



Deployment of Cisco ISE

- [Cisco ISE Deployment Terminology, on page 1](#)
- [Personas in Distributed Cisco ISE Deployments, on page 2](#)
- [Configure a Cisco ISE Node, on page 2](#)
- [Support for Multiple Deployment Scenarios, on page 5](#)
- [Cisco ISE Distributed Deployment, on page 5](#)
- [Deployment and Node Settings, on page 9](#)
- [Logging Settings, on page 17](#)
- [Admin Access Settings, on page 19](#)
- [Administration Node, on page 22](#)
- [Support for Automatic Failover for the Administration Node, on page 30](#)
- [Policy Service Node, on page 30](#)
- [Monitoring Node, on page 33](#)
- [Monitoring Database, on page 37](#)
- [Configure MnT Nodes for Automatic Failover, on page 39](#)
- [Cisco pxGrid Node, on page 40](#)
- [Cisco pxGrid Cloud Overview, on page 46](#)
- [View Nodes in a Deployment, on page 53](#)
- [Download Endpoint Statistical Data from MnT Nodes, on page 53](#)
- [Database Crash or File Corruption Issues, on page 53](#)
- [Device Configuration for Monitoring, on page 54](#)
- [Synchronize Primary and Secondary Cisco ISE Nodes, on page 54](#)
- [Change Node Personas and Services, on page 54](#)
- [Effects of Modifying Nodes in Cisco ISE , on page 55](#)
- [Create a Policy Service Node Group, on page 55](#)
- [Remove a Node from Deployment, on page 56](#)
- [Shut Down a Cisco ISE Node, on page 57](#)
- [Scenarios In Which Need to Reregister a Node, on page 57](#)
- [Change the Hostname or IP Address of a Standalone Cisco ISE Node, on page 58](#)

Cisco ISE Deployment Terminology

The following terms are commonly used when discussing Cisco ISE deployment scenarios:

- **Service:** A service is a specific feature that a persona provides, such as network access, profiler, posture, security group access, monitoring and troubleshooting, and so on.
- **Node:** A node is an individual instance that runs the Cisco ISE software. Cisco ISE is available as an appliance and also as a software that can be run on VMware. Each instance (appliance or VMware) that runs the Cisco ISE software is called a node.
- **Persona:** The persona of a node determines the services provided by the node. A Cisco ISE node can assume any of the following personas: Administration, Policy Service, Monitoring, and pxGrid. The menu options that are available through the Admin portal are dependent on the role and personas that a Cisco ISE node assumes.
- **Deployment Model:** Determines if your deployment is distributed, standalone, or high availability in standalone, which is a basic two-node deployment.

Personas in Distributed Cisco ISE Deployments

A Cisco ISE node can assume the Administration, Policy Service, or Monitoring personas.

A Cisco ISE node can provide various services based on the persona that it assumes. Each node in a deployment can assume the Administration, Policy Service, and Monitoring personas. In a distributed deployment, you can have the following combination of nodes in your network:

- Primary Policy Administration Node (primary PAN) and secondary Policy Administration Node (secondary PAN) for high availability
- Primary Monitoring Node (primary MnT node) and Secondary Monitoring Node (secondary MnT node) for high availability
- A pair of health check nodes or a single health check node for the primary PAN automatic failover
- One or more Policy Service Nodes (PSNs) for the session failover

Configure a Cisco ISE Node


After you install a Cisco ISE node, all the default services provided by the Administration, Policy Service, and Monitoring personas run on it. This node is in a standalone state. You must log in to the Admin portal of the Cisco ISE node to configure it. You cannot edit the personas or services of a standalone Cisco ISE node. You can, however, edit the personas and services of the primary and secondary Cisco ISE nodes. You must first configure a primary ISE node and then register secondary ISE nodes to the primary ISE node.

If you are logging in to the node for the first time, you must change the default administrator password and install a valid license.

We recommend that you do not change the host name and the domain name configured on Cisco ISE in production. If required, reimagine the appliance, make changes, and configure the details during the initial deployment.


Before you begin

You should have a basic understanding of how distributed deployments are set up in Cisco ISE. See [Guidelines for Setting Up a Distributed Deployment](#).

-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**.
- Step 2** Check the check box next to the Cisco ISE node that you want to configure, and click **Edit**.
- Step 3** Enter the values, as required, and click **Save**.
-

Configure a Primary Policy Administration Node

To set up a distributed deployment, you must first configure a Cisco ISE node as your primary PAN.

-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**. The **Register** button is disabled initially. To enable this button, configure a Primary PAN.
- Step 2** Check the check box next to the current node, and click **Edit**.
- Step 3** Click **Make Primary** to configure your primary PAN.
- Step 4** Click **Save** to save the node configuration.
-

What to do next

1. Add secondary nodes to your deployment.
2. Enable the profiler service and configure the probes, if required.

Register a Secondary Cisco ISE Node

You can register Cisco ISE nodes to the primary PAN to form a multinode deployment. Nodes in a deployment other than the primary PAN are referred to as secondary nodes. While registering a node, you can select the personas and services that must be enabled on the node. Registered nodes can be managed from the primary PAN (for example, managing the node personas, services, certificates, licenses, applying patches, and so on).

When a node is registered, the primary PAN pushes the configuration data to the secondary node, and the application server on the secondary node restarts. After the complete data replication, further configuration changes done on the primary PAN are replicated to the secondary node. The time taken for the changes to be replicated on the secondary node depends on various factors, such as network latency, load on the system, and so on.


Before you begin

Ensure that the primary PAN and the node being registered are DNS resolvable to each other. If the node that is being registered uses an untrusted self-signed certificate, you are prompted with a certificate warning along with details of the certificate. If you accept the certificate, it is added to the trusted certificate store of the primary PAN to enable TLS communication with the node.

If the node uses a certificate that is not self-signed (for example, signed by an external CA), you must manually import the relevant certificate chain of that node to the trusted certificate store of the primary PAN. When you import the secondary node's certificate to the trusted certificate store, check the **Trust for Authentication**

within ISE check box in the **Trusted Certificates** window for the PAN to validate the secondary node's certificate.

While registering a node with session services enabled (such as Network Access, Guest, Posture, and so on), you can add it to a node group. See [Create a Policy Service Node Group, on page 55](#) for more details.

-
- Step 1** Log in to the primary PAN.
- Step 2** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Deployment**.
- Step 3** Click **Register** to initiate registration of a secondary node.
- Step 4** Enter the DNS-resolvable fully qualified domain name (FQDN) of the standalone node that you are going to register (in the format hostname.domain-name, for example, abc.xyz.com). The FQDN of the primary PAN and the node being registered must be resolvable from each other.
- Step 5** Enter the GUI-based administrator credentials for the secondary node in the **Username** and **Password** fields.
- Step 6** Click **Next**.

The primary PAN tries to establish TLS communication (for the first time) after the node is registered.

- If the node uses a certificate that is trusted, you can proceed to Step 7.
- If the node uses a self-signed certificate that is not trusted, a certificate warning message is displayed with details about the certificate (such as, Issued-to, Issued-by, Serial number, and so on), which can be verified against the actual certificate on the node. Click the **Import Certificate and Proceed** option to trust this certificate and proceed with registration. Cisco ISE imports the default self-signed certificate of that node to the trusted certificate store of the primary PAN. If you do not want to use the default self-signed certificate, click **Cancel Registration** and manually import the relevant certificate chain of that node to the trusted certificate store of the primary PAN. When you import the secondary node's certificate to the trusted certificate store, check the **Trust for Authentication within ISE** check box adjacent to the corresponding PAN to validate the secondary node's certificate.
- If the node uses a CA-signed certificate, an error message is displayed, stating that the registration cannot proceed until certificate trust is set up.

- Step 7** Check the check boxes to select the personas and services to be enabled on the node, and then click **Save**.

When a node is registered, an alarm (which confirms that a node has been added to the deployment) is generated on the primary PAN. You can view this alarm in the **Alarms** dashlet in the Cisco ISE GUI **Dashboard**. After the registered node is synchronized and restarted, you can log in to the secondary node GUI using the same credentials used on the primary PAN.

What to do next

- For time-sensitive tasks such as guest user access and authorization, logging, and so on, ensure that the system time on your nodes is synchronized.
- If you registered a secondary PAN, and are using the internal Cisco ISE CA service, you must back up the Cisco ISE CA certificates and keys from the primary PAN and restore them on the secondary PAN.

Support for Multiple Deployment Scenarios

Cisco ISE can be deployed across an enterprise infrastructure, supporting 802.1X wired, wireless, and Virtual Private Networks (VPNs).

The Cisco ISE architecture supports both standalone and distributed (also known as *high availability* or *redundancy*) deployments, where one machine assumes the primary role, and another *backup* machine assumes the secondary role. Cisco ISE features distinct configurable personas, services, and roles, which allow you to create and apply Cisco ISE services where needed in the network. The result is a comprehensive Cisco ISE deployment that operates as a fully functional and integrated system.

Cisco ISE nodes can be deployed with one or more of the Administration, Monitoring, and Policy Service personas. Each persona performs a different, but vital, part in your overall network policy management topology. Installing Cisco ISE with an administration persona allows you to configure and manage your network from a centralized portal to promote efficiency and ease of use.

Cisco ISE Distributed Deployment

A deployment that has more than one Cisco ISE node is called a distributed deployment. To support failover and to improve performance, you can set up your deployment with multiple Cisco ISE nodes in a distributed fashion. In a Cisco ISE distributed deployment, the administration and monitoring activities are centralized, and processing is distributed across the PSNs. Depending on your performance needs, you can scale your deployment. Each Cisco ISE node in a deployment can assume any of these personas—Administration, Policy Service, and Monitoring.

Cisco ISE Deployment Setup

After you install Cisco ISE on all your nodes, as described in the [Cisco Identity Services Engine Hardware Installation Guide](#), the nodes come up in a standalone state. You must then define one node as your primary PAN. While defining your primary PAN, you must enable the administration and monitoring personas on that node. You can optionally enable the policy service persona on the primary PAN. After you complete the task of defining personas on the primary PAN, you can register other secondary nodes to the primary PAN and define personas for the secondary nodes.

All Cisco ISE system and functionality-related configurations should be done only on the primary PAN. The configuration changes that you perform on the primary PAN are replicated to all the secondary nodes in your deployment.

There must be at least one MnT in a distributed deployment. At the time of configuring your primary PAN, you must enable the Monitoring persona. After you register an MnT node in your deployment, you can edit the primary PAN and disable the Monitoring persona, if required.

Data Replication from Primary to Secondary Cisco ISE Nodes

When you register a Cisco ISE node as a secondary node, Cisco ISE immediately creates a data replication channel from the primary to the secondary node and begins the process of replication. Replication is the process of sharing Cisco ISE configuration data from the primary to the secondary nodes. Replication ensures consistency among the configuration data available in all the Cisco ISE nodes that are part of your deployment.

, click the corresponding radio button to enable or disable the replication of the dynamically discovered endpoints across all the nodes in your Cisco ISE deployment:

To enable or disable endpoint replication using OpenAPI, see the [Cisco ISE API Reference Guide](#).

A full replication typically occurs when you first register a Cisco ISE node as a secondary node. Incremental replication occurs after a full replication and ensures that any new changes, such as additions, modifications, or deletions to the configuration data in the PAN are reflected in the secondary nodes. The process of replication ensures that all the Cisco ISE nodes in a deployment are in sync. You can view the status of replication in the **Node Status** column in the **Deployment** window of the Cisco ISE Admin portal. When you register a Cisco ISE node as a secondary node or perform a manual synchronization with the PAN, the node status shows an orange icon, indicating that the requested action is in progress. After the synchronization is complete, the node status turns green, indicating that the secondary node is synchronized with the PAN.



Note In Cisco ISE replication, if the primary node is unable to publish replication events to the secondary nodes and the number of un-published events exceeds 1.5 million, the system marks all the secondary nodes OUT OF SYNC.

On the other hand, if the delta between the published sequence and the secondary nodes consuming the replication events exceeds 2 million, the system marks those specific PSNs OUT OF SYNC.

Cisco ISE Node Deregistration

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

When a PSN is deregistered, the endpoint data is lost. If you want the PSN to retain the endpoint data after it becomes a standalone node, you can do one of the following:

- Obtain a backup from the primary PAN, and when the PSN becomes a standalone node, restore this data backup on it.
- Change the persona of the PSN to administration (secondary PAN), synchronize the data from the **Deployment** window of the Admin portal, and then deregister the node. This node will now have all the data. You can then add a secondary PAN to the existing deployment.



Note You cannot deregister a primary PAN.

