, page 5-1](#)
- [Best Practices for Inline Posture Deployment, page 5-9](#)
- [Inline Posture Node Guidelines, page 5-10](#)
- [Inline Posture Node Authorization, page 5-11](#)
- [Deploying an Inline Posture Node, page 5-13](#)
- [Configuring a High-Availability Pair, page 5-19](#)
- [Configuring Inline Posture Node as RADIUS Client in Administration Node, page 5-20](#)
- [Removing an Inline Posture Node from Deployment, page 5-21](#)
- [Health of an Inline Posture Node, page 5-22](#)
- [Remote Access VPN Use Case, page 5-22](#)
- [Collecting Inline Posture Node Logs, page 5-24](#)
- [Kclick process in Inline Posture Node, page 5-24](#)

Role of Inline Posture Node in a Cisco ISE Deployment

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 requests, such as wireless LAN controllers (WLCs) and VPN devices.

After the initial authentication of a client using the EAP/802.1x and RADIUS protocols, the client must 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, an Inline Posture node 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, posture assessed, and remediation completed. This restriction is accomplished by downloading and installing Downloadable Access Control Lists (DACLS) that are tailored for specific client flows. See [Downloadable ACLs, page 21-11](#).

When the client is compliant, a CoA is sent to the Inline Posture node by the Policy Service 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, and therefore duplicate downloads are not necessary as long as the DACL content does not change in the Cisco ISE servers.

Inline Posture Policy Enforcement Flow

Figure 5-1 illustrates the Inline Posture policy enforcement process and shows the flow for WLC enforcement for traffic to the Policy Service node. The access steps are similar for an inline deployment with VPN gateways.

Figure 5-1 Inline Posture Policy Enforcement Flow



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, which is usually the Policy Service node (in this illustration, the RADIUS Access-Request message is sent to the Inline Posture node).

3. The 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 the Inline Posture node to send an Authorization-Only request to the Policy Service node to retrieve the profile for the session.
7. The Policy Service node returns an Access-Accept message, along with the necessary Inline Posture node profile.
8. If the access control list (ACL) that is defined in the profile is not already available on the Inline Posture node, the Inline Posture node downloads it from the Policy Service 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 node 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, which is 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 access only 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 node initiates a RADIUS (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 the Inline Posture node.

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 user 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 during 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 Inline Posture deployment:

- **Trusted**—The interface that talks to the Policy Service node and other trusted devices *inside* the Cisco ISE network. The trusted interface is always designated to Eth0 interface.
- **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 interface.

Dedicated Nodes Required for Inline Posture

Unlike other personas, 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 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. You can also have as many active-standby pairs or standalone Inline Posture nodes as needed, up to the limit of supported endpoints for your deployment.

For information on Cisco ISE distributed deployments, see [Chapter 4, “Setting Up Cisco ISE in a Distributed Environment.”](#)

Standalone Inline Posture Node in a Cisco ISE Deployment

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.

[Figure 5-2](#) illustrates a simple Inline Posture standalone configuration, with client access through WLC and VPN devices.

Inline Posture High Availability

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 and forwards all network packets until it fails and 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.

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. The secondary Inline Posture node is read-only, and cannot be used for configuration of any kind, even high availability.

When you configure an Inline Posture high-availability pair, 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.

[Figure 5-4](#) illustrates a routed mode high-availability Inline Posture configuration.

An Inline Posture high-availability pair consists of two physical Inline Posture nodes configured as a cluster that have heartbeat links on the eth2 and eth3 interfaces, and are connected by dedicated cables.

The eth2 and eth3 interfaces of both nodes communicate with heartbeat protocol exchanges to determine the health of the nodes. 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 service IP address to both the trusted and untrusted interfaces for both nodes of the active-standby pair, thus making the service IP address the address of the cluster, representing it as a single entity to the rest of the network.

Automatic Failover in Inline Posture Nodes

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.

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.

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.

In addition to the heartbeat monitor, an optional (but highly recommended) link-detect mechanism is available. With the use of this mechanism, 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.

When failover occurs:

1. The standby Inline Posture node takes over the service IP address.
2. 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 [Synchronizing an Inline Posture Node](#), page 5-20.

