Cisco ISE 2.7 Admin Guide: Deployment

Deployment

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 or personas of a node determine the services provided by a 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 an 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 on your network:

- Primary and secondary Administration nodes for high availability
- A single or a pair of non-administration nodes for health check of Administration nodes for automatic failover
- A pair of health check nodes or a single health check node for Primary Administration Node (PAN) automatic failover
- One or more Policy Service Nodes (PSN) for 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 will be 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.

It is recommended not to change the host name and the domain name on Cisco ISE that have been configured or in production. If it is required, then reimage 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

Procedure

Step 1 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 PAN

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

Procedure

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

The Register button will be disabled initially. To enable this button, you must 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 ISE nodes to Primary PAN to form a multi-node 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, Primary PAN will push the configuration data to the secondary node and the application server on the secondary node will restart. After this is complete, further configuration changes done on 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 the 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 being registered uses an untrusted self-signed certificate, you will be prompted with a certificate warning with the details of the certificate. If you accept the certificate, it will be added to the trusted certificate store of Primary PAN to enable TLS communication with the node.

If the node uses a certificate that is not self-signed (for example, signed by external CA), you must manually import the relevant certificate chain of that node to the trusted certificate store of Primary PAN. When you import the secondary node's certificate in to the trusted certificate store, check the Trust for Authentication within ISE check box 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 section Create a Policy Service Node Group, on page 49 for more details.

Procedure

Step 1 Log in to the Primary PAN.

Step 2 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 UI-based administrator credentials for the secondary node in the Username and Password fields.

Step 6 Click **Next**.

Primary PAN tries to establish TLS communication (for the first time) with the node being 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. The certificate warning message displays 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. You can select 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 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 Primary PAN. When you import the secondary node's certificate in to the trusted certificate store, check the **Trust for Authentication within ISE** check box for the PAN to validate the secondary node's certificate.
- If the node uses a CA signed certificate, an error message is displayed that the registration cannot proceed until certificate trust is set up.

Step 7 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) will be generated on the Primary PAN. You can view this alarm on the Alarms page. After the registered node is synchronized and restarted, you can log in to the secondary node GUI using the same credentials used on 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 will be 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.

  See "Backup and Restore of Cisco ISE CA Certificates and Keys" in Chapter "Basic Setup" of the *Cisco ISE Administrator Guide*.

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 “redundant”) 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 they are 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 one performing a different 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 Cisco ISE distributed deployment, administration and monitoring activities are centralized, and processing is distributed across the Policy Service nodes. Depending on your performance needs, you can scale your deployment. Each Cisco ISE node in a deployment can assume any of the following 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 then 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 Monitoring node in a distributed deployment. At the time of configuring your Primary PAN, you must enable the Monitoring persona. After you register a Monitoring node in your deployment, you can edit the Primary PAN and disable the Monitoring persona, if required.

Data Replication from Primary to Secondary ISE Nodes

When you register an 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 present in all Cisco ISE nodes that are part of your deployment.

A full replication typically occurs when you first register an 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 Cisco ISE nodes in a deployment are in sync. You can view the status of replication in the Node Status column from the deployment pages 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. Once it is complete, the node status turns green indicating that the secondary node is synchronized with the PAN.

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 will be 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 page of the Admin portal, and then deregister the node. This node will now have all the data. You can then add a secondary Admin node to the existing deployment.

Note You cannot deregister a Primary PAN.

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, ISE node. For Administration, Policy Service, and Monitoring capabilities, you must choose an ISE node.
- 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 during the setup of 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 Admin password when you install Cisco ISE. The previous Cisco ISE Admin 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 Domain Name System (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 will fail.
- Configure the forward and reverse DNS lookup for all 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 Monitoring node, and restore the data to the new secondary Monitoring node. This ensures that the history of the primary Monitoring node is in sync with the new secondary node as 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.
- 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 will then be replicated to the new node.

Menu Options Available on Primary and Secondary Nodes

The menu options available in Cisco ISE nodes that are 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 are enabled on them.

If a node assumes more than one persona, for example, the Policy Service persona, and a Monitoring persona with an Active role, then the menu options listed for Policy Service nodes and Active Monitoring node will be available on that node.

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

Table 1: Cisco ISE Nodes and Available Menu Options

Cisco ISE Node	Available Menu Options
All Nodes	• View and configure system time and NTP server settings. • Install server certificate, manage certificate signing request. You can perform server certificate operations, for all the nodes in the deployment, via the Primary PAN that centrally manages all server certificates. **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.
Primary PAN	All menus and sub-menus.
Active Monitoring Node	• Provides access to monitoring data (on both Primary and Active Monitoring nodes). **Note** The Operations menu can be viewed only from the Primary PAN. The Operations menu does not appear in the Monitoring nodes in Cisco ISE 2.1 and above.
Policy Service Nodes	Option to join, leave, and test Active Directory connection. Each Policy Service node 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 Policy Service nodes to the Active Directory domain individually.
Secondary PAN	Option to promote the Secondary PAN to become the Primary PAN. **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.

Deployment and Node Settings

Deployment and Node Settings

The Deployment Nodes page enables you to configure Cisco ISE (Administration, Policy Service, and Monitoring) nodes and to set up a deployment.

Deployment Nodes List Window

The following table describes the fields on the **Deployment Nodes Lis**t page, which you can use to configure Cisco ISE nodes in a deployment. The navigation path for this page is **Administration** > **System** > **Deployment**.

Field Name	Usage Guidelines
Hostname	Displays the hostname of the node.
Node Type	Displays the node type. It can be one of the following: • Cisco ISE (Administration, Policy Service, and Monitoring) nodes
Personas	(Only appears if the node type is Cisco ISE) Lists the personas that an Cisco ISE node has assumed. For example, Administration, Policy Service.
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: • PRI(A): Refers to the Primary PAN • SEC(A): Refers to the Secondary PAN • PRI(M): Refers to the Primary Monitoring Node • SEC(M): Refers to the Secondary Monitoring Node
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: • Session • Profiling • All

Field Name	Usage Guidelines
Node Status	Indicates the status of each ISE node in a deployment for data replication. • Green (Connected): Indicates that an ISE node, which is already registered in the deployment is in sync with the Primary PAN. • Red (Disconnected): Indicates that an ISE node is not reachable or is down or data replication is not happening. • Orange (In Progress): Indicates that an ISE node is newly registered with the Primary PAN or you have performed a manual sync operation or the ISE node is not in sync (out of sync) with the Primary PAN. For more details, click the quick view icon for each ISE node in the Node Status column.

Related Topics

Cisco ISE Distributed Deployment, on page 4
Cisco ISE Deployment Terminology, on page 1
Configure a Cisco ISE Node, on page 1
Register a Secondary Cisco ISE Node

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. The navigation path for this window is: **Administration** > **System** > **Deployment** > **Deployment Node** > **Edit** > **General Settings**.

Table 2: 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	

Field Name	Usage Guidelines
Administration	Check this check box 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. Role: Displays the role that the Administration persona has assumed in the deployment. Could take on any one of the following values: Standalone, Primary, Secondary Make Primary: Click this button to make this node your primary Cisco ISE node. You can have only one primary Cisco ISE node in a deployment. The other options on this page 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, a Make Primary button appears next to it.If the node has a Secondary role, a Promote to Primary button appears next to it.If the node has a Primary role and there are no other nodes registered with it, a Make Standalone button appears next to it. You can click this button to make your primary node a standalone node.

Field Name	Usage Guidelines
Monitoring	Check this check box 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. 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 2.5 MB per Cisco ISE node in your network, per day. 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 for you to configure the Primary-Secondary roles. To configure these roles, choose one of the following: • 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. 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 following window: **Administration** > **System** > **Logging** > **Remote Logging Targets**.