Node Status in a Cisco ISE Deployment

Table 1: Cisco ISE Node Status

Node Status	Description	Guidelines
Connected	Node is connected and replication works fine.	-

Node Status	Description	Guidelines
Replication Stopped	Node is connected but replication has stopped. The node is reachable on ports 443 and 12001. This is a temporary status that will change when the underlying replication issue is resolved.	The underlying issue may resolve by itself. If not, log in to the Cisco ISE GUI and perform a manual synchronization from the Deployment window.
Node Not Reachable	Node is not reachable on port 443 but replication works fine.	The node may come back up by itself when the underlying issue resolves. If not, log in to the Cisco ISE GUI and perform a manual synchronization from the Deployment window.
Disconnected	Node is not reachable and replication has stopped.	This status is set if the node is down for more than five minutes. The node may come back up by itself when the underlying issue resolves. If not, log in to the Cisco ISE GUI and perform a manual synchronization from the Deployment window.
Registration failed	Node registration has failed.	Perform a manual synchronization. Deregister and register the affected Cisco ISE node again.
In Progress	Node registration or manual sync is in progress. Currently, the In Progress status has a timeout value of 300 minutes, after which the status changes to: <ul style="list-style-type: none"> • Connected if registration or manual sync is successful. • Registration Failed if node registration fails. • Not in Sync if manual sync fails. 	No action is required if the status changes to Connected. If it changes to Registration Failed or Not in Sync, check the respective rows and perform the necessary actions.
Not in Sync	Node is not in sync.	Perform a manual synchronization. Deregister and register the affected Cisco ISE node again.
Upgrading	Node upgrade is in progress.	-

Guidelines for Setting Up a Distributed Deployment

Read the following statements carefully before you set up Cisco ISE in a distributed environment:

- Choose a node type for the Cisco ISE server. You must choose a Cisco ISE node for administration, policy service, and monitoring capabilities.
- Choose the same Network Time Protocol (NTP) server for all the nodes. To avoid timezone issues among the nodes, you must provide the same NTP server name when setting up each node. This setting ensures that the reports and logs from the various nodes in your deployment are always synchronized with timestamps.
- Configure the Cisco ISE administrator password when you install Cisco ISE. The previous Cisco ISE administrator default login credentials (admin/cisco) are no longer valid. Use the username and password that was created during the initial setup, or the current password if it was changed later.
- Configure the DNS server. Enter the IP addresses and fully qualified domain names (FQDNs) of all the Cisco ISE nodes that are part of your distributed deployment in the DNS server. Otherwise, node registration fails.
- Configure the forward and the reverse DNS lookup for all the Cisco ISE nodes in your distributed deployment in the DNS server. Otherwise, you may run into deployment-related issues when registering and restarting Cisco ISE nodes. Performance might be degraded if reverse DNS lookup is not configured for all the nodes.
- (Optional) Deregister a secondary Cisco ISE node from the primary PAN to uninstall Cisco ISE from it.
- Back up the primary MnT, and restore the data to the new secondary MnT. This ensures that the history of the primary MnT is in sync with the new MnT because the new changes are replicated.
- Ensure that the primary PAN and the standalone node that you are about to register as a secondary node are running the same version of Cisco ISE.
- Enable **Internal CA Settings** on your Cisco ISE primary PAN before you add another node to your deployment to ensure that the Cisco ISE certificate services function as expected. To enable Internal CA Settings, click the **Menu** icon (☰) and choose **Administration > System > Certificates > Certificate Authority > Internal CA Settings**.
- While adding a new node to the deployment, make sure that the issuer certificate chain of wildcard certificates is part of the trusted certificates of the new node. When the new node is added to the deployment, the wildcard certificates are replicated to the new node.
- When configuring your Cisco ISE deployment to support Cisco TrustSec, or when Cisco ISE is integrated with Cisco Catalyst Center, do not configure a PSN as SXP-only. SXP is an interface between Cisco TrustSec and non-Cisco TrustSec devices. SXP does not communicate with the Cisco TrustSec-enabled network devices.

Menu Options Available on Primary and Secondary Nodes

The menu options that are available in Cisco ISE nodes that are a part of a distributed deployment depend on the personas that are enabled on them. You must perform all administration and monitoring activities through the primary PAN. For other tasks, you must use the secondary nodes. Therefore, the user interface of the secondary nodes provides limited menu options based on the persona that is enabled on them.

If a node assumes more than one persona, for example, the Policy Service persona, and a Monitoring persona with a primary role, the menu options listed for the PSNs and the primary MnT are available on that node.

The following table lists the menu options that are available on the Cisco ISE nodes that assume different personas.

Table 2: Cisco ISE Nodes and Available Menu Options

Cisco ISE Node	Available Menu Options
All Nodes	<ul style="list-style-type: none"> View and configure the system time and the NTP server settings. Install the server certificate and manage certificate signing request. You can perform server certificate operations for all the nodes in the deployment through the primary PAN that centrally manages all the server certificates. <p>Note The private keys are not stored in the local database and are not copied from the relevant node. The private keys are stored in the local file system.</p>
Primary Policy Administration node (primary PAN)	All menus and submenus.
Primary Monitoring node (primary MnT node)	<ul style="list-style-type: none"> Provides access to monitoring data. <p>Note The Operations menu can be viewed only from the primary PAN. The Operations menu does not appear in the monitoring nodes.</p>
PSNs (Policy Service nodes)	Options to join, leave, and test the Active Directory connection are available. Each PSN must be separately joined to the Active Directory domain. You must first define the domain information and join the PAN to the Active Directory domain. Then, join the other PSNs to the Active Directory domain individually.
Secondary Policy Administration node (secondary PAN)	<p>Option to promote the secondary PAN to primary PAN.</p> <p>Note After you have registered the secondary nodes to the primary PAN, while logging in to the Admin portal of any of the secondary nodes, you must use the login credentials of the primary PAN.</p>

Deployment and Node Settings

The **Deployment Nodes** window enables you to configure the Cisco ISE (PAN, PSN, and MnT) nodes and to set up a deployment.

Deployment Nodes List Window


The following table describes the fields in the **Deployment Nodes List** window, which you can use to configure Cisco ISE nodes in a deployment. To view this window, click the **Menu** icon () and choose **Administration > System > Deployment**.

Table 3: Deployment Nodes List

Field Name	Usage Guidelines
Hostname	Displays the hostname of the node.
Personas	(Only appears if the node type is Cisco ISE) Lists the personas that a Cisco ISE node has assumed, for example, Administration, Policy Service, Monitoring, or pxGrid. For example, Administration , Policy Service , Monitoring , or pxGrid .
Role	Indicates the role (primary, secondary, or standalone) that the Administration and Monitoring personas have assumed, if these personas are enabled on this node. The role can be any one or more of the following: <ul style="list-style-type: none"> • PRI(A): Refers to the primary PAN. • SEC(A): Refers to the secondary PAN. • PRI(M): Refers to the primary MnT. • SEC(M): Refers to the secondary MnT.
Services	(Only appears if the Policy Service persona is enabled) Lists the services that run on this Cisco ISE node. Services can include any one of the following: <ul style="list-style-type: none"> • Identity Mapping • Session • Profiling • All
Node Status	Indicates the status of each Cisco ISE node in a deployment for data replication: <ul style="list-style-type: none"> • Green (Connected): Indicates that a Cisco ISE node, which is already registered in the deployment, is in sync with the primary PAN. • Red (Disconnected): Indicates that a Cisco ISE node is not reachable, is down, or data replication is not happening. • Orange (In Progress): Indicates that a Cisco ISE node is newly registered with the primary PAN, you have performed a manual sync operation, or the Cisco ISE node is not in sync (out of sync) with the primary PAN. For more information, click the quick view icon for each Cisco ISE node in the Node Status column.

Related Topics

[Cisco ISE Distributed Deployment](#), on page 5

[Cisco ISE Deployment Terminology](#), on page 1

[Configure a Cisco ISE Node](#), on page 2

[Register a Secondary Cisco ISE Node](#), on page 3

General Node Settings



The following table describes the fields on the **General Settings** window of a Cisco ISE node. In this window, you can assign a persona to a node and configure the services to be run on it. To view this window, click the **Menu** icon () and choose **Administration > System > Deployment > Deployment Node > Edit > General Settings**.

Table 4: General Node Settings

Field Name	Usage Guidelines
Hostname	Displays the hostname of the Cisco ISE node.
FQDN	Displays the fully qualified domain name of the Cisco ISE node, for example, ise1.cisco.com.
IP Address	Displays the IP address of the Cisco ISE node.
Node Type	Displays the node type.
Personas	
Administration	<p>Enable this toggle button if you want a Cisco ISE node to assume the Administration persona. You can enable the Administration persona only on nodes that are licensed to provide the administrative services.</p> <p>Role: Displays the role that the Administration persona has assumed in the deployment. The persona can take one of these values—Standalone, Primary, or Secondary.</p> <p>Make Primary: Click this to make this node your primary Cisco ISE node. You can have only one primary Cisco ISE node in a deployment. The other options in this window will become active only after you make this node primary. You can have only two Administration nodes in a deployment. If the node has a Standalone role, the Make Primary button appears next to it. If the node has a Secondary role, the Promote to Primary button appears next to it. If the node has a Primary role, and there are no other nodes registered with it, the Make Standalone button appears next to it. Click the Make Standalone button to make your primary node a standalone node.</p>

Field Name	Usage Guidelines
Monitoring	<p>Click this toggle button if you want a Cisco ISE node to assume the Monitoring persona and function as your log collector. There must be at least one Monitoring node in a distributed deployment. At the time of configuring your primary PAN, you must enable the Monitoring persona. After you register a secondary Monitoring node in your deployment, you can edit the primary PAN and disable the Monitoring persona, if required.</p> <p>To configure a Cisco ISE node on a VMware platform as your log collector, use the following guidelines to determine the minimum amount of disk space that you need: 180 KB per endpoint in your network per day and 2.5 MB per Cisco ISE node in your network per day.</p> <p>You can calculate the maximum disk space that you need based on how many months of data you want to have in your Monitoring node. If there is only one Monitoring node in your deployment, it assumes the standalone role. If you have two Monitoring nodes in your deployment, Cisco ISE displays the name of the other Monitoring node too for you to configure the primary-secondary roles. To configure these roles, choose one of the following:</p> <ul style="list-style-type: none"> • Primary: For the current node to be the primary Monitoring node. • Secondary: For the current node to be the secondary Monitoring node. • None: If you do not want the Monitoring nodes to assume the primary-secondary roles. <p>If you configure one of your Monitoring nodes as primary or secondary, the other Monitoring node automatically becomes the secondary or primary node, respectively. Both the primary and secondary Monitoring nodes receive Administration and Policy Service logs. If you change the role for one Monitoring node to None, the role of the other Monitoring node also becomes None, thereby cancelling the high availability pair after you designate a node as a Monitoring node. You will find this node listed as a syslog target in the Remote Logging Targets window. To view this window, click the Menu icon () and choose Administration > System > Logging > Remote Logging Targets.</p>

Field Name	Usage Guidelines
Policy Service	

Field Name	Usage Guidelines
	<p>Click this toggle button to enable any one or all of the following services:</p> <ul style="list-style-type: none"> Enable Session Services: Check this check box to enable network access, posture, guest, and client-provisioning services. From the Include Node in Node Group drop-down list, choose the group to which this Policy Service node belongs. Note that Certificate Authority (CA) and Enrollment over Secure Transport (EST) services can only run on a Policy Service node that has session services enabled on it. <p>For Include Node in Node Group, choose None if you do not want this Policy Service node to be a part of a group.</p> <p>All the nodes within the same node group should be configured on the network access device (NAD) as RADIUS clients and authorized for CoA, because any one of them can issue a CoA request for the sessions that are established through any node in the node group. If you are not using a load balancer, the nodes in a node group should be the same as, or a subset of the RADIUS servers and clients configured on the NAD. These nodes would also be configured as RADIUS servers.</p> <p>While a single NAD can be configured with many Cisco ISE nodes as RADIUS servers and dynamic-authorization clients, it is not necessary for all the nodes to be in the same node group.</p> <p>The members of a node group should be connected to each other using high-speed LAN connection such as Gigabit Ethernet. The node group members need not be L2 adjacent, but L2 adjacency is highly recommended to ensure sufficient bandwidth and reachability. See the Create a Policy Service Node Group, on page 55 for more details.</p> <ul style="list-style-type: none"> Enable Profiling Service: Check this check box to enable the Profiling service. If you enable the Profiling service, you must click the Profiling Configuration tab and enter the details, as required. When you enable or disable any of the services that run on the Policy Service node or make any changes to this node, you will be restarting the application server processes on which these services run. Expect a delay while these services restart. You can determine when the application server has restarted on a node by using the show application status ise command from the CLI. Enable Threat-Centric NAC Service: Check this check box to enable the Threat-Centric Network Access Control (TC-NAC) feature. This feature allows you to create authorization policies based on the threat and vulnerability attributes received from the threat and vulnerability adapters. Threat severity levels and vulnerability assessment results can be used to dynamically control the access level of an endpoint or a user. Enable SXP Service: Check this check box to enable SXP service on the node. You must also specify the interface to be used for SXP service. <p>If you have configured NIC bonding or teaming, the bonded interfaces are also listed along with the physical interfaces in the Use Interface drop-down list.</p> <ul style="list-style-type: none"> Enable Device Admin Service: Check this check box to create TACACS policy sets, policy results, and so on, to control and audit the configuration of network

Field Name	Usage Guidelines
	<p>devices.</p> <ul style="list-style-type: none"> • Enable Passive Identity Service: Check this check box to enable the Identity Mapping feature. This feature enables you to monitor users who are authenticated by a Domain Controller and not by Cisco ISE. In networks where Cisco ISE does not actively authenticate users for network access, you can use the Identity Mapping feature to collect user authentication information from the Active Directory Domain Controller.
pxGrid	Check this check box to enable the pxGrid persona. Cisco pxGrid is used to share the context-sensitive information from the Cisco ISE session directory to other policy network systems such as Cisco Adaptive Security Appliance (ASA). The pxGrid framework can also be used to exchange policy and configuration data between nodes, for example, sharing tags and policy objects between Cisco ISE and third-party vendors, and for non-Cisco ISE-related information exchanges such as threat information.

Related Topics

[Personas in Distributed Cisco ISE Deployments](#), on page 2

[Administration Node](#), on page 22

[Policy Service Node](#), on page 30

[Monitoring Node](#), on page 33

[Cisco pxGrid Node](#), on page 40

[Synchronize Primary and Secondary Cisco ISE Nodes](#), on page 54

[Create a Policy Service Node Group](#), on page 55

[Deploy Cisco pxGrid Node](#), on page 41

[Change Node Personas and Services](#), on page 54

[Configure MnT Nodes for Automatic Failover](#), on page 39

Profiling Node Settings

The following table describes the fields in the **Profiling Configuration** window, that you can use to configure the probes for the profiler service. To access this window, click the **Administration > System > Deployment > ISE Node > Edit > Profiling Configuration**.

Table 5: Profiling Node Settings

Field Name	Usage Guidelines
NetFlow	<p>Click this toggle button to enable NetFlow for each Cisco ISE node that has assumed the Policy Service persona to receive NetFlow packets sent from the routers. Enter the required values for the following options:</p> <ul style="list-style-type: none"> • Interface: Choose the interface on the Cisco ISE node. • Port: Enter the NetFlow listener port number on which NetFlow exports are received from the routers. The default port is 9996.