3. Active sessions are automatically reauthenticated and authorized.

Inline Posture Operating Modes

The Inline Posture operating mode that you choose depends largely on the architecture of your existing network. Cisco ISE supports the following operating modes:

- [Inline Posture Routed Mode, page 5-6](#)
- [Inline Posture Bridged Mode, page 5-7](#)
- [Inline Posture Maintenance Mode, page 5-8](#)

Inline Posture Routed Mode

The Inline Posture routed 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.

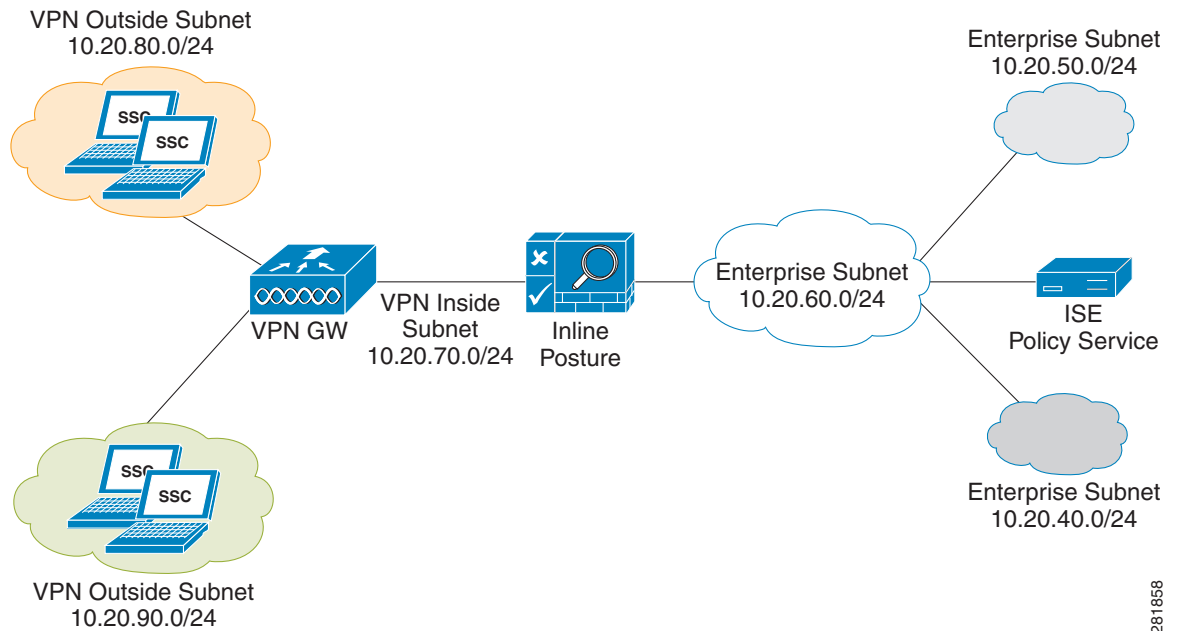
In routed mode, the Inline Posture node operates as a Layer 3 router, and becomes the default gateway for the untrusted network with its managed clients. All traffic between the untrusted and trusted networks passes through the Inline Posture node, which applies the IP filtering rules, access policies, 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 5-2](#) illustrates an Inline Posture routed mode configuration. In this example, Inline Posture is a hop for the client traffic from the VPN gateway (GW) en route to the Policy Service 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 5-2 *Inline Posture Routed Mode Configuration*