Field Name	Usage Guidelines
Policy Service	

Field Name	Usage Guidelines
	Check this check box to enable any one or all of the following services: • **Enable Session Services**: Check this check box to enable network access, posture, guest, and client provisioning services. Choose the group to which this Policy Service node belongs from the Include Node in Node Group drop-down list. Note that CA and EST services can only run on a Policy Service node that has session services enabled on it. For **Include Node in Node Group**, choose **None** if you do not want this Policy Service node to be part of any group. 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. While a single NAD can be configured with many ISE nodes as RADIUS servers and dynamic-authorization 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 the Create Policy Service Node Group section in *Cisco ISE Admin Guide: Deployment* for more details. • Enable Profiling Service: Check this check box to enable the Profiler 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. You must expect a delay while these services restart. You can

Field Name	Usage Guidelines
	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. 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. • 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 devices. • Enable Passive Identity Service: Check this check box to enable the Identity Mapping feature. This feature enables you to monitor users that are authenticated by a Domain Controller (DC) 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 (AD) Domain Controller.
pxGrid	Check this check box to enable pxGrid persona. Cisco pxGrid is used to share the context-sensitive information from 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 like sharing tags and policy objects between ISE and third party vendors, and for non-ISE related information exchanges such as threat information.

Related Topics

Personas in Distributed Cisco ISE Deployments, on page 1

Administration Node, on page 23
Policy Service Node, on page 31
Monitoring Node, on page 33
pxGrid Node, on page 40
Synchronize Primary and Secondary Cisco ISE Nodes, on page 47
Create a Policy Service Node Group, on page 49
Deploy pxGrid Node, on page 42
Change Node Personas and Services, on page 48
Configure Monitoring Nodes for Automatic Failover, on page 39

Profiling Node Settings

The following table describes the fields on the Profiling Configuration page, which you can use to configure the probes for the profiler service. The navigation path for this page is: **Administration** > **System** > **Deployment** > **ISE Node** > **Edit** > **Profiling Configuration**.

Table 3: Profiling Node Settings

Field Name	Usage Guidelines
NetFlow	Check this check box if you want to enable NetFlow per Cisco ISE node that has assumed the Policy Service persona to receive Netflow packets sent from the routers.Choose these options: • **Interface**: Choose the interface on the ISE node. • **Port**: Enter the NetFlow listener port number on which NetFlow exports are received from the routers. The default port is 9996.
DHCP	Check this check box if you want to enable DHCP per Cisco ISE node that has assumed the Policy Service persona to listen for DHCP packets from IP helper.Choose these options: • **Port**: Enter the DHCP server UDP port number. The default port is 67. • **Interface**: Choose the interface on the ISE node. • **Port**: Enter the DHCP server UDP port number. The default port is 67.
DHCP SPAN	Check this check box if you want to enable DHCP SPAN per Cisco ISE node that has assumed the Policy Service persona to collect DHCP packets. • **Interface**: Choose the interface on the ISE node.

Field Name	Usage Guidelines
HTTP	Check this check box if you want to enable HTTP per Cisco ISE node that has assumed the Policy Service persona to receive and parse HTTP packets. • **Interface**: Choose the interface on the ISE node.
RADIUS	Check this check box if you want to enable RADIUS per ISE node that has assumed the Policy Service persona to collect RADIUS session attributes as well as CDP, LLDP attributes from the IOS Sensor enabled devices.
Network Scan (NMAP)	Check this box to enable the NMAP probe.
DNS	Check this check box if you want to enable DNS per ISE node that has assumed the Policy Service persona to perform a DNS lookup for the FQDN.Enter the timeout period in seconds. **Note** For the DNS probe to work on a particular Cisco ISE node in a distributed deployment, you must enable any one of the following probes: DHCP, DHCP SPAN, HTTP, RADIUS, or SNMP. For DNS lookup, one of the probes mentioned above must be started along with the DNS probe.
SNMP Query	Check this check box if you want to enable SNMP Query per ISE node that has assumed the Policy Service persona to poll network devices at specified intervals.Enter values for the following fields: Retries, Timeout, Event Timeout, and an optional Description. **Note** In addition to configuring the SNMP Query probe, you must also configure other SNMP settings in the following location: Administration > Network Resources > Network Devices. When you configure SNMP settings on the network devices, ensure that you enable the Cisco Device Protocol (CDP) and Link Layer Discovery Protocol (LLDP) globally on your network devices.

Field Name	Usage Guidelines
SNMP Trap	Check this check box if you want to enable SNMP Trap probe per ISE node that has assumed the Policy Service Persona to receive linkUp, linkDown, and MAC notification traps from the network devices.Choose any of the following: • **Link Trap Query**: Check this check box to receive and interpret linkup and linkdown notifications received through the SNMP Trap. • **MAC Trap Query**: Check this check box to receive and interpret MAC notifications received through the SNMP Trap. • **Interface**: Choose an interface on the ISE node. • **Port**: Enter the UDP port of the host to use. The default port is 162.
Active Directory	Scans the defined Active Directory servers for information about Windows users.
pxGrid	Allows ISE to collect (profile) endpoint attributes over pxGrid.

Related Topics

Cisco ISE Profiling Service

Network Probes Used by Profiling Service

Configure Profiling Service in Cisco ISE Nodes

Logging Settings

These pages allow you 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 on the Remote Logging Targets page, which you can use to create external locations (syslog servers) to store logging messages. The navigation path for this page is: **Administration** > **System** > **Logging** > **Remote Logging Targets**.

Table 4: Remote Logging Target Settings

Fields	Usage Guidelines
Name	Enter the name of the new target.
Target Type	Select the target type. By default it is set to UDP Syslog.
Description	Enter a brief description of the new target.

Fields	Usage Guidelines
IP Address	Enter the IP address or hostname of the destination machine where you want to store the logs. ISE supports IPv4 and IPv6 formats for logging.
Port	Enter the port number of the destination machine.
Facility Code	Choose the syslog facility code to be used for logging. Valid options are Local0 through Local7.
Maximum Length	Enter the maximum length of the remote log target messages. Valid options are from 200 to 1024 bytes.
Buffer Message When Server Down	Check this check-box if you want Cisco ISE to buffer the syslog messages when TCP syslog targets and secure syslog targets are unavailable. ISE retries sending the messages to the target when the connection resumes. After the connection resumes, messages are sent by the order from oldest to newest and 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 existing buffered messages for the specific target are lost.
Reconnect Timeout (Sec)	Give in seconds how long will the TCP and secure syslogs be kept before being discarded, when the server is down.
Select CA Certificate	Select a client certificate.
Ignore Server Certificate Validation	Check this check-box if you want ISE to ignore server certificate authentication and accept any syslog server.

Related Topics

Cisco ISE Logging Mechanism
Cisco ISE System Logs
Remote Syslog Message Format
Cisco ISE Message Catalogs
Collection Filters
Event Suppression Bypass Filter
Configure Remote Syslog Collection Locations
Configure Collection Filters

Logging Category Settings

The following table describes the fields on the Logging Categories page, which you can use to configure the log severity level and choose logging targets for the logs of selected categories to be stored. The navigation path for this page is **Administration** > **System** > **Logging** > **Logging Categories**.