Field Name	Usage Guidelines
DHCP	<p>Click this toggle button to enable DHCP for each Cisco ISE node that has assumed the Policy Service persona to listen for DHCP packets from the IP helper. Provide values for the following options:</p> <ul style="list-style-type: none"> • Interface: Choose the interface on the Cisco ISE node. • Port: Enter the DHCP server UDP port number. The default port is 67.
DHCP SPAN	<p>Click this toggle button to enable DHCP SPAN for each Cisco ISE node that has assumed the Policy Service persona to collect DHCP packets.</p> <ul style="list-style-type: none"> • Interface: Choose the interface on the Cisco ISE node.
HTTP	<p>Click this toggle button to enable HTTP per Cisco ISE node that has assumed the Policy Service persona to receive and parse HTTP packets.</p> <ul style="list-style-type: none"> • Interface: Choose the interface on the Cisco ISE node.
RADIUS	<p>Click this toggle button to enable the RADIUS server for each Cisco ISE node that has assumed the Policy Service persona to collect RADIUS session attributes as well as Cisco Device Protocol (CDP) and Link Layer Discovery Protocol (LLDP) attributes from the Cisco IOS Sensor-enabled devices.</p>
Network Scan (NMAP)	<p>Click this toggle button to enable the NMAP probe.</p>
DNS	<p>Click this toggle button to enable DNS for each Cisco ISE node that has assumed the Policy Service persona to perform a DNS lookup for the FQDN. Enter the Timeout period in seconds.</p> <p>Note For the DNS probe to work on a particular Cisco ISE node in a distributed deployment, you must enable one of these probes—DHCP, DHCP SPAN, HTTP, RADIUS, or SNMP. For DNS lookup, one of these probes must be started along with the DNS probe.</p>
SNMP Query	<p>Click this toggle button to enable SNMP query for each Cisco ISE node that has assumed the Policy Service persona to poll network devices at specified intervals. Enter values in Retries, Timeout, Event Timeout (mandatory), and Description (optional) fields.</p> <p>Note In addition to configuring the SNMP Query probe, you must also configure other SNMP settings in Administration > Network Resources > Network Devices. When you configure SNMP settings on the network devices, ensure that you enable CDP and LLDP globally on your network devices.</p>

Field Name	Usage Guidelines
SNMP Trap	<p>Click this toggle button to enable an SNMP Trap probe for each Cisco ISE node that has assumed the Policy Service Persona to receive linkUp, linkDown, and MAC notification traps from the network devices. Provide or enable the following information:</p> <ul style="list-style-type: none"> • Link Trap Query: Enable this toggle button to receive and interpret the notifications received through the SNMP trap. • MAC Trap Query: Enable this toggle button to receive and interpret the MAC notifications received through the SNMP trap. • Interface: Choose an interface on the Cisco ISE node. • Port: Enter the UDP port of the host to use. The default port is 162.
Active Directory	<p>Click this toggle button to scan the defined Active Directory servers for information about Windows users.</p> <ul style="list-style-type: none"> • Days before rescan: Choose the days after which you want the scan to run again.
pxGrid	<p>Click this toggle button to allow Cisco ISE to collect (profile) endpoint attributes over pxGrid.</p>

Related Topics

[Cisco ISE Profiling Service](#)

[Network Probes Used by Profiling Service](#)

[Configure Profiling Service in Cisco ISE Nodes](#)

Logging Settings

The following sections explain how to configure the severity of debug logs, create an external log target, and enable Cisco ISE to send log messages to these external log targets.

Remote Logging Target Settings

The following table describes the fields in the **Remote Logging Targets** window that you can use to create external locations (syslog servers) to store logging messages. To access this window, **Administration > System > Logging > Remote Logging Targets**, and click **Add**.

Table 6: Remote Logging Target Settings

Field Name	Usage Guidelines
Name	Enter a name for the new syslog target.
Target Type	Select the target type from the drop-down list. The default value is UDP Syslog .
Description	Enter a brief description of the new target.

Field Name	Usage Guidelines
IP Address	Enter the IP address or hostname of the destination machine that will store the logs. Cisco ISE supports IPv4 and IPv6 formats for logging.
Port	Enter the port number of the destination machine.
Facility Code	Choose the syslog facility code that must be used for logging, from the drop-down list. Valid options are Local0 through Local7.
Maximum Length	Enter the maximum length of the remote log target messages. Valid values are from 200 through 1024 bytes.
Include Alarms For this Target	When you check this check box, alarm messages are sent to the remote server as well.
Comply to RFC 3164	When you check this check box, the delimiters (, ; { } \ \) in the syslog messages sent to the remote servers are not escaped even if a backslash (\) is used.
Buffer Message When Server Down	This check box is displayed when you choose TCP Syslog or Secure Syslog from the Target Type drop-down list. Check this check box to allow Cisco ISE to buffer the syslog messages when a TCP syslog target or secure syslog target is unavailable. Cisco ISE retries sending messages to the target when the connection to the target resumes. After the connection resumes, messages are sent sequentially, starting with the oldest, and proceeding to the newest. Buffered messages are always sent before new messages. If the buffer is full, old messages are discarded.
Buffer Size (MB)	Set the buffer size for each target. By default, it is set to 100 MB. Changing the buffer size clears the buffer, and all the existing buffered messages for the specific target are lost.
Reconnect Timeout (Sec)	Enter the time (in seconds) to configure how long the TCP and secure syslogs are stored for before being discarded when the server is down.
Select CA Certificate	This drop-down list is displayed when you choose Secure Syslog from the Target Type drop-down list. Choose a client certificate from the drop-down list.
Ignore Server Certificate Validation	This check box is displayed when you choose Secure Syslog from the Target Type drop-down list. Check this check box for Cisco ISE to ignore server certificate authentication and accept any syslog server. By default, this option is set to Off unless the system is in FIPS mode when this is disabled.

Related Topics

- [Cisco ISE Logging Mechanism](#)
- [Cisco ISE System Logs](#)
- [Cisco ISE Message Catalogs](#)
- [Collection Filters](#)
- [Event Suppression Bypass Filter](#)
- [Configure Remote Syslog Collection Locations](#)
- [Configure Collection Filters](#)

Configure Logging Categories

The following table describes the fields that you can use to configure a logging category. Set a log severity level and choose the logging targets for the logs of a logging category. To access this window, choose **Administration > System > Logging > Logging Categories**.

Click the radio button next to the logging category that you want to view, and click **Edit**. The following table describes the fields that are displayed in the edit window of the logging categories.

Table 7: Logging Category Settings

Field Name	Usage Guidelines
Name	Displays the name of the logging category.
Log Severity Level	For some logging categories, this value is set by default, and you cannot edit it. For some logging categories, you can choose one of the following severity levels from a drop-down list: <ul style="list-style-type: none"> • FATAL: Emergency level. This level means that you cannot use Cisco ISE and you must immediately take the necessary action. • ERROR: This level indicates a critical error condition. • WARN: This level indicates a normal but significant condition. This is the default level set for many logging categories. • INFO: This level indicates an informational message. • DEBUG: This level indicates a diagnostic bug message.
Local Logging	Check this check box to enable logging events for a category on the local node.
Targets	This area allows you to choose the targets for a logging category by transferring the targets between the Available and the Selected areas using the left and right arrow icons. The Available area contains the existing logging targets, both local (predefined) and external (user-defined). The Selected area, which is initially empty, then displays the targets that have been chosen for the category.

Related Topics

[Cisco ISE Message Codes](#)

[Configure Remote Syslog Collection Locations](#)

[Set Severity Levels for Message Codes](#)

Admin Access Settings

These sections enable you to configure access settings for administrators.

Administrator Password Policy Settings

The following table describes the fields in the **Password Policy** tab that you can use to define a criteria that administrator passwords should meet. To view this window, click the **Menu** icon (☰) and choose **Administration > System > Admin Access > Authentication > Password Policy**.

Table 8: Administrator Password Policy Settings

Field Name	Usage Guidelines
Minimum Length	Specify the minimum length of the password (in characters). The default is six characters.
Password must not contain	Admin name or its characters in reverse order: Check this check box to restrict the use of the administrator username or its characters in reverse order as the password.
	Cisco or its characters in reverse order: Check this check box to restrict the use of the word "Cisco" or its characters in the reverse order as the password.
	This word or its characters in reverse order: Check this check box to restrict the use of any word that you define or its characters in the reverse order as the password.
	Repeated characters four or more times consecutively: Check this check box to restrict the use of repeated characters four or more times consecutively as the password.
	Dictionary words, their characters in reverse order, or their letters replaced with other characters: Check this check box to restrict the use of dictionary words, their characters in reverse order, or their letters replaced with other characters, as the password. Substitution of \$ for s, @ for a, 0 for o, 1 for l, ! for i, 3 for e, and so on, is not permitted. For example, Pa\$\$w0rd is not permitted. <ul style="list-style-type: none"> • Default Dictionary: Choose this option to use the default Linux dictionary in Cisco ISE. The default dictionary contains approximately 480,000 English words. This option is selected by default. • Custom Dictionary: Choose this option to use your customized dictionary. Click Choose File to select a custom dictionary file. The text file must comprise newline-delimited (JSON format) words, .dic extension, and a size less than 20 MB.
Password must contain at least one character of each of the selected types	Check the check box for the type of characters an administrator's password must contain. Choose one or more of the following options: <ul style="list-style-type: none"> • Lowercase alphabetic characters • Uppercase alphabetic characters • Numeric characters • Non-alphanumeric characters

Field Name	Usage Guidelines
Password History	<p>Specify the number of previous passwords from which the new password must be different, to prevent the repeated use of the same password. Check the Password must be different from the previous <i>n</i> versions check box, and enter the number in the corresponding field.</p> <p>Enter the number of days before which you cannot reuse a password. Check the Cannot reuse password within <i>n</i> days check box, and enter the number in the corresponding field.</p>
Password Lifetime	<p>Check the check boxes for the following options to force users to change passwords after a specified time period:</p> <ul style="list-style-type: none"> • Administrator passwords expire <i>n</i> days after creation or last change: Time (in days) before the administrator account is disabled if the password is not changed. The valid range is 1 to 3650 days. • Send an email reminder to administrators <i>n</i> days prior to password expiration: Time (in days) before which administrators are reminded that their password will expire. The valid range is 1 to 3650 days.
Display Network Device-Sensitive Data	
Require Admin Password	Check this check box if you want the admin user to enter the login password to view network device-sensitive data such as shared secrets and passwords.
Password cached for <i>n</i> Minutes	The password that is entered by the admin user is cached for this time period. The admin user will not be prompted to enter the password again during this period to view the network device-sensitive data. The valid range is from 1 to 60 minutes.

Related Topics

[Cisco ISE Administrators](#)

[Create a New Administrator](#)

Session Timeout and Session Information Settings

The following table describes the fields in the **Session** window that you can use to define session timeout and terminate an active administrative session. To access this window, click the **Menu** icon (☰) and choose **Administration > System > Admin Access > Settings > Session**.

Table 9: Session Timeout and Session Information Settings

Field Name	Usage Guidelines
Session Timeout	
Session Idle Timeout	Enter the time, in minutes, that you want Cisco ISE to wait for, before it logs out the administrator if there is no activity. The default value is 60 minutes. The valid range is from 6 to 100 minutes.
Session Info	

Field Name	Usage Guidelines
Invalidate	Check the check box adjacent to the session ID that you want to terminate and click Invalidate .

Related Topics

- [Administrator Access Settings](#)
- [Configure Session Timeout for Administrators](#)
- [Terminate an Active Administrative Session](#)

Administration Node

A Cisco ISE node with the Administration persona allows you to perform all administrative operations on Cisco ISE. It handles all the system-related configurations that are related to functionalities such as authentication, authorization, auditing, and so on. In a distributed environment, you can have a maximum of two nodes running the Administration persona. The Administration persona can take on of these following roles—Standalone, Primary, or Secondary.

High Availability for Administrative Node

In a high-availability configuration, the primary Policy Administration Node (PAN) is in the Active state. The secondary PAN is in the Standby state, which means it receives all configuration updates from the primary PAN, but is not active in the Cisco ISE network.

Cisco ISE supports manual and automatic failover. With automatic failover, when the primary PAN goes down, an automatic promotion of the secondary PAN is initiated. Automatic failover requires a nonadministration secondary node, which is called a health check node. The health check node checks the health of the primary PAN. If the health check node detects that the primary PAN is down or unreachable, it initiates the promotion of the secondary PAN to take over the primary role.

To deploy the Automatic Failover feature, you must have at least three nodes, with two of them assuming the Administration persona, and one acting as the health check node. A health check node is a nonadministration node and can be a PSN, MnT, or pxGrid node, or a combination of these. If the primary and secondary PANs are in different data centers, you must have a health check node for each PAN.

The following table lists the features that are affected when the primary PAN goes down and the secondary PAN is yet to take over.

Table 10: Availability of Features

Feature Name	Available When Primary PAN is Down? (Yes/No)
Existing internal user RADIUS authentication	Yes
Existing or new AD user RADIUS authentication	Yes

Feature Name	Available When Primary PAN is Down? (Yes/No)
Existing endpoint with no profile change	Yes
Existing endpoint with profile change	No
New endpoint learned through profiling.	No
Existing guest: Local Web Authentication (LWA)	Yes
Existing guest: Central Web Authentication (CWA)	Yes (apart from flows enabled for device registration, such as Hotspot, BYOD, and CWA with automatic device registration)
Guest change password	No
Guest: AUP	No
Guest: Max Failed Login Enforcement	No
New Guest (Sponsored or Self-registered)	No
Posture	Yes
BYOD with Internal CA	No
Existing Registered Devices	Yes
MDM on-boarding	No
pxGrid Service	No
Log in to GUI of secondary nodes	Yes (The login process is delayed because a blocking call to the PAN is attempted to update the last login details. Login proceeds after this call times out.)



Note To support certificate provisioning with the internal CA, you must to import the root certificate of the original primary PAN and its key into the new primary node, after promotion. Certificate provisioning does not work after automatic failover for the PSN nodes that are added after the promotion of the secondary node to primary PAN.