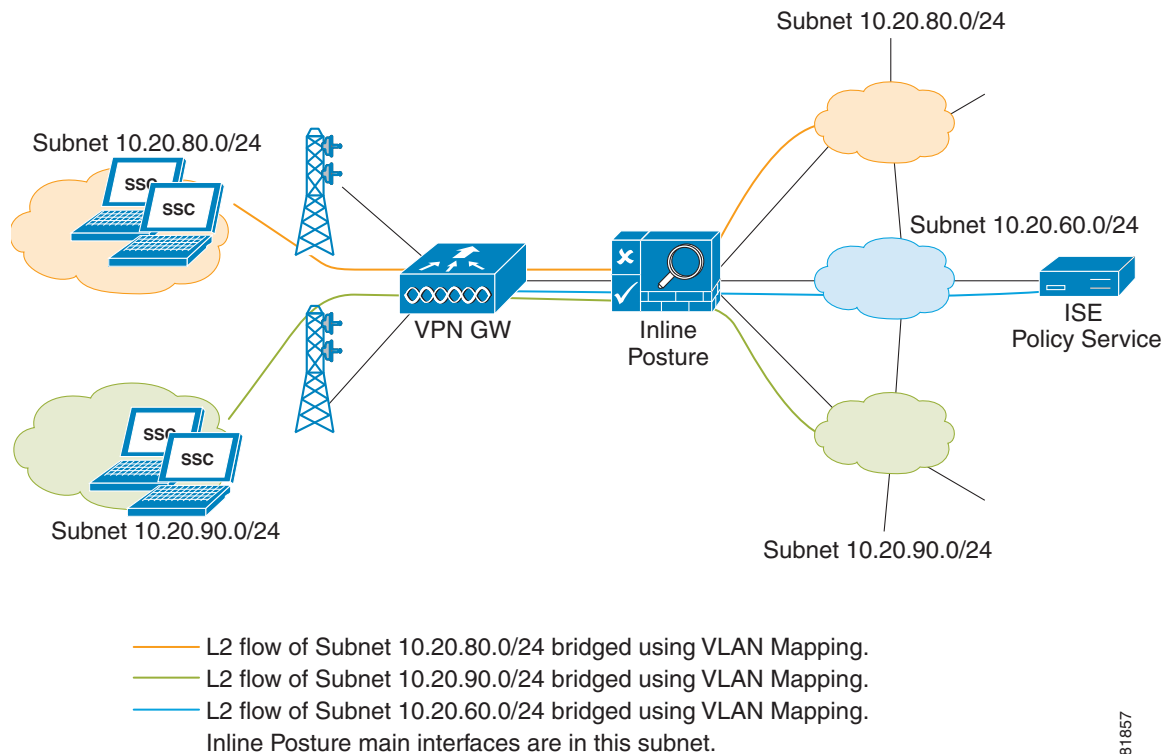
281858

Inline Posture Bridged Mode

The Inline Posture bridged mode acts as a Layer 2 “bump” in the wire, forwarding packets without regard to the destination address.

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 5-3 shows the Inline Posture node acting as a bridge for the Layer 2 client traffic from the WLC to the Cisco ISE network, managed by the Policy Service 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 5-3 *Inline Posture Bridged Mode Configuration*

281857

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

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

Inline Posture Maintenance Mode

The Inline Posture maintenance 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, and before you perform other configurations.

Inline Posture High Availability in Routed and Bridged Modes

Figure 5-4 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.

In this example, the untrusted IP address for Inline Posture 1 can be 10.20.70.101, and the untrusted IP address for Inline Posture 2 can be 10.20.70.102. However, the service IP address for both nodes on the untrusted side of the network would be 10.20.70.100. 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 5-4 *Inline Posture High-Availability Routed Mode Configuration*

In a bridged mode, Inline Posture eth0 and eth1 interfaces 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.

Best Practices for Inline Posture Deployment

You can follow the best practices listed here to manage your Inline Posture deployment efficiently:

- [Use Filters to Define Access Privileges, page 5-9](#)
- [Configure Managed Subnets and Static Routes, page 5-9](#)
- [Configure High-Availability Pair, page 5-10](#)

Use Filters to Define Access Privileges

Consider the following when configuring filters for Inline Posture:

- In a typical implementation, 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, MAC and IP, or 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.



Caution

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 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.

Configure Managed Subnets and Static Routes

Consider the following when configuring managed subnets for Inline Posture:

- Configure a managed subnet for Inline Posture. A managed subnet 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.
- Configure managed subnets for endpoints in Layer 2 proximity of the Inline Posture node, such as, a WLC that delivers packets directly to the untrusted interface of the Inline Posture node.
- 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.
- Ensure that subnets on the trusted side of the Inline Posture node are different from the subnets on the untrusted side.

- Ensure that an Administration node, Policy Service node, and Monitoring node are not on the same subnet as the Inline Posture node, 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.
- Configure static routes for all downstream host networks that are typical of VPN address pools.