Table 5: Logging Category Settings

Fields	Usage Guidelines
Name	Displays the name of the logging category.
Log Severity Level	Allows you to choose the severity level for the diagnostic logging categories from the following options: • **FATAL**—Emergency. This option means that Cisco ISE cannot be used and you must take action immediately • **ERROR**—This option indicates a critical or error condition. • **WARN**—This option indicates a normal but significant condition. This is the default condition. • **INFO**—This option indicates an informational message. • **DEBUG**—This option indicates a diagnostic bug message.
Local Logging	Check this check box to enable logging event for the category on the local node.
Target	Allows you to change the targets for a category by transferring the targets between the Available and the Selected boxes using the left and right icons. The Available box contains the existing logging targets, both local (predefined) and external (user-defined). The Selected box, which is initially empty, contains the selected targets for the specific category.

Related Topics

Remote Syslog Message Format

Cisco ISE Message Codes

Configure Remote Syslog Collection Locations

Set Severity Levels for Message Codes

Admin Access Settings

These pages enable you to configure access settings for administrators.

Administrator Password Policy Settings

The following table describes the fields on the Administrator Password Policy page, which you can use to define a criteria that administrator passwords should meet. The navigation path for this page is:**Administration** > **System** > **Admin Access** > **Authentication** > **Password Policy**.

Table 6: Administrator Password Policy Settings

Fields	Usage Guidelines
Minimum Length	Specifies the minimum length of the password (in characters). The default is six characters.

Fields	Usage Guidelines
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.
	"cisco" or its characters in reverse order—Check this check box to restrict the use of the word "cisco" or its characters in reverse order.
	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 reverse order.
	Repeated characters four or more times consecutively—Check this check box to restrict the use of repeated characters four or more times consecutively.
	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. Substitution of "$" for "s", "@" for "a", "0" for "o", "1" for "l", "!" for "i", "3" for "e" is not permitted. For example, Pa$$w0rd • Default Dictionary—Choose this option to use the default Linux dictionary in Cisco ISE. The default dictionary contains approximately 480,000 English words. By default, this option is selected. • Custom Dictionary—Choose this option to use your customized dictionary. Click **Choose File** to select the custom dictionary file. The text file must be of newline-delimited words, .dic extension, and size less than 20 MB.
Required Characters	Specifies that the administrator password must contain at least one character of the type that you choose from the following choices: • Lowercase alphabetic characters • Uppercase alphabetic characters • Numeric characters • Non-alphanumeric characters

Fields	Usage Guidelines
Password History	Specifies the number of previous passwords from which the new password must be different to prevent the repeated use of the same password. Also, specifies the number of characters that must be different from the previous password. Enter the number of days before which you cannot reuse a password.
Password Lifetime	Specifies the following options to force users to change passwords after a specified time period: • Time (in days) before the administrator account is disabled if the password is not changed. (The allowable range is 0 to 2,147,483,647 days.) • Reminder (in days) before the administrator account is disabled.
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	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 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 on the Session page, which you can use to define session timeout and terminate an active administrative session. The navigation path for this page is:**Administration** > **System** > **Admin Access** > **Settings** > **Session**.

Table 7: Session Timeout and Session Info Settings

Fields	Usage Guidelines
Session Timeout	
Session Idle Timeout	Enter the time in minutes that you want Cisco ISE to wait 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.

Fields	Usage Guidelines
Session Info	
Invalidate	Check the check box next 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 system-related configurations that are related to functionality 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 any one of the following roles: Standalone, Primary, or Secondary.

High Availability for the Administrative Node

In a high availability configuration, the Primary Administration Node (PAN) is in the active state. The Secondary PAN (backup PAN) is in the standby state, which means it receives all configuration updates from the Primary PAN, but is not active in the 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 non-administration secondary node, which is called a health check node. The health check node checks the health of Primary PAN. If the health detects that the Primary PAN is down or unreachable, the health check node initiates the promotion of the Secondary PAN to take over the primary role.

To deploy the auto-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 non-administration node and can be a Policy Service, Monitoring, 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.

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

Features	Available When Primary PAN is Down (Yes/No)
Existing internal user RADIUS authentication	Yes
Existing or New AD user RADIUS authentication	Yes
Existing endpoint with no profile change	Yes
Existing endpoint with profile change	No
New endpoint learned through profiling.	No
Existing guest – LWA	Yes

Features	Available When Primary PAN is Down (Yes/No)
Existing guest – 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

To support certificate provisioning with the internal certificate authority, 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 auto-failover for 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 Primary PAN is called the active health check node. The health check node for Secondary PAN is called the passive health check node. The active health check node is responsible for checking status of Primary PAN, and managing the automatic failover of Administration nodes. We recommended using two non-administration ISE nodes as health check nodes, one for the Primary 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 PANs are in the same data center, you can use a single non-administration 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 non-administration node, which means it can be a Policy Service, Monitoring, or pxGrid node, or a combination of those. We recommend that you designate PSN nodes in the same data center as the Administration nodes as health check 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/MnT) not having the Administration persona can be used as health check node.

If you chose not to enable automatic failover, and rely on manually promoting the secondary node when the primary PAN fails, then 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 would monitor 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 health of Administration nodes again. Designating or removing the high-availability health check role on 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 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, then the health check node considers the Primary PAN to be in good health. Otherwise, the health check node considers the Primary PAN to be in bad health. If the health of the Primary PAN is bad continuously for more than the configured failover period, then the health check node initiates failover to the Secondary PAN.

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

Response from 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 would raise an alarm. However, a promotion request will be made to the Secondary PAN.

Changing Health Nodes

You can change the 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 some other node (H2) is made the health check node of the Primary PAN. In such a case, once the Primary PAN goes down, there is no way for N1 to know that there is another node (H2) checking the same Primary PAN. Later, if H2 goes down or out of network, an actual failover is required. The Secondary PAN, however, retains the right to reject the promotion request. So, once the Secondary PAN has been 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 Primary PAN.

Automatic Failover to the Secondary PAN

You can configure ISE to automatically the promote the secondary PAN when the primary PAN becomes unavailable. The configuration is done on the primary administrative node (Primary PAN) on the **Administration** > **System** > **Deployment** page. 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**. With the default configuration, that time is 10 minutes. Promotion of the secondary PAN to primary 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 isn’t available in network.
- The failover request came from a valid health check node.
- The failover request is for this 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 would be attempted.

- Health of 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 primary PAN is down (network port shut or network service down), or it’s 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 primary PAN and will check the health status afresh.

Bring Your Own Device In Case Of Automatic Failover To Secondary PAN

When the primary PAN is down, authentication isn’t 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 won’t be onboarded. BYOD onboarding requires an active primary PAN.

Once the original primary PAN is brought back up or the secondary PAN is promoted, new BYOD endpoints will be onboarded without any issues.

If the primary PAN that failed can’t be re-joined 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, endpoints 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 would be avoided or promotion request to the secondary node would be rejected.

- Node receiving the promotion request is not the secondary node.

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

Functionalities Affected by the PAN Auto-Failover Feature

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

Functionality	Affect Details
Operations that are Blocked	
Upgrade	Upgrade via the CLI is blocked. The PAN auto-failover feature will be available for configuration after you upgrade from a previous version of Cisco ISE to release 1.4. By default, this feature is disabled. To deploy the auto-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 non-administration node and can be a Policy Service, Monitoring, 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.
Restore of Backup	Restore via the CLI and user interface will be blocked. If PAN auto-failover configuration was enabled prior to restore, you must reconfigure it after a successful restore.
Change Node Persona	Change of the following node personas via the user interface will be blocked: • Admin persona in both the Administration nodes. • Persona of the PAN. • Deregistration of health check node after enabling the PAN auto-failover feature.