High-Availability Health Check Nodes

The health check node for the Primary PAN is called the active health check node. The health check node for the Secondary PAN is called the passive health check node. The active health check node is responsible for checking the status of the Primary PAN, and managing the automatic failover of Administration nodes. We recommend that you use two nonadministrative ISE nodes as health check nodes, one for the Primary PAN and one for the Secondary PAN. If you use only one health check node, and that node goes down, automatic failover will not happen.

When both the PANs are in the same data center, you can use a single nonadministrative ISE node as the health check node for both the Primary PAN and the Secondary PAN. When a single health check node checks the health of both the Primary PAN and the Secondary PAN, it assumes both the active and passive roles.

A health check node is a nonadministration node, which means it can be a Policy Service, Monitoring, or pxGrid node, or a combination of these. We recommend that you designate PSN nodes as health check nodes in the same data center as the Administration nodes. However, in a small or a centralized deployment, where the two Administration nodes are not in the same location (LAN or data center), any node (PSN, pxGrid, or MnT) not having the Administration persona can be used as the health check node.



Note If you chose to not enable automatic failover, and rely on manually promoting the secondary node when the primary PAN fails, you do not need any check nodes.

Health Check Node for the Secondary PAN

The health check node for the Secondary PAN is a passive monitor. It does not take any action until the Secondary PAN has been promoted as the Primary PAN. When the Secondary PAN takes over the primary role, its associated health check node takes the active role for managing automatic failover of Administration nodes. The health check node of the previous Primary PAN becomes the health check node for the Secondary PAN now and monitors it passively.

Disabling and Restarting Health Check

When a node is removed from the health check role or auto failover configuration is disabled, the health check service is stopped on that node. When the auto failover configuration is enabled on the designated high-availability health check node, the node starts checking the health of Administration nodes again. Designating or removing the high-availability health check role of a node does not involve any application restart on that node; only the health check activities are started or stopped.

If the high-availability health check node is restarted, it ignores the previous downtimes of the Primary PAN and starts checking the health status afresh.

Health Check Nodes

The active health check node checks the health status of the primary PAN at a configured polling interval. It sends a request to the primary PAN, and if the response that it receives matches the configuration, the health check node considers the primary PAN to be in good health. If the health of the primary PAN is continuously poor for more than the configured failover period, the health check node initiates failover to the secondary PAN.

If, at any time during a health check, the health status is found to be good after being reported as poor previously within the failover period, the health check node marks the primary PAN status as good, and resets the health check cycle.


The response from the health check of the primary PAN is validated against the configuration values available on its health check node. If the response does not match, it raises an alarm. However, a promotion request is made to the secondary PAN.

Changing Health Nodes

You can change the Cisco ISE node that you are using for a health check, but there are some things to consider.

For example, assume that the health check node (H1) goes out-of-sync, and another node (H2) is made the health check node of the primary PAN. In such a case, after the primary PAN goes down, there is no way for H1 to know that another node (H2) is checking the same primary PAN. Later, if H2 goes down or goes out of the network, an actual failover is required. The secondary PAN, however, retains the right to reject the promotion request. So, after the secondary PAN is promoted to the primary role, a promotion request from H2 is rejected with an error. Even if a health check node for the primary PAN is out of sync, it continues to check the health of the primary PAN.

Automatic Failover to the Secondary PAN

You can configure Cisco ISE to automatically promote the Secondary PAN when the Primary PAN becomes unavailable. The configuration is done on the Primary PAN in the **Deployment** window. To view this window, click the **Menu** icon () and choose **Administration > System > Deployment**. The failover period is defined as the number of times configured in **Number of Failure Polls Before Failover** times the number of seconds configured in **Polling Interval**. In the default configuration, that time is 10 minutes. Promotion of the Secondary PAN to Primary PAN takes another 10 minutes. So, by default, the total time from Primary PAN failure to secondary PAN working is 20 minutes.

When the Secondary PAN receives the failover call, it carries out the following validations before proceeding with the actual failover:

- The Primary PAN is not available in the network.
- The failover request came from a valid health check node.
- The failover request is for the Secondary PAN.

If all the validations pass, the Secondary PAN promotes itself to the primary role.

The following are some sample (but not limited to) scenarios where automatic failover of the Secondary PAN can be attempted:

- Health of the Primary PAN is consistently not good for the **Number of failure polls before failover** value during the polling period.

- Cisco ISE services on the Primary PAN are manually stopped, and remain stopped for the failover period.
- The Primary PAN is shut down using soft halt or reboot option, and remains shut down for the configured failover period.
- The Primary PAN goes down abruptly (power down), and remains down for the failover period.
- The network interface of the Primary PAN is down (network port shut or network service down), or it is not reachable by the health check node for any other reason, and remains down for the configured failover period.

Health Check Node Restarts

Upon restart, the high-availability health check node ignores the previous downtimes of the Primary PAN and checks the health status afresh.

Bring Your Own Device in Case of Automatic Failover to Secondary PAN

When the Primary PAN is down, authentication is not interrupted for the endpoints that already have certificates issued by the Primary PAN root CA chain. This is because all the nodes in the deployment have the entire certificate chain for trust and validation purposes.

However, until the Secondary PAN is promoted to Primary, new BYOD devices will not be onboarded. BYOD onboarding requires an active Primary PAN.

After the original primary PAN is brought back up or the Secondary PAN is promoted, new BYOD endpoints are onboarded without any issues.

If the Primary PAN that failed can not be rejoined as the Primary PAN, regenerate the root CA certificate on the newly promoted Primary PAN (the original secondary PAN).

For existing certificate chains, triggering a new root CA certificate results in the automatic generation of the subordinate CA certificates. Even when new subordinate certificates are generated, endpoint certificates that were generated by the previous chain continue to be valid.

Sample Scenarios when Automatic Failover is Avoided

The following are some sample scenarios that depict cases where automatic failover by the health check node might be avoided or a promotion request to the secondary node rejected:

- The node receiving the promotion request is not the secondary node.
- The promotion request received by the Secondary PAN does not have the correct Primary PAN information.
- The promotion request is received from an incorrect health check node.
- The promotion request is received, but the Primary PAN is up and in good health.
- The node receiving the promotion request goes out-of-sync.

Functionalities Affected by the PAN Automatic Failover Feature

The following table lists the functionalities that are blocked or require additional configuration changes if the PAN automatic failover configuration is enabled in your deployment.


Functionality	Affected Details
Operations that are Blocked	
Upgrade	<p>Upgrade through the CLI is blocked.</p> <p>By default, this feature is disabled.</p> <p>To deploy the Automatic Failover feature, you must have at least three nodes, where two of the nodes assume the Administration persona, and one node acts as the health check node. (A health check node is a nonadministration node and can be a PSN, MnT, or pxGrid node, or a combination of these). If the PANs are in different data centers, you must have a health check node for each PAN.</p>
Restore of Backup	<p>Restore action through the CLI and user interface is blocked.</p> <p>If the PAN automatic failover configuration was enabled prior to restore, you must reconfigure it after a successful restore.</p>
Change Node Persona	<p>Change of the following node personas through the GUI is blocked:</p> <ul style="list-style-type: none"> • Administration persona in both the Primary and Secondary PANs • Persona of the PAN • Deregistration of health check node after enabling the PAN Automatic Failover feature
Other CLI Operations	<p>The following admin operations through the CLI is blocked:</p> <ul style="list-style-type: none"> • Patch installation and rollback • DNS server change • IP address change of eth1, eth2, and eth3 interfaces • Host alias change of eth1, eth2, and eth3 interfaces • Time zone change
Other Administration Portal Operations	<p>The following administrative operations through the GUI is blocked:</p> <ul style="list-style-type: none"> • Patch installation and rollback • Change of HTTPS certificate • Change of admin authentication type from password-based authentication to certificate-based authentication and vice versa
Users with maximum connected devices cannot connect.	Some session data is stored on the failed PAN, and cannot be updated by the PSN.
Operations that Require PAN Automatic Failover to be Disabled	

Functionality	Affected Details
CLI Operations	<p>The following administrative operations through the CLI display a warning message if the PAN automatic failover configuration is enabled. These operations may trigger automatic failover if a service or system is not restarted within the failover window. Hence, while performing the following operations, we recommend that you to disable the PAN automatic failover configuration:</p> <ul style="list-style-type: none"> • Manually stopping the Cisco ISE service • Soft reload (reboot) of Cisco ISE using the admin CLI

Configure Primary PAN for Automatic Failover

Before you begin

To deploy the Automatic Failover feature, you must have at least three nodes, of which two nodes assume the Administration persona, and one node acts as the health check node. A health check node is a nonadministration node and can be a PSN, MnT, or pxGrid node, or a combination of these. If the PANs are in different data centers, you must have a health check node for each PAN.

-
- Step 1** Log in to the Primary PAN GUI.
- Step 2** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Deployment > PAN Failover**.
- Step 3** Check the **Enable PAN Auto Failover** check box to enable automatic failover of the primary PAN.
- Note** You can only promote a Secondary PAN to become the Primary PAN. Cisco ISE nodes that assume only the PSN, MnT, or pxGrid node, or a combination of these, cannot be promoted to become the Primary PAN.
- Step 4** Choose the health check node for the primary PAN from the **Primary Health Check Node** drop-down list containing all the available secondary nodes.
- We recommend that you have this node in the same location or data center as the primary PAN.
- Step 5** Choose the health check node for the secondary PAN, from the **Secondary Health Check Node** drop-down list containing all the available secondary nodes.
- We recommend that you have this node in the same location or data center as the secondary PAN.
- Step 6** Provide the **Polling Interval** time after which the PAN status is checked. The valid range is 30 to 300 seconds.
- Step 7** Provide the count for **Number of Failure Polls before Failover**.
- Failover occurs if the status of the PAN is not good for the specified number of failure polls. The valid range is 2 to 60 counts.
- Step 8** Click **Save**.
-

What to do next

After the promotion of the Secondary PAN to the Primary PAN, do the following:


- Manually sync the old Primary PAN to bring it back into the deployment.
- Manually sync any other secondary node that is out of sync, to bring it back into the deployment.

Manually Promote Secondary PAN to Primary

If the Primary PAN fails and you have not configured PAN automatic failover, you must manually promote the Secondary PAN to become the new Primary PAN.

Before you begin

Ensure that you have a second Cisco ISE node configured with the Administration persona to promote as your Primary PAN.

-
- Step 1** Log in to the Secondary PAN GUI.
- Step 2** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**.
- Step 3** In the **Edit Node** window, click **Promote to Primary**.

Note You can only promote a secondary PAN to become the primary PAN. Cisco ISE nodes that assume only the Policy Service or Monitoring persona, or both, cannot be promoted to become the primary PAN.

If the node that was originally the Primary PAN, comes back up, it will be demoted automatically and become the Secondary PAN. You must perform a manual synchronization on this node (that was originally the Primary PAN) to bring it back into the deployment.

In the **Edit Node** window of a secondary node, you cannot modify the personas or services because the options are disabled. You have to log in to the Admin portal to make changes.

- Step 4** Click **Save**.
-

Reusing a Node of an Existing Cisco ISE Deployment as a Primary PAN for a New Cisco ISE Deployment

If you want to repurpose a node of an existing Cisco ISE deployment to the primary PAN of a new Cisco ISE deployment you must, perform these steps:

-
- Step 1** Run the Cisco ISE Perform System Erase utility, as described in the *Cisco ISE Installation Guide* for your version of Cisco ISE. This document is available at: <https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-guides-list.html>
- Step 2** Perform a fresh install of Cisco ISE, as described in the *Cisco ISE Installation Guide*.
- Step 3** Configure the standalone node as a primary Policy Administration node, as described in [Configure a Primary Policy Administration Node, on page 3](#).
-

Restoring Service to the Primary PAN

Cisco ISE does not support automatic fallback to the original primary PAN. After the automatic failover to the secondary PAN is initiated, if you bring the original primary PAN back into the network, you should configure it as the secondary PAN.

Support for Automatic Failover for the Administration Node

Cisco ISE supports automatic failover for the Administration persona. To enable the Automatic Failover feature, at least two nodes in your distributed setup should assume the Administration persona and one node should assume the nonadministration persona. If the Primary PAN goes down, an automatic promotion of the Secondary PAN is initiated. For this, a nonadministration secondary node is designated as the health check node for each of the PANs. The health check node checks the health of the primary PAN at configured intervals. If the health check response received for the primary PAN is not good for reasons, such as the device being down or unreachable, the health check node initiates the promotion of the secondary PAN to take over the primary role after waiting for the configured threshold value. Some features are unavailable after automatic failover of the secondary PAN. Cisco ISE does not support fallback to the original primary PAN. See [High Availability for Administrative Node](#).

Policy Service Node

A Policy Service node (PSN) is a Cisco ISE node with the Policy Service persona, and provides network access, posture, guest access, client provisioning, and profiling services.

At least one node in your distributed setup should assume the Policy Service persona. This persona evaluates the policies and makes all the decisions. Typically, there is more than one PSN in a distributed deployment.

All the PSNs that reside in the same high-speed Local Area Network (LAN) or behind a load balancer can be grouped together to form a node group. If one of the nodes in a node group fails, the other nodes detect the failure and reset URL-redirectioned sessions, if any.

High Availability in Policy Service Nodes

To detect node failure and to reset all URL-redirectioned sessions on the failed node, two or more PSNs can be placed in the same node group. When a node that belongs to a node group fails, another node in the same node group issues a Change of Authorization (CoA) for all URL-redirectioned sessions on the failed node.

All the nodes within the same node group should be configured on the network access device (NAD) as RADIUS clients and authorized for CoA, because any one of them can issue a CoA request for the sessions that are established through any node in the node group. If you are not using a load balancer, the nodes in a node group should be the same as, or a subset of, the RADIUS servers and clients configured on the NAD. These nodes should also be configured as RADIUS servers.



Note While a single NAD can be configured with many Cisco ISE nodes as RADIUS servers and dynamicauthorization clients, it is not necessary for all the nodes to be in the same node group.

The members of a node group should be connected to each other using high-speed LAN connection such as Gigabit Ethernet. The node group members need not be L2 adjacent, but L2 adjacency is highly recommended to ensure sufficient bandwidth and reachability. See [Create a Policy Service Node Group, on page 55](#) for more details.