Configure High-Availability Pair

Consider the following when configuring Inline Posture for high availability:

- Assign a service IP address (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.

Inline Posture Node Guidelines

Before you configure an Inline Posture node in a distributed deployment, read and understand the following statements:

1. The Inline Posture node is supported only on Cisco ISE-3300 series and SNS-3415 appliances. It is not currently supported on Cisco SNS-3495 appliance or as a virtual appliance.
2. Inline Posture is unable to run concurrently with Administration, Policy Service, or Monitoring personas and, therefore, is a dedicated node.
3. An Inline Posture node must be registered to the primary Administration node on your network.
4. For each deployment instance of an Inline Posture node, you can deploy a standalone node, or an active-standby pair.
5. At any network entry point, like VPN headend using ASA or group of ASAs in an HA cluster, a maximum of 2 Inline Posture nodes can be deployed as active-standby pair for high-availability. You can have several HA pairs in a deployment.
6. Inline Posture nodes are similar to network access devices (NAD) in function from the perspective of Cisco ISE node. Inline Posture nodes can serve as multiple NADs like switches, Wireless Lan Controllers, and VPN devices. Based on the deployment needs, you can deploy multiple instances of Inline Posture nodes. To determine the maximum number of deployment instances, treat the Inline Posture nodes as access devices.
7. For an Inline Posture high-availability, two nodes are configured as an active-standby pair. One node is designated as the primary node and the other as the secondary node. The primary node becomes the active node when both nodes come up at the same time.
8. For an Inline Posture active-standby pair configuration, all configuration must be applied from the ISE administrative user interface. The standby node configuration displays only basic tables when viewed from the ISE administrative user interface.
9. You can synchronize an Inline Posture active node configuration to its peer standby node from the Failover tab of the active node. For more information, see [Synchronizing an Inline Posture Node](#), page 5-20.

**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.

10. Registering an Inline Posture node results in system restart. High-availability changes and changes to infrastructure configurations such as the eth1 IP address or Inline Posture mode 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.
11. Do not change the CLI password for Inline Posture node once it has been added to the deployment. If the password is changed, when you access the Inline Posture node through the Administration node, a Java exception error is displayed and the CLI gets locked. You need to recover the password by using the installation DVD and rebooting the Inline Posture node. Or, you can set the password to the original one.

If you need to change the password, then deregister the Inline Posture node from the deployment, modify the password, and then add the node to the deployment with the new credentials.

12. After you register an Inline Posture node to the Administration node, you are not allowed to change the eth0 (Trusted) IP address through the Admin portal. The reason for this is that, if you change the eth0 IP address of a registered Inline Posture node, it cannot communicate with the Administration node. Any attempted communication between the Inline Posture node and Administration node then fails, leading to a potential exception.

**Note**

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 when 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 starts 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 this time is in Maintenance mode.

Inline Posture Node Authorization

The following images illustrate the client authorization flow and session recovery using Lazy Fetch mechanism for Inline Posture node.