Functionality	Affect Details
Other CLI Operations	The following admin operations via the CLI will be blocked: • Patch Installation and Roll back • DNS Server change • IP address change of eth1, eth2, and eth3 interfaces • Host alias change of eth1, eth2, and eth3 interfaces • Timezone change
Other Administration Portal Operations	The following admin operations via the user interface will be blocked: • Patch Installation and Roll back • Change HTTPS certificate. • Change admin authentication type from password-based authentication to certificate-based authentication and viceversa.
Users with maximum connected devices cannot connect.	Some session data is stored on the failed PAN, and can't be updated by the PSN.
Operations that Require PAN Auto-Failover to be Disabled	
CLI Operations	The following admin operations via the CLI will display a warning message if PAN auto-failover configuration is enabled. These operations may trigger auto-failover if service/system is not restarted within failover window. Hence, while performing the below operations it is recommended to disable PAN auto-failover configuration: • Manual ISE service stop • Soft reload (reboot) using admin CLI

Configure Primary PAN for Automatic Failover

Before you begin

To deploy the auto-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 non-administration node and can be a Policy Service, Monitoring, 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.

Procedure

Step 1 Log in to the user interface of the Primary PAN.

Step 2 Choose **Administration** > **System** > **Deployment** > **PAN Failover**.

Step 3 Check the **Enable PAN Auto Failover** check box, to enable automatic failover of the Primary PAN.

You can only promote a Secondary PAN to become the Primary PAN. Cisco ISE nodes that assume only the Policy Service, Monitoring, or pxGrid persona, or a combination of these, cannot be promoted to become the Primary PAN.

Step 4 Select the health check node for Primary PAN from the **Primary Health Check Node** drop down list containing all the available secondary nodes.

It is recommended to have this node in the same location or data center as the Primary PAN.

Step 5 Select the health check node for Secondary PAN, from the **Secondary Health Check Node** drop down list containing all the available secondary nodes.

It is recommended to have this node in the same location or data center as the Secondary PAN.

Step 6 Provide the **Polling Interval** time after which the Administration node status will be checked . The valid range is from 30 to 300 seconds.

Step 7 Provide the count for **Number of Failure Polls before Failover**.

The failover will occur if the status of the Administration node is not good for the specified number of failure polls. The valid range is from 2 to 60 counts.

Step 8 Click **Save**.

What to do next

After the promotion of 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 auto-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.

Procedure

Step 1 Log in to the user interface of the Secondary PAN.

Step 2 Choose **Administration** > **System** > **Deployment**.

Step 3 In the Edit Node page, click **Promote to Primary**.

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.

Step 4 Click **Save**.

What to do next

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 page 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.

Reusing a node of an existing ISE Deployment as a Primary PAN for a new ISE Deployment

If you want to repurpose a node of an existing ISE Deployment to the Primary PAN of a new ISE Deployment you must follow these steps:

Procedure

Step 1 First run the ISE utility "Perform System Erase", as described in the ISE Installation guide for your version of ISE https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-guides-list.html

Step 2 Perform a fresh install of ISE, as described in the ISE installation guides.

Step 3 Configure the stadalone node as a Primary Administration Node by referring Configure a Primary PAN, on page 2.

Restoring Service to the Primary PAN

Cisco ISE does not support automatic fallback to 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 auto-failover feature, at least two nodes in your distributed setup should assume the Administration persona and one node should assume the non-Administration persona. If the Primary Administration Node (PAN) goes down, an automatic promotion of the Secondary Administration Node is initiated. For this, a non-administration secondary node is designated as the health check node for each of the administration nodes. The health check node checks the health of PAN at configured intervals. If the health check response received for the PAN health is not good due to being down or not reachable, health check node initiates the promotion of the Secondary Administration

Node to take over the primary role after waiting for the configured threshold value. There are some features that are unavailable after auto-failover of the Secondary Administrative Node. Cisco ISE does not support fallback to the original PAN. Refer to the High Availability for the Administrative Node section for more information.

Policy Service Node

A Policy Service Node (PSN) is a Cisco ISE node with the Policy Service persona 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 would be more than one Policy Service node in a distributed deployment.

All Policy Service nodes 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 any URL-redirected sessions.

High Availability in Policy Service Nodes

To detect node failure and to reset all URL-redirected sessions on the failed node, two or more Policy Service nodes 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-redirected 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 would also be configured as RADIUS servers.

While a single NAD can be configured with many ISE nodes as RADIUS servers and dynamic-authorization 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 49 section for more details.

Load Balancer To Distribute Requests Evenly Among PSNs

When you have multiple Policy Service nodes 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. Refer to the Cisco and F5 Deployment Guide: ISE Load Balancing using BIG-IP for information on and best practices about deploying PSNs behind a load balancer.

Session Failover in Policy Service Nodes

When a Policy Service node that has active URL-redirected sessions fails, the endpoints are stuck in an intermediate state. Even if the redirect endpoint detects that the Policy Service node that it has been communicating with has failed, it cannot re-initiate authorization.

If the Policy Service nodes are part of a node group, the nodes within a node group exchange heartbeat messages to detect node failures. If a node fails, one of its peers from the node group learns about the active URL-redirected sessions on the failed node and issues a CoA to disconnect those sessions.

As a result, the sessions are handled by another Policy Service node that is available in the same node group. The session failover does not automatically move the sessions over from a Policy Service node that has gone down to one that is available, but issues a CoA to achieve that.

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 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 Policy Service node failures.

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

Light Data Distribution

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

Light Data Distribution consists of the following two 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.
- **TTL**—This value specifies the maximum time a session will wait for the batch to complete before updating the Light Data Distribution. The default value is 1000 milliseconds.

In case of connectivity issues between 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, 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 ISE Messaging Services for inter-node communication. ISE messaging service uses different certificate (signed by internal-CA chain). In case you are facing issues with ISE messaging service, you will have to regenerate ISE messaging service certificate. Choose **Administration > system > Certificates > Certificate management >Certificate Signing request**. Select **ISE Messaging service** in the **Certificate(s) will be used for** section. 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. The **RADIUS Session Directory** stores only the session attributes that are required for Change of Authorization(CoA).

This feature is enabled by default from Cisco ISE Release 2.7. To enable or disable this feature, choose **Administration** > **System** > **Settings** > **Light Data Distribution** and check or uncheck the **RADIUS Session Directory** check box.

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. To enable or disable this feature, choose **Administration** > **System** > **Settings** > **Light Data Distribution** and check or uncheck the **Enable Endpoint Owner Directory** check box.

Monitoring Node

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

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

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

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

Manually Modify MnT Role

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

Procedure

Step 1 Log in to the user interface of the Primary PAN.

Step 2 Choose **Administration** > **System** > **Deployment**.

Step 3 Select MnT node you want to change the role from list of nodes.

Step 4 Click **Edit**.

Step 5 In the **Monitoring** section, change the role to **Primary/ Secondary**.

Step 6 Click **Save**.

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

Syslog over Cisco ISE Messaging

Cisco ISE 2.6 offers MnT WAN Survivability for the built-in UDP syslog collection targets (LogCollector and LogCollector2) by the option **Use ISE Messaging Service for UDP Syslogs delivery to MnT**. This option is disabled by default in Cisco ISE 2.6 First Customer Ship (FCS). This option is enabled by default with Cisco ISE Release 2.6 cumulative Patch 2 onwards.

Using the ISE Messaging Service for UDP syslogs, this feature retains 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 ISE node from all other ISE nodes in the deployment. The following features also use ISE Messaging Service: Light Session Directory (see the section "Light Session Directory" in Chapter "Set Up Cisco ISE in a Distributed Environment" in *Cisco Identity Service Engine Administrator Guide* , and Profiler Persistence Queue. .

Note If your deployment is using TCP/Secure syslogs for ISE deployment, the functionality remains same as the earlier releases.