Load Balancer to Distribute Requests Evenly Among PSNs

When you have multiple PSNs in the deployment, you can use a load balancer to distribute the requests evenly. The load balancer distributes the requests to the functional nodes behind it. See [Cisco and F5 Deployment Guide: ISE Load Balancing using BIG-IP](#) for more information, and to know about best practices when deploying PSNs behind a load balancer.

Session Failover in Policy Service Nodes

PSNs in a node group share session information. The nodes exchange heartbeat messages to detect node failures. If a node fails, one of its peers from the node group knows which sessions were on the failed PSN, and issues a CoA to disconnect those sessions. Most clients automatically reconnect, and establish a new session.

Some clients don't automatically reconnect. For example, if a client connects through a VPN, then that client may not see the CoA. Clients that are IP phones, multihost 802.1X ports, or virtual machines may also not see or be able to respond to a CoA. URL-redirectioned clients (webauth) also can't connect automatically. Those clients must manually reconnect.

Timing issues can also prevent reconnection, for example, if the posture state is pending at the time of PSN failover.

For more information about PSN session sharing, see [Light Data Distribution, on page 31](#).

Number of Nodes in a Policy Service Node Group

The number of nodes that you can have in a node group depends on your deployment requirements. Node groups ensure that node failures are detected and that a peer issues a CoA for sessions that are authorized, but not yet postured. The size of the node group does not have to be very large.

If the size of the node group increases, the number of messages and heartbeats that are exchanged between the nodes increases significantly. As a result, traffic also increases. Having fewer nodes in a node group helps reduce the traffic and at the same time provides sufficient redundancy to detect PSN failures.

There is no hard limit on the number of PSNs that you can have in a node group cluster.

Light Data Distribution

Light Data Distribution is used to store user session information and replicate it across the PSNs in a deployment, thereby eliminating the need to be dependent on the PAN or MnT nodes for user session details.

Light Data Distribution consists of the following directories:

- [RADIUS Session Directory](#)
- [Endpoint Owner Directory](#)

In addition, you can configure the following options under **Advanced Settings**:


- **Batch Size:** The session updates can be sent in batches. This value specifies the number of records sent in each batch from a Light Data Distribution instance to the other PSNs in the deployment. If this field is set to 1, the session updates are *not* sent in batches. The default value is 10 records.
- **TTL:** This value specifies the maximum time a session will wait for a batch to complete before updating the Light Data Distribution. The default value is 1000 milliseconds.

In case of connectivity issues between the PSNs, for example, when a PSN is down, the session details are retrieved from the MnT session directory and stored for future use.

Large deployments can hold up to 2,000,000 session records. Small deployments can store 1,000,000 session records. When an accounting stop request is received for a session, the corresponding session data is deleted from all Light Data Distribution instances. When the number of stored records exceeds the maximum limit, the oldest sessions are deleted based on the timestamp.



Note

- If the IPv6 prefix length of a session is less than 128 bits and the interface ID is not specified, the IPv6 prefix is rejected, thereby preventing multiple sessions from having the same key.
- Light Data Distribution uses Cisco ISE messaging services for inter-node communication. Cisco ISE Release 3.0 and later support Certificate Signing Request generation for Cisco ISE Messaging Service. Thus, Cisco ISE Release 3.0 and later have both internal and external CA support for ISE Messaging Service. If you face issues with the Cisco ISE messaging service, you have to regenerate Cisco ISE messaging service certificate.
- **1.** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Certificates > Certificate Management > Certificate Signing Requests**
- **2.** In the **Certificate(s) will be used for** section, select **ISE Messaging service**.
- **3.** Click **Generate ISE messaging service certificate**.

RADIUS Session Directory

The **RADIUS Session Directory** is used to store the user session information and replicate it across the PSNs in a deployment. This directory stores only the session attributes that are required for CoA.

This functionality is enabled by default from Cisco ISE Release 2.7. You can enable or disable this functionality, by checking or unchecking the **RADIUS Session Directory** check box in the **Light Data Distribution** window. To view this window, click the **Menu** icon () and choose **Administration > System > Settings > Light Data Distribution**.

Endpoint Owner Directory

Until Cisco ISE Release 2.6, when an endpoint probe is received on a Policy Service node (PSN) that is different from the one that originally handled the requests for that specific endpoint, the endpoint owner is changed to the new PSN. This results in endpoint ownership flapping.

From Cisco ISE Release 2.7, the **Endpoint Owner Directory** is used to store the PSN FQDN of each MAC address connecting to Cisco ISE and to replicate this data across the PSNs in a deployment. This avoids

endpoint ownership flapping because all the PSNs are now aware of all the endpoint owners. The endpoint ownership now changes only in case of a successful RADIUS authentication of that endpoint on another PSN.

In addition, the static endpoint assignments are prioritized over the attributes received by an incoming probe for the same endpoint, avoiding attribute override issues.

This feature is enabled by default from Cisco ISE Release 2.7. If required, you can disable it to fall back to the old mechanism of not using the endpoint owner directory. The **Endpoint Owner Directory** is also used in profiling, and disabling this option will use the legacy profiler owner's directory. You can enable or disable this feature by checking or unchecking the **Enable Endpoint Owner Directory** check box in the **Light Data Distribution** window. (To view this window, click the **Menu** icon (☰) and choose **Administration > System > Settings > Light Data Distribution**).

Monitoring Node

A Cisco ISE node with the Monitoring persona functions as the log collector and stores log messages from the PANs and PSNs in your network. This persona provides advanced monitoring and troubleshooting tools that you can use to effectively manage your network and resources. A node with this persona aggregates and correlates the data that it collects to provide you with meaningful information in the form of reports.

Cisco ISE allows you to have a maximum of two nodes with this persona that can take on primary or secondary roles for high availability. Both the primary and secondary MnT nodes collect log messages. If the primary MnT goes down, the primary PAN points to the secondary node to gather monitoring data. But the secondary node will not be promoted to primary automatically. This should be done by following the procedure described in [Manually Modify the MnT Role](#).

At least one node in your distributed setup should assume the Monitoring persona. We recommend that you do not have the Monitoring and Policy Service personas enabled on the same Cisco ISE node, and that the node be dedicated solely to monitoring, for optimum performance.

You can access the Monitoring menu from the PAN in your deployment.



Note If you have enabled pxGrid, you must create a new certificate for the pxGrid node. Create the certificate template with digital signature usage and generate a new PxGrid certificate.

Manually Modify the MnT Role

You can manually modify MnT roles (both from primary to secondary and from secondary to primary) from the primary PAN.

-
- Step 1** Log in to the primary PAN GUI.
 - Step 2** In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **Administration > System > Deployment**.
 - Step 3** From the list of nodes, check the check box next to the MnT node for which you want to change the role.
 - Step 4** Click **Edit**.
 - Step 5** In the **Monitoring** section, change the role to **Primary** or **Secondary**.

You can enable the **Dedicated MnT** option if you want to disable all the other personas and services enabled on that node. When this option is enabled, the configuration data replication process is stopped on that node. This helps to improve the performance of the MnT node. When you disable this option, manual synchronization is triggered.

Step 6 Click **Save**.

Syslog over Cisco ISE Messaging Service

Cisco ISE, Release 2.6, offers MnT WAN survivability for the default, built-in UDP syslog collection targets, LogCollector and LogCollector2. This survivability can be enabled by the option **Use "ISE Messaging Service" for UDP Syslogs delivery to MnT** (In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **System > Logging > Log Settings**). After you enable this option, the UDP syslogs are protected by Transport Layer Security (TLS).

The **Use "ISE Messaging Service" for UDP Syslogs delivery to MnT** option is disabled by default in Cisco ISE, Release 2.6, First Customer Ship (FCS). This option is enabled by default in Cisco ISE, Release 2.6, Cumulative Patch 2 and later releases.

Using the Cisco ISE messaging service for UDP syslogs retains the operational data for a finite duration even when the MnT node is unreachable. The MnT WAN survivability period is approximately 2 hours and 30 mins.

This service uses TCP port 8671. Please configure your network accordingly and allow the connections to TCP port 8671 on each Cisco ISE node from all other Cisco ISE nodes in the deployment.



Note If your deployment uses TCP or secure syslogs for Cisco ISE deployment, the functionality remains same as the earlier releases.

We recommend defining the UDP syslog target using the IP address. This avoids the latencies encountered while defining UDP syslog targets using the DNS name or DNS server.

Queue Link Alarm

The Cisco ISE messaging service uses a different certificate, signed by the internal CA chain. You will get a `queue-link alarm` message in the **Alarms** dashlet in the Cisco ISE GUI dashboard. Ensure that the following are in place to resolve the alarm:

- All the nodes are connected and synced.
- All the nodes and Cisco ISE messaging services are functional.
- The Cisco ISE messaging services ports are not blocked by external entities such as firewalls.
- The Cisco ISE messaging certificate chain on each node is not broken, and the state of the certificate is good.


If the prerequisites listed above are met, the queue-link alarm might be triggered due to the upgrade process.

To resolve the queue link alarm, regenerate the Cisco ISE root CA chain:

1. In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **Administration > System > Certificates > Certificate Management > Certificate Signing Requests**.

2. Click **Generate Certificate Signing Request (CSR)** and then select the **ISE Root CA** in the **Certificate(s) will be used for** drop-down list.
3. Click **Replace ISE root CA Certificate Chain**.

A **Queue Link Error** alarm is generated in the following scenarios:


- **Timeout**: A **Queue Link Error** alarm with the cause as **Timeout** occurs when there is a network issue between two nodes in the Cisco ISE deployment. Check the connectivity on port 8671 to troubleshoot this error.
- **Unknown CA**: A **Queue Link Error** alarm with the cause as **Unknown CA** occurs when there is a broken Cisco ISE Messaging Certificate present in the **System Certificates** window (To view this window, click the **Menu** icon () and choose **Administration > System > Certificates > System Certificates**.) This issue can be resolved by regenerating the Cisco ISE Messaging Certificate by choosing **Administration > System > Certificates > Certificate Management > Certificate Signing Requests** and then clicking on **Generate Certificate Signing Request (CSR)** on the Cisco ISE GUI.



Note Regeneration is not required if you have already replaced the Cisco ISE root CA certificate chain.

When you replace the Cisco ISE Root CA chain, the Cisco ISE Messaging Service certificate is also replaced. This is followed by the restart of the Cisco ISE Messaging service with a downtime of about two minutes. The syslogs are lost during this downtime. To avoid losing the syslogs during the downtime, the Cisco ISE Messaging Service can be disabled for a short period.

To enable or disable the Cisco ISE messaging service for UDP syslogs delivery to MnT, perform this procedure:

-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **System > Logging > Log Settings**.
 - Step 2** Check or uncheck the **Use “ISE Messaging Service” for UDP Syslogs delivery to MnT** check box to enable or disable the use of the Cisco ISE Messaging Service.
 - Step 3** Click **Save**.

[Cisco ISE Community Resource](#)

For more information about the Queue Link alarm, see [Queue Link Error](#).

Automatic Failover in MnT Nodes

MnT nodes do not offer high availability, but do offer active standby. The PSN copies operational audit data to both the primary and secondary MnT nodes.

Automatic Failover Process

When a primary MnT node goes down, the secondary MnT node takes over all the monitoring and troubleshooting information.

To manually convert the secondary node to a primary node, see [Manually Modify the MnT Role](#). If the primary node comes up again after the secondary node is promoted, the primary node takes on the secondary role. If the secondary node is not promoted, the primary MnT node resumes the primary role after it comes up again.



Caution When the primary node comes back up after a failover, back up of the secondary and restore the data to update the primary node.

Guidelines for Setting Up an Active Standby Pair of MnT Nodes

You can specify two MnT nodes on a Cisco ISE network, and configure them to be an active standby pair. We recommend that you back up the primary MnT node, and restore the data to the new secondary MnT node. This ensures that the history of the primary MnT node is synchronized with the new secondary node because the primary replicates new data. The following rules apply to an active standby pair:

- All changes are logged to the primary MnT node. The secondary node is read-only.
- Changes made to the primary node are automatically replicated on the secondary node
- Both the primary and secondary nodes are listed as log collectors, to which all other nodes send logs.
- The Cisco ISE dashboard is the main entry point for monitoring and troubleshooting. Monitoring information is displayed on the dashboard from the PAN . If the primary node goes down, monitoring information is available on the secondary node.
- Backing up and purging MnT data is not a part of a standard Cisco ISE node backup process. You must configure repositories for backup and data purging on both the primary and secondary MnT nodes, and use the same repositories for each.

MnT Node Failover Scenarios

The following scenarios apply to the active-standby or single-node configurations corresponding to the MnT nodes:

- In an active-standby configuration of the MnT nodes, the primary PAN always points to the primary MnT node to collect the monitoring data. After the primary MnT node fails, the PAN points to the standby MnT node. The failover from the primary node to the secondary node takes place after it is down for more than five minutes.

However, after the primary node fails, the secondary node does not become the primary node. If the primary node comes up, the PAN starts collecting the monitoring data again from the resumed primary node.

- If the primary MnT node is down, and you want to promote the standby MnT node to active status, you can do so by following the procedure provided in [Manually Modify the MnT Role](#) or by deregistering the existing primary MnT node. When you deregister the existing primary MnT node, the standby node becomes the primary MnT node, and the PAN automatically points to the newly promoted primary node.
- In an active-standby pair, if you deregister the secondary MnT node, or if the secondary MnT node goes down, the existing primary MnT node remains the current primary node.
- If there is only one MnT node in the Cisco ISE deployment, then that node acts as the primary MnT node, and provides monitoring data to the PAN. However, when you register a new MnT node, and make

it the primary node in the deployment, the existing primary MnT node automatically becomes the standby node. The PAN points to the newly registered primary MnT node to collect monitoring data.

Monitoring Database

The rate and amount of data that is utilized by the monitoring functions requires a separate database on a dedicated node that is used for these purposes.

Like PSN, the MnT node has a dedicated database that requires you to perform maintenance tasks, as described in the topics covered in this section.

Back Up and Restore the Monitoring Database

The Monitoring database handles large volumes of data. Over time, the performance and efficiency of the MnT node depends on how well you manage that data. To increase efficiency, we recommend that you back up the data and transfer it to a remote repository on a regular basis. You can automate this task by scheduling automatic backups.