Figure 5-5 Inline Posture Node Client Authorization Flow

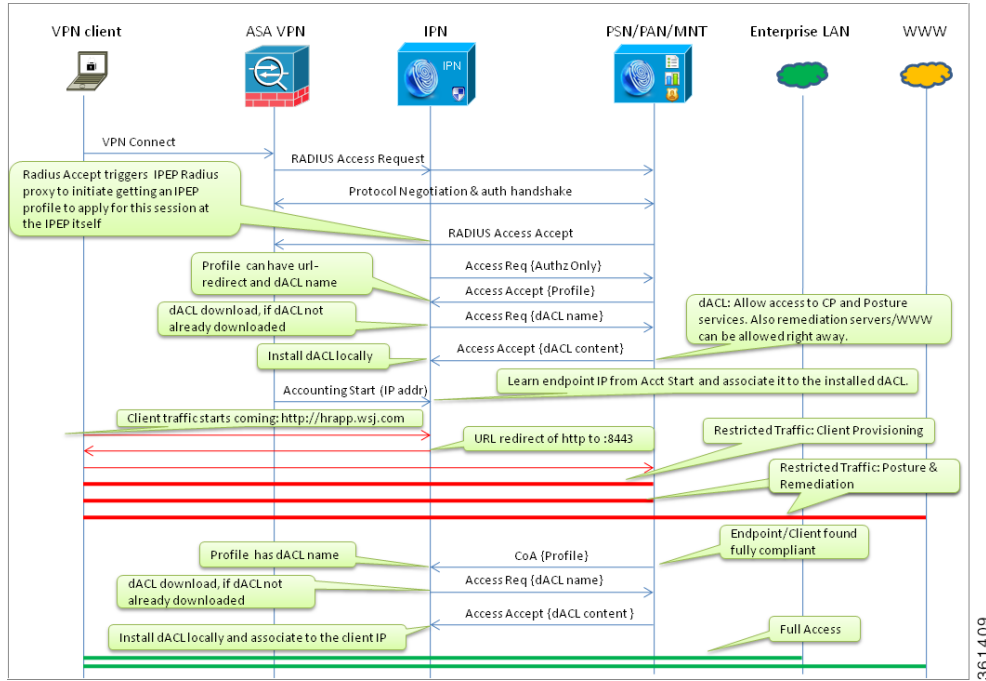
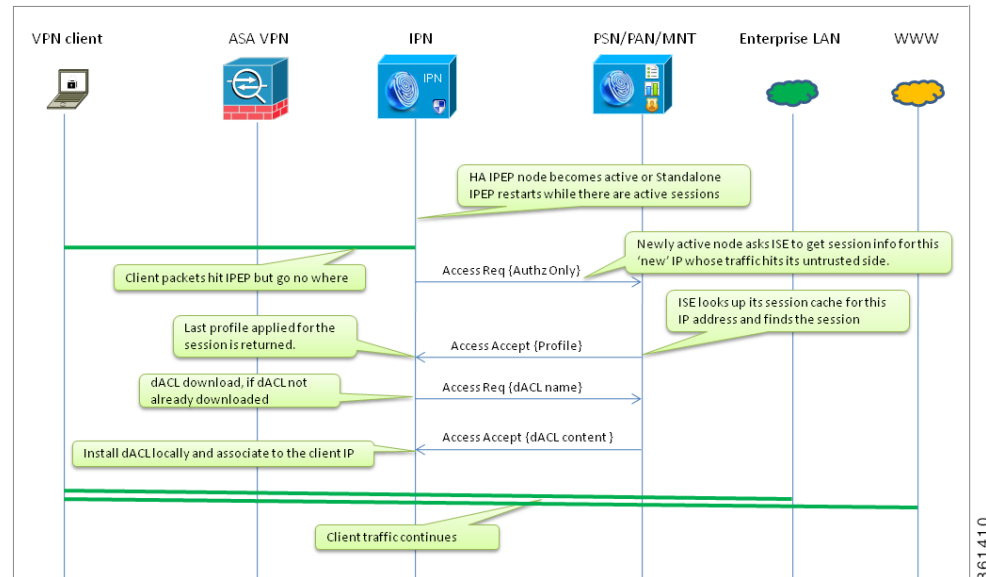


Figure 5-6 Inline Posture Node Session Recovery Using Lazy Fetch Mechanism



Inline Posture Node Session Removal due to Client Disconnect

When a wireless client is wandering off from the WLC control, the WLC is required to send a RADIUS Accounting Stop similar to the VPN gateway to ensure that the Inline Posture node cleans up the session corresponding to the client.

Deploying an Inline Posture Node

The initial process for deploying an Inline Posture node is the same, whether it is intended to be a standalone node or part of an active-standby pair.



Note

Inline Posture is supported on the Cisco ISE 3415, ISE 3315, ISE 3355, and ISE 3395 platforms.

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

1. [Configuring an Inline Posture Node, page 5-13](#)
2. [Creating Inline Posture Downloadable Access Control Lists, page 5-16](#)
3. [Creating Inline Posture Node Profiles, page 5-17](#)
4. [Creating an Inline Posture Authorization Policy, page 5-18](#)

Configuring an Inline Posture Node

Inline Posture is a dedicated node registered to the Administration node. You configure Inline Posture from the administration console, and that configuration is then replicated to the Inline Posture node. A copy of the configuration is stored locally in the administration database. After an Inline Posture node is registered, it is rebooted.

