



Illumio Core[®]

Version 23.5

PCE Administration Guide

June 2024

30000-100-23.5

Legal Notices

Copyright © 2024 Illumio 920 De Guigne Drive, Sunnyvale, CA 94085. All rights reserved.

The content in this documentation is provided for informational purposes only and is provided "as is," without warranty of any kind, expressed or implied of Illumio. The content in this documentation is subject to change without notice.

Product Version

PCE Version: 23.5

For the complete list of Illumio Core components compatible with Core PCE, see the Illumio Support portal (login required).

For information on Illumio software support for Standard and LTS releases, see [Versions and Releases](#) on the Illumio Support portal.

Resources

Legal information, see <https://www.illumio.com/legal-information>

Trademarks statements, see <https://www.illumio.com/trademarks>

Patent statements, see <https://www.illumio.com/patents>

License statements, see <https://www.illumio.com/eula>

Open source software utilized by the Illumio Core and their licenses, see [Open Source Licensing Disclosures](#)

Contact Information

To contact Illumio, go to <https://www.illumio.com/contact-us>

To contact the Illumio legal team, email us at legal@illumio.com

To contact the Illumio documentation team, email us at doc-feedback@illumio.com

Contents

Chapter 1 Overview of PCE Administration	10
About This Administration Guide	10
How to Use This Guide	10
Before Reading This Guide	10
Notational Conventions in This Guide	11
PCE Architecture and Components	11
About the PCE Architecture	11
Description of PCE Components	12
Management Interfaces for PCE and VEN	14
PCE Control Interface and Commands	15
PCE Organization and Users	16
RBAC Users Roles and Permissions	16
Invite Users to Your Organization	16
Chapter 2 Manage PCE Nodes and Clusters	18
Manage Data and Disk Capacity	18
Identify Data Management Strategies	19
Detect Disk Usage	19
Respond to Disk Capacity Issues	20
Recover Disk Usage	21
Cluster Nodes and Command-Line Operations	21
PCE Control Commands	21
Database Commands	22
Start and Stop Nodes and Cluster	23
Start Individual PCE Node	23
Stop a PCE Node or Entire Cluster	23
Restart a PCE Node or Entire Cluster	23
Check Node and Cluster Status	24
Check Node Environment	24
Check PCE Node Status	24
Check Services on a PCE Node	25
Check PCE Cluster Status	26
Check PCE Version	27
Check PCE Cluster Members	27
Update PCE Configuration	27

Back up PCE Runtime File	27
Update Runtime Configuration	27
Get Current PCE Runlevel	28
Set PCE Runlevel	28
Update PCE Certificates	28
Change the PCE FQDN	29
Upgrade the OS on a Running PCE	30
Firewall Coexistence	35
Firewall Tampering Protection	36
Firewall Coexistence Modes	36
Prerequisites and Recommendations	38
Set Firewall Coexistence	39
PCE Listen Only Mode	40
About PCE Listen Only Mode	40
Enable PCE Listen Only Mode	41
Determine if PCE Is in Listen Only Mode	41
VEN Heartbeat and Listen Only Mode	41
Disable PCE Listen Only Mode	43
Expand 2x2 Cluster to 4x2	43
Prepare Environment for Cluster Expansion	43
Back Up PCE Database	45
Configure Existing Nodes for Expansion	45
Install and Configure PCE on Nodes	46
Verify Cluster Expansion	48
Replace PCE Nodes or Uninstall Cluster	49
Replace a Failed Node	49
Replace a Running Node	50
Uninstall the PCE Cluster	51
Chapter 3 PCE Database Management	53
About the PCE Databases	53
Policy and Traffic Data Databases	53
Data Retention of Traffic Flow Summaries	53
Determine the Primary Database	54
Show Database Replication Information	54
Rotate Database Passwords and Other Secrets	54
Anonymize Database Export	55

View Events Using PCE Command Line	56
PCE Database Backup	57
About PCE Database Backup	58
Back Up the Policy Database	58
Back Up the Traffic Database	59
Back Up the PCE Runtime Environment File	61
Database Migration, Failover, and Restore	62
Migrate PCE Databases	62
Manage Automatic Database Failover	62
Manual Database Failover	63
Restore from Data Backup	64
Manage Multi-Node Traffic Database	66
Expand Existing Traffic Database to Multiple Nodes	66
Add or Remove a Worker Node	69
Back Up and Restore Multi-Node Traffic Database	70
Database Management Commands for Multi-Node Traffic Database	70
PCE Default Object Limits	70
Types of Object Limits	71
Check Object Limits and Usage	72
Object Limits During Bulk Create	73
Object Limits and Concurrent Transactions	73
PCE Object Limits	73
Chapter 4 Monitor and Diagnose PCE Health	79
<hr/>	
PCE Logs	79
Log Files for PCE Services	79
Log Files (Non-syslog)	80
Password-related Event Logging	82
Search the PCE Log Files	82
Monitor PCE Health	84
PCE Health Monitoring Techniques	84
Minimum Required Monitoring	85
Health Monitoring Using PCE Web Console	86
Health Monitoring Using Health REST API	87
Health Monitoring Using Syslog	87
Health Monitoring Using PCE Command Line	88
PCE Health Troubleshooting	90

PCE Health Metrics Reference	96
Support Reports for PCE	109
Generate PCE Support Bundle in Web Console	109
Generate PCE Support Report at Command Line	109
View Host and System Inventory	111
Chapter 5 PCE HA and DR	113
<hr/>	
PCE HA and DR Concepts	113
Overview of PCE HA and DR	113
Design Goals for PCE HA	113
PCE HA and DR Requirements	115
PCE Cluster Front End Load Balancing	115
Traffic Load Balancer Requirements	116
DNS Load Balancing	116
Network Latency Between Nodes	116
PCE Replication and Failover	117
Standby PCE Prerequisites	117
Set Up a Standby PCE	121
Failover to Standby PCE	123
Monitoring Replication	124
Limitations and Constraints	125
PCE Failures and Recoveries	126
Types of PCE Failures	126
PCE-VEN Network Partition	127
Service Failure	129
Core Node Failure	131
Data Node Failure	132
Site Failure (Split Clusters)	136
Cluster Network Partition	143
Multi-Node Traffic Database Failure	145
Complete Cluster Failure	146
Complete Cluster Recovery	148
PCE-Based VEN Distribution Recovery	149
Restore VENs Paired to Failed PCE	150
Chapter 6 Connectivity Configuration for PCE	151
<hr/>	
Connectivity Settings	151
Private Data Centers	151

Offline Timers	152
Set the IP Version for Workloads	155
Manage Security Settings	156
Enable IP Forwarding	158
SecureConnect Setup	159
Features of SecureConnect	159
Use Pre-Shared Keys with SecureConnect	160
Use PKI Certificates with SecureConnect	161
Prerequisites, Limitations, and Caveats	161
Configure SecureConnect to Use Pre-Shared Keys	163
Configure SecureConnect to Use Certificates	163
Requirements for Certificate Setup on Workloads	164
AdminConnect Setup	165
Features of AdminConnect	166
Prerequisites and Limitations	167
Certificates for AdminConnect	167
Secure Laptops with AdminConnect	168
Chapter 7 Access Configuration for PCE	171
Role-based Access Control	171
Overview of Role-based Access Control	171
Use Cases	172
Features of Role-based Access Control	173
About Roles, Scopes, and Granted Access	174
Prerequisites and Limitations	181
Setup for Role-based Access Control	182
Add a Scoped Role	182
Manage a Local User	182
Manage a Service Account	185
Add or Remove an External User	186
Add or Remove an External Group	187
Change Users and Groups Added to Roles	189
View User Activity	189
Change Your Profile Settings	190
Role-based Access for Application Owners	192
Overview	192
Updates to Roles	193

Configuration	195
Facet Searches for Scoped Roles	196
Ruleset Viewer	196
Scoped Roles and Permissions	197
Scoped Users and PCE	200
Labeled Objects	204
Rulesets and Rules	205
App Group Map	206
Policy Generator and Explorer	207
My Roles	208
Configure Access Restrictions and Trusted Proxy IPs	208
Configure Access Restrictions	208
Configure Trusted Proxy IPs	209
Password Policy Configuration	211
About Password Policy for the PCE	211
Password Requirements	212
Password Expiration and Reuse	212
Change Password Policy Settings	213
Authentication	215
SAML SSO Authentication	216
Signing for SAML Requests	217
LDAP Authentication	219
Active Directory Single Sign-on	225
Overview of AD FS SSO Configuration	225
Configure AD Users to Use Different UPN Suffixes	225
Initial AD FS SSO Configuration	228
Create a Relying Party Trust	235
Create Claim Rules	244
Obtain ADFS SSO Information for the PCE	255
Configure the PCE for AD FS SSO	257
Azure AD Single Sign-on	258
Prerequisites	259
STEP 1: Obtain URLs from the Illumio PCE Web Console	259
STEP 2: Configure SSO settings in Azure AD	259
STEP 3: Obtain SAML certificate and URLs from Azure AD	263
STEP 4: Configure SAML SSO settings in the Illumio PCE	264
STEP 5: Create App Roles in Azure AD	265

STEP 6: Assign users and groups to app roles in Azure AD	266
STEP 7: Add External Groups and assign roles in the PCE Web Console	267
STEP 8: Turn on SAML authentication in the PCE Web Console	269
STEP 9: Test SSO	269
Okta Single Sign-on	270
Prerequisite for Okta SSO	270
Configure the PCE for Okta SSO	270
OneLogin Single Sign-on	272
Configure SSO for OneLogin	272
Ping Identity Single Sign-on	274
Configure SSO for Ping Identity	274
Chapter 8 PCE Administration Troubleshooting	279
PCE Administration Troubleshooting Scenarios	279
Transaction ID Wraparound in PostgreSQL Database	279

Overview of PCE Administration

This chapter contains the following topics:

About This Administration Guide	10
PCE Architecture and Components	11
PCE Control Interface and Commands	15
PCE Organization and Users	16

This section explains concepts that will help you with ongoing PCE operations and administration.

About This Administration Guide

The following sections give useful information to help you get the most out of this guide.

How to Use This Guide

This guide describes how to maintain and operate the Policy Compute Engine (PCE). It also includes other important tasks required to manage your PCE deployment.

Before Reading This Guide

Before attempting the procedures in this guide, you should be familiar with the following technology:

- Your organization's security goals
- Illumio Core

- General computer system administration of Linux and Windows operating systems, including startup/shutdown, common processes or services
- Linux shell (bash) and Windows PowerShell
- TCP/IP networks, including protocols and well-known ports
- PKI certificates

Notational Conventions in This Guide

- Newly introduced terminology is italicized. Example: *activation code* (also known as pairing key)
- Command-line examples are monospace. Example: `illumio-ven-ctl --activate`
- Arguments on command lines are monospace italics. Example: `illumio-ven-ctl -
-activate activation_code`
- In some examples, the output might be shown across several lines but is actually on one single line.
- Command input or output lines not essential to an example are sometimes omitted, as indicated by three periods in a row. Example:

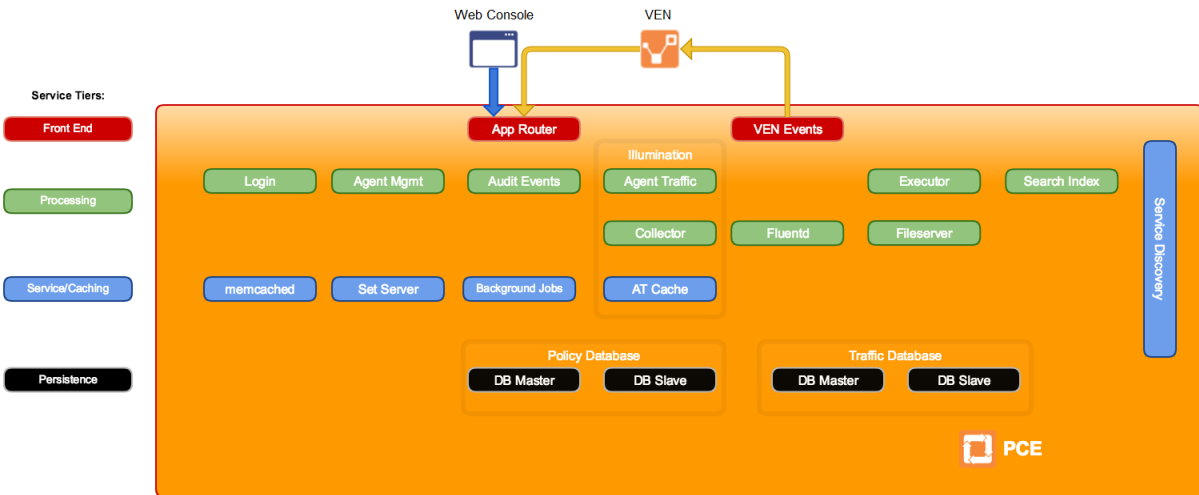
```
...  
some command or command output  
...
```

PCE Architecture and Components

This section describes how the PCE functions, and provides an overview of its components and how they function together.

About the PCE Architecture

The PCE has service tiers responsible for various functions.




Description of PCE Components

Tier	PCE component	Description
Front-end	Management interfaces: PCE web console and VEN	Management interfaces include: <ul style="list-style-type: none"> • PCE web console • REST API • PCE command line • VEN command line
	VEN events	For information, see VEN Architecture and Components in the <i>VEN Administration Guide</i> .
	App Router	Directs requests to the proper service.

Tier	PCE component	Description
Processing	Login	Central server for authentication.
	Agent Manager	Manages data in the policy domain, such as workload context and policy definitions. Also, manages data for all user and organization authentication and authorization, such as users, organizations, API keys, and roles.
	Agent Traffic	Provides information about traffic to and from VENs. Serves as the service underlying Illumination.
	Collector	Aggregates packet and traffic flow information sent from the VEN. Serves as the service underlying Illumination.
	Audit Events	Creates an overview of auditable system events across the PCE and VENs.
	Fluentd	Log forwarder service that forwards the flow log files received from VENs.
	Executor	Backbone for asynchronous job execution, such as report generation and background jobs.
	Fileserver	Central storage and retrieval for large data files.
	Search Index	Supports auto-completion in the PCE web console.
Service	memcached	Open source component: in-memory cache.
	Background Jobs	Backbone for asynchronous job execution, such as report generation and background jobs.
	Set Server	In-memory cache to aid in policy calculations.
	Agent Traffic cache	Stores the traffic flow data and graphs for Illumination. See Agent Traffic. In the PCE architecture diagram, labeled “AT Cache.”
Persistence	Policy primary database and replica	Postgres database that contains all the policy and agent related data. The primary and replica databases run on separate data nodes.
	Traffic database primary and replica	Postgres database that contains all the historical traffic flow data. Traffic Explorer is backed by this datastore. The primary and replica databases run on separate data nodes.

Management Interfaces for PCE and VEN

The following diagram illustrates the logical view of the management interfaces to the PCE and VEN.