Note You should not perform a backup when a purge operation is in progress. If you start a backup during a purge operation, the purge operation stops or fails.

If you register a secondary MnT node, we recommend that you first back up the primary MnT node and then restore the data to the new secondary MnT node. This ensures that the history of the primary MnT node is in sync with the new secondary node when the new changes are replicated.


Monitoring Database Purge

The purging process allows you to manage the size of the Monitoring database by specifying the number of months to retain the data during a purge. The default is three months. This value is utilized when the disk space usage threshold for purging (80 percentage of the total disk space) is met. For this option, each month consists of 30 days. A default of three months equals 90 days.

Guidelines for Purging the Monitoring Database

Follow these guidelines for optimal Monitoring database disk usage:

- If the Monitoring database disk usage is greater than 80 percent of the threshold setting, that is 60 percent of total disk space, a critical alarm is generated, indicating that the database size is about to exceed the maximum amount of allocated disk size. If the disk usage is greater than 90 percent of the threshold setting, that is 70 percent of total disk space, another alarm is generated, indicating that the database size has exceeded the maximum amount of allocated disk size.

A purge process runs, creating a status history report that you can view in the **Data Purging Audit** window. To view this window, click the **Menu** icon () and choose **Operations > Reports > Reports > Audit > Data Purging Audit**. An information (INFO) alarm is generated after the purge is completed.

- Purging is also based on the percentage of consumed disk space for the database. When the consumed disk space for the Monitoring database is equal to or exceeds the threshold (the default is 80 percentage of the total disk space), the purge process starts. This process deletes only the oldest seven days' monitoring data, irrespective of what is configured in the Admin portal. It continues this process in a loop until the disk space is below 80 percent. Purging always checks the Monitoring database disk space limit before proceeding.


Operational Data Purging

Cisco ISE Monitoring Operational database contains information that is generated as Cisco ISE reports. Recent Cisco ISE (Cisco ISE Release 2.4 and above) releases have options to purge the monitoring operational data and reset the monitoring database when the **application configure ise** command is run.

The purge option is used to clean up the data and prompts you to enter the number of days for which to retain the data. The reset option is used to reset the database to the factory default, so that all the data that is backed up is permanently deleted. Specify the database if the files are consuming too much file system space.



Note The reset option causes Cisco ISE services to be temporarily unavailable.

The **Operational Data Purging** window contains the **Database Utilization** and **Purge Data Now** areas. To view this window, click the **Menu** icon () and choose **Administration > System > Maintenance > Operational Data Purging**. You can view the total available database space and the RADIUS and TACACS data stored in the **Database Utilization** area. Hover the mouse over the status bar to display the available disk space and the number of days the existing data is stored for in the database. Specify the period for which the RADIUS and TACACS data is supposed to be retained in the **Data Retention Period** area. Data is purged at 4 a.m. every day, and you can configure the export of data to a repository before it is purged, by specifying the number of retention days. Check the **Enable Export Repository** check box to select and create a repository, and specify an **Encryption Key**.

In the **Purge Data Now** area, you can purge all the RADIUS and TACACS data or specify the number of days beyond which data is supposed to be purged.



Note You must export RADIUS authentication and accounting, TACACS authorization and accounting, RADIUS errors, and misconfigured supplicants tables to a repository before purging.

Related Topics


[Purge Older Operational Data](#), on page 38

Purge Older Operational Data

The operational data is collected in the server over a period of time. It can be purged either instantly or periodically. You can verify the success of the data purge by viewing the **Data Purging Audit** report.

Before you begin

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

Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Maintenance > Operational Data Purging**.

Step 2 Do one of the following:

- In the **Data Retention Period** area:
 - a. Specify the time period, in days, for which RADIUS and TACACS data should be retained. All the data prior to the specified time period is exported to a repository.
 - b. In the **Repository** area, check the **Enable Export Repository** check box to choose the repository to save data.
 - c. In the **Encryption Key** field, enter the required password.
 - d. Click **Save**.

Note If the configured retention period is less than the existing retention thresholds corresponding to the diagnostics data, the configured value overrides the existing threshold values. For example, if you configure the retention period as three days and this value is less than the existing thresholds in the diagnostics tables (for example, a default of five days), the data is purged according to the value that you configure (three days) in this window.

- In the **Purge Data Now** area:
 - a. Choose to purge all the data or to purge the data that is older than the specified number of days. Data is not saved in any repository.
 - b. Click **Purge**.
-

Configure MnT Nodes for Automatic Failover

If you have two MnT nodes in a deployment, you can configure a primary-secondary pair for automatic failover to avoid downtime in the Cisco ISE Monitoring service. A primary-secondary pair ensures that a secondary MnT node automatically provides monitoring if the primary node fails.

Before you begin


- Before you configure MnT nodes for automatic failover, they must be registered as Cisco ISE nodes.
- Configure monitoring roles and services on both the nodes and name them for their primary and secondary roles, as appropriate.
- Configure repositories for backup and data purging on both the primary and secondary MnT nodes. For the backup and purging features to work properly, use the same repositories for both the nodes. Purging takes place on both the primary and secondary nodes of a redundant pair. For example, if the primary MnT node uses two repositories for backup and purging, you must specify the same repositories for the secondary node.

Configure a data repository for a MnT node using the **repository** command in the system CLI.



Note For scheduled backup and purge to work properly on the nodes of a monitoring redundant pair, configure the same repository, or repositories, on both the primary and secondary nodes using the CLI. The repositories are not automatically synced between the two nodes.

From the Cisco ISE dashboard, verify that the MnT nodes are ready. The **System Summary** dashlet shows the MnT nodes with a green check mark to the left when their services are ready.

-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Deployment**.
- Step 2** In the **Deployment Nodes** window, check the check box next to the MnT node that you want to specify as primary, and click **Edit**.
- Step 3** Click the **General Settings** tab and choose **Primary** from the **Role** drop-down list.
- When you choose an MnT node as primary, the other MnT node automatically becomes secondary. In the case of a standalone deployment, primary and secondary role configuration is disabled.
- Step 4** Click **Save**. Both the primary and secondary nodes restart.
-

Cisco pxGrid Node

You can use Cisco pxGrid to share the context-sensitive information from Cisco ISE session directory with other network systems such as Cisco ISE ecosystem, partner systems, and other Cisco platforms. The pxGrid framework can also be used to exchange policy and configuration data between nodes, such as sharing tags and policy objects between Cisco ISE and third-party vendors, and for other information exchanges. Cisco pxGrid also allows third-party systems to invoke adaptive network control actions (ANC) to quarantine users or devices or both in response to a network or security event. Cisco TrustSec information, such as tag definition, value, and description can be passed from Cisco ISE through the Cisco TrustSec topic to other networks. The endpoint profiles with Fully Qualified Names (FQNs) can be passed from Cisco ISE to other networks through an endpoint profile meta topic. Cisco pxGrid also supports bulk download of tags and endpoint profiles.

You can publish and subscribe to SXP bindings (IP-SGT mappings) through Cisco pxGrid. For more information about SXP bindings, see [e](#).

The following logs are available for the Cisco pxGrid node:

- pxgrid.log: Provides state change notifications.
- pxgrid-cm.log: Displays updates on publisher or subscriber or both and data exchange activity between the client and the server.
- pxgrid-controller.log: Displays the details of client capabilities, groups, and client authorization.
- pxgrid-jabberd.log: Displays all the logs related to system state and authentication.
- pxgrid-pubsub.log: Displays all the information related to publisher and subscriber events.

**Note**

- You can enable Cisco pxGrid and Cisco pxGrid persona with the Cisco ISE Advantage license.
- Cisco pxGrid should be defined in order to work with the Passive ID Work Center. For more information, see [PassiveID Work Center](#)

High Availability for pxGrid 2.0

pxGrid 2.0 nodes operate in an Active/Active configuration. For high availability, there should be at least two pxGrid nodes in the deployment. Large deployments can have up to four nodes for increased scale and redundancy. We recommend that you configure IP addresses for all the nodes, so that if one node goes down, that node's clients connect to the working node. When the PAN goes down, the pxGrid server stops handling the activations. Manually promote the PAN to activate the pxGrid server. For more information about pxGrid deployments, see [Performance and Scalability Guide for Cisco Identity Services Engine](#)

All the pxGrid service provider clients periodically reregister themselves with the pxGrid controller within a span of 7.5 minutes. If the client does not reregister, the PAN node assumes that the client is inactive and deletes the client. If the PAN node goes down for more than 7.5 minutes, when it comes back up, it deletes all the clients with timestamp values older than 7.5 minutes. All those clients must then register again with the pxGrid controller.

pxGrid 2.0 clients use WebSocket and REST-based APIs for pub/sub and query. These APIs are served by the ISE application server on port 8910. The pxGrid processes shown by `show logging application pxgrid` don't apply to pxGrid 2.0.

**Note**


All the references to pxGrid 1.0 processes in the GUI and the CLI have been removed.

Deploy Cisco pxGrid Node

You can enable Cisco pxGrid persona both on a standalone node and distributed deployment node.

Before you begin

- You must have a Cisco ISE Advantage license to enable the Cisco pxGrid persona. For licensing requirements, see [ISE Licensing / Ordering](#).
- All nodes use the CA certificate for Cisco pxGrid service usage. If you used the default certificate for Cisco pxGrid service before the upgrade, the upgrade replaces that certificate with the internal CA certificate.


-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Deployment**.
- Step 2** In the **Deployment Nodes** window, check the check box next to the node for which you want to enable the Cisco pxGrid services, and click **Edit**.
- Step 3** Click the **General Settings** tab and enable the **pxGrid** toggle button.
- Step 4** Click **Save**.

Note When you upgrade from the previous version, the **Save** option might be disabled. This happens when the browser cache refers to the old files from the previous version of Cisco ISE. Clear the browser cache to enable the **Save** option.

Configure Cisco pxGrid Settings

Before you begin

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

Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > pxGrid Services > Settings**.

Step 2 Check one of the following check boxes based on your requirements:

- **Automatically approve new certificate-based accounts:** Check this check box to automatically approve the connection requests from new Cisco pxGrid clients.
- **Allow password-based account creation:** Check this check box to enable username or password-based authentication for Cisco pxGrid clients. When this option is enabled, Cisco pxGrid clients cannot be automatically approved.

Step 3 Click **Save**.

Use the **Test** option in the Cisco pxGrid **Settings** window to run a health check on the Cisco pxGrid node. View the details in the pxgrid or pxgrid-test.log file.

Use the **PxGrid Client Auto Approval** API to:

- Enable automatic approval of certificate-based connection requests from new pxGrid clients. Enable this option only when you trust all the clients in your environment.
- Enable username or password-based authentication for the pxGrid clients. When this option is enabled, pxGrid clients cannot be automatically approved. A pxGrid client can register itself with the pxGrid controller by sending the username through a REST API. The pxGrid controller generates a password for the pxGrid client during client registration. An administrator can approve or deny the connection request.

For more information about the PxGrid Client Auto Approval API, see the “pxGrid Settings” section in the ERS SDK. You can access the ERS SDK at the following URL:


<https://<ISE-Admin-Node>:9060/ers/sdk>

Only users with ERS Admin role can access the ERS SDK.

Generate Cisco pxGrid Certificate

Before you begin

- You must not use the same certificate for Cisco ISE pxGrid server and pxGrid clients. You must use client certificates for the pxGrid clients. To generate client certificates, choose **Administration > System > Certificates**.
- Some versions of Cisco ISE have a certificate for Cisco pxGrid that uses NetscapeCertType. We recommend that you generate a new certificate.
- To perform the following task, you must be a Super Admin or System Admin.
- A Cisco pxGrid certificate must be generated from the primary PAN.
- If the Cisco pxGrid certificate uses the subject alternative name (SAN) extension, be sure to include the FQDN of the subject identity as a DNS name entry.
- Create a certificate template with digital signature usage and use that to generate a new Cisco pxGrid certificate.

-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > pxGrid Services > Client Management > Certificates**.
- Step 2** From the **I want to** drop-down list, choose one of the following options:
- **Generate a single certificate (without a certificate signing request)**: You must enter the Common Name (CN) if you select this option.
 - **Generate a single certificate (with a certificate signing request)**: You must enter the Certificate Signing Request details if you select this option.
- Step 3** (Optional) Enter a description for this certificate.
- Step 4** Click the **pxGrid_Certificate_Template** link to download and edit the certificate template based on your requirements.
- Step 5** Enter the **Subject Alternative Name (SAN)**. You can add multiple SANs. The following options are available:
- **IP address**: Enter the IP address of the Cisco pxGrid client to be associated with the certificate.
 - **FQDN**: Enter the FQDN of the pxGrid client.
- Step 6** From the **Certificate Download Format** drop-down list, choose one of the following options:
- **Certificate in Private Enhanced Electronic Mail (PEM) format, key in PKCS8 PEM format (including certificate chain)**: The root certificate, the intermediate CA certificates, and the end entity certificate are represented in the PEM format. PEM-formatted certificates are BASE64-encoded ASCII files. Each certificate starts with the "-----BEGIN CERTIFICATE-----" tag and ends with the "-----END CERTIFICATE-----" tag. The end entity's private key is stored using PKCS* PEM. It starts with the "-----BEGIN ENCRYPTED PRIVATE KEY-----" tag and ends with the "-----END ENCRYPTED PRIVATE KEY-----" tag.
 - **PKCS12 format (including certificate chain; one file for both the certificate chain and key)**: A binary format to store the root CA certificate, the intermediate CA certificate, and the end entity's certificate and private key in one encrypted file.

Step 7 Enter the password for the certificate.

Step 8 Click **Create**.

You can view the certificate that you created in the **Issued Certificates** window. To view this window, click the **Menu** icon () and choose **Administration > System > Certificates > Certificate Authority > Issued Certificates**.

Note From Cisco ISE 2.4 patch 13 onwards, the certificate requirements have become stricter for the pxGrid service. If you are using the Cisco ISE default self-signed certificate as the pxGrid certificate, Cisco ISE might reject that certificate after applying Cisco ISE 2.4 patch 13 or later. This is because the earlier versions of that certificate have the **Netscape Cert Type** extension specified as **SSL Server**, which now fails (a client certificate is also required now).