Queue-link Alarm

The ISE Messaging Service uses a different certificate, signed by internal-CA chain.You might get a `queue-link alarm` in the **Administrations** > **Alarms** window. This alarm is expected in case you are performing any deployment operations such as registering a node to deployment, manually syncing a node from PPAN, a node being in out-of-sync state or in nodes application service is getting restarted. Ensure the following to resolve the alarm:

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

If the prerequisites listed above are met, the queue-link alarm is triggered due to the following actions:

- Changing the domain name or hostname of your PAN or PSN.
- Restoring a backup on a new deployment.
- Promoting the old Primary PAN to new Primary PAN post upgrade.

To resolve the queue-link alarm, regenerate the ISE Root CA chain. Choose **Administration > System > Certificates > Certificate Management > Certificate Signing Request**. Click on **Generate Certificate Signing Request (CSR)**. Select the **ISE Root CA** in the **Certificate(s) will be used for** drop-down list. Click on **Replace ISE root CA Certificate Chain**.

To enable or disable the ISE Messaging Service for UDP Syslogs delivery to MnT:

Procedure

Step 1 Choose **Administrations> System> Logging> Log Settings.ISE root CA**

Step 2 Toggle the **Use “ISE Messaging Service” for UDP Syslogs delivery to MnT** option to use or unuse the ISE Messaging Service for UDP syslog delivery.

Step 3 Click **Save**.

Automatic Failover in Monitoring Nodes

Monitoring nodes do not offer high availablity, but do offer active standby. The Policy Service Node (PSN) copies operational audit data to both the primary and secondary Monitoring nodes.

Automatic Failover Process

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

To manually convert the secondary monitoring node to a primary node, Manually Modify MnT Role. If the primary node comes back up after the secondary node was promoted, it takes the secondary role. If the secondary node was not promoted, the primary Monitoring node resumes the primary role, after it comes back up.

⚠

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

Guidelines for Setting Up an Active Standby Pair of Monitoring Nodes

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

- All changes are logged to the primary Monitoring 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 monitoring data is not part of a standard Cisco ISE node backup process. You must configure repositories for backup and data purging on both the primary and secondary Monitoring nodes, and use the same repositories for each.

Monitoring Node Failover Scenarios

The following scenarios apply to the active/standby or single node configurations corresponding to the monitoring nodes:

- In an active/standby configuration of the monitoring nodes, the Primary Administration Node (PAN) always points to the primary monitoring node to collect the monitoring data. After the primary monitoring node fails, the PAN points to the standby monitoring node. The failover from the primary monitoring node to the standby monitoring node happens after it is down for more than 5 minutes.

  However, after the primary node fails, the standby node does not become the primary node. In case the primary node comes up, the Administration node starts collecting the monitoring data again from the resumed primary node.
- If the primary monitoring node is down, and you want to promote the standby monitoring node to active status, you can promote the standby monitoring node to primary by Manually Modify MnT Role or by de-registering the existing primary monitoring node. When you de-register the existing primary monitoring node, the standby node becomes the primary monitoring node, and the PAN automatically points to the newly promoted primary node.
- In an active/standby pair, if you de-register the secondary monitoring node or if the secondary monitoring node goes down, the existing primary monitoring node remains current the primary node.

- If there is only one monitoring node in the ISE deployment, then that node acts as the primary monitoring node, and provides monitoring data to the PAN. However, when you register a new monitoring node, and make it the primary node in the deployment, the existing primary monitoring node automatically becomes the standby node. The PAN points to the newly registered primary monitoring node to collect monitoring data.

Monitoring Database

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

Like Policy Service, Monitoring has a dedicated database that requires you to perform maintenance tasks, such as the topics covered in this section:

Back Up and Restore of the Monitoring Database

Monitoring database handles large volumes of data. Over time, the performance and efficiency of the monitoring 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 Monitoring node, we recommend that you first back up the primary Monitoring node and then restore the data to the new secondary Monitoring node. This ensures that the history of the primary Monitoring node is in sync with the new secondary node as 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 data during a purge. The default is three months. This value is utilized when the disk space usage threshold for purging (percentage of 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

The following are some guidelines to follow relating to Monitoring database disk usage:

- If the Monitoring database disk usage is greater than 80 percent of the threshold setting, critical alarm is generated indicating that the database size has exceeded the allocated disk size. If the disk usage is greater than 90 percent another alarm is generated.

  A purge process runs, creating a status history report that you can view by choosing **Operations** > **Reports** > **Deployment Status** > **Data Purging Audit**. An information (INFO) alarm is generated when the purge completes.

- 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 percent), the purge process starts. This process deletes only the last seven days of monitoring data, irrespective of what is configured in the Admin portal. It will continue 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

ISE MnT Operational (OPS) database contains information that is generated as ISE reports. Recent ISE releases have options to Purge M&T Operational Data and Reset M&T Database after running the ISE admin CLI command **application configure ise.**

The purge option is used to clean up the data and will prompt to ask the number of days to be retained. The reset option is used to reset the database to the factory default, so that all the data that is backed up will be permanently deleted. You can reset the database if the files are consuming too much file system space.

Note The reset option will cause ISE services to be temporarily unavailable until it restarts.

The **Operational Data Purging** page (Administration > System > Maintenance > Operational Data Purging) contains the **Database Utilization** and **Purge Data Now** areas. You can view the total available database space and the RADIUS and TACACS data stored in the **Database Utilization** area. You can hover the mouse over the status bar to display the available disk space and the number of days the existing data is stored in the database. You can specify the period during 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 to export data to a repository before it is purged by specifying the number of retention days. You can 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 RADIUS and TACACS data or specify the number of days beyond which data is supposed to be purged..

Note You can export the 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 gets 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.

Procedure

Step 1 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 will be exported to a repository.

  b. In the Repository area, check the Enable Export Repository check box to choose the repository to save data. Refer to the Create Repositories section for more information.

  c. In the Encryption Key text box, 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, then the configured value overrides the existing threshold values. For ample, if you configure the retention period as 3 days and this value is less than the existing thresholds in the diagnostics tables (for example, a default of 5 days), then data is purged according to the value that you configure (3 days) in this page.

- In the Purge Data Now area:

  a. Choose to purge all data or to purge data that is older than the specified number of days. Data is not saved in any repository.

  b. Click Purge.

Configure Monitoring Nodes for Automatic Failover

If you have two Monitoring 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 Monitoring node automatically provides monitoring should the primary node fail.

Before you begin

- Before you can configure Monitoring nodes for automatic failover, they must be registered as Cisco ISE nodes.
- Configure monitoring roles and services on both 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 Monitoring 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 Monitoring node uses two repositories for backup and purging, you must specify the same repositories for the secondary node.

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

⚠ **Caution** 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 Monitoring nodes are ready. The System Summary dashlet shows the Monitoring nodes with a green check mark to the left when their services are ready.

Procedure

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

Step 2 In the Deployment Nodes page, check the check box next to the Monitoring node that you want to specify as active, and click **Edit**.

Step 3 Click the **General Settings** tab and choose **Primary** from the **Role** drop-down list.

When you choose a Monitoring node as primary, the other Monitoring node automatically becomes secondary. In the case of a standalone deployment, primary and secondary role configuration is disabled.

Step 4 Click **Save**. The active and standby nodes restart.

pxGrid Node

You can use Cisco pxGrid to share the context-sensitive information from Cisco ISE session directory with other network systems such as ISE Eco system partner systems and other Cisco platforms. The pxGrid framework can also be used to exchange policy and configuration data between nodes like sharing tags and policy objects between Cisco ISE and third party vendors, and for other information exchanges. pxGrid also allows 3rd party systems to invoke adaptive network control actions (EPS) to quarantine users/devices in response to a network or security event. The TrustSec information like tag definition, value, and description can be passed from Cisco ISE via TrustSec topic to other networks. The endpoint profiles with Fully Qualified Names (FQNs) can be passed from Cisco ISE to other networks through a 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 pxGrid. For more information about SXP bindings, see the Security Group Tag Exchange Protocol section in .