This guide focuses on the use of the `illumio-pce-ctl` control script and related administrative programs on the PCE itself.

Interface	Notes	See...
PCE web console	With the PCE web console, you can perform many common tasks for managing the Illumio Core.	<i>Visualization Guide</i>
PCE command line	Use of the command line directly on the PCE. The <code>illumio-pce-ctl</code> command-line tool is the primary management tool on the PCE. You can perform many common tasks for managing the Illumio Core, including installing and updating the VEN.	<i>PCE Administration Guide</i>
REST API	With the Illumio Core REST API, you can perform many common management tasks, such as automate the management of large groups of workloads, rather than each workload individually. The endpoint for REST API requests is the PCE itself, not the workload. The REST API does not communicate directly with the VEN.	<i>REST API Developer Guide</i>
VEN com-	The <code>illumio-ven-ctl</code> command-line tool is the	<i>VEN Administration</i>

Interface	Notes	See...
mand line	primary management tool for the VEN.	<i>Guide</i>

PCE Control Interface and Commands

The Illumio PCE control interface `illumio-pce-ctl` is a command-line tool for performing key tasks for operating your PCE cluster, such as starting and stopping nodes, setting cluster runlevels, and checking the cluster status.

IMPORTANT:

In this guide, all command-line examples based on an RPM installation. When you install the PCE using the tarball, you must modify the commands based on your PCE user account and the directory where you installed the software.

The PCE includes other command-line utilities used to set up and operate your PCE:

- `illumio-pce-env`: Verify and collect information about the PCE runtime environment.
- `illumio-pce-db-management`: Manage the PCE database.
- `supercluster-sub-command`: Manage specific Supercluster operations.

The PCE control interface can only be executed by the PCE runtime user (`ilo-pce`), which is created during the PCE RPM installation.

Control Command Access with `/usr/bin`

For easier command execution, PCE installation creates softlinks in `/usr/bin` by default for the Illumio PCE control commands. The `/usr/bin` directory is usually included by default in the `PATH` environment variable in most Linux systems. When your `PATH` does not include `/usr/bin`, add it to your `PATH` with the following command. You might want to add this command to your login files (`$HOME/.bashrc` or `$HOME/.cshrc`).

```
export PATH=$PATH:/usr/bin
```

Syntax of `illumio-pce-ctl`

To make it simpler to run the PCE command-line tools, you can run the following Linux softlink commands or add them to your `PATH` environment variable.

```
$ cd /usr/bin
$ sudo ln -s /opt/illumio-pce/illumio-pce-ctl ./illumio-pce-ctl
$ sudo ln -s /opt/illumio-pce/illumio-pce-db-management ./illumio-pce-db-
management
$ sudo ln -s /opt/illumio-pce/illumio-pce-env ./illumio-pce-env
```

After these commands are executed, you can run the PCE command-line tools using the following syntax:

```
$ sudo -u ilo-pce illumio-pce-ctl sub-command --option
```

Where:

sub-command is an argument displayed by `illumio-pce-ctl --help`.

PCE Organization and Users

A PCE organization is a group of policies and users targeted toward a specific business group or unit, including all the networking security rules and people who are associated with the policy. An organization can contain any number of users, workloads, policy objects (rulesets, IP lists, services, and security settings), and labels.

Organizations are initially set up by your Illumio administrator. When an organization is created, an email is sent that contains a user login for the organization. When this user logs in, the organization is created, and users can now be invited to join.

RBAC Users Roles and Permissions

For information on creating local or external users and assigning PCE permissions to those users, see [Role-based Access Control](#).

Invite Users to Your Organization

When you are an organization owner, you can invite other users to your organization and grant roles to specify permissions for those users.

When you invite a user to your organization, the user receives an email at the specified address that contains a link for their account setup. The link in invitation email is valid only for 7 days, after which it expires. If you invited a user who did not receive their email or did not sign up using that email, you can re-invite them.

External Users and Non-Email Usernames

When you use an external corporate Identity Provider (IdP) to authenticate users with the PCE, but your IdP usernames do not use email addresses, the PCE cannot send email invitations to those users when you add them to the PCE. When you add this type of user, send them a login URL that they can use to set up their Illumio Core accounts and log in to the PCE web console.

Invitation Emails Are Not Sent

When users you invite do not receive their invitation emails, the SMTP server might not be configured correctly with the PCE.

- Make sure that your PCE's IP address is allowed to relay messages and that its emails are not blocked by any anti-spam protection.
- Check your PCE's `runtime_env.yml` file to make sure that the `smtp_relay_address` value is correct.

Manage PCE Nodes and Clusters

This chapter contains the following topics:

Manage Data and Disk Capacity	18
Cluster Nodes and Command-Line Operations	21
Start and Stop Nodes and Cluster	23
Check Node and Cluster Status	24
Update PCE Configuration	27
Firewall Coexistence	35
PCE Listen Only Mode	40
Expand 2x2 Cluster to 4x2	43
Replace PCE Nodes or Uninstall Cluster	49

This section describes how to manage PCE infrastructure, which is made up of core and data nodes organized into one or more clusters.

Manage Data and Disk Capacity

The amount of data collected and stored by the PCE can be large. Events, Explorer, and the internal syslog all generate data that is stored in PCE databases and log files. When the amount of stored data is not managed carefully, disks can become overfull. This occurrence can cause a variety of symptoms: inability to take backups, failing API calls, and general PCE functionality issues. Even when these issues do not occur, a large amount of stored data creates larger database backups, and it takes longer to back up and restore the database.

To successfully manage these issues, consider the following recommendations:

- **Identify:** Know your organization’s policies, backup strategies, and monitoring strategies.
- **Detect:** Monitor ongoing disk usage.
- **Respond:** Know how to troubleshoot and fix issues related to data storage.
- **Recover:** Set up your PCE deployment to reduce disk usage.

Identify Data Management Strategies

Identify your organization's policies and strategies related to data storage and retention, backups, and monitoring. This knowledge forms the basis for any ongoing data management activities. You'll need the following information:

- **Records retention policy:** How many days of events data must be available at all times? When your policy requires fewer days of events data than the PCE's default, you can decrease the PCE's events retention period, which helps avoid filling up disk space.
- **System backup policy:** Are full backups always necessary, or would weekly full backups be sufficient, supplemented by smaller daily backups that do not include events data?
- **Disk usage trends:** How fast is data usage growing in your Illumio Core deployment? What is the additional data usage each day?
- **Monitoring tools:** What disk monitoring tools are in place? If none, is there a useful tool that could be added? Do the monitoring tools integrate with the PCE Health API?

Detect Disk Usage

Monitor disk usage to be sure you are aware of status and trends, especially any unusual activity, such as sudden spikes or other anomalies.

- Watch the PCE Health page. For information, see [Monitor PCE Health](#).
 - Check the Disk Usage figures.
 - When disk usage is too high, the PCE displays warnings, such as "Disk Critical."
 - You can call the page's underlying PCE Health API with external monitoring tools.
- Check the system health messages that are sent to syslog from each node in the cluster.
- Use the command `illumio-pce-ctl events-db disk-usage-show` to get the number of events in the database, the amount of disk used by the Events database, and the average number of events per day. For more information, see [View Events Using PCE Command Line](#) in the *Events Administration Guide*.
- Run your own disk monitoring tools or use standard Linux commands, such as `df` and `du`.

Respond to Disk Capacity Issues

You can prevent many disk capacity issues by deploying the PCE with sufficient resources. Be sure your disk meets the recommendations in [PCE Capacity Planning](#) in the *PCE Installation and Upgrade Guide*.

When you are running out of storage space, use Linux tools to find the parts of the disk that are being utilized heavily. Then, depending on your findings, try some of these techniques:

- Are the PCE log files taking up disk space? Look for extra, older files you can move or delete from the log directory (usually `/var/logs/illumio-pce`).
- Are other system logs taking too much space? Rotate and compress them, or delete them.
- After a PCE successfully joins a Supercluster, a directory called `postgres1.bak` is sometimes left behind in the `<postgres1 directory>`, especially on the database master node. You can delete the directory `postgres1.bak` and all its contents. This file directory is kept in case the `cluster-join` command fails and you need to recover, but once the `cluster-join` is complete, and your disk space needs become the higher priority, the directory can be removed.
- Delete any large or unnecessary files in the `/tmp` directory; for example, core files.
- Remove copies of backups stored on PCE nodes. In general, don't use the PCE as a place to store backup files.
- Reduce the retention period for events data, making sure it is still acceptable according to your organization's record retention policy. The PCE automatically deletes excess older records from the database. The default data retention period for events is 30 days. You can decrease the retention period to as little as 1 day. However, exercise caution; balance the need to minimize disk usage against your company's data retention policies and your need to retain data for analysis. For information about how to change the data retention period, see [Configure Events Settings in PCE Web Console](#) in the *Events Administration Guide*.
- The PCE provides short-term storage of events data. Consider forwarding events data to Splunk or other SIEM software for long-term storage in accordance with your organization's data retention policies.
- Consider excluding events from most database dumps. Use the option `--no-include-events` for the `illumio-pce-db-management dump` command. When your organization's policies permit it, perform a full database dump (which includes events data) once during each events data retention period.

Recover Disk Usage

- **Extend the disk:** When the current disk or partition is smaller than the recommended size, increase the partition size. The file `runtime_env.yml` can be configured with different local partition settings.
- **Add a partition or slice for logs or backups:** Copy the old files in `/var/logs/illumio-pce` to a new disk. Mount the new disk to the same location on the PCE with the same permissions as the original disk.
- **Create a new disk or partition:** Mount a new disk or partition to a suitable location for saving backup files.
- **Move the Explorer database to its own disk:** Mount a new dedicated disk and move files from the existing traffic datastore to this dedicated disk. For information, see [How to Move an Existing Explorer Database to a Separate Disk](#) in the Illumio Knowledge Base (login required).

Cluster Nodes and Command-Line Operations

The PCE control interface commands are restricted to the type of node they can be executed on. For example, the command to set a cluster’s runlevel can be run on any core or data node. Database-specific commands must only be run on specific data nodes. The following tables list the command-line operations you can perform and the specific nodes the commands must be run on.

PCE Control Commands

The following table shows commands you can use to control various aspects of PCE behavior. Some of the commands affect a single node and others affect the entire PCE cluster. The commands have the following general syntax:

```
# sudo -u ilo-pce illumio-pce-ctl sub-command --option
```

Sub-Command	Description	Run on Node
Single-node commands		
start	Start PCE software on a single node.	Any
start --run-level n	Start PCE software at a specified runlevel on a single node.	
stop	Stop PCE software on a single node.	Any
restart	Restart the PCE software on a single node.	Any

Sub-Command	Description	Run on Node
status	Show status of the PCE software on a single node.	Any
check-env	Check the runtime_env.yml file on a single node.	Any
service-discovery-status	Get status of service-discovery services on a single node.	Any
check-consul-status	Get status of the consul service on a single node.	Any
Cluster-wide commands		
set-runlevel	Set the software runlevel for the PCE software on all nodes.	Any
get-runlevel	Get the runlevel of the PCE software on all nodes.	Any
cluster-status	Get the status of the PCE software across the cluster.	Any
cluster-stop	Shut down the cluster.	An
cluster-restart	Restart the cluster.	Any
cluster-leave	Force the current node or the node defined by the IP address to be removed from the cluster.	Any
cluster-members	Show all cluster members.	Any

Database Commands

The following table shows commands you can use to control various aspects of PCE database behavior. The commands have the following general syntax:

```
# sudo -u ilo-pce illumio-pce-db-management sub-command --option
```

Sub-Command	Description	Run on Node
setup	Begin initial setup of the PCE database.	Any
migrate	Migrate the database to the latest schema.	Any
dump	Dump the database to a file.	Data node where agent_traffic_redis_server service is running.
restore	Restore the database from a file.	Any data node
create-domain	Create the first organization and user in the system.	Any data node

Sub-Command	Description	Run on Node
show-master	Show which node is the primary database.	Any
show-replication-info	Show replication lag between the replica and primary databases.	Any

Start and Stop Nodes and Cluster

This section describes how to stop and start the PCE.

Start Individual PCE Node

This command starts the node where it is run:

```
$ sudo -u ilo-pce illumio-pce-ctl start
```

Stop a PCE Node or Entire Cluster

This command stops the node where it is run:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

This command stops the entire cluster and can be run on *any node* in the cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-stop
```

Restart a PCE Node or Entire Cluster

This command restarts the node where it is run:

```
$ sudo -u ilo-pce illumio-pce-ctl restart
```

This command restarts the entire cluster and can be run on *any node* in the cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-restart
```

When the PCE is restarted, the UI can become available before all the required PCE services are running. In this case, an informative message is displayed in the UI, like "PCE is Unavailable."

Check Node and Cluster Status

This section describes several ways you can check the status of PCE nodes and clusters.

Check Node Environment

Run this command to examine the main PCE configuration file `runtime_env.yml` and validate it for syntax and basic structure:

```
$ sudo -u ilo-pce illumio-pce-env check
```

Check PCE Node Status

Run this command to display the status of the PCE node:

```
$ sudo -u ilo-pce illumio-pce-ctl status
```

Node Status Codes:

- 0 - Stopped
- 1 - All required processes running
- 2 - Partial, not all required processes running

For example, when you run the following status command (with semicolon) and echo `$?`, you receive the following output:

```
$ sudo -u ilo-pce illumio-pce-ctl status; echo $?
Checking Illumio Runtime                RUNNING 0.29s
1
```

To see the PCE node status with standard Linux statuses, you have two options:

Run the status command with the `--stdexit` option to see the following node status:

- 0 - Running
- 1 - Running at runlevel 1
- 2 - Error
- 3 - Stopped

For example:


```
$ sudo -u ilo-pce illumio-pce-ctl status --stdexit
```

Run the PCE service script, which calls the `illumio-pce-ctl` command and provides standard Linux status codes.

For example:

```
$ service illumio-pce status
```

NOTE:

Running the service script to retrieve status automatically returns the `--stdexit` status values. However, running the `service illumio-pce ctl status` command does not insert the `--stdexit` option.

Check Services on a PCE Node

Run the following command and the `-v` (verbose) option to display the status of individual services on a PCE node:

```
$ sudo -u ilo-pce illumio-pce-ctl status -v
```

Example output:

```
$ sudo -u ilo-pce illumio-pce-ctl status -v
```

```
Checking Illumio Runtime csaefh iimntttttt RUNNING 0.75s
```

The colored string represents the status of the PCE services as described by the following table. Use the characters to determine whether services are in the steady state.

For more information about the services, enter `status -s`.