Any client with a noncompliant certificate fails to integrate with Cisco ISE. Use a certificate issued by the internal CA, or generate a new certificate with proper usage extensions:


- The **Key Usage** extension in the certificate must contain the **Digital Signature** and **Key Encipherment** fields.
- The **Extended Key Usage** extension in the certificate must contain the **Client Authentication** and **Server Authentication** fields.
- The **Netscape Certificate Type** extension is not required. If you want to include that extension, add both **SSL Client** and **SSL Server** in the extension.
- If you are using a self-signed certificate, the **Basic Constraints CA** field must be set to **True**, and the **Key Usage** extension must contain the **Key Cert Sign** field.

Control Permissions for Cisco pxGrid Clients

You can create Cisco pxGrid authorization rules for controlling the permissions for the Cisco pxGrid clients. Use these rules to control the services that are provided to the Cisco pxGrid clients.

You can create different types of groups and map the services provided to the Cisco pxGrid clients to these groups. Use the **Groups** option in the **Client Management** window to add new groups. You can view the example authorization rules in the **Client Management > Policies** window. Note that you can update only the **Custom Operations** field for the predefined rules.

To create an authorization rule for pxGrid clients:

Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > pxGrid Services > Client Management > Policy**.

Step 2 From the **Service** drop-down list, choose one of the following options:

- **com.cisco.ise.pubsub**
- **com.cisco.ise.config.anc**
- **com.cisco.ise.config.profiler**
- **com.cisco.ise.config.trustsec**
- **com.cisco.ise.service**

- **com.cisco.ise.system**
- **com.cisco.ise.radius**
- **com.cisco.ise.sxp**
- **com.cisco.ise.trustsec**
- **com.cisco.ise.mdm**

Step 3 From the **Operation** drop-down list, choose one of the following options:

- **<ANY>**
- **publish**
- **publish /topic/com.cisco.ise.session**
- **publish /topic/com.cisco.ise.session.group**
- **publish /topic/com.cisco.ise.anc**
- **<CUSTOM>**: You can specify a custom operation if you select this option.

See [pxGrid Operations and Services Use Cases](#) for more information.

Note The following additional attributes are published to the /topic/com.cisco.ise.session for REST ID store for Cisco ISE 3.0 patch 5 and later releases:

- identityProvider:Azure
- oid
- tenantID
- preferredUsername

Step 4 From the **Groups** drop-down list, choose the groups that you want to map to this service.

ANC and manually added groups are listed in this drop-down list.

Note Only the clients that belong to the groups included in the policy can subscribe to the service specified in that policy. For example, if you define a pxGrid policy for com.cisco.ise.pubsub service and assign the ANC group to this policy, only the clients that belong to the ANC group can subscribe to the com.cisco.ise.pubsub service.

pxGrid Operations and Services Use Cases

When creating a new pxGrid policy, note that some pxGrid operations are only applicable to specific services.

You can find the following pxGrid operations in the Cisco ISE GUI.

Operation <ANY>

When you use the <ANY> operation with a service and a particular user group, any operation related to that service is only accessible to the users in the chosen user group.

Consider the following example.

Service: com.cisco.ise.session; Operation: <ANY>; Groups: SessionUsers.

In this example, only pxGrid clients that are part of the 'SessionUsers' group will be able to perform any operation related to session topic (like subscribe/gets operations).

Operation publish

All the publish related operations are applicable only when com.cisco.ise.pubsub is chosen as the service. You can use the publish operation to create a pxGrid policy specifying that only a pxGrid client of a particular user group can publish a chosen topic or can publish all topics.

Operation <Custom>

You can use the <Custom> operation to specify an operation that is not provided in the Operation drop-down list. Currently, pxGrid supports the following operations but not all of them are listed in the Operation drop-down list:

1. 'sets' (applicable on all services and topics except pubsub) – You can use this to restrict access to REST API calls that perform a set operation.
2. 'gets' (applicable on all services and topics except pubsub) – You can use this to restrict access to REST API calls that perform a get operation.
3. 'publish' followed by a particular topic name (only applicable on pubsub service) – You can use this to restrict access to users who can publish a particular topic.

For example, Service: com.cisco.ise.pubsub, Operation: publish/topic/com.cisco.ise.session.

However, some rules with the same operation, service, and topic are incomprehensible and must be avoided. For example, Service: com.cisco.ise.session, Operation: publish /topic/com.cisco.ise.session.

4. 'subscribe' followed by topic name (only applicable on pubsub service) – You can use this to restrict access to users who can subscribe to a particular topic.

For example, Service: com.cisco.ise.pubsub, Operation: publish /topic/com.cisco.ise.session

Cisco pxGrid Cloud Overview

Cisco pxGrid Cloud is a new Cisco cloud offer that extends pxGrid, ERS, and Open API access to cloud-based applications. Cisco ISE 3.1 patch 3 and later releases support Cisco pxGrid Cloud.

To allow connectivity between a Cisco ISE deployment and Cisco pxGrid Cloud, the **pxGrid Cloud** option must be enabled on one or more pxGrid nodes in the Cisco ISE deployment. If you have configured high availability for pxGrid nodes, one of the nodes acts as the Active node and the other one will be the Standby node. The Standby node takes over when the Active node goes down.

Only the Active node establishes connection to Cisco pxGrid Cloud and handles the traffic between the Cisco ISE deployment and Cisco pxGrid Cloud. No other Cisco ISE node interacts with Cisco pxGrid Cloud.

The pxGrid Cloud agent resides in Cisco ISE and serves as the bridge between Cisco ISE and Cisco pxGrid Cloud. A pxGrid Cloud application can subscribe to a pxGrid topic. The pxGrid Cloud agent in Cisco ISE learns about this subscription from Cisco pxGrid Cloud and establishes the actual subscription to the pxGrid service in Cisco ISE. When the agent receives a notification on the pxGrid topic, it forwards the notification

to Cisco pxGrid Cloud over a logical channel dedicated to the pxGrid service. The pxGrid Cloud application can invoke pxGrid, ERS, and Open APIs in the Cisco ISE deployment. The pxGrid Cloud agent proxies a REST request from Cisco pxGrid Cloud to Cisco ISE, and returns the response back to Cisco pxGrid Cloud.

Cisco, its partners, and its customers can develop pxGrid Cloud-based applications and register them with the pxGrid Cloud offer. These applications use the pxGrid, ERS, and Open APIs to exchange information with Cisco ISE.

Cisco ISE customers who have a pxGrid Cloud subscription can register their Cisco ISE deployment with Cisco pxGrid Cloud and use the applications listed in the offer. To do this, they must:


1. Acquire and activate the pxGrid Cloud subscription.
2. Enable the pxGrid Cloud service on one or two pxGrid nodes in the Cisco ISE deployment.
3. Register the Cisco ISE deployment with Cisco pxGrid Cloud (associating it with the subscription) and receive an authentication token.
4. Enter the authentication token in the **Setup Connection** window in Cisco ISE (under **Administration > pxGrid Services > Client Management > pxGrid Cloud Connection**).

This activates the pxGrid Cloud agent on the Active pxGrid Cloud node and establish a connection between the Cisco ISE deployment and Cisco pxGrid Cloud.

5. Select a pxGrid Cloud application from the offer and associate it with the subscription. The application will then have access to the Cisco ISE deployment.

Enable pxGrid Cloud Service in Cisco ISE

Before you begin

- Ensure that you install and activate the Advantage license in your Cisco ISE deployment.
- The pxGrid Cloud agent creates an outbound HTTPS connection to Cisco pxGrid Cloud. Therefore, you must configure Cisco ISE proxy settings if the customer network uses a proxy to reach the internet. To configure proxy settings in Cisco ISE, click the **Menu** icon () and choose **Administration > System > Settings > Proxy**.
- The Cisco ISE Trusted Certificates Store must include the root CA certificate required to validate the server certificate presented by Cisco pxGrid Cloud. Ensure that the **Trust for Authentication of Cisco Services** option is enabled for this root CA certificate.

Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Deployment**.

Step 2 Click the node in which you want to enable the pxGrid Cloud service.

Step 3 In the **General Settings** tab, click the toggle button to enable the **pxGrid** service.


Step 4 Check the **Enable pxGrid Cloud** check box.

The pxGrid Cloud service can be enabled on two nodes to enable high availability.

Note You can enable the **pxGrid Cloud** option only when the **pxGrid** service is enabled on that node.

Connect Cisco ISE to Cisco pxGrid Cloud

After the pxGrid Cloud service is enabled, you must connect the Cisco ISE deployment to Cisco pxGrid Cloud. You must register your Cisco ISE deployment in Cisco pxGrid Cloud and generate an authentication token.

Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > pxGrid Services > Client Management > pxGrid Cloud Connection**.

Step 2 Click **Setup Connection**.

Step 3 Enter the authentication token in the **Setup Connection** window, and then click **Connect**.

The connection setup includes the following steps:

- a. **Enrollment:** A request is sent to Cisco pxGrid Cloud to enroll the Cisco ISE deployment using the authentication token. When this step is successfully completed, the pxGrid Cloud agent is started on the Active node in the Cisco ISE deployment.
- b. **pxGrid Connection:** The pxGrid Cloud agent establishes a persistent connection to the pxGrid component running locally on the same Cisco ISE node. All pxGrid notifications from Cisco ISE are sent to the pxGrid Cloud agent using this connection.
- c. **Cloud Connection:** The pxGrid Cloud agent establishes a persistent connection to Cisco pxGrid Cloud and sets up the logical channels. These logical channels are used to receive the pxGrid, ERS, and Open API requests from Cisco pxGrid Cloud, and to send the pxGrid notifications to Cisco pxGrid Cloud.

You can view the connection setup progress in the **pxGrid Cloud Connection** window. After all these steps are completed, the status is displayed as **Connected**, and the name of the Active pxGrid node is displayed.

To terminate the pxGrid Cloud connection, click **Disconnect** in the **pxGrid Cloud Connection** window. This disconnects the Cisco ISE deployment from Cisco pxGrid Cloud and terminates the pxGrid Cloud agent in the Active node.

After the Cisco ISE deployment is connected to Cisco pxGrid Cloud, the pxGrid Cloud agent (called HERMES process) is listed in the output of the **show application status ise** CLI command.

Disable pxGrid Cloud Service on Cisco ISE

Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Deployment**.

Step 2 Check the check box next to the corresponding pxGrid node and click **Edit**.


Step 3 Uncheck the **Enable pxGrid Cloud** check box.

This stops the pxGrid Cloud agent in the Cisco ISE deployment. You can re-enable the pxGrid Cloud service later when needed.

Configure a pxGrid Cloud Policy

By default, pxGrid Cloud applications are not permitted to access any pxGrid, ERS, or Open APIs in the Cisco ISE deployment. Access must be explicitly granted by configuring policies in Cisco ISE.

You can create a policy to specify what is allowed or denied between your Cisco ISE deployment and the pxGrid Cloud service. Authorization policies specific to each partner environment can be configured in the cloud portal. You will need the Cisco ISE Advantage license to configure a pxGrid Cloud policy.


Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > pxGrid Services > Client Management > pxGrid Cloud Policy**.

Step 2 In the **pxGrid Services** area, choose the required services from the list. You can enable one or more pxGrid services by clicking their names.

Step 3 In the **ERS APIs** area, click the toggle button enable the **ERS APIs** option to provide ERS API access to pxGrid Cloud applications.

Note The **ERS APIs** option is disabled here if the ERS service is disabled in Cisco ISE.

To enable this service in Cisco ISE, perform these steps:


a. In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Settings > API Settings > API Service Settings**.

b. Click the toggle button to enable the **ERS (Read/Write)** option.

Step 4 In the **Open APIs** area, click the toggle button to enable the **Open APIs** option to provide Open API access to pxGrid Cloud applications.

Note The **Open APIs** option is disabled here if the **Open API** option is disabled in Cisco ISE.

To enable this service in Cisco ISE, perform these steps:

a. In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration > System > Settings > API Settings > API Service Settings**.

b. Click the toggle button to enable the **Open API (Read/Write)** option.

Note

- By default, both ERS and Open API options are disabled in the **pxGrid Cloud Policy** window.
- When you enable the ERS or Open APIs option, the pxGrid Cloud applications are granted **Read Only** access by default (only HTTP GET operations can be performed). Click the toggle button to enable the **Read/Write** option in the **pxGrid Cloud Policy** window if you want to allow POST, PUT, and DELETE operations as well.

Cisco pxGrid Cloud Clients

To view the pxGrid Cloud applications, choose **Administration > pxGrid Services > Client Management > Clients > pxGrid Cloud Clients**.

The pxGrid Cloud offer provides a collection of registered applications that pxGrid Cloud subscribers can select and use. For example, if a subscriber registers Cisco ISE deployment in Cisco pxGrid Cloud and uses two applications, those two applications are listed in the **pxGrid Cloud Clients** tab. Note that you can only view the pxGrid Cloud applications in this tab. You cannot make any changes from this tab.

You can view the total number of pxGrid Cloud applications that are currently running on this deployment in the **Total Clients** pane in the **Summary** window (**Administration > pxGrid Services > Summary**).

High Availability for pxGrid Nodes

You can configure two or more pxGrid nodes to enable high availability. When the Cisco ISE deployment is successfully connected to Cisco pxGrid Cloud, one of the nodes is selected as the Active node and the pxGrid Cloud agent is started on that node. If the Active node is down, or if the network connectivity to the Active node is lost, the Standby node is moved to the Active state. The pxGrid Cloud agent is started on that node and the connectivity to Cisco pxGrid Cloud is established again.



Note The failover process might take around 30 seconds.