In a high-availability configuration, Cisco pxGrid servers replicate information between the nodes through the PAN. When the PAN goes down, pxGrid server stops handling the client registration and subscription. You need to manually promote the PAN for the pxGrid server to become active. You can check the pxGrid Services page (Administration > pxGrid Services) to verify whether a pxGrid node is currently in active or standby state.

For XMPP (Extensible Messaging and Presence Protocol) clients, pxGrid nodes work in Active/Standby high availability mode which means that the pxGrid Service is in "running" state on the active node and in "disabled" state on the standby node.

After the automatic failover to the secondary pxGrid node is initiated, if the original primary pxGrid node is brought back into the network, the original primary pxGrid node will continue to have the secondary role and will not be promoted back to the primary role unless the current primary node goes down.

Note At times, the original primary pxGrid node might be automatically promoted back to the primary role.

In a high availability deployment, when the primary pxGrid node goes down, it might take around 3 to 5 minutes to switchover to the secondary pxGrid node. It is recommended that the client waits for the switchover to complete, before clearing the cache data in case of primary pxGrid node failure.

The following logs are available for pxGrid node:

- pxgrid.log—State change notifications.
- pxgrid-cm.log—Updates on publisher/subscriber and data exchange activity between client and server.
- pxgrid-controller.log—Displays the details of client capabilities, groups, and client authorization.
- pxgrid-jabberd.log—All logs related to system state and authentication.
- pxgrid-pubsub.log—Information related to publisher and subscriber events.

Note If pxGrid service is disabled on a node, port 5222 will be down, but port 8910 (used by Web Clients) will be functional and will continue to respond to the requests.

Note You can enable pxGrid with Base license, but you must have a Plus license to enable pxGrid persona. In addition, certain extended pxGrid services may be available in your Base installation if you have recently installed an upgrade license for .

Note pxGrid should be defined in order to work with the Passive ID Work Center. For more information, see the PassiveID Work Center section in *Cisco ISE Admin Guide: Asset Visibility* .

pxGrid Client and Capability Management

Clients connecting to Cisco ISE must register and recieve account approved before using pxGrid services. pxGrid clients use the pxGrid Client Library available from Cisco through the pxGrid SDK to become the clients.Cisco ISE supports both auto and manual approvals. A client can log in to pxGrid using a unique name and certificate-based mutual authentication. Similar to the AAA setting on a switch, clients can connect to either a configured pxGrid server hostname or an IP Address.

pxGrid "Capabilities" are information topics or channels on pxGrid for clients to publish and subscribe. In Cisco ISE, only capabilities such as Identity, adaptive network control, and SGA are supported. When a client creates a new capability, it appears in. **Administration** > **pxGrid Services** > **View by Capabilities**. You can enable or disable capabilities individually. Capability information is available from the publisher through publish, directed query, or bulk download query.

Note Users that are assigned to EPS user group can perform actions in Session group, because pxGrid Session group is part of EPS group. If a user is assigned to EPS group, the user will be able to subscribe to Session group on pxGrid client.

Related Topics

Generate pxGrid Certificate, on page 44

Enable pxGrid Clients

Before you begin

- Enable the pxGrid persona on at least one node to view the requests from the Cisco pxGrid clients.
- Enable Passive Identity Services. Choose **Administration** > **Deployment**, checkmark the desired node, click **Edit** and from the settings screen, checkmark **Enable Passive Identity Service**.

Procedure

Step 1 Choose **Administration** > **pxGrid Services**.

Step 2 Check the checkbox next to the client and click **Approve**.

Step 3 Click **Refresh** to view the latest status.

Enable pxGrid Capabilities

Before you begin

- Enable the pxGrid persona on at least one node to view the requests from the Cisco pxGrid clients.
- Enable a pxGrid client.

Procedure

Step 1 Choose **Administration** > **pxGrid Services**.

Step 2 Click **View by Capabilities** at the top-right.

Step 3 Select the capability you want to enable and click **Enable**.

Step 4 Click **Refresh** to view the latest status.

Deploy pxGrid Node

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

Before you begin

- You can enable pxGrid with Base license, but you must have a Plus license to enable pxGrid persona.
- Cisco pxGrid services running on a Cisco ISE SNS 3415/3495 Appliance or in VMWare.
- All nodes are configured to use the CA certificate for pxGrid usage. If default certificate is used for pxGrid before upgrade, it will be replaced by the internal CA certificate after upgrade.
- If you are using a distributed deployment or upgrading from Cisco ISE 1.2, then you need to enable the pxGrid Usage option for the certificates. To enable the pxGrid Usage option, go to **Administration** > **Certificates** > **System Certificates**. Choose the certificate being used in the deployment and click **Edit**. Check the pxGrid: use certificate for the pxGrid Controller checkbox.

Procedure

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

Step 2 In the Deployment Nodes page, check the check box next to the node to which you want to enable the pxGrid services, and click **Edit**.

Step 3 Click the **General Settings** tab and check the pxGrid checkbox.

Step 4 Click **Save**.

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.

Cisco pxGrid Live Logs

The Live Logs page displays all the pxGrid management events. Event info includes the client and capability names along with the event type and timestamp.

Navigate to **Administration** > **pxGrid Services** > **Live Log** to view the list of events. You can also clear the logs and resynchronize or refresh the list.

Configure pxGrid Settings

Before you begin

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

Procedure

Step 1 Choose **Administration > pxGrid Services > Settings**.

Step 2 Select the following options based on your requirements:

- Automatically Approve New Accounts—Check this check box to automatically approve the connection requests from new pxGrid clients.

- Allow Password Based Account Creation—Check this check box to enable username/password based authentication for pxGrid clients. If this option is enabled, the pxGrid clients cannot be automatically approved.

  A pxGrid client can register itself with the pxGrid controller by sending the username via REST API. The pxGrid controller generates a password for the pxGrid client during client registration. The administrator can approve or deny the connection request.

Step 3 Click **Save**.

You can use the **Test** option on the pxGrid Settings page to run a health check on the pxGrid node. You can view the details in the pxgrid/pxgrid-test.log file.

Generate pxGrid Certificate

Before you begin

- To perform the following task, you must be a Super Admin or System Admin.
- pxGrid certificate must be generated from the Primary PAN.
- If the pxGrid certificate uses the subject alternative name (SAN) extension, be sure to include the FQDN of the subject identity as a DNS name entry.

Procedure

Step 1 Choose **Administration > pxGrid Services > Certificates**.

Step 2 Select one of the following options from the **I want to** drop-down list:

- 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.
- Generate bulk certificates—You can upload a CSV file that contains the required details.
- Download root certificate chain—You can download the root certificates and add them to the trusted certificate store. You must specify the host name and the certificate download format.

You can download the certificate template from the Certificate Template link and edit the template based on your requirements.

Step 3 (Required if you choose to Generate a single certificate (without a certificate signing request) option) Enter the FQDN of the pxGrid client.

Step 4 (optional) You can enter a description for this certificate.

Step 5 Specify the Subject Alternative Name (SAN). You can add multiple SANs. The following options are available:

- IP address—Enter the IP address of the pxGrid client to be associated with the certificate.
- FQDN—Enter the fully qualified domain name of the pxGrid client.

Note This field is not displayed if you have selected the Generate Bulk Certificate option.

Step 6 Select one of the following options from the **Certificate Download Format** drop-down list:

- 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 certificate 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 a certificate password.

Step 8 Click **Create**.