To introduce an Inline Posture node in your Cisco ISE network, you must first register the Inline Posture node with the primary Administration 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 node as a client making the Policy Service node its RADIUS server.

**Note**

After 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 a high-availability pair, the service IP address for the active node is used.

Before You Begin

To perform the following task, you must be a Super Admin or System Admin.

Inline Posture is not supported on the Cisco ISE 3495 platform. Ensure that you install Inline Posture on any one of the following supported platforms: ISE 3315, ISE 3355, ISE 3395, or ISE 3415.

Follow and apply the guidelines for configuring certificates for Inline Posture. Refer to *Cisco Identity Services Engine Hardware Installation Guide, Release 1.2* for details.

Register the Inline Posture node with the primary Administration node. All nodes must be registered with the primary Administration node to function as a member of the Cisco ISE distributed system.

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, IP addresses for other Cisco ISE nodes, shared secret information for the 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.

**Caution**

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.

Step 1 Choose **Administration > System > Deployment**.

Step 2 Check the **Inline Posture node** check box in the Deployment Nodes page and click **Edit**.

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

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 following tabs and enter the appropriate information for the fields in the tabs. See [Inline Posture Node Settings](#) for more information on the fields.

- **Basic Information**
- **Deployment Modes**—A newly registered Inline Posture node comes up in maintenance mode. For production purposes, you must choose the Routed or Bridged mode.
- **Filters**—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.
- **Radius Config**—RADIUS configuration is mandatory. At least one client and one server configuration is necessary for Inline Posture.
- **Managed Subnets**—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. You must enter the following values: IP Address, Subnet Mask, VLAN ID, and Description.
- **Static Routes**—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.

- **Logging**—Click the **Logging** tab and enter the IP address and port number for the logging server, which is typically the Monitoring 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.

- **Failover**—This tab is for Inline Posture High Availability configuration.

Step 5 Click **Save**. The Inline Posture node restarts automatically.

Step 6 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 a high-availability pair, the service IP address for the active node is used.

What To Do Next

To complete the deployment of the Inline Posture node, you must create DACLs, authorization profiles, and authorization policy rules: unknown, compliant, and noncompliant.



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.

Related Topics

- [Inline Posture Node Settings, page A-6](#)
- [Registering an Inline Posture Node, page 4-14](#)
- [Configuring a High-Availability Pair, page 5-19](#)
- [Creating Inline Posture Downloadable Access Control Lists, page 5-16](#)
- [Creating Inline Posture Node Profiles, page 5-17](#)
- [Creating an Inline Posture Authorization Policy, page 5-18](#)

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.

Before You Begin

To perform the following task, you must be a Super Admin or System Admin.

-
- Step 1** Choose **Policy > Policy Elements > Results > Authorization > Downloadable ACLs**.
- Step 2** Click **Add**.
- Step 3** Enter the name of the DACL and its description.
- Step 4** Create the following DACLS:
- ipn-compliant (Permit All): Use the following syntax: permit ip any any
 - ipn-noncompliant (Deny All): Use the following syntax: deny ip any any
 - ipn-unknown (Pre-Posture): Use at least one ACL to allow supplicants and the Policy Service node to have access to each other for posture evaluation. This DACL can be used to block or quarantine users until they pass authentication. Here is an example syntax:

```
deny tcp any any eq 80
deny tcp any any eq 443
permit ip any 10.1.2.4 0.0.0.0
permit udp any any eq 53
deny ip any any
```
- Step 5** Save the DACLS.
-

What To Do Next

[Creating Inline Posture Node Profiles, page 5-17](#)

Related Topics

[Cisco ISE Authorization Profiles, page 21-2](#)

Creating Inline Posture Node Profiles

You must create three Inline Posture authorization profiles, as well as an authorization profile for a NAD.

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

Before You Begin

To perform the following task, you must be a Super Admin, System Admin, or Policy Admin.

-
- Step 1** Choose **Policy > Policy Elements > Results > Authorization > Inline Posture Node Profiles**.
- Step 2** Click **Add**.
- Step 3** Enter a name and description for the authorization profile. Supported characters for the name field are: space, ! # \$ % & ' () * + , - . / ; = ? @ _ { .