Character	Service
a	Agent background worker
c	PCE web console
e	Event service
f	Fluentd
h	HAproxy
i	ilo_monitor or ilocron, in that order
m	memcached

Character	Service
n	nginx
s	Console discovery
t	Various “thin” services

Check PCE Cluster Status

Run this command to display the PCE cluster status:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

For example:

```
$ sudo -u ilo-pc illumio-pce-ctl cluster-status
Reading /var/illumio-pce-data/runtime_env.yml.

SERVICES (runlevel: 5)          NODES (Reachable: 4 of 4)
=====
agent_service                  10.6.31.18 10.6.31.17
agent_traffic_redis_cache     10.6.31.20 10.6.31.19
agent_traffic_redis_server    10.6.31.20
agent_traffic_service         10.6.31.18 10.6.31.17
auditable_events_service     10.6.31.18 10.6.31.17
collector_service             10.6.31.18 10.6.31.18 10.6.31.17 10.6.31.17
database_service              10.6.31.20
database_slave_service        10.6.31.19
ev_service                    10.6.31.18 10.6.31.17
executor_service              10.6.31.18 10.6.31.17
fileserver_service            10.6.31.20
fluentd_source_service        10.6.31.17 10.6.31.18
login_service                 10.6.31.18 10.6.31.17
memcached                     10.6.31.17 10.6.31.18
node_monitor                  10.6.31.18 10.6.31.18 10.6.31.17 10.6.31.17
pg_listener_service           10.6.31.20
search_index_service          10.6.31.17 10.6.31.18
server_load_balancer          10.6.31.17 10.6.31.18
service_discovery_agent       10.6.31.31
service_discovery_server      10.6.31.19 10.6.31.20 10.6.31.32
set_server_redis_server       10.6.31.19
```

```
Cluster status: RUNNING
```

Check PCE Version

Run this command to display the version of the installed PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl version
```

Check PCE Cluster Members

Run this command to display the members of the PCE cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-members
```

Update PCE Configuration

This section describes how to change the configuration of a PCE at any time after the initial configuration is set during PCE installation.

Back up PCE Runtime File

Store a copy of each node's `runtime_env.yml` file on a system that is not part of the Supercluster. The default location of the PCE Runtime Environment File is `/etc/illumio-pce/runtime_env.yml`.

Update Runtime Configuration

Update the `runtime_env.yml` file with the configuration changes.

Run the following command to validate the `runtime_env.yml` file:

```
$ sudo -u ilo-pce illumio-pce-env check
```

Run the following command to restart the node with the configuration changes:

```
$ sudo -u ilo-pce illumio-pce-ctl restart
```

Get Current PCE Runlevel

When you first install the PCE software and start the PCE application, the runlevel is set to 1 by default. At runlevel 1, only the database services are running. This setting allows you to set up the database before the entire PCE application starts running.

Runlevel 1 is also used for upgrading the PCE software. When upgrade the PCE, you need to set the PCE runlevel to 1 before you migrate the PCE database. After database migration finishes, you can set the PCE runlevel back to 5 to start the entire PCE application.

When the PCE software is already at runlevel 5, setting the runlevel to 1 takes effect the next time the software is started.

For more information about upgrading the PCE software, see the *PCE Installation and Upgrade Guide*.

Run this command to display the current Illumio PCE runlevel:

```
$ sudo -u ilo-pce illumio-pce-ctl get-runlevel
```

Set PCE Runlevel

Run this command to start the PCE cluster at one of the following runlevels:

- Runlevel 1, which only starts the PCE database
- Runlevel 5, which starts the entire PCE cluster

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel [1 or 5]
```

Update PCE Certificates

Whenever the PCE certificates are updated, you must obtain the new certificate and update it on all PCE nodes. Use the following steps.

1. Obtain the new certificate. The certificate must meet certificate requirements described in the *PCE Installation and Upgrade Guide*.
2. Stop *all nodes* in your deployment:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

3. On *every node*, load the certificate into the correct directory.

For example:

```
/var/lib/illumio_pce/cert
```

4. When the name of the new certificate is different from the name of the old certificate, update the file names in your `runtime_env.yml` file on *every node*.
5. On *every node*, validate the certificate:

```
$ sudo -u ilo-pce illumio-pce-env check
```

6. Start *all nodes* in your deployment:

```
$ sudo -u ilo-pce illumio-pce-ctl start
```

Change the PCE FQDN

WARNING: Before starting this process, add or generate another certificate with a new FQDN. If you skip this step, your cluster will stay down with old certificates.

You can change the fully-qualified domain name (FQDN) of a PCE as long as the PCE is not part of a Supercluster.

1. On *any node*, shut down all PCE nodes:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-stop
```

2. Open the file `runtime_env.yml`.
3. Modify the parameter `pce_fqdn` and save the file.
4. Validate the `runtime_env.yml` file:

```
$ sudo -u ilo-pce illumio-pce-env check
```

5. On *any node*, restart the PCE:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-restart
```

Workloads that were paired with the old FQDN automatically detect and pair with the new FQDN as long as the PCE was stopped long enough for each VEN to attempt and fail at least one heartbeat.

Upgrade the OS on a Running PCE

You can upgrade the operating system on a running PCE cluster without stopping the entire cluster. Isolate one node at a time, wipe its disk, and install the new operating system while the other nodes in the PCE cluster continue to operate. The PCE can function with a mix of operating system versions on the different nodes.

Use this procedure when upgrading from one operating system version to another. If you are merely installing an operating system patch, you do not need to wipe the disk.

The general steps are as follows:

1. Back up the PCE databases.
2. Remove one node from the cluster.
3. Wipe the disk and install the new operating system version.
4. Install and configure the PCE software.
5. Restore the node to the cluster.
6. Repeat this procedure for the other nodes in the PCE cluster.

Back Up the PCE

1. Back up the PCE policy and traffic databases and `runtime_env.yml` file. Follow the steps in [PCE Database Backup](#). For a Supercluster, follow the steps in [Back Up Supercluster](#) in the *PCE Supercluster Deployment Guide*.
2. Save a copy of the PCE certificate in a safe location (not on the PCE node). Take note of the directory path where the certificate was stored. You will need to replace the certificate in the same location later.
3. Save a copy of the private key in a safe location. Take note of the directory path where the key file was stored. You will need to replace the key in the same location later.

Remove a Node From the Cluster

Remove one node from the PCE cluster so you can update its operating system. The cluster will continue to operate using the remaining nodes.

Remove and upgrade the nodes in this order:

- Core nodes
- Replica data node
- Primary data node

CAUTION:

Remove and upgrade the policy database primary data node last to avoid unnecessary failover. To find the primary data node, run the following command on any node in the PCE cluster:

```
$ sudo -u ilo-pce illumio-pce-db-management show-master
```

1. Verify that the cluster is running and healthy. If you remove a node from a PCE that is not in a healthy state, it can cause downtime. There are several ways to check the health of the PCE cluster; see [Monitor PCE Health](#).

One way to check PCE health is to run the following command:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

2. On the node that is to be removed, stop the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

Stopping the PCE software causes PCE services to fail over to their backup node.

3. Check to be sure the PCE node is stopped.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

Expected output:

```
Checking Illumio Runtime                STOPPED 1.76s
```

4. When you are removing the *leader node*, wait until the PCE has promoted another node to the leader before proceeding. Run the following command to determine the new leader node:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leader
```

5. On the *leader node*, run the following command to be sure the data nodes are synchronized.

CAUTION:
To avoid data loss, the data nodes must be synchronized before removing the node from the PCE cluster. Be sure the output from this command shows that the nodes are synchronized.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

Expected output is similar to the following:

```
Reading /etc/illumio-pce/runtime_env.yml.
SERVICES (runlevel: 5)          NODES (Reachable: 3 of 4)
=====
agent_background_worker_service 192.0.2.241    192.0.2.242
agent_service                   192.0.2.241    192.0.2.242
agent_traffic_redis_cache       192.0.2.240
agent_traffic_redis_server      192.0.2.240
agent_traffic_service           192.0.2.241    192.0.2.241
192.0.2.242    192.0.2.242
app_gateway_service            192.0.2.240    192.0.2.241
192.0.2.242
auditable_events_service       192.0.2.241    192.0.2.242
citus_coordinator_replica_service NOT RUNNING
citus_coordinator_service      192.0.2.240
cluster_management_service     192.0.2.241    192.0.2.241
collector_service              192.0.2.241    192.0.2.241
192.0.2.242    192.0.2.242
data_job_queue_redis_replica_service NOT RUNNING
data_job_queue_redis_service   192.0.2.240
data_job_queue_service         192.0.2.241    192.0.2.241
192.0.2.242    192.0.2.242
database_monitor               192.0.2.240
database_service               192.0.2.240
```



```

database_slave_service          NOT RUNNING
db_cache_manager_service        192.0.2.240
ev_service                      192.0.2.241      192.0.2.242
events_background_worker_service 192.0.2.241      192.0.2.242
executor_service                192.0.2.241      192.0.2.242
fileserver_service              192.0.2.240
fileserver_slave_service        NOT RUNNING
flow_analytics_monitor_service   192.0.2.240
flow_analytics_service          192.0.2.240      192.0.2.240
fluentd_data_service            192.0.2.240
fluentd_source_service          192.0.2.241      192.0.2.242
fluentd_sys_event_fwd_service   192.0.2.240      192.0.2.241
192.0.2.242
login_service                   192.0.2.241      192.0.2.242
memcached                       192.0.2.241      192.0.2.242
network_device_service          192.0.2.241      192.0.2.242
node_monitor                    192.0.2.240      192.0.2.241
192.0.2.242
report_generator_service         192.0.2.241      192.0.2.242
report_monitor_service          192.0.2.240
reporting_database_monitor       192.0.2.240
reporting_database_replica_service NOT RUNNING
reporting_database_service       192.0.2.240
reporting_etl_service            192.0.2.241
reporting_management_service     192.0.2.241      192.0.2.242
search_index_service            192.0.2.241      192.0.2.242
server_load_balancer            192.0.2.241      192.0.2.242
service_discovery_agent         NOT RUNNING
service_discovery_server        192.0.2.240      192.0.2.241
192.0.2.242
set_server_redis_server         192.0.2.240
traffic_database_monitor         192.0.2.240
traffic_query_service            192.0.2.240
traffic_worker_service           192.0.2.241      192.0.2.241
192.0.2.242      192.0.2.242
web_server                      192.0.2.241      192.0.2.242
    
```

```
Cluster status: RUNNING
```

6. Wait until the cluster status has returned to RUNNING.
7. On the *leader node*, remove the node. For *ip_address*, substitute the IP address of the node you are removing:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leave ip_address
```

Expected output:

```
Removed node successfully.
```

8. Check the status of the PCE again to confirm it is still running normally:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

Expected output is similar to that shown in step 5.

Remove OS and Install New

Remove the old operating system version. Then install the new version. Use the documentation provided by your operating system vendor.

Reinstall the PCE

1. Install the PCE software and configure its runtime parameters. Follow the steps in the [PCE Installation](#) section in the *PCE Installation and Upgrade Guide*.

IMPORTANT:
Do not start the PCE yet.

- Be sure the PCE FQDN (hostname) is the same as before the upgrade.
- Be sure the and IP addresses for all NICs are the same as before the upgrade.
- Set up NTP and IPTables as described in [OS Setup and Package Dependencies](#).

Restore PCE Files

1. Copy the `runtime_env.yml` file to the same location where it was before.
2. Replace the certificate and key files in the same directory path where they were before.
3. Compare the certificate and key file locations to the specified locations in the `runtime_env.yml` file to be sure they match.

Restore Node to Cluster

Restore the node to the cluster.

1. On the node where you just upgraded the OS, run the following command. For `ip_address`, substitute the IP address of any running node in the PCE cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-join ip_address
```

After the node successfully joins the PCE cluster, the PCE software is started.

2. Verify that the cluster is functional and data has been synchronized to all data nodes.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status -w
```

Wait until this command returns output that shows all services are running. The output concludes with this line:

```
Cluster status: RUNNING
```

Upgrade and Restore Remaining Nodes

Repeat this procedure for the other nodes in the PCE cluster. Reminder: Upgrade the primary database node last.

Firewall Coexistence

To provide additional security, you can supplement Illumio's firewall with your organization's firewalls using Firewall Coexistence. The Illumio firewall can be set to either **Exclusive** mode or **Coexistence** mode via the PCE web console or the Illumio REST API. In both modes, the Illumio firewall is always separate from other firewalls.

IMPORTANT:

The Firewall Coexistence feature deprecates the following features:

- Windows FAS VEN coexistence
- Linux VEN NAT ignore
- Linux VEN container mode

Firewall Tampering Protection

- *When coexistence is turned on in primary or secondary mode*

The VEN only monitors its own firewall rules against tampering. When the VEN detects tampering of Illumio firewall rules, an alert is raised, and the VEN reconfigures its firewall rules to its pre-tampered state in order to protect the workload. You can program non-Illumio rules in any table without generating any tampering alerts.

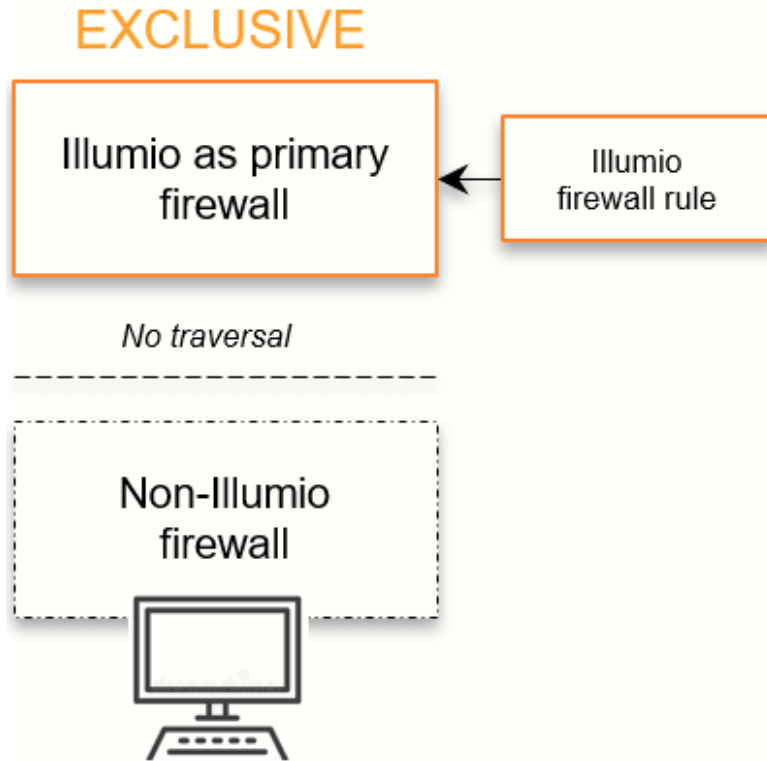
- *When coexistence is turned on in primary mode*

The VEN also monitors that the Illumio rule in the main tables “stay on the top” when you choose Illumio to be the primary firewall. When the VEN detects that the Illumio rule is not on the top, an alert is raised, and the VEN moves the Illumio rule back to the top.

Firewall Coexistence Modes

Exclusive Mode

The default mode is Exclusive, in which Illumio is the only firewall. In this mode, any non-Illumio firewall is not traversed. This behavior applies to all tables in iptables, such as filter, NAT, Raw, or Mangle.




Coexistence Mode

With a set of labels and policy states, you can enable Firewall Coexistence for a set of workloads. You can configure coexistence in two ways:

- A configuration in which Illumio is the primary firewall.
- A configuration in which Illumio is *not* the primary firewall.

NOTE:

The Coexistence mode applies to all tables of the Linux firewall. Therefore, when you have your own NAT rules and use the “ignore nat” option on the VEN with the Coexistence mode, you do not see any change in behavior.



Prerequisites and Recommendations

This release of the Firewall Coexistence feature requires that you upgrade the VEN to 18.3.1 or later. The older versions of Illumio Firewall Coexistence are deprecated.

Windows VEN version 18.3.x ignores the older `limited_wfas_coexistence` and `full_wfas_coexistence` VEN settings for coexistence located in the VEN `runtime_env.yml` file. Linux VEN version 18.3.x ignores settings in `/etc/default/illumio-agent` for NAT table coexistence (container mode).

The following upgrade sequence is required. You must upgrade the VEN last and only after configuring firewall coexistence in the PCE:

Recommended Firewall Setting

For better security, Illumio strongly recommends setting the Illumio firewall as the primary firewall.

When you select Illumio to be the primary firewall, the VEN ensures that the Illumio rule in the main tables “stay on the top” only when you choose Illumio to be the primary firewall. The VEN does not enforce the Illumio rules to be on the top when Illumio is not the primary firewall. This behavior applies to all tables in iptables, such as filter, NAT, Raw, or Mangle.

When the Illumio firewall is set as primary, non-Illumio firewalls are traversed only when the Illumio firewall rules allow the traversal, in which case, packets are passed to non-Illumio firewalls.

IMPORTANT:

When the Illumio firewall is not set as primary, packets passed by non-Illumio firewalls are seen by the Illumio firewall; however, packets accepted by the non-Illumio firewall are not seen by the Illumio firewall.

Example

When the Illumio firewall is not set as primary, and the non-Illumio firewall logs and accepts all traffic on port 22, the Illumio firewall does not see the traffic on port 22.

When packets are allowed by the Illumio firewall, they are passed to other firewalls. Illumio's firewall does not monitor packets dropped by other firewalls. Packets dropped by the Illumio firewall are not passed to non-Illumio firewalls.

Set Firewall Coexistence

You can set firewall coexistence using either interface:

- PCE web console
- Illumio REST API

To view firewall coexistence settings in the PCE web console:

From the PCE web console menu, choose **Settings > Security > Firewall Coexistence**.

The PCE web console displays the following settings:

- **Default:** Illumio Core is the exclusive firewall by default. You can configure firewall coexistence as needed for all workloads and/or specific labels.
- **Firewall Coexistence:**

To add the scope for firewall coexistence:

1. Click **Add**.



The screenshot shows a dialog box titled "Add Scope for Firewall Coexistence". It contains three rows of configuration options:

- Scope:** A dropdown menu with the text "Select Labels" and a refresh icon to its right.
- Enforcement:** A dropdown menu with the text "All".
- Illumio Core is Primary Firewall:** A dropdown menu with the text "Yes".

At the bottom right of the dialog are two buttons: "Cancel" and "Add".

Start adding and configuring the Scope, Enforcement, and whether it is the Primary Firewall.

2. From the Scope drop-down list, select the labels.
3. From the Enforcement drop-down list, select All, Enforced, or Illuminated.
4. In the Illumio Core is Primary Firewall, select either Yes or No.
5. Once the selections are made, click on **Add**.

PCE Listen Only Mode

This section describes how to use Listen Only mode when you want to temporarily stop the PCE from sending policy updates to your VENs.

About PCE Listen Only Mode

Enabling Listen Only mode for the PCE is typically used in these situations:

- During PCE maintenance windows, such as PCE backup or maintenance on parts of your network.
- After restoring the PCE from a backup. See [PCE Database Backup](#) for information.

In Listen Only mode, VENs still report updated workload information to the PCE; however, the PCE does not modify the firewall rules on any workloads or send any updates to the VENs. The PCE does not mark workloads as offline or remove them from policy when Listen Only mode is enabled.

When this mode is enabled, you can still write policy, pair new workloads, provision policy changes, assign or change workload labels; however, changes are not be sent

to the VENs until you disable Listen Only mode. You can disable Listen Only mode when you are ready to resume normal policy operations.

Enable PCE Listen Only Mode

1. On *all nodes* in the cluster, stop the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

2. Set *all nodes* in the PCE cluster at runlevel 1:

```
$ sudo -u ilo-pce illumio-pce-ctl start --runlevel 1
```

3. On *any node* in the cluster, enable Listen Only mode:

```
$ sudo -u ilo-pce illumio-pce-ctl listen-only-mode enable
```

4. Set the PCE runlevel to 5:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 5
```

Determine if PCE Is in Listen Only Mode

On a *data node* in the cluster, determine whether the PCE is in Listen Only mode :

```
$ sudo -u ilo-pce illumio-pce-ctl listen-only-mode status
```

Additionally, when the PCE is in Listen Only mode, the PCE web console displays a banner that indicates how long the PCE has been in Listen Only mode.

When Listen Only mode is enabled, the Workloads list page and Workload detail pages indicate the VEN connectivity status is **Syncing** and Policy Sync is **Verified**.

After you disable Listen Only mode and set the PCE runlevel to 5, the PCE receives each VEN's heartbeat and begins applying any changes. After the changes have been synchronized, the VEN connectivity status is **Online** and Policy Sync is **Active**.

VEN Heartbeat and Listen Only Mode

Before you disable Listen Only mode, determine whether your VENs sent recent heartbeats to the PCE while Listen Only mode was enabled. When a VEN hasn't sent a

heartbeat to the PCE within the last hour, the PCE will remove that VEN from policy after you disable Listen Only mode. Large numbers of VENs that haven't heartbeat with the PCE might indicate a problem in the environment that is preventing the VENs from communicating with the PCE. To prevent large numbers of workloads from being marked as offline and removed from policy, investigate and resolve any problems before disabling Listen Only mode.

To determine a VEN's most recent heartbeat, use the Illumio Core REST API. Use the Workloads API with the `last_heartbeat_on` property to GET a workload collection or individual workload.

Examples:

```
GET [api_version][org_href]/workloads
```

```
GET [api_version][workload_href]
```

To determine the last heartbeat for each workload, check the `last_heartbeat_on` property in the agent section (the REST API name for the VEN) of the response.

```
},
"agent": {
"status": {
"last_heartbeat_on": "2017-11-30T01:30:04.734Z",
...
}
},
```

Additionally, use the REST API to query workloads for a VEN heartbeat time that occurred *before* you enabled PCE Listen Only mode. Before you disable Listen Only mode, investigate any workloads with a heartbeat timestamp prior to when you enabled it.

See [Workload Operations](#) in the *REST API Developer Guide* for more information.

Query Parameters

Parameter	Description	Data Type	Required
<code>last_heartbeat_on [lte]</code>	Allows you to search for workloads whose last heartbeat occurred before a specific time.	String (timestamp_in_rfc3339)	No

Parameter	Description	Data Type	Required
	lte: Less than or equal to.		
last_heartbeat_on [gte]	Allows you to search for workloads whose last heartbeat occurred after a specific time. gte: Greater than or equal to.	String (timestamp_in_rfc3339)	No

Example Query

You enabled PCE Listen Only mode on February 23, 2020 at 7:20 PM. Use the following query parameter to return only those workloads whose last heartbeat occurred before this time. Any workloads that are returned should be checked for connectivity before you disable Listen Only mode.

```
GET [api_version][org_href]/workloads?last_heartbeat_on[lte]=2020-02-23T19:20:29+02:00
```

Disable PCE Listen Only Mode

NOTE:

You must run the command to disable PCE Listen Only mode at runlevel 1 or 5.

1. From *one of the data nodes*, disable Listen Only node:

```
$ sudo -u ilo-pce illumio-pce-ctl listen-only-mode disable
```

2. Verify that PCE Listen Only mode is disabled:

```
$ sudo -u ilo-pce illumio-pce-ctl listen-only-mode status
```

Expand 2x2 Cluster to 4x2

This section describes how to expand an existing PCE 2x2 cluster to a 4x2 cluster by adding two core nodes.

Prepare Environment for Cluster Expansion

This section helps you prepare your PCE cluster environment for the new core nodes.

Prepare Server Load Balancer or DNS

Add the new core node information for a server load balancer (SLB) or DNS:

- **Server load balancer (SLB)**

Before installing the PCE software on the two new core nodes, perform the following tasks:

- Add the IP addresses of the two new nodes to your load balancer configuration.
- Configure your load balancer to check the health of the new core nodes.
- Run a health check and verify that the two new core nodes are down.
- Verify that traffic is *not* being forwarded to the new nodes.

- **DNS**

Perform the following tasks:

- Add the two new nodes to your DNS configuration.
- When TCP connectivity from the VENs to the PCE is direct and not routed through a virtual IP (VIP), modify the `runtime_env.yml` on all four nodes in the existing cluster and change the `cluster_public_ip > cluster_fqdn` to include the two new core nodes.

Define this parameter as a list of IP addresses that the VENs can connect to, which is the load balancing VIP or a list of all core nodes in the cluster.

For example:

```
cluster_public_ips:
  cluster_fqdn:
    - <existing_core_node_ip_address>
    - <existing_core_node_ip_address>
    - <new_core_ip_node_address>
    - <new_core_ip_node_address>
```

Ensure Connectivity from VENs to New Nodes

Ensure that connectivity from existing VENs to the new core nodes is allowed and working; for example, you might need to update your network's firewall policies to permit access from existing VENs to the new core nodes.

Prepare the Cluster for New Nodes

Before you install the PCE software on the new core nodes, perform the following tasks.

1. Stop the cluster by running this command:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

2. Validate the cluster's configuration by running this command:

```
$ sudo -u ilo-pce illumio-pce-ctl check-env
```

3. Start the cluster by running this command:

```
$ sudo -u ilo-pce illumio-pce-ctl start
```

The PCE configures all VENs to include access to the new core nodes. When complete, all your VENs should be listed as online.

Back Up PCE Database

Before you expand your 2x2 cluster, create a backup of your PCE database. See [Back Up the PCE](#) for information.

Configure Existing Nodes for Expansion

1. On *all nodes* in the existing cluster, stop the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

2. Before you modify the `runtime_env.yml` file on the existing nodes, create a file backup in case you need to revert back to the last known configuration.

For example, on *all nodes*, run this command:

```
cp /etc/illumio-pce/runtime_env.yml /etc/illumio-pce/runtime_env.yml.bak
```

3. Modify both new core nodes' `runtime_env.yml` file so that the `node_type` parameter is defined as `core`. For example, change the parameter from `core0` or `core1` to `core`.

4. On *all nodes*, modify the `runtime_env.yml` file to define the `cluster_type` parameter as `6node_v0` and save the file. Your `runtime_env.yml` file might not have this parameter; you only need to add it when it does not already exist.

For example:

```
cluster_type: 6node_v0
```

5. On *all nodes* in the existing cluster, check the syntax of the `runtime_env.yml` configuration:

```
$ sudo -u ilo-pce illumio-pce-env check
```

6. On *all nodes* in the existing cluster, restart the PCE with the configuration changes:

```
$ sudo -u ilo-pce install_root/illumio-pce-ctl restart
```

7. On *any node* in the cluster, check the cluster status:

```
$ sudo -u ilo-pce install_root/illumio-pce-ctl cluster-status
```

The status of the cluster should return as `RUNNING`.

Install and Configure PCE on Nodes

Install the PCE software and configure the new core nodes using the same RPM used to install the existing nodes, and use the same system and environmental configuration as the existing two core nodes. This configuration includes all `runtime_env.yml` settings, kernel performance modifications, syslog configurations, DNS, and NTP. See [PCE Installation](#) for information.

CAUTION:

Use the same RPM you used to install the existing PCE nodes to install the PCE software on the new nodes.

After you have installed the PCE software, perform these steps:

1. For layer 4 load balancer implementations, confirm that two of the core nodes are present and UP on the load balancer. These nodes should match with those

shown in `cluster-status` with the role of `server_load_balancer`. When nodes in the cluster fail, the nodes that own the `server_load_balancer` role can change.

2. Ensure that the TLS certificate is valid for the new nodes as well as the existing nodes. The certificate might contain only the cluster name, or might include each of the core node names in the SAN field. When the SAN field is used, ensure that both of the new core nodes are included.
3. Copy the certificate and key from the existing core nodes to the new core nodes in `/var/lib/illumio-pce/cert` (or wherever you defined this location in the `runtime_env.yml` file).
4. Copy the `runtime_env.yml` file from an existing core node to the new core nodes. Ensure that when nodes have a specific configuration, such as `internal_service_ip`, you configure this parameter on the new core nodes to correctly reflect the configuration on the two new nodes.
5. Verify that the new nodes have the correct `node_type` (`core`) and `cluster_type` (`6node_v0`) and, when using a DNS load balancer, verify that all four core nodes are defined in the runtime parameter named `cluster_public_ips > cluster_fqdn`.
6. On *all new core nodes*, verify that the new core nodes were configured correctly:

```
$ sudo -u ilo-pce illumio-pce-ctl check-env
```

7. Find the IP address of the cluster leader node:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leader
```

8. On any existing node in the cluster (not the new node you are about to add), run the following command. For `ip_address`, substitute the IP address of the first new node.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-nodes allow ip_address
```

9. On the *first new node*, insert the first new core node into the cluster. Use the cluster leader node IP address that you found in the earlier step.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-join ip_address_of_Leader_node
```

This command should confirm the node is added and report that there are 5 nodes in the cluster.

10. On any existing node in the cluster (not the second new node you are about to add), run the following command. For *ip_address*, substitute the IP address of the second new node.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-nodes allow ip_address
```

11. On the *second new node*, insert the second new core node into the cluster:

```
sudo -u ilo-pce illumio-pce-ctl cluster-join ip_address_of_Leader_node
```

This command should confirm the node is added and report that there are 6 nodes in the cluster.

12. On *all nodes*, restart the PCE software with the configuration changes:

```
$ sudo -u ilo-pce illumio-pce-ctl restart
```

Verify Cluster Expansion

Perform these steps to ensure that you have successfully expanded your PCE 2x2 to a 4x2 cluster.