Table 11: Events that Trigger High-Availability Response

Event	High-Availability Response
pxGrid Cloud service disabled on Active node	Standby node immediately becomes the Active node
Active node restarted because of a crash or user-initiated sequence	Standby node becomes Active. When the restarted node comes up, this node becomes the Standby node and monitors the Active node.
Deployment upgrade (or standalone Cisco ISE node upgrade) with one pxGrid node	After the upgrade, the node functions as the Active node.
Upgrade deployment with Active and Standby nodes	After the upgrade, the Standby node remains as Standby and continues to monitor the Active node. When the Active node is upgraded, the Standby node takes over as the Active node. When the upgraded node comes up, it becomes Standby and monitors the Active node.
Network issue occurs between the Active and Standby nodes	Both the nodes operate in Active mode. When this occurs, the names of both the nodes are displayed in the pxGrid Cloud Connection window. After the connectivity between the nodes is restored, one of the nodes is selected as the Active node and the other node acts as the Standby node.
Add a new pxGrid node to the deployment	The new node initially acts as the Active node. After the node is fully synchronized and is able to communicate with its peer, one of the nodes is selected as the Active node.

The following configuration changes restart the pxGrid Cloud agent:

- Replacing the pxGrid system certificate
- Replacing the Admin system certificate
- Enabling or disabling the **Trust for authentication within ISE** or **Trust for authentication of Cisco Services** option for any trust certificate
- Changing Cisco ISE proxy settings
- Enabling or disabling the ERS service for Cisco pxGrid Cloud
- Enabling or disabling any pxGrid service in the **pxGrid Cloud Policy** window

Log Files Specific to pxGrid Cloud Service

You can check the following log files in the active pxGrid node if there is any issue related to pxGrid Cloud service:


Log File	Contents	Available Location
pxcloudlog	<ul style="list-style-type: none"> • pxGrid Cloud service configuration changes • pxGrid Cloud service connection status • High-availability status (selection of Active node, detection of failures, and so on) 	Cisco ISE nodes where the pxGrid Cloud service is enabled.

hermes.log	<p>All activities logged by the pxGrid Cloud agent including:</p> <ul style="list-style-type: none"> • Cisco ISE and Cisco pxGrid Cloud connection status • pxGrid topic subscription status • Handling pxGrid, ERS, and Open API requests from Cisco pxGrid Cloud • Configuration changes made in Cisco ISE 	<p>Active pxGrid node</p> <p>Note If the Standby node was previously Active, hermes.log is retained in that node, but the log file is not updated after it moves to the Standby state.</p>
------------	--	---

These log files are included in the Cisco ISE support bundle when the **Include Debug Logs** option is enabled. To download these logs, choose **Operations > Troubleshoot > Download Logs > Debug Logs > Application Logs**.

Configure Debug Log Level for pxGrid Cloud Service

Perform this procedure to configure the level of detail included in the pxcloud.log and hermes.log files.

-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **Operations > Troubleshoot > Debug Wizard > Debug Log Configuration**.
- Step 2** Click the pxGrid node.
- Step 3** Click **pxGrid Cloud**.
- Step 4** Choose one of the following options from the **Log Level** drop-down list:
- **Trace**
 - **Debug**
 - **Info**
 - **Warn**
 - **Error**
 - **Fatal**

The selected log level applies to both the pxcloud.log and the hermes.log.

Note The hermes.log supports only the **Debug**, **Info**, **Warn**, and **Error** log levels. Hence, if you choose **Trace**, the log level is set as **Debug** for the hermes.log. If you choose **Fatal**, the log level is set as **Error** for the hermes.log.

View Nodes in a Deployment

In the **Deployment Nodes** window, you can view all the Cisco ISE nodes, primary and secondary, that are a part of your deployment.

Step 1 Log in to the primary Cisco ISE Admin portal.

Step 2 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**.

Step 3 Click **Deployment** in the navigation pane on the left.

All the Cisco ISE nodes that are a part of your deployment are listed.

Download Endpoint Statistical Data from MnT Nodes

You can download statistical data about the endpoints that connect to your network from the MnT nodes. Key Performance Metrics (KPM), which include the load, CPU usage, and authentication traffic data are available. You can use this data to monitor and troubleshoot issues in your network. From the Cisco ISE CLI, run the **application configure ise** command and choose Option 12 or Option 13 to download the daily KPM statistics or the KPM statistics for the last eight weeks.

The output of this command provides the following data about endpoints:

- Total endpoints in your network
- Number of endpoints that established a successful connection
- Number of endpoints that failed authentication
- Total number of new endpoints that have connected each day
- Total number of endpoints onboarded each day

The output also includes time stamp details, the total number of endpoints that connected through each of the Policy Service nodes (PSNs) in the deployment, total number of endpoints, active endpoints, load, and authentication traffic details.

See the [Cisco Identity Services Engine CLI Reference Guide](#) for more information on this command.

Database Crash or File Corruption Issues

Cisco ISE may crash if the Oracle database files are corrupted because of a power outage or other reasons, resulting in data loss. Based on the incident, follow these instructions to recover from data loss:

- In case of PAN corruption in the deployment, you should [promote the Secondary PAN to Primary PAN](#). If the secondary PAN's promotion is not possible because the deployment is small or any other reason, [restore](#) the most recent available backup as described in [Cisco Identity Services Engine CLI Reference Guide](#).
- In case of PSN corruption, follow the steps to deregister, reset config, and register, as described in the [Cisco Identity Services Engine CLI Reference Guide](#).
- In case of a standalone device, restore the most recent backup that is available, as described in the [Cisco Identity Services Engine CLI Reference Guide](#).




Note Obtain the backup from the standalone device regularly to avoid loss in the latest configuration changes.

Device Configuration for Monitoring

The MnT node receives and uses data from the devices on a network to populate the dashboard display. To enable communication between the MnT node and the network devices, the switches and NADs must be configured properly.

Synchronize Primary and Secondary Cisco ISE Nodes

You can make configuration changes to Cisco ISE only through the primary PAN. The configuration changes get replicated to all the secondary nodes. If, for some reason, this replication does not occur properly, you can manually synchronize the secondary PAN with the primary PAN.


-
- Step 1** Log in to the primary PAN.
- Step 2** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**.
- Step 3** Check the check box next to the node that you want to synchronize with the primary PAN, and click **Syncup** to force a full database replication.
-

Change Node Personas and Services



Note When you enable or disable any of the services that run on a PSN or make any changes to a PSN, you will be restarting the application server processes on which these services run. Expect a delay while these services restart. Because this delay in restarting services, automatic failover, if enabled in your deployment, might get initiated. To avoid this, make sure that the automatic failover configuration is turned off.

You can edit the Cisco ISE node configuration to change the personas and services that run on the node.

-
- Step 1** Log in to the primary PAN.
- Step 2** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**.
- Step 3** Check the check box next to the node whose personas or services you want to change, and then click **Edit**.
- Step 4** Choose the personas and services that you want to modify.
- Step 5** Click **Save**.
- Step 6** Verify the receipt of an alarm on your primary PAN to confirm the persona or service change. If the persona or service change is not saved successfully, an alarm is not generated.
-

Effects of Modifying Nodes in Cisco ISE

When you make any of the following changes to a node in a Cisco ISE, that node restarts, which causes a delay:

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



Note When you promote the secondary Administration node to the primary PAN position, the primary node will assume a secondary role. This causes both the primary and secondary nodes to restart, causing a delay.

Create a Policy Service Node Group

When two or more Policy Service nodes (PSNs) are connected to the same high-speed Local Area Network (LAN), we recommend that you place them in the same node group. This design optimizes the replication of endpoint profiling data by retaining less significant attributes that are local to the group and reducing the information that is replicated to the remote nodes in the network. Node group members also check on the availability of peer group members. If the group detects that a member has failed, it attempts to reset and recover all URL-redirectioned sessions on the failed node.

The node groups are used for the PSN failover in the sessions on which URL redirect (posture services, guest services, and MDM) is imposed.



Note We recommend that you put all the PSNs in the same local network and as a part of the same node group. PSNs need not be a part of a load-balanced cluster to join the same node group. However, each local PSN in a load-balanced cluster should typically be part of the same node group.

Node group members can communicate over TCP/7800.

Before you add PSNs as members to a node group, you must create the node group. You can create, edit, and delete PSN groups from the **Deployment** window of the Admin portal.

Step 1 In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**.

Step 2 Click the **Settings** icon at the top of the left navigation pane.

Step 3 Click **Create Node Group**.

Step 4 Enter a unique name for your node group.

Note We recommend that you do not configure a node group with the name **None**, because it may cause issues during node registration.

Step 5 (Optional) Enter a description for your node group.

Step 6 (Optional) Check the **Enable MAR Cache Distribution** check box and fill in the other options. Ensure that the MAR is enabled in the **Active Directory** window before checking this check box.

Step 7 Click **Submit** to save the node group.

After you save the node group, it should appear in the left navigation pane. If you do not see the node group in the left pane, it may be hidden. Click the **Expand** button on the navigation pane to view the hidden objects.

Add a node to a node group, or edit a node by choosing the corresponding node group from the **Include node in node group** drop-down list in the **Policy Service** area.

Remove a Node from Deployment

To remove a node from a deployment, you must deregister it. The deregistered node becomes a standalone Cisco ISE node.

It retains the last configuration that it received from the primary PAN and assumes the default personas of a standalone node, that is, Administration, Policy Service, or Monitoring. If you deregister an MnT node, this node will no longer be a syslog target.


When a Primary PSN is deregistered, the endpoint data is lost. If you want the PSN to retain the endpoint data after it becomes a standalone node, do one of the following:

- Obtain a backup from the primary PAN, and when the PSN becomes a standalone node, restore this data backup on it.
- Change the persona of the PSN to Administration (secondary PAN), synchronize the data in the **Deployment** window of the Admin portal, and then deregister the node. This node will now have all the data. You can then add a secondary PAN to the existing deployment.

You can view these changes in the **Deployment** window of the primary PAN. However, expect a delay of five minutes for the changes to take effect and appear in the **Deployment** window.

Before you begin

Before you remove a secondary node from a deployment, perform a backup of Cisco ISE configuration, which you can then restore later, if needed.

-
- Step 1** In the Cisco ISE GUI, click the **Menu** icon () and choose **Administration** > **System** > **Deployment**.
 - Step 2** Check the check box next to the secondary node that you want to remove, and click **Deregister**.
 - Step 3** Click **OK**.
 - Step 4** Verify the receipt of an alarm on your primary PAN to confirm that the secondary node is deregistered successfully. If the secondary node fails to deregister from the primary PAN, it means the alarm is not generated.
-

Shut Down a Cisco ISE Node

Before you run the **halt** command from the Cisco ISE CLI, we recommend that you stop the Cisco ISE application service and ensure that it is not performing a backup, restore, installation, upgrade, or remove operation. If you run the **halt** command while the Cisco ISE is performing any of these operations, you will get one of the following warning messages:

```
WARNING: A backup or restore is currently in progress! Continue with halt?
```

```
WARNING: An install/upgrade/remove is currently in progress! Continue with halt?
```

If no processes are running when you use the **halt** command, or if you enter **Yes** in response to the warning message displayed, then you must respond to the following question:

```
Do you want to save the current configuration?
```

If you enter **Yes** to save the existing Cisco ISE configuration, the following message is displayed:

```
Saved the running configuration to startup successfully.
```



Note We recommend that you stop the application process before rebooting the appliance.

We also recommend that you stop the application process before rebooting Cisco ISE. For more information, see the [Cisco Identity Services Engine CLI Reference Guide](#).

Scenarios In Which Need to Reregister a Node

The following table summarizes some of the scenarios where you need to reregister a node when it is corrupted:

Scenarios	What needs to be done
If any of the nodes other than the Primary PAN is corrupted	<ol style="list-style-type: none"> 1. Deregister the failed node from the deployment. 2. Reinstall Cisco ISE on the failed node. 3. Reregister the node in the existing deployment. <p>Note You must import the old certificates to the node before or after the registration.</p>
If the Primary PAN is corrupted	<p>If, for example, there are two nodes, N1 (Primary PAN) and N2 (Secondary PAN):</p> <ol style="list-style-type: none"> 1. Promote secondary PAN (N2) to Primary PAN. 2. Remove the failed node (N1) from the deployment. 3. Reinstall Cisco ISE on the failed node (N1). 4. Register the node (N1) as Secondary PAN to deployment. 5. Import the old certificates to the node (N1) after the registration is completed. 6. Promote the node (N1) back to Primary PAN to have similar deployment as earlier.
If both Primary PAN and Secondary PAN are corrupted	<p>If, for example, there are two nodes, N1 (Primary PAN) and N2 (Secondary PAN):</p> <ol style="list-style-type: none"> 1. Reinstall Cisco ISE on Primary PAN node (N1) and Secondary PAN node (N2). 2. Restore configuration backup in Primary PAN node (N1). 3. Import old certificates in Primary PAN node (N1). 4. Register the other node (N2) as Secondary PAN in the deployment. 5. Perform reset-config on other nodes and register the nodes in the deployment. 6. Import certificates to all the nodes. <p>Note If the Primary PAN and Secondary PANs are VMs, reinstalling Cisco ISE might change the UDI. Hence, you must reinstall the licenses with the new UDIs.</p>

Change the Hostname or IP Address of a Standalone Cisco ISE Node

You can change the hostname, IP address, or domain name of standalone Cisco ISE nodes. However, you cannot use **localhost** as the hostname for a node.

Before you begin

If a Cisco ISE node is a part of a distributed deployment, you must first remove it from the deployment and ensure that it is a standalone node.

-
- Step 1** Change the hostname or IP address of the Cisco ISE node using the **hostname**, **ip address**, or **ip domain-name** command from the Cisco ISE CLI.
- Step 2** Reset the Cisco ISE application configuration using the **application stop ise** command from the Cisco ISE CLI to restart all the services.
- Step 3** Register the Cisco ISE node to the primary PAN if it is a part of a distributed deployment.

Note If you are using the hostname while registering the Cisco ISE node, the fully qualified domain name (FQDN) of the standalone node that you are going to register, for example, *abc.xyz.com*, must be DNS-resolvable from the primary PAN. Otherwise, node registration fails. You must enter the IP addresses and FQDNs of the Cisco ISE nodes that are a part of your distributed deployment in the DNS server.

After you register the Cisco ISE node as a secondary node, the primary PAN replicates the change in the IP address, hostname, or domain name to the other Cisco ISE nodes in your deployment.