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 4** Create the following authorization profiles for Inline Posture that correspond to the DACLs you created. Specify the appropriate DACL for each of the following authorization profiles:
- IPN-Unknown-Compliant (Pre-Posture): This profile requires that you enter a URL redirect. To do this, check the URL Redirect check box.
The URL redirect appears in the Attributes Details field.
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.
 - IPN-Compliant (Permit All)
 - IPN-Noncompliant (Deny All).
- Step 5** Click **Submit**.
-

What To Do Next

[Creating an Inline Posture Authorization Policy, page 5-18](#)

Related Topics

- [Cisco ISE Authorization Policies, page 21-1](#)
- [Cisco ISE Authorization Profiles, page 21-2](#)

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 rules when creating 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.

Before You Begin

To perform the following task, you must be a Super Admin or System Admin.

-
- Step 1** Choose **Policy > Authorization**.
- Step 2** Leave the default rules as is.
- Step 3** Create the following Unknown Posture Status Rule:
- Identity Group: Any
 - Condition: Session:PostureStatus EQUALS = Unknown
 - Permissions: IPN-Unknown-Compliant + nad-authorization-profile
- Step 4** Create the following Compliant Posture Rule:
- Identity Group: Any
 - Condition: Session:PostureStatus EQUALS = Compliant
 - Permissions: IPN-Compliant + nad-authorization-profile
- Step 5** Create the following Noncompliant Posture Rule:
- Identity Group: Any
 - Condition: Session:PostureStatus EQUALS = Noncompliant
 - Permissions: IPN-Noncompliant + nad-authorization-profile
- Step 6** Save the policy. The Inline Posture node deployment is now complete.
-

What To Do Next

[Configuring Inline Posture Node as RADIUS Client in Administration Node, page 5-20.](#)

Related Topics

- [Cisco ISE Authorization Policies, page 21-1](#)
- [Cisco ISE Authorization Profiles, page 21-2](#)

Configuring a High-Availability Pair

When you configure two Inline Posture nodes for high availability, you specify one node as the primary unit in the pair and it 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.

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.

**Note**

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.

Before You Begin

- To perform the following task, you must be a Super Admin or System Admin.
- You should have successfully configured two (2) Inline Posture nodes, and registered them on the Cisco ISE network.
- 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 5-4](#) 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 the course of this procedure.
- Have all necessary network configuration information for your installation on hand. Check with your system architect for a complete list of information you will need.

-
- Step 1** Choose **Administration > System > Deployment**.
- Step 2** 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** Enter the appropriate information in the fields.
- Step 6** 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.
- Step 7** Verify the node status by checking the check box next to it, and then clicking the **Failover** tab. Ensure that your primary and secondary Inline Posture nodes are configured correctly.
-

What To Do Next

Complete the following task: [Configuring Inline Posture Node as RADIUS Client in Administration Node, page 5-20](#).

Related Topics

- [Inline Posture Node Settings, page A-6](#)
- [Configuring an Inline Posture Node, page 5-13](#)
- [Synchronizing an Inline Posture Node, page 5-20](#)

Synchronizing an Inline Posture Node

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 synchronize the standby node with the latest Inline Posture database from the active node.

Before You Begin

- You must be a Super Admin or System Admin.
- You must configure two Inline Posture nodes.
- You must establish a relationship between the two nodes.

-
- Step 1** Choose **Administration > System > Deployment**.
- Step 2** Check the check box next to the Inline Posture node that you want to sync with the other node (usually the active node), and click the **Edit** icon.
- Step 3** Click the **Failover** tab. See [Deployment Settings](#) for a description of the fields in this tab.
- Step 4** Click **Sync Peer Node**. Data from the selected node is automatically transferred to its peer node.
-

Related Topics

- [Configuring an Inline Posture Node, page 5-13](#)
- [Configuring a High-Availability Pair, page 5-19](#)

Configuring Inline Posture Node as RADIUS Client in Administration Node

For an Inline Posture node to act as a RADIUS proxy, you must add it as a RADIUS client in the Administration node.

Before You Begin

- You must be a Super Admin or System Admin.
- You must deploy Inline Posture in your Cisco ISE deployment.