1. To verify that the cluster is fully up and running and all PCE services are at run-level 5, run the status command:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

2. Confirm that the cluster contains 6 nodes:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-members
```

3. When you are using a server load balancer to manager PCE traffic, confirm on the load balancer that two of the core nodes are present and listed as UP. These nodes should match those shown from the `cluster-status` command with the role of `server_load_balancer`. When nodes in the cluster fail, the nodes that own the `server_load_balancer` role can change.

4. Verify that you can log into the PCE web console and navigate the interface successfully.
5. Verify that logs are being populated in the logging directory of the new nodes, and (when configured) logs are being forwarded to external log destinations.
6. Verify that your workload VENS are online in the Workloads page of the PCE web console. Be aware that VENS might be offline occasionally for unrelated reasons; therefore, compare the VEN connectivity status to your baseline.

NOTE:

Large numbers of VENS remaining in Syncing state can indicate that one of the core nodes is not reachable due to a network firewall, load balancer, or `runtime_env.yml` misconfiguration.

Replace PCE Nodes or Uninstall Cluster

This section describes how to add a new node to take the place of one that has failed. It also describes how to uninstall the PCE.

NOTE:

You can replace only one PCE node at a time.

Replace a Failed Node

1. Determine which node is the cluster leader:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leader
```

2. On the *cluster leader node*, remove the failed node:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leave ip_address
```

Where `ip_address` is the IP address of the failed node.

3. Before adding the new replacement node, ensure that:
 - The new node has a valid `runtime_env.yml` file configured.
 - The PCE software is not running.
4. On any existing node in the cluster (not the new node you are about to add), run the following command. For `ip_address`, substitute the IP address of the new

node.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-nodes allow ip_address
```

5. On the *new node*, run the following command to add the new node to the cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-join ip_address
```

Where *ip_address* is the IP address of any existing running node within the cluster.

After the new node successfully joins the PCE cluster, the PCE software is started.

Replace a Running Node

Perform this procedure to take offline or replace a running node in the cluster; for example, when you need to upgrade the host hardware.

NOTE:

Performing these steps on a *data node* can result in the loss of your Illumination data and existing VEN Support Reports.

1. Stop the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

Stopping the PCE software causes PCE services to fail over to their backup node.

2. Wait for the node to enter the FAILED state. To check this status, run the following command on any other node:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-members
```

3. When you are removing the *leader node*, wait until the PCE has promoted another node to the leader before proceeding. Run the following command to determine the new leader node:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leader
```

4. On the *leader node*, remove the failed node:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leave ip_address
```

5. Before adding the new replacement node, ensure that:
 - The node has a valid `runtime_env.yml` file configured.
 - The PCE system software is not running.
6. On any existing node in the cluster (not the new node you are about to add), run the following command. For *ip_address*, substitute the IP address of the new node.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-nodes allow ip_address
```

7. On the *new node*, run the following command to add the new node to the cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-join ip_address
```

Where *ip_address* is the IP address of any existing running node within the cluster.

After the new node successfully joins the PCE cluster, the PCE software is started.

Uninstall the PCE Cluster

To completely uninstall and remove the PCE for your system, perform the following steps:

1. Remove the PCE UI package:

```
$ rpm -e illumio-pce-ui
```

2. Remove the main PCE package:

```
$ rpm -e illumio-pce
```

3. Manually delete these directories:

```
/var/lib/illumio-pce  
/var/log/illumio-pce  
/etc/illumio-pce
```

PCE Database Management

This chapter contains the following topics:

About the PCE Databases	53
PCE Database Backup	57
Database Migration, Failover, and Restore	62
Manage Multi-Node Traffic Database	66
PCE Default Object Limits	70

This section describes how to manage the PCE databases, backups, failover and restore.

About the PCE Databases

This section describes concepts you need to know to successfully administer the PCE databases.

Policy and Traffic Data Databases

The PCE uses two databases: one for policies and the other for traffic flow data. Both databases need to be backed up or restored.

Database	Summary of Command	Notes
Policy	<code>illumio-pce-db-management dump --file <i>backup_filename</i></code>	Backs up the policy database.
Traffic	<code>illumio-pce-db-management traffic dump --file <i>traffic_backup_filename</i></code>	Back up the traffic database by adding the traffic parameter.

Data Retention of Traffic Flow Summaries

The PCE removes traffic flow data summaries (used by the Explore features in the PCE web console) when these conditions occur:

- The disk size of the traffic flow summaries exceeds the disk space allocated for the data. See [PCE Capacity Planning](#) in the *PCE Installation and Upgrade Guide* for information.
- The traffic data database has been inactive for 90 days.

When FlowLink is used, the following limits apply on traffic data:

- The default storage limit on traffic data from all of an organization's FlowLink servers is 500MB.
- The default storage size limit is based on the number of server VENS, endpoints, and container VENS. Kubelink flows (from container VENS) are grouped with server and endpoint flows.
- When the storage limit or the 90-day limit is reached, traffic flow data is pruned. The order of pruning is first data from endpoints, then Kubelink, and lastly Server VENS.

Determine the Primary Database

Policy Database

Run the following command to determine the primary policy database:

```
sudo -u ilo-pce illumio-pce-db-management show-master
```

Traffic Database

Run the following command to determine the primary traffic database:

```
sudo -u ilo-pce illumio-pce-db-management traffic show-master
```

Show Database Replication Information

Run the following command to view information about data replication between the primary and replica databases:

```
sudo -u ilo-pce illumio-pce-db-management show-replication-info
```

Rotate Database Passwords and Other Secrets

At any time, an Illumio Administrator can rotate the PCE database passwords and other auto-generated secrets used within the PCE. The new secrets take effect when the PCE is restarted. To rotate secrets, run the following command on any node:

```
sudo -u ilo-pce illumio-pce-ctl rotate-secrets
```

In a Supercluster, run this command once for each region.

Anonymize Database Export

You can anonymize the database dump file to protect confidential data before sending it to Illumio Customer Support for troubleshooting purposes. You can safely share policy and configuration data with Illumio for support requests. Sensitive data, such as usernames, passwords, and IP addresses, are masked.

1. Dump the policy or traffic database by running one of the following commands.

Policy database

```
sudo -u ilo-pce /opt/illumio_pce/illumio-pce-db-management dump --file  
backup_filename
```

Traffic database

```
sudo -u ilo-pce /opt/illumio_pce/illumio-pce-db-management traffic dump --  
for-masking --file traffic_backup_filename
```

2. Anonymize the policy or traffic dump file by running one of the following commands.

Policy dump file

```
sudo -u ilo-pce /opt/illumio_pce/illumio-pce-db-management mask-db-dump --in-  
file backup_filename --out-file masked_filename --dict-file dictionary.txt --  
tmpdir path_to_alternate_tmp_dir;
```

Traffic dump file (add the `--traffic` flag)

```
sudo -u ilo-pce /opt/illumio_pce/illumio-pce-db-management mask-db-dump --  
traffic --in-file backup_filename --out-file masked_filename --dict-file  
dictionary.txt --tmpdir path_to_alternate_tmp_dir;
```

Optional `--tmpdir` parameter

The `/tmp` directory stores intermediate files and can sometimes run out of space. Use `--tmpdir` to specify an alternate temporary directory with adequate space.

Example command output

```
Dictionary file /home/pce/dictionary.txt will be created
Reading /home/pce/backup.july.11.2019.tar.bz2
Processing avenger_fileserver_dev.sql
Processing avenger_executor_dev.sql
Processing avenger_ops_dev.sql
Processing avenger_events_dev.sql
Processing avenger_agent_dev.sql
Processing avenger_login_dev.sql
Processing dump-info
Processing avenger_node.uuid
Processing avenger_cluster.uuid
Writing /home/pce/masked_backup.july.11.2019.tar.bz2
Writing dictionary file /home/pce/dictionary.txt
Done
```

3. Send the anonymized output file named in `--out-file` to Illumio Customer Support.

CAUTION:

Do not send the dictionary file to Illumio (`dictionary.txt` in the command above). Retain it at your own site. It contains the mapping from the unmasked data to the masked data.

Illumio recommends consistently using the same dictionary file. This approach ensures that the same value is consistently masked and you can compare changes between different masked database dumps.

View Events Using PCE Command Line

You can view events using the PCE command line. For more details about viewing events, see [View and Export Events](#).

Run the following command at any runlevel to display:

- The total number of events
- The average number of events per day

```
sudo -u ilo-pce illumio-pce-db-management events-db events-db-show
```

Run the following command at any runlevel to display:

- The amount of disk space used by events
- The total number of events
- The disk usage based on type of event

```
sudo -u ilo-pce illumio-pce-db-management events-db disk-usage-show
```

Example

```
illumio-pce-db-management events-db disk-usage-show
Reading /opt/pce_config/etc/runtime_env.yml.
INSTALL_ROOT=/var/illumio_pce
RENV=development
```

Events database disk usage summary:

Number of events: 6

Average number of events per day: 6

Total disk usage: 0.539 MB (565248.0 bytes)

Disk usage by event_type:

Event Type	Count	Disk Usage
system_task.prune_old_log_events	1	0.090 MB
user.login	1	0.090 MB
user.logout	1	0.090 MB
user.sign_in	1	0.090 MB
user.sign_out	2	0.180 MB

PCE Database Backup

This section provides step-by-step instructions for backing up the PCE databases. Before you start, be sure you understand the technical details of the two PCE databases; see [About the PCE Databases](#) for information.

NOTE:

The PCE runtime configuration file, `runtime_env.yml`, is not included in database backups. You must back up this important file separately. See [Back Up the PCE Runtime Environment File](#).

About PCE Database Backup

You use the PCE database command line utility `illumio-pce-db-management` to back up, migrate, manage failover, and restore the PCE databases.

IMPORTANT:

You must run the PCE database commands as the PCE runtime user `ilo-pce`

When to Back Up

Follow your organization's backup policies and procedures, including frequency (such as, hourly, daily, or weekly) and retention location (namely, offsite or on a system other than the PCE cluster nodes).

Illumio recommends backing up the PCE databases in the following situations:

- Before and after a PCE version upgrade
- After pairing a large number of VENs
- After updating a large number of workloads (such as, changing workload policy state or applying labels)
- After provisioning major policy changes
- After making major changes in your environment that affect workload information (such as, IP address changes)
- On-demand backups before performing the procedures in this guide

Back Up the Policy Database

Perform these steps to back up all PCE data, such as before upgrading the PCE.

1. (On an SNC, skip this step.) Before you back up the PCE, determine which data node is running the `agent_traffic_redis_server` service:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

You see the following output:

```
SERVICES (runlevel: 5) NODES (Reachable: 1 of 1)
=====
agent_background_worker_service 192.168.33.90
agent_service NOT RUNNING
```

```
agent_slony_service 192.168.33.90
agent_traffic_redis_cache 192.168.33.90
agent_traffic_redis_server 192.168.33.90      <=== run the dump command
from this node
agent_traffic_service NOT RUNNING
...
```

2. On the *data node* that is running the `agent_traffic_redis_server` service, run the following commands:

```
$ sudo -u ilo-pce illumio-pce-db-management dump --file <location-of-db-dump-
file>
$ sudo -u ilo-pce illumio-pce-db-management traffic dump --file <location-of-
traffic-dump-file>
```

In *location-of-db-dump-file* and *location-of-traffic-dump-file* enter a file name for the policy database dump and the traffic database dump files, respectively.

NOTE:

On an SNC, run these commands on the single node.

3. After the dump commands finish, copy the backup files to a fault-tolerant storage location.

Back Up the Traffic Database

The traffic database dump can be very large, depending on the traffic datastore size. Therefore, the Supercluster database dump on leader and member PCEs does not include the traffic database dump. The following procedure is provided to back up the traffic data separately.

NOTE:

If you have a multi-node traffic database, do not use this procedure for routine backups. In a multi-node traffic database, the procedure in this section is used only for the initial installation of the multi-node database or when adding or removing worker nodes. For routine backups in a multi-node traffic database, use `pgbackrest` instead. See [Using pgbackrest for Traffic Data Backups](#).

Perform these steps to back up the traffic database only. If you need to back up the traffic flow data, perform this procedure on every region; traffic flow information is unique to every (region) PCE.

1. On *any data node*, run the following command:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic dump --file <path_to_
traffic_backup_file.tar.gz>
```

In `path_to_traffic_backup_file.tar.gz`, include the filename extension `.tar.gz`.

2. After the command finishes, copy the backup file to a fault-tolerant storage location.

Using pgbackrest for Traffic Data Backups

Instead of using the built-in PCE backup commands, you can use the `pgbackrest` tool. For example, `pgbackrest` can be useful if you have dedicated storage for backups, such as NFS network shared storage. If you have a multi-node traffic database, you must use `pgbackrest` for backups to ensure adequate space and performance.

Hardware Requirements

A shared filesystem such as NFS mount which is mounted on all the PCE nodes is required for `pgbackrest` to work. Make sure the NFS disk has enough space to store multiple backups. Specify the root location of this mount with the `backup_root` key in the `runtime_env.yaml`, shown below in "Enabling `pgbackrest`."

The NFS mount can be used to store other data in addition to the traffic data. For example, it could store the policy database and `runtime_env.yaml` file. The NFS mount must be a solid-state drive (SSD) disk. Rotational disks cannot be used, because they are too slow for the amount of data involved.

To calculate the size of the NFS mount needed for a multi-node traffic database, use the following formula: Number of worker node pairs x 150 GB x number of days retained + storage needed when occasionally adding or removing a node, which is 400 GB x number of worker node pairs. Optionally, add the amount of storage needed for any additional uses, such as the policy database.

Enabling `pgbackrest`

To enable the `pgbackrest` tool, add the following commands to the server `runtime_env.yaml`, with your cluster values specified where needed:

```
traffic_datastore_backup_service:
  pgbackrest_enabled: true
  backup_destination_type: 'filesystem'
  backup_root: '<Location of NFS root>'
  backup_encryption_key: '<Location of file that contains the backup encryption
key>'
  max_full_backups: '<max number of full backups to retain>' # Defaults to 2
```

Back Up the Traffic Database (pgbackrest)

Use the following command to take a backup of the traffic database cluster. In a multi-node traffic database, you can run this command on any coordinator or worker node:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic cluster-backup
```

List Available Backups (pgbackrest)

Use the following command to get the list of backups available, in the order in which they were taken:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic cluster-backup-list
```

Restore a Backup (pgbackrest)

Use the following commands to restore data from a given backup. For *backupLabel*, substitute the label of the backup to restore:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 1
$ sudo -u ilo-pce illumio-pce-db-management traffic cluster-restore --backup-
label backupLabel
```

Back Up the PCE Runtime Environment File

The PCE runtime configuration file, `runtime_env.yml`, is not included in automatic PCE backups. You must manually back up this file to a secure location.

Store a copy of each node's `runtime_env.yml` file on a system that is not part of the PCE cluster. By default, the PCE Runtime Environment File is located at the following location on each node:

```
/etc/illumio-pce/runtime_env.yml
```

If the file is not found there, it has been moved to a custom location. To find the file, check the `ILLUMIO_RUNTIME_ENV` environment variable.

IMPORTANT:

The `runtime_env.yml` file contains sensitive information that should be kept secret, such as encryption keys. Take steps to ensure the confidentiality of this file.