The certificate that you created is visible in ISE under **Administration** > **System** > **Certificates** > **Certificate Authority** > **Issued Certificates**, and downloaded to your browser's downloads directory.

Control Permissions for pxGrid Clients

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

You can create different types of groups and map the services provided to the pxGrid clients to these groups. Use the **Manage Groups** option in the **Permissions** window to add new groups. You can view the predefined authorization rules that use predefined groups (such as EPS and ANC) in the Permissions window. Note that you can update only the **Operations** field for the predefined rules.

To create an authorization rule for pxGrid clients:

Procedure

Step 1 From the **Administration** tab, choose **pxGrid Services > Permissions**.

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>**

  Note You can specify a custom operation if you select this option.

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

Predefined groups (such as EPS and ANC) and manually added groups (using the **Manage Groups** option in the **Permissions** window) are listed in this drop-down list.

View Nodes in a Deployment

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

Procedure

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

Step 2 Choose **Administration** > **System** > **Deployment**.

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

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

Download Endpoint Statistical Data From Monitoring Nodes

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

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

- Total endpoints on 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.

Refer to 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 due to power outage or other reasons resulting in data loss. Based on the incident, follow the steps below to recover from data loss.

- In case of PAN corruption in deployment, you should promote the Secondary PAN to Primary PAN.
- If SPAN is promotion is not possible due to small deployment or any other reason, restore the most recent available backup.
- In case of PSN corruption, follow the steps to de-register, reset config and register the node again.
- In case of Standalone box, restore most recent available backup.

Note Obtain backup from the standalone box regularaly to avoid loss in the latest configuration changes.

Device Configuration for Monitoring

The monitoring node receives and uses data from the devices on a network to populate the dashboard display. To enable communication between the monitoring 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.

Before you begin

You must click the Syncup button to force a full replication if the Sync Status is set to Out of Sync or if the Replication Status is Failed or Disabled.

Procedure

Step 1 Log in to the Primary PAN.

Step 2 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

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

Before you begin

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

Procedure

Step 1 Log in to the Primary PAN.

Step 2 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.

Step 5 Click **Save**.

Step 6 Verify 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 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

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 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-redirected sessions on the failed node.

Note We recommend that you make all PSNs in the same local network part of the same node group. PSNs need not be 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.

Before you can add PSNs as members to a node group, you must create the node group first. You can create, edit, and delete Policy Service node groups from the Deployment pages of the Admin portal.

Before you begin

Node group members can communicate over TCP/7800.

Procedure

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

Step 2 Click the **action** icon, and then click **Create Node Group**.

Step 3 Enter a unique name for your node group.

Step 4 (Optional) Enter a description for your node group.

Step 5 (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 page before enabling this option.

Step 6 Click **Submit** to save the node group.

After you save the node group, it should appear in the navigation pane on the left. 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.

What to do next

Add a node to a node group. Edit the node by choosing the node group from the Member of Node Group drop-down list.

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 are Administration, Policy Service, and Monitoring. If you deregister a Monitoring 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, 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 page 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 from the Deployment page of the Primary PAN. However, expect a delay of 5 minutes for the changes to take effect and appear on the Deployment page.

Before you begin

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

Procedure

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

Step 2 Check the check box next to the secondary node that you want to remove, and then click **Deregister**.

Step 3 Click **OK**.

Step 4 Verify 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, the alarm is not generated.

Shut Down an ISE Node

Before you issue the halt command, it is advised that you stop Cisco ISE application service and ensure that it is not performing any backup, restore, installation, upgrade, or remove operation. If you issue 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 click 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 click Yes to save the existing Cisco ISE configuration, the following message is displayed:

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

Note It is recommended that you stop the application process before rebooting the appliance.

This is also applicable to reboot ISE. For more information: see Cisco Identity Services Engine CLI Reference Guide

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. You cannot use "localhost" as the hostname for a node.

Before you begin

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

Procedure

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 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 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.

Cisco ISE Deployment Upgrade

Cisco ISE offers a GUI-based centralized upgrade from the Admin portal. The upgrade process is much simplified and the progress of the upgrade and the status of the nodes are displayed on screen. Refer to the *Cisco Identity Services Engine Upgrade Guide* for a list of pre and post upgrade tasks.

The Upgrade Overview page lists all the nodes in your deployment, the personas that are enabled on them, the version of ISE installed, and the status (indicates whether a node is active or inactive) of the node. You can begin upgrade only if the nodes are in the Active state.

Different Types of Deployment

- Standalone Node—A single Cisco ISE node assuming the Administration, Policy Service, and Monitoring persona.
- Multi-Node Deployment—A distributed deployment with several ISE nodes. The procedure to upgrade a distributed deployment is discussed in the following listed references.

ISE Community Resource

For information on how to assess the network for ISE deployment readiness, see ISE Deployment Assistant (IDA).

Upgrade a Distributed Deployment

You can upgrade all the nodes in a Cisco ISE deployment using the Admin portal from Release 2.0 onwards, you can also upgarde a Limited Availability Release of Cisco ISE 2.0 or later to the General Availability Release.

Before you begin

Ensure that you have performed the following tasks before you upgrade:

- Obtain a backup of the ISE configuration and operational data.
- Obtain a backup of the system logs.
- Disable scheduled backups. Reconfigure the backup schedules after deployment upgrade is complete.
- Export the certificates and private keys.
- Configure a repository. Download the upgrade bundle and place it in the repository.
- Make a note of Active Directory join credentials and RSA SecurID node secret, if applicable. You need this information to connect to Active Directory or RSA SecurID server after upgrade.
- Purge the operational data to improve upgrade performance.
- Ensure that your Internet connection to the repository is good.

Note When you download an upgrade bundle from a repository to a node, the download times out if it takes more than 35 minutes to complete. This issue occurs because of poor Internet bandwidth.

Procedure

Step 1 Click the **Upgrade** tab in the Admin portal.

Step 2 Click **Proceed**.

The **Review Checklist** window appears. Read the given instructions carefully.

Step 3 Check the **I have reviewed the checklist** check box, and click **Continue**.

The **Download Bundle to Nodes** window appears.

Step 4 Download the upgrade bundle from the repository to the nodes:

a) Check the check box next to the nodes to which you want to download the upgrade bundle.

b) Click **Download**.

The **Select Repository and Bundle** window appears.

c) Select the repository.

You can select the same repository or different repositories on different nodes, but you must select the same upgrade bundle on all the nodes.

d) Check the check box next to the bundle that you want to use for the upgrade.

e) Click **Confirm**.

Once the bundle is downloaded to the node, the node status changes to **Ready for Upgrade**.

Step 5 Click **Continue**.

The **Upgrade Nodes** window appears.