-
- Step 1** Choose **Administration > Network Resources > Network Devices**.
- Step 2** In the Network Devices navigation panel, click **Network Devices**.
- Step 3** Enter a Name and Description for the device.
- Step 4** Enter the IP address of the Inline Posture node.
- 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 RADIUS shared secret information.
- Step 8** Click **Save**.
-

Removing an Inline Posture Node from Deployment

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

Before You Begin

To perform the following task, you must be a Super Admin or System Admin.


-
- Step 1** Choose **Administration > System > Deployment**.
- Step 2** 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.
- Step 6** Click **Deregister**.
- Step 7** Click **OK**.
-

Health of 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 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.

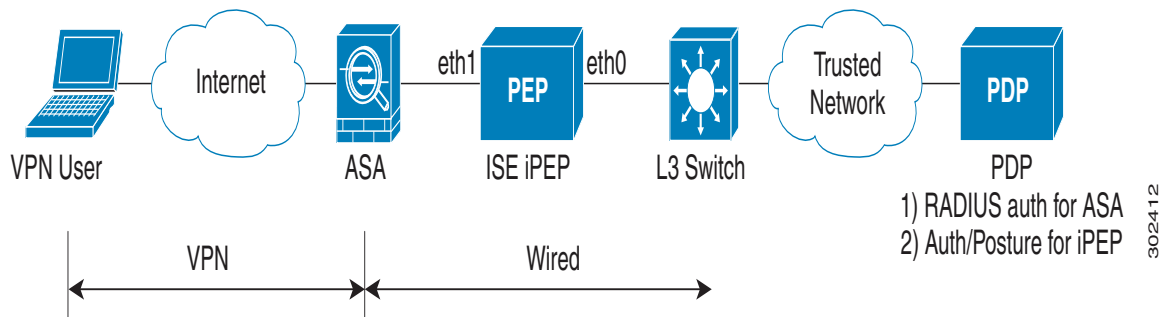
Figure 5-7 System Summary Quick View Status



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 5-8](#) 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 5-8 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 access control lists (DACLS) for policy enforcement by the Inline Posture node.
 - b. Standard authorization profile: Specifies any RADIUS attributes intended for the NAD, which is 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/HTTPS 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.

Related Topics

[Configuring an Inline Posture Node with a VPN Device, page 5-23](#)

Configuring an Inline Posture Node with a VPN Device

Before You Begin

Ensure that your network infrastructure is configured correctly to route or switch traffic to and from the Inline Posture node and its downstream networks.

-
- Step 1** Configure a standalone Cisco ISE node. For more information, refer to [Configuring a Cisco ISE Node, page 4-11](#).
 - Step 2** Register the standalone Cisco ISE node as an Inline Posture node to an existing primary Administration node, and configure the Inline Posture node from the primary Administration node. For more information, refer to [Deploying an Inline Posture Node, page 5-13](#).
 - Step 3** Optionally, you can configure a second Inline Posture node and configure an Active/Standby pair. For more information, refer to [Inline Posture High Availability, page 5-4](#).
 - Step 4** Set up a Policy Service node to be the RADIUS server for the Inline Posture node. Configure the Policy Service 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.
 - Step 6** (Optional) You can configure standard authorization profiles for the NAD's use. For more information, refer to [Creating Inline Posture Node Profiles, page 5-17](#) and [Creating Inline Posture Downloadable Access Control Lists, page 5-16](#).
 - Step 7** Configure an 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 5-18](#).

- Step 8** 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 9** 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**.

Collecting Inline Posture Node Logs

From the Inline Posture node CLI, all the logs can be archived and collected using the backup-logs command.

```
PEP/admin# config terminal
PEP/admin# repository remoteloc
PEP/admin# url ftp://myremoteserver/store
PEP/admin# user <myremoteuser> password plain <myremotepasswd>
PEP/admin# end
PEP/admin# backup-logs myipeplogs repository remoteloc
% Creating log backup with timestamped filename: myipeplogs-110317-1836.tar.gz
```



Note Collecting Inline Posture node logs remotely from the Primary Administration UI is not supported.

Kclick process in Inline Posture Node

Click kernel module process, called as kclick owns CPU scheduling in Inline Posture node. Kclick provides the CPU cycles for other processes that request it. Due to this the 'top' output at an Inline Posture Node displays the kclick using all the CPU cycles in the system including idle cycles.