Database Migration, Failover, and Restore

This section describes how to perform database management tasks.

Migrate PCE Databases

These steps explain how to migrate the database from a previous version to the current version. You must run this command at runlevel 1 in the following cases:

- After you have upgraded to a newer version of the PCE software.
- After you have restored a backup file that is from a previous version of the PCE software.

To migrate the PCE database:

1. On any node, migrate the PCE database:

```
$ sudo -u ilo-pce illumio-pce-db-management migrate
```

2. On the *primary database*, set the cluster to runlevel 5:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 5
```

Setting runlevel might take some time to complete.

3. Check the progress to see when the status is `RUNNING`:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status -w
```

Manage Automatic Database Failover

When the primary database experiences a failure event lasting more than 2 minutes, the PCE automatically fails over to the backup database. Failing over the database

causes other PCE services to restart. During the database failover period, REST API requests might fail and the PCE web console might become unresponsive.

When the primary database node comes back online and rejoins the cluster, it will detect it is no longer the primary and become the backup database.

Determine Which Node Is Primary

NOTE:

When you install the PCE software, the first data node you install becomes the primary database. Upgrading the PCE does not change the primary database to another data node.

```
$ sudo -u ilo-pce illumio-pce-db-management show-master
```

View Auto Failover Mode

```
$ sudo -u ilo-pce illumio-pce-db-management get-auto-failover
```

Example output:

```
$ sudo -u ilo-pce illumio-pce-db-management get-auto-failover
```

```
Database Failover mode: 'off'
```

Turn Auto Failover Off or On

Automatic failover is enabled by default. To disable it, run the following command:

```
$ sudo -u ilo-pce illumio-pce-db-management set-auto-failover off
```

Manual Database Failover

1. Determine which node that is running as the primary database:

```
$ sudo -u ilo-pce illumio-pce-db-management show-master
```

2. On the *primary database node*, stop the PCE software on the node:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

Wait roughly two minutes for the new node to take over.

3. On the *new database node*, verify that the database service is running:

```
$ sudo -u ilo-pce illumio-pce-db-management show-master
```

4. On the *previous primary database node* in the PCE cluster, restart the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl start
```

After the node starts, the PCE recognizes it as the replica database node and will sync it with the primary database node.

Restore from Data Backup

This task describes how to restore a PCE cluster from a data backup.

We can restore to a different FQDN using the `--update-fqdn` option on the `restore` for the policy DB. This requires the `runtime_env.yml` to have the `pce_fqdn` option set to the new PCE FQDN before running starting the PCE in runlevel 1.

NOTE:

Illumio recommends you wait at least 15 minutes to restore a backup of the policy database after taking the backup. When you restore a policy database backup sooner than 15 minutes, the PCE might not apply policy correctly to all workloads.

1. On *all nodes* in the PCE cluster, stop the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

2. On *all nodes* in the PCE cluster, start the PCE at runlevel 1:

```
$ sudo -u ilo-pce illumio-pce-ctl start --runlevel 1
```

3. On *any node*, verify the runlevel:


```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status -w
```

4. Restore the policy database to the *data node* that is running the `agent_traffic_redis_server` service. (For information about how to determine which node this is, see [Back Up the Policy Database](#).)

```
$ sudo -u ilo-pce illumio-pce-db-management restore --file /path/to/policy_db_dump_file
$ sudo -u ilo-pce illumio-pce-db-management migrate
```

5. Copy the Illumination data file from the primary *data node* that is running the `agent_traffic_redis_server` service to the replica data node. The file is located in the following directory on both nodes.

```
persistent_data_root/redis/redis_traffic_0_master.rdb
```

6. Restore the traffic database. Run this command on the same node where you took the traffic database backup.

```
$ sudo -u ilo-pce illumio-pce-db-management traffic restore --file /path/to/traffic_db_dump_file
```

When prompted to bring the PCE to runlevel 5, reply “yes” if you want the PCE to automatically finish migrating the traffic database and bring the PCE to fully operational status. Reply “no” if you don’t want to migrate the traffic database.

7. If you chose “no” in the previous step:
 - a. Return the PCE cluster to runlevel 5:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 5
```

8. On *any node*, verify the runlevel is 5:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status -w
```

9. Take the PCE out of Listen Only mode:

```
$ sudo -u ilo-pce /opt/illumio-pce/illumio-pce-ctl listen-only-mode disable
```

NOTE:

Explorer will be in maintenance mode for some time after the restore commands complete. The PCE is made available immediately, but the Explorer database restore continues in the background.

Manage Multi-Node Traffic Database

You can scale traffic data by sharding it across multiple PCE data nodes. This can be done when first installing the PCE, as described in [Scale Traffic Database to Multiple Nodes](#). You can also expand an existing traffic database to multiple nodes and change the number of nodes as needed. Reasons for doing so include:

- If you experience performance problems with ingestion or Explorer with a single-node traffic database, these performance issues could be solved by migrating to a multi-node traffic database.
- If you need to store more data than the single-node traffic database can handle (for example, if you want to store 90 days of data), a multi-node traffic database may be required.

Expand Existing Traffic Database to Multiple Nodes

To reconfigure an existing PCE cluster to scale the traffic database to multiple nodes, use the following steps. The PCE will have to be taken offline for a maintenance window. The duration of this maintenance window depends on the amount of data in the traffic database. For a database of 400GB, the downtime is up to approximately 3 hours.

1. On *any data node*, run the following command to back up the traffic database:

```
$ sudo -u ilo-pc e illumio-pce-db-management traffic dump --file trafficdb-backup.tar.gz
```

2. On *any data node*, run the following command to back up the reporting database:

```
$ sudo -u ilo-pc e illumio-pce-db-management report dump --file reportdb-backup.tar.gz
```

3. On *all new nodes*, run the following command to allow multi-node traffic, where the address is the IP address of each new node:

```
illumio-pce-ctl cluster-nodes allow <address>
```

4. On *all nodes*, stop the PCE:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

5. Install the PCE software on the new coordinator and worker nodes, using the same version of the PCE that is present on the existing nodes in the cluster. There must be exactly two (2) coordinator nodes. There must be two (2) or more pairs of worker nodes.
6. Update the `runtime_env.yml` configuration on every node (the new ones you just added as well as the ones that were already in the PCE cluster) as follows. For examples, see [Example Configurations for Multi-Node Traffic Database](#).
 - Set the cluster type to `4node_dx` for a 2x2 PCE or `6node_dx` for a 4x2 PCE.
 - In the `traffic_datastore` section, set `num_worker_nodes` to the number of worker node pairs. For example, if the PCE cluster has 4 worker nodes, set this parameter to 2.
 - On each coordinator node, in addition to the settings already described, set `node_type` to `citus_coordinator`.
 - On each worker node, in addition to the settings already described, set `node_type` to `citus_worker`.
 - If you are using a split-datacenter deployment, set the `datacenter` parameter on each node to an arbitrary value that indicates what part of the datacenter the node is in.

7. Check the runtime configuration:

```
$ sudo -u ilo-pce illumio-pce-env check
```

8. On *all nodes*, start the PCE at runlevel 1:

```
$ sudo -u ilo-pce illumio-pce-ctl start --runlevel 1
```

9. When the PCE is up and running at level 1, restore the reporting database backup. Run this command on the node where you took the backup.

```
$ sudo -u ilo-pce illumio-pce-db-management report restore --file pce-reportdb-dump.tar.gz
```

10. On *one of the coordinator nodes*, migrate the traffic database. This will create the database on the coordinator node.

```
$ sudo -u ilo-pce illumio-pce-db-management traffic migrate
```

11. On the *node where you took the backup*, restore the traffic database backup that you made in step 1:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic restore --file trafficdb-backup.tar.gz
```

When prompted, reply Y if you want to bring the PCE up to runlevel 5 while the database restore continues in the background. This makes all PCE features except Explorer available immediately, without having to wait for the restore to complete.

If you do not choose to go to runlevel 5 at this time, you can do so later by running the following command on *any node*:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 5
```

12. On *any node*, check the cluster status:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status -w
```

13. When the cluster status is UP and RUNNING, verify successful setup. Log in to the PCE web console and verify that the health of the PCE is good. Check Explorer by running a few queries.

Add or Remove a Worker Node

To add or remove a worker node in a multi-node traffic database, use the following steps. The PCE will have to be taken offline for a maintenance window. The duration of this maintenance window depends on the amount of data in the traffic database.

WARNING:

Be sure that the final number of worker nodes is an even number. Worker nodes can only function in groups of two.

1. On *any data node*, run the following command to back up the traffic database:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic dump --file trafficdb_
backup.tar.gz
```

2. On *any node*, set the PCE to runlevel 1:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 1
```

3. When removing a node, run the following command on the node you are removing:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-leave
```

4. On *all nodes*, stop the PCE cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-stop
```

5. On *every PCE node*, update the value of `traffic_datastore.num_worker_nodes` in `runtime_env.yml`. The value should always be twice as large as the number of individual worker nodes, because the worker nodes are configured in pairs.

6. On *all nodes*, start the PCE at runlevel 1:

```
$ sudo -u ilo-pce illumio-pce-ctl start --runlevel 1
```

7. On the *data node where you took the backup*, restore the traffic database backup that you made in step 1:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic restore --file  
trafficdb_backup.tar.gz
```

8. On *any node*, set the PCE to runlevel 5:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 5
```

9. Verify successful setup. Log in to the PCE web console and verify that the health of the PCE is good. Check Explorer by running a few queries.

Back Up and Restore Multi-Node Traffic Database

When your PCE cluster includes a multi-node traffic database, the data size increases, and the standard PCE backup and restore commands consume too much time and resources. To back up and restore multi-node traffic data, use `pgbackrest` instead. For more information, see [Using pgbackrest for Traffic Data Backups](#).

Database Management Commands for Multi-Node Traffic Database

Following are some useful commands to get information about a cluster where the traffic database is distributed to multiple nodes.

To show the worker node configuration:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic citus-worker-metadata
```

To show worker primary nodes:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic show-citus-worker-primaries
```

To show worker replication information:

```
$ sudo -u ilo-pce illumio-pce-db-management traffic show-citus-worker-  
replication-info
```

PCE Default Object Limits

The PCE enforces certain soft and hard limits to restrict the total number of system objects that you can create. These limits are set based on the tested performance and

capacity limits of the PCE.

Types of Object Limits

This section describes the difference between soft and hard limits.

Soft Limits

Soft limits serve as an early warning for potential PCE scale and performance issues. When you see a soft limit warning, contact Illumio Customer Support to discuss the potential impact of this alert on your deployment.

When the PCE reaches a soft limit, it logs an organization (audit) event that indicates the soft limit for that object has been reached:

```
soft_limit_exceeded
```

You should investigate soft limit alerts on a non-emergency basis. When PCE services are functioning normally, but the PCE is generating a lot of soft limit alerts, consult Illumio Customer Support about altering or suppressing the soft limit alerts.

NOTE:

When you lower a soft limit below the current actual usage, the PCE does not generate an event.

Hard Limits

Hard limits protect the PCE from usage and performance overloads, such as creating too many workloads, or too large a security policy. When you receive a hard limit warning, Illumio recommends that you investigate it immediately. When a hard limit is reached in conjunction with a service outage, a PCE core capacity might be overloaded.

When a hard limit is reached, any attempt to create more objects of that type will fail and result in an error message in the PCE web console or a HTTP 406 error returned in REST API. In addition, the PCE logs this event:

```
hard_limit_exceeded
```

When you reach a hard limit, contact Illumio Customer Support to discuss your PCE deployment.

Check Object Limits and Usage

To check the status and usage of the current object limits, run the following command:

```
$ sudo -u ilo-pce <install_root>/illumio-pce-ctl obj-limits list
```

WARNING:When your current usage for any object type shows that you are approaching a soft or hard object limit, contact Illumio Customer Support for assistance.

The CLI commands `illumio-pce-db-management events-storage` and `illumio-pce-env` show information about hard and soft limits and related events.

- `illumio-pce-db-management events-storage` CLI commands list when the soft-cap reached, hard-cap reached, and hard-cap exited conditions were last observed.
- `illumio-pce-db-management events-storage` CLI commands list the current soft-cap and hard-cap limits.
- `illumio-pce-env` command displays a warning if a hard cap condition exists, but the command does not fail.

Example:

```
$ illumio-pce-db-management events-storage

Reading /opt/pce_config/etc/runtime_env.yml.
INSTALL_ROOT=/var/illumio_pce
RENV=development

Event limit conditions status
Current events soft_limit, hard_limit (in MB): [7132, 8915]
Events soft limit last exceeded at:
Events hard limit last exceeded at:
Last recovered from events hard limit exceeded condition at:

Done.
```


Object Limits During Bulk Create

When you use the Illumio REST API to perform an asynchronous job, such as bulk creation of multiple workloads, and you reach the workload object limit during the job, the job will successfully create as many workloads within the limit, and fail to create more workloads.

The HTTP response shows that some workloads were successfully created, and includes a failure message for each workload that was not created due to the hard limit.

For example:

```
[
  {
    "token": "object_limit_hard_limit_reached",
    "message": "Object limit hard limit reached"
  }
]
```

Object Limits and Concurrent Transactions

When multiple users create the same type of object simultaneously, the PCE can reach the hard object limit for that object concurrently during the parallel transactions. This type of “race” condition is atypical but can occur.

For example, a PCE has 900 rules. Two users each simultaneously add 100 rules in a single transaction. After their two transactions, the rule object count is 1100. When the two transactions occur simultaneously and the PCE reaches a hard limit for that object, both transaction can return an error after the PCE reaches the limit.

PCE Object Limits

The following table lists all PCE object limits, identified by each object name followed by the object’s keyname in parentheses. The object keyname is displayed when you run the `illumio-pce-ctl obj-limits list` command on one of the nodes in your cluster.