Figure 1: Upgrade Window Showing the Current Deployment and the New Deployment

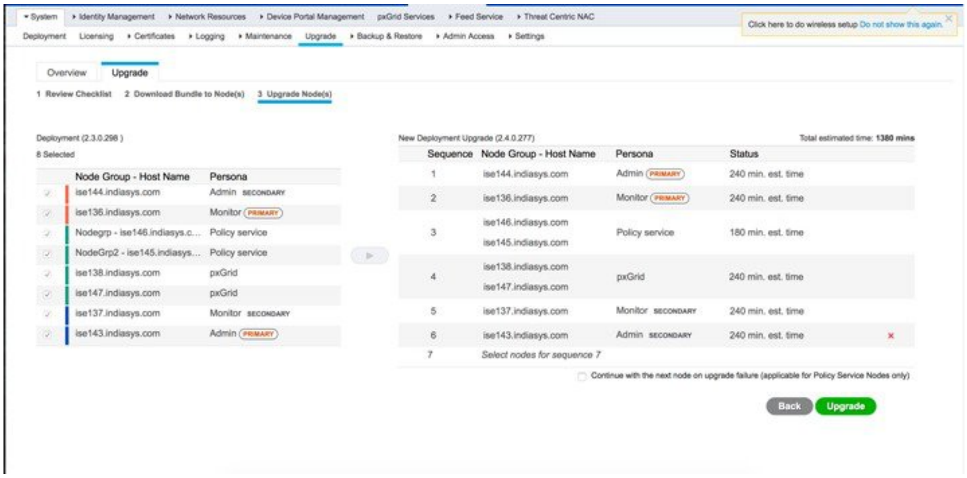

Deployment (2.3.0.298)

8 Selected

	Node Group - Host Name	Persona
☑	ise144.indiasys.com	Admin SECONDARY
☑	ise136.indiasys.com	Monitor PRIMARY
☑	Nodegrp - ise146.indiasys.c...	Policy service
☑	NodeGrp2 - ise145.indiasys...	Policy service
☑	ise138.indiasys.com	pxGrid
☑	ise147.indiasys.com	pxGrid
☑	ise137.indiasys.com	Monitor SECONDARY
☑	ise143.indiasys.com	Admin PRIMARY

New Deployment Upgrade (2.4.0.277)

Total estimated time: **1380 mins**

Sequence	Node Group - Host Name	Persona	Status	
1	ise144.indiasys.com	Admin PRIMARY	240 min. est. time	
2	ise136.indiasys.com	Monitor PRIMARY	240 min. est. time	
3	ise146.indiasys.com ise145.indiasys.com	Policy service	180 min. est. time	
4	ise138.indiasys.com ise147.indiasys.com	pxGrid	240 min. est. time	
5	ise137.indiasys.com	Monitor SECONDARY	240 min. est. time	
6	ise143.indiasys.com	Admin SECONDARY	240 min. est. time	×
7	*Select nodes for sequence 7*			

☐ Continue with the next node on upgrade failure (applicable for Policy Service Nodes only)

Step 6 Choose the upgrade sequence.

When you move a node to the new deployment, a time estimate for the upgrade is displayed on the **Upgrade Nodes** window. You can use this information to plan for upgrade and minimize downtime. Use the sequence given below if you have a pair of Administration and Monitoring Nodes, and several Policy Service Nodes.

a) By default, the Secondary Administration Node is listed first in the upgrade sequence. After upgrade, this node becomes the Primary Administration Node in the new deployment.
b) The Primary Monitoring Node is the next one in the sequence to be upgraded to the new deployment.
c) Select the Policy Service Nodes and move them to the new deployment. You can alter the sequence in which the Policy Service Nodes are upgraded.

   You can upgrade the Policy Service Nodes in sequence or in parallel. You can select a set of Policy Service Nodes and upgrade them in parallel.

d) Select the Secondary Monitoring Node and move it to the new deployment.
e) Finally, select the Primary Administration Node and move it to the new deployment.

Step 7 Check the **Continue with upgrade on failure** check box if you want to continue with the upgrade even if the upgrade fails on any of the Policy Service Nodes in the upgrade sequence.

This option is not applicable for the Secondary Administration Node and the Primary Monitoring Node. If any one of these nodes fail, the upgrade process is rolled back. If any of the Policy Service Nodes fail, the Secondary Monitoring Node and the Primary Administration Node are not upgraded and remain in the old deployment.

Step 8 Click **Upgrade** to begin the deployment upgrade.

Figure 2: Upgrade Window Showing the Upgrade Progress

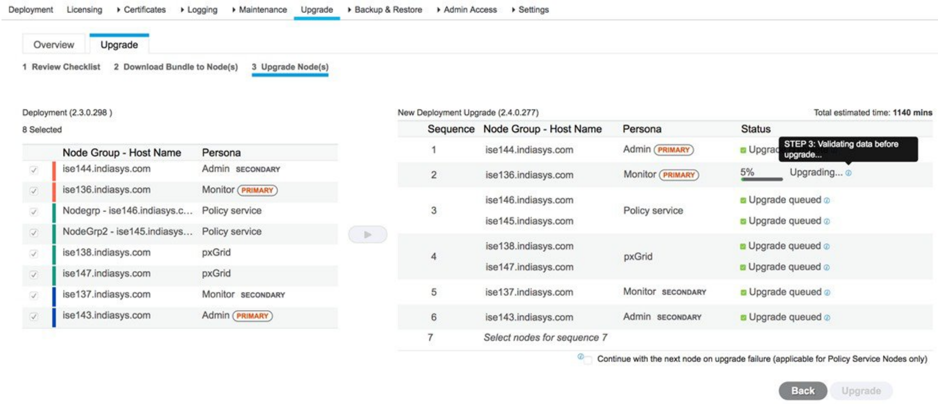

Deployment (2.3.0.298)

8 Selected

	Node Group - Host Name	Persona
☑	ise144.indiasys.com	Admin SECONDARY
☑	ise136.indiasys.com	Monitor PRIMARY
☑	Nodegrp - ise146.indiasys.c...	Policy service
☑	NodeGrp2 - ise145.indiasys...	Policy service
☑	ise138.indiasys.com	pxGrid
☑	ise147.indiasys.com	pxGrid
☑	ise137.indiasys.com	Monitor SECONDARY
☑	ise143.indiasys.com	Admin PRIMARY

New Deployment Upgrade (2.4.0.277)

Total estimated time: **1140 mins**

Sequence	Node Group - Host Name	Persona	Status
1	ise144.indiasys.com	Admin PRIMARY	Upgrad
2	ise136.indiasys.com	Monitor PRIMARY	5% Upgrading...
3	ise146.indiasys.com ise145.indiasys.com	Policy service	Upgrade queued Upgrade queued
4	ise138.indiasys.com ise147.indiasys.com	pxGrid	Upgrade queued Upgrade queued
5	ise137.indiasys.com	Monitor SECONDARY	Upgrade queued
6	ise143.indiasys.com	Admin SECONDARY	Upgrade queued
7	*Select nodes for sequence 7*		

STEP 3: Validating data before upgrade...

☐ Continue with the next node on upgrade failure (applicable for Policy Service Nodes only)

The upgrade progress is displayed for each node. On successful completion, the node status changes to **Upgrade Complete**.

Note When you upgrade a node from the Admin portal, if the status does not change for a long time (and remains at 80%), you can check the upgrade logs from the CLI or the status of the upgrade from the console. Log in to the CLI or view the console of the Cisco ISE node to view the progress of upgrade. You can use the **show logging application** command to view the *upgrade-uibackend-cliconsole.log* and *upgrade-postosupgrade-yyyymmdd-xxxxxx.log*.

You can view the following upgrade logs from the CLI using the show logging application command:

- DB Data Upgrade Log
- DB Schema Log
- Post OS Upgrade Log

In case you get a warning message: **The node has been reverted back to its pre-upgrade state**, go to the **Upgrade** window, click the **Details** link. Address the issues that are listed in the **Upgrade Failure Details** window. After you fix all the issues, click **Upgrade** to reinitiate the upgrade.

Note If the posture data update process is running on the Primary Administration Node in the new deployment, you cannot register a node to the Primary Administration Node. You can either wait till the posture update process is over (which might take approximately 20 minutes) or disable the posture auto-update feature from the **Administration** > **System** > **Settings** > **Posture** > **Updates** page while upgrading or registering a node to the new deployment.

Note When you upgrade from Cisco ISE release 2.2 or later to release 2.7, MAC SPW bundle is not listed in **Policy** > **Results** > **Client Provisioning** > **Resources**. Download *mac-spw-dmg-2.7.0.1-isebundle* from cisco.com and upload to the resources to provision MAC OS X 10.15 release .