Object	Description	Soft Limit	Hard Limit
VENS per PCE (active_agents_per_pce)	Total number of VENS that have been installed on managed workloads	SNC: 8,000 2x2 (small): 2,000 2x2: 8,000 4x2: 20,000	SNC:10,000 2x2 (small): 2,500 2x2: 10,000 4x2: 25,000
Labels (total_labels)	Total number of labels	20,000	25,000
Label Groups (total_label_groups)	Total number of label groups	8,000	10,000
Label Group members (label_group_members)	Total number of labels in a label group, including nested label groups For example, you have label groups A and B, and each group contains 1000 labels. Label group C contains label groups A and B. The total number of label_group_members in C is 2002 (1000 + 1000 + 2). Every nested label group and all its members are counted in the object limit.	8,000	10,000
IP List entries (total_ip_list_entries)	Total number of all IP list entries in all IP lists in the system	8K	10K
Interfaces per Unmanaged Workload (interfaces_per_unmanaged_workload)	Total number of network interfaces supported per unmanaged workload An unmanaged workload does not have a VEN installed on it.	102	128
Interfaces per VEN (interfaces_per_agent)	Total number of interfaces supported per managed workload A managed workload has a VEN installed on it.	32	None (-1)
Items per Rule	Total number of items allowed per	50	200

Object	Description	Soft Limit	Hard Limit
(total_actors_per_rule)	<p>rule in the Providers and Consumers fields.</p> <p>A rule contains labels, workloads, and IP lists. When you have a rule that has two Provider items and two Consumer items, the rule has 4 items.</p>		
Pairing Keys (active) (total_active_pairing_keys)	<p>Total number of active pairing keys</p> <p>A pairing key is active when you create a pairing profile, click Start Pairing, and generate the key.</p> <p>When you click Stop Pairing, the pairing key becomes inactive and is no longer counted in the object limit.</p>	1200	5K
Pairing Profiles (total_pairing_profiles)	Total number of pairing profiles	1200	5K
RBAC Permissions (total_org_permissions)	<p>Total number of RBAC permissions</p> <p>Each RBAC permission is a three tuple of an RBAC user or user group, role, and scope.</p>	10K	35K
Policy Services (total_policy_services)	Total number of services that you have added to the PCE and provisioned to use in rules	10K	None (-1)
Port ranges per Policy Service (port_ranges_per_policy_service)	Total number of port ranges per service	50	None (-1)
Services per Rule (total_services_per_rule)	Total number of services that can be associated with a single rule	40	50
Ports per Rule	Total number of ports that can be	400	500

Object	Description	Soft Limit	Hard Limit
(total_service_ports_per_rule)	associated with a single rule. Each service has a certain number of ports or port ranges. Note that in this instance, "service" refers not to a proper service or virtual service as such, but to a port representing a service. This means that this object limit governs your adding a distinct port or port range to a rule.		
Rules (total_rules)	Total number of all rules in all rulesets	40K	50K
Scopes and Rules (total_scopes_rules)	Sum of the total number of rules times the total number of scopes in all rulesets For example, you have two rulesets: RuleSet1 (2 rules, 3 scopes) and RuleSet2 (2 rules, 1 scope). In this example, the total number of scopes and rules is $(2 \times 3) + (2 \times 1) = 8$.	40K	50K
Total stateless Rules (total_stateless_rules)	The total number of stateless rules in your organization	80	100
Total selective enforcement rules total_selective_enforcement_rules	Total number of selective enforcement rules	400	500
RBAC Users and Groups (total_org_auth_security_principals)	Total number of all RBAC users and groups	1600	2000
Adaptive User Segmentation	Total number of Adaptive User Segmentation (AUS) users used in rules	45K	50K

Object	Description	Soft Limit	Hard Limit
(AUS) users (total_security_principals)			
Service Bindings (total_service_bindings)	Total number of service bindings created between workloads and virtual services	90K	100K
Services per VEN (services_per_agent)	Total number of services on a managed workload that the VEN reports to the PCE When you add more than 200 services to a managed workload, the PCE ignores any services over the 200 limit.	160	200
Workloads (total_workloads)	Total number of managed and unmanaged workloads A managed workload has a VEN installed on it, while an unmanaged workload does not.	SNC: 2,000 2x2 (small): 10,000 2x2: 40,000 4x2: 100,000	SNC: 2,500 2x2(small): 12,500 2x2: 50,000 4x2: 125,000
Container workloads (total_container_workloads)	Total number of container workloads. The term <i>container workloads</i> refers to containerized workloads in a container cluster that is managed by a Kubelink that is not in Cluster Local Actor Store (CLAS) mode.	8K	10K
Kubernetes workloads (total_kubernetes_workloads)	Total number of Kubernetes workloads. The term <i>Kubernetes workloads</i> refers to containerized workloads in a container cluster that is managed by a Kubelink that is in Cluster Local Actor Store (CLAS) mode.	8K	10K
Container workload pro-	Total number of Container Workload Profiles in each container cluster.	800	1K

Object	Description	Soft Limit	Hard Limit
files (container_workload_profiles_per_container_cluster)			
Container clusters (total_container_clusters)	Total number of container clusters.	80	100
User sessions (total_active_sessions)	Maximum number of user sessions on a single PCE cluster at the same time. This limit includes only actual logged-in user sessions, and omits impersonated sessions, such as scheduled jobs that log in to access PCE data. When the limit is exceeded, anyone who tries to log in is refused with an explanatory message.	100	125

Monitor and Diagnose PCE Health

This chapter contains the following topics:

PCE Logs	79
Monitor PCE Health	84
PCE Health Metrics Reference	96
Support Reports for PCE	109

This section describes how to monitor the PCE to ensure it is operating correctly. For example, you can view events that are generated by the PCE, read PCE logs, and generate reports about PCE activity.

PCE Logs

Most PCE logs are written to syslog, but some logs are written directly to a file in the directory you specify with the `log_dir` parameter in the PCE `runtime_env.yml` file.

Log Files for PCE Services

This table lists the main PCE services and the log file name or the syslog filter for the service.

PCE Service	Syslog Filter Rule or Log File Name
agent_service	program("illumio_pce/agent")
agent_background_worker_service	
agent_traffic_redis_cache	program("illumio_pce/agent_traffic")
agent_traffic_redis_server	
agent_traffic_service	
auditable_events_service	message('"category":"auditable"');
collector_service	program("illumio_pce/collector");
database_mon-	program("illumio_pce/database_monitor");

PCE Service	Syslog Filter Rule or Log File Name
itor	
database_service database_slave_service	program("illumio_pce/postgresql");
ev_service	program("EventService");
executor_service	program("illumio_pce/executor");
fileserver_service	program("illumio_pce/fileserver");
fluentd_source_service	program("illumio_pce/fluentd");
ilocron	program("illumio_pce/ilocron");
login_service	program("illumio_pce/login");
memcached	program("illumio_pce/memcached");
node_monitor	program("illumio_pce/system_health");
redis	program("redis");
search_index_service	program("illumio_pce/search_index");
server_load_balancer	program("haproxy"); <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>NOTE: HAProxy logs to /dev/log using a datagram socket. When using syslog-ng, you might need to update your syslog-ng configuration to listen on /dev/log on a datagram socket.</p> </div>
service_discovery_service	program("illumio_pce/service_discovery"); program("consul");
web_server	match("nginx;" value("MESSAGE"));

Log Files (Non-syslog)

The following PCE log files are written to the value defined in the `log_dir` parameter of the runtime configuration file.

- `agent_background_worker_0.log`
- `cache_0_master.log`
- `consul.log`
- `config_manager`

- fileserver.3400.log
- fluentd-source.log
- ilo_node_monitor.log
- nginx_error.log
- passenger.log
- pce_error.log
- pg_listener.log
- set_server_0_master.log
- system_history.log
- thin_agent_traffic.3200.log
- thin_collector.3100.log
- thin_login.3300.log
- thin_search_engine.3500.log
- tmessenger/compact.log
- tmessenger/heartbeat.log
- tmessenger/relay.log
- traffic_0_master.log
- traffic_worker_0.log
- traffic_worker.log

In addition, the PCE software writes system stats to the following two files in the `log_dir/systats` directory every 10 minutes:

- `perflog`
- `app_stats`

CAUTION:

Do not delete these files. They contain useful system and application statistics that can help Illumio Customer Support in troubleshooting PCE operational issues.

System Upgrade Log

On each PCE node, the `log_dir` directory contains a log file called `system_history` that records the following information:

- Initial PCE version
- PCE version upgrades (old version and new version)

- PCE backups (how many times the PCE software on the node has been backed up)
- PCE restores with the timestamp of the backup that was restored
- A timestamp for each log entry indicating when the operation occurred

Example system upgrade log:

```
2016-09-24 05:04:11.216: Change in PCE software version detected. Previous:
16.6.0-4114, Current: 16.9.0-4121.
2016-09-24 05:04:39.583: Data dump to file '/tmp/illumio_pce_data/db_
backup.tar.gzip' started.
2016-09-24 05:04:47.950: Data dump to file '/tmp/illumio_pce_data/db_
backup.tar.gzip' completed. MD5 checksum: 02cef311e9657710a1900d8c5deb49d9
```

Password-related Event Logging

The system records auditable events for the following occurrences:

- When an Illumio administrator changes the password requirements
- When users successfully change their passwords as required by password policy
- When users fail to change their passwords according to required password policy

Search the PCE Log Files

The PCE Support Report search function allows you to search PCE log files (log files written to `/var/log/illumio-pce`) based on the following criteria:

- **From (timestamp) & To (timestamp):** Search between two specific points in time.
- **From (timestamp) & Duration (hours):** Search a duration of time starting at a specific point in time.
- **Duration (hours) & To (timestamp):** Search for a duration of time up to a specific point in time.
- **Duration (hours) & At (timestamp):** Search for something that occurred during a general time frame and gather logs from before and after the event; (timestamp) is the midpoint.
- **From (timestamp) + Search term:** Search from a starting time for specific types of information using the standard search terms.

Examples of Searching

The following examples use questions to frame the log search goals and help formulate your searches.

From/To Dates

Question:

“I want to search 12 hours worth of PCE logs starting on February 1, 2020 and ending 12 hours after (from midnight on 2/1 to noon 2/1).”

Search syntax:

```
$ sudo -u ilo-pce ./support_report logs from=02/01/2020 to=02/02/2020
$ sudo -u ilo-pce ./support_report list
```

Duration/To

Question:

“I want to search for 6 hours worth of PCE logs ending on midnight of February 2, 2020. Effectively from 1800 on February 1 through 0000 on February 2, 2020.”

The default value of hours in a date is midnight.

Search syntax:

```
$ sudo -u ilo-pce ./support_report logs duration=6 to=02/02/2020
$ sudo -u ilo-pce ./support_report list
```

At/Duration

Explanation: Use the “at” operator in conjunction with the “duration” operator in the following example. To find details for a specific event that occurred at a known time, use “at.” “At” is the approximate time at which an event of interest occurs. The duration is the time range on either side of this timestamp. In this example, duration returns all messages between 10:00:15 and 12:00:15 on February 2, 2020 and “at” narrows the search to a more specific time, in this case, 11:00AM.

Question:

“I want to search a time window between the hours of 10:00AM and 12:00PM on February 2, 2020, for a specific event that occurred at 11:00AM.”

Search syntax:

```
$ sudo -u ilo-pce ./support_report logs at=02/02/2020T11:00:15 duration=2
$ sudo -u ilo-pce ./support_report list
```

From + Search Term Included

Question:

“I want to see all PCE logs entries starting from February 2, 2020, until the present that refer to JOB_STORE.”

Search syntax:

```
$ sudo -u ilo-pce ./support_report from=02/02/2020 include=JOB_STORE
$ sudo -u ilo-pce ./support_report list
```

From + Search Term Included and Excluded

Question:

“I want to see all PCE logs entries starting from February 2, 2020, until the present that refer to JOB_STORE and timed_work but for all servers excluding core0.”

Search syntax:

```
$ sudo -u ilo-pce ./support_report from=02/02/2020 include=JOB_STORE
include=timed_work exclude=core0
$ sudo -u ilo-pce ./support_report list
```

Monitor PCE Health

This section describes how to monitor the health of the PCE.

PCE Health Monitoring Techniques

You can monitor the PCE software health using the following methods:

- **PCE web console:** The Health page in the PCE web console provides health information about your on-premises PCE, whether you deployed a 2x2 cluster, 4x2 cluster, or SNC.
- **REST API:** The PCE Health API can be used to obtain health information.
- **Syslog:** When you configure syslog with the PCE software, the PCE reports `system_health` messages to syslog for all nodes in the PCE cluster.

- **PCE command-line interface:** Run commands to obtain health status for the entire PCE cluster and each node in the cluster.

Minimum Required Monitoring

The PCE provides several different methods you can use to monitor PCE health, as described in [PCE Health Monitoring Techniques](#).

No matter which technique you use, there is one main signal that it is important to watch for: the overall system status. You must monitor it as follows:

- If you are using the PCE web console, keep an eye on the **PCE Health** status near the top of the page. It indicates whether the PCE is in a Normal, Warning, or Critical state of health. For details, see [Health Monitoring Using PCE Web Console](#).
- If you are using the API, similarly, monitor the status field. For details, see [Health Monitoring Using Health REST API](#).
- If you are using the PCE syslog to monitor PCE health, watch for any messages that contain the text `sev=WARN` or `sev=ERR`. In such messages, check the other fields for details. For more details, see [Health Monitoring Using Syslog](#).

The rest of this section provides details about the meaning of the various PCE health metrics and what to do if a warning or error state is seen.

PCE Health Status Codes

The following table lists the status shown in the PCE web console (or PCE Health API), the severity code shown in syslog, the corresponding color code in the PCE web console, and the most commonly encountered causes for each level of health.

Status/Severity	Color	Typical Meaning
Normal (healthy) or sev= v=INFO	Green	<ul style="list-style-type: none"> • All required nodes and services are running. • CPU usage, memory usage, and disk usage of all nodes is less than 95%, and all other metrics are below their thresholds. • Database replication lag is less than or equal to 30 seconds. • (In a PCE Supercluster only) Supercluster replication lag is less than or equal to 120 seconds.
Warning or sev= v=WARN	Yellow	<ul style="list-style-type: none"> • One or more nodes are unreachable. • One or more optional services are missing, or one or more required services have been degraded.

Status/Severity	Color	Typical Meaning
		<ul style="list-style-type: none"> The CPU usage, memory usage, or disk usage of any node is greater than or equal to 95%, or another health metric has exceeded its warning threshold. For more information, see PCE Health Metrics Reference. Database replication lag is greater than 30 seconds. (In a PCE Supercluster only) Supercluster replication lag is greater than 120 seconds.
Critical or sev- v=ERR	Red	<ul style="list-style-type: none"> One or more required services are missing. A health metric has exceeded its critical/error threshold. For more information, see PCE Health Metrics Reference.

If a warning threshold has been exceeded, a warning icon appears in three places in the PCE web console: the upper right of the PCE Health dashboard, the General summary area of the dashboard, and next to the appropriate tab.

Health Monitoring Using PCE Web Console

Click the Health icon at the top of the PCE web console to see the general health of the PCE.

Tabs categorize the health information by Node, Application, Database Replication, and Supercluster.

The Node tab shows node information, including the health metric Disk Latency. It also displays a hardware requirements message for each node, to tell whether the hardware provisioned meets the requirements as documented in the [Capacity Planning](#) topic. If a node is found to have sufficient resources to meet specifications, the message "Node Specs Meet requirements" appears with a green checkmark. If the node does not have sufficient resources to meet the required specifications, the alert "Node Specs Do not meet requirements" appears with a yellow triangle. The requirements vary depending on the type of PCE cluster (single-node, 2x2 multi-node, 4x2 multi-node, etc.). This is determined based on the `cluster_type` runtime parameter, which is set for every node. The hardware requirements check needs to know the cluster type so it can use the right set of hardware requirements.

The Application tab shows a variety of information, including database health metrics. (For details, see [PCE Health Metrics Reference](#).) The tab is divided into sections:

- Collector Summary (flow rate, success vs. failure rates)
- Traffic Summary (ingestion, backlog, database utilization)
- Policy Database Summary (database size, transaction ID age, vacuum backlog)
- VEN Heartbeat (success vs. failure, latency)
- VEN Policy (request rate, latency)

The Database Replication tab shows the database replication lag.

The Supercluster tab shows the Supercluster replication lag (applicable only in a PCE Supercluster).

PCE Health Status Indicator

The PCE web console provides an indicator that reflects overall status. Near the top of the PCE Health page in the PCE web console, a warning indicator labeled **PCE Health** shows normal, warning, or critical. You can find more details on the tab that corresponds to the issue.

Health Monitoring Using Health REST API

With the PCE Health API, you can display PCE health information using the following syntax:

```
GET [api_version]/health
```

For details, see [PCE Health](#) in the *REST API Developer Guide*.

Health Monitoring Using Syslog

Each PCE node reports its status to the local syslog daemon once every minute. The PCE uses the program name `illumio_pce/system_health` for these messages.

Example Syslog Messages

Example syslog message from a non-leader PCE node:

```
2015-12-17T00:40:31+00:00 level=info host=ip-10-0-0-26 ip=127.0.0.1
program=illumio_pce/system_health| sec=312831.757 sev=INFO pid=9231 tid=12334020
rid=0 leader=10.0.24.26 database_replication_lag=3.869344 cpu=2% disk=11%
memory=19%
```

Example syslog message from a leader PCE node for a healthy PCE cluster:

```
2015-12-23T22:52:59+00:00 level=info host=ip-10-0-24-26 ip=127.0.0.1
program=illumio_pce/system_health| sec=911179.836 sev=INFO pid=5633 tid=10752960
rid=0 cluster=healthy cpu=2% disk=10% memory=37%
```

Example syslog message from a leader PCE node for a degraded PCE cluster with one node missing:

```
2015-12-23T22:56:00+00:00 level=notice host=ip-10-0-24-26 ip=127.0.0.1
program=illumio_pce/system_health| sec=911360.719 sev=WARN pid=5633 tid=10752960
rid=0 cluster=degraded missing=1 cpu=34% disk=10% memory=23%
```

Health Monitoring Using PCE Command Line

This section gives several techniques you can use at the command line to monitor PCE health.

Monitor a PCE Cluster

The following command displays the status of the PCE cluster, including where each individual service is running:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

Return codes:

- 0 - NOT RUNNING
- 1 - RUNNING
- 2 - PARTIAL (not all required services running)

For example:

```
$ ./illumio-pce-ctl cluster-status

SERVICES (runlevel: 5)          NODES (Reachable: 4 of 4)
=====
agent_background_worker_service 10.0.26.49    10.0.6.171
agent_service                   10.0.26.49    10.0.6.171
agent_traffic_redis_cache       10.0.11.96    10.0.25.197
agent_traffic_redis_server      10.0.25.197
agent_traffic_service           10.0.26.49    10.0.26.49    10.0.6.171
```



```

10.0.6.171
auditable_events_service      10.0.26.49      10.0.6.171
collector_service             10.0.26.49      10.0.26.49      10.0.6.171
10.0.6.171
database_monitor              10.0.11.96      10.0.25.197
database_service              10.0.25.197
database_slave_service        10.0.11.96
ev_service                    10.0.26.49      10.0.6.171
executor_service              10.0.26.49      10.0.6.171
filesaver_service             10.0.25.197
fluentd_source_service        10.0.26.49      10.0.6.171
login_service                 10.0.26.49      10.0.6.171
memcached                     10.0.26.49      10.0.6.171
node_monitor                  10.0.11.96      10.0.25.197      10.0.26.49
10.0.6.171
pg_listener_service           10.0.11.96
search_index_service          10.0.26.49      10.0.6.171
server_load_balancer          10.0.26.49      10.0.6.171
service_discovery_agent       10.0.25.197
service_discovery_server      10.0.11.96      10.0.26.49      10.0.6.171
set_server_redis_server       10.0.11.96
traffic_worker_service         10.0.26.49      10.0.6.171
web_server                    10.0.26.49      10.0.6.171

```

This command displays the members of the PCE cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-members
```

For example:

```

[illumio@core0 illumio-pce]$ ./illumio-pce-ctl cluster-members
Reading /var/illumio-pce/data/runtime_env.yml.
Node                Address             Status  Type
core0.mycompany.com 10.6.1.19:8301     alive  server
data0.mycompany.com 10.6.1.20:8301     alive  server
core1.mycompany.com 10.6.1.32:8301     alive  server
data1.mycompany.com 10.6.1.31:8301     alive  client

```

Monitor Database Replication

On *either data node*, run the following command to display the status of replication between the primary database and replica:

```
$ sudo -u ilo-pce illumio-pce-db-management show-replication-info
```

The PCE updates this information every two minutes.

IMPORTANT:

To prevent data loss during a database failover operation, monitor the PCE databases for excessive database replication lag.

For example:

```
$ ./illumio-pce-db-management show-replication-info
Reading /var/illumio/data/runtime_env.yml.
INSTALL_ROOT=/var/illumio/software
RENV=development

Current Time: 2016-02-16 22:42:03 UTC

Master: (10.6.1.73)
Last Sampling Time : 2016-02-16 22:41:14 UTC
Transaction Log location : 0/41881E8

Slave(s):
IP Address: 10.6.1.72
Last Sampling Time : 2016-02-16 22:41:16 UTC
Streaming : true
Receive Log Location : 0/41881E8
Replay Log Location : 0/4099048
Receive Lag (bytes) : 0
Replay Lag (bytes) : 979360
Transaction Lag (secs) : 4.633377
Last Transaction Replayed Time: 2016-02-16 22:37:12.920179 UTC
```

PCE Health Troubleshooting

This section tells what action to take if you see a non-normal status when monitoring PCE health. The recommended response depends on which metric has departed from

the Normal state. If you are not able to diagnose and fix it yourself, contact Illumio Support.

The health metrics may occur in the PCE web console, API response status field, or in the syslog severity field. When multiple conditions result in differing levels of severity, the more critical level is reported. If you receive a non-normal level for any of the following, here are the suggested actions to take. For additional details, see [PCE Health Metrics Reference](#).

Name	Troubleshoot
Disk Latency	Warning/Critical: Disk latency on data nodes is an indication that DB/Traffic service needs to be investigated further for possible performance issues. Typically higher disk latency numbers indicate Disk I/O bottlenecks.
CPU	When the PCE is under heavy load, CPU usage increases, and the Warning status is reported. Typically, the load should decrease without intervention in less than 20 minutes. If the Warning condition persists for 30 minutes or more, decrease the load on the CPU or increase capacity.
Memory	When the PCE is under heavy load, memory usage increases, and the Warning status is reported. Typically, the load should decrease without intervention in less than 20 minutes. If the Warning condition persists for 30 minutes or more, increase the available memory.
Disk Space	The PCE manages disk space using log rotation, and this is usually sufficient to address any Warning condition. If the Warning level persists for more than one day, and the amount of disk space consumed keeps increasing, notify Illumio Support.
Policy Database Summary	<ul style="list-style-type: none"> disk_usage (database disk utilization): Warning: Plan to increase the capacity of the disk partition holding the Policy DB or make more room by deleting unnecessary data as soon as possible. Critical: Immediately increase the disk partition holding the Policy DB or make more room by deleting unnecessary data. txid_max_age (transaction ID maximum age): Warning: Contact Illumio Support and plan a manual full vacuum as soon as possible. Critical: Immediately contact Illumio Support. vacuum_backlog (vacuum backlog):

Name	Troubleshoot
	Warning, Critical: If the situation persists, contact Illumio Support so that the reason for the underperformance of the auto-vacuum can be investigated.
VEN heart-beat performance	<ul style="list-style-type: none"> <li data-bbox="428 384 946 415">• avg_latency, hi_latency (latency): If the VEN heartbeat latency is high, examine the application logs on core nodes and system resource utilization across the entire PCE cluster. IOPS-related issues may often be diagnosed by examining database logs and observing long wait times for committing database transactions to disk. <li data-bbox="428 674 862 705">• rate, result (response stats): Warning/Critical: Examine the application logs on core nodes for more information about the precise cause of the failure.
Policy performance	<ul style="list-style-type: none"> <li data-bbox="428 825 946 856">• avg_latency, hi_latency (latency): If latency is abnormally high, investigate the cause. For example, examine the logs to try to find out why the policy is changing. <li data-bbox="428 982 760 1014">• rate (request count): If abnormally large, investigate the cause (see latency). The default threshold is conservative by design. Each organization has its own expected rate of change of VEN policy, so there is no universal correct warning threshold. You can modify the threshold to better match expectations (see Configurable Thresholds for Health Metrics). If the number of VEN policy requests is too high, examine application logs to find the reasons for the policy changes, and determine whether the policy changes are expected.
Collector summary	<ul style="list-style-type: none"> <li data-bbox="428 1392 857 1423">• Flow summaries rate, node: A 4x2 PCE cluster is configured to handle approximately 10,000 flow summaries per second by default. If fewer posts are reported and you see a large number of failed posts, the collector count can be increased with help from Illumio Support. <li data-bbox="428 1640 740 1671">• Success rate, node: This metric is informational. However, if counts differ across core machines, ensure intra-PCE latency is within the 10ms limit. <li data-bbox="428 1797 902 1829">• Failure percentage ratio, node:

Name	Troubleshoot
	<p>On startup, or when connections are reestablished, VEN post rates can overwhelm the PCE, causing it to reject posts. This is normal unless persistent. If this ratio is large, or if the value is consistent and large (0.1), it means VENs may not be able to upload flow data, and they will start dropping after 24 hrs. The solution is usually to add more collectors.</p>
Traffic summary	<ul style="list-style-type: none"> <p>• Ingest rate, node:</p> <p>A 4x2 PCE cluster is configured to handle approximately 10,000 flows per second by default. If this rate is exceeded, and a backlog begins to grow, the PCE will eventually prune the backlog and lose data. Adding additional <code>flow_analytics</code> daemons will distribute the work, but eventually PostgreSQL itself could become the bottleneck, requiring the use of DX.</p> <p>• Backlog size, node:</p> <p>If the size of the backlog increases continuously, this indicates performance issues with the flow analytics service which processes the flows in the backlog. Contact Illumio support if the backlog exceeds the safe threshold.</p> <p>• Backlog size percentage:</p> <p>Increasing values indicate that the buffered new flow data is growing, meaning the PCE is unable to keep up with the rate of data posted. The PCE collector flow summary rate and PCE traffic summary ingest rate need to be to be roughly equal, or this buffered backlog will grow.</p>
Database Replication Lag	<p>Warning: Check whether the PCE is running properly, and verify that there is no network issue between the nodes. If the replication lag keeps increasing, contact Illumio Support.</p>
Supercluster Replication Lag	<p>Warning: Check whether all PCEs are running properly, and verify that there is no network issue between the lagging PCEs. If the replication lag keeps increasing, contact Illumio Support.</p>

Configurable Thresholds for Health Metrics

You can configure the thresholds that define the normal, warning, and critical status for each health metric. Each health metric has predefined thresholds for normal (green), warning (yellow), and critical (red). You can use the command `illumio-pce-`

`env metrics --write` to adjust these thresholds. This command can be used to modify any Boolean, number, float, or string, or array of these types (no nested arrays). For example:

```
illumio-pce-env metrics --write CollectorHealth:failure_warning_percent=15.0
```

After setting the desired threshold values, copy `/var/lib/illumio-pce/data/illumio/metrics.conf` to every node in the cluster to ensure consistent application of the thresholds.

Examples of when you might want to use this feature:

- At a larger installation, the default memory threshold is set to 80%, but memory usage routinely spikes to 95%. Every time the memory utilization exceeds the threshold, the PCE Health page displays a warning. By configuring a higher threshold, you can reduce the frequency of warnings.
- Database replication lag can exceed a threshold for a brief time, raising a warning, but the system will catch up with replication after some time. To reduce these warnings, you can configure a longer time period for database replication lag to be tolerated. Note: This is not the same as configuring the threshold of the replication lag itself, but the permissible period of time for the lag to be non-zero.
- The default thresholds might be acceptable when the PCE is first installed, but as more VENs are paired to the PCE over time, the default thresholds might need adjustment.

To set health metrics thresholds:

1. Run the following command to get a list of the available metrics, their current settings, and the thresholds you can modify:

```
illumio-pce-env metrics --list
```

Example output:

Engine	Param	Value	Default
CollectorHealth	failure_warning_percent		10.0
	failure_critical_percent		20.0
	summary_warning_rate		12000

	summary_critical_rate	15000
DiskLatencyMetric		
FlowAnalyticsHealth	backlog_warning_percent	10.0
	backlog_critical_percent	50.0
	summary_warning_rate	12000
	summary_critical_rate	15000
PolicyDBDiskHealthMetric		
PolicyDBTxidHealthMetric		
PolicyDBVacuumHealthMetric		
PolicyHealth		
TrafficDBMigrateProgress		

If nothing appears in the Param column for a given metric, you can't modify the thresholds for that metric. This example output shows that the Collector Health metric has four thresholds you can modify.

2. Run the following command:

```
illumio-pce-env metrics --write MetricName:threshold_name=value
```

For *MetricName*, *threshold_name*, and *value*, substitute the desired values. For example:

```
illumio-pce-env metrics --write CollectorHealth:failure_warning_percent=15.0
```

NOTE: Do not insert any space characters around the equals sign (=).

3. Copy `/var/lib/illumio-pce/data/illumio/metrics.conf` to every node in the cluster. The path to `metrics.conf` might be different if you have customized `persistent_data_root` in `runtime_env.yml`.
4. Restart the PCE.
5. When a metrics configuration is detected, the PCE loads and applies it. In `ilo_node_monitor.log`, you should see a message like "Loaded metric configuration for *MetricName*."

The metrics command provides other options as well. This section discussed only the most useful ones. For complete information, run the command with the `-h` option to see the help text:

```
illumio-pce-env metrics -h
```

PCE Health Metrics Reference

The health metrics consist of a set of key value pairs. The following table describes the possible keys that can appear.

Category	Key	Description	Severity Levels
Disk Space	disk, disk_space_percent_thresholds, disk_inode_percent_thresholds	<p>The PCE node reports disk space for the PCE application directories (configured in the runtime_env.yml file):</p> <ul style="list-style-type: none"> ephemeral_data_root runtime_data_root log_dir persistent_data_root directories <p>When all these directories are on a single mount point, the node reports: disk=n%</p> <p>When multiple mount points exist, the node reports the first discovered path by name, such as:</p> <pre>ephemeral_data_root=n% log_dir=n%</pre> <p>When the PCE encounters an error determining this information, the node reports: disk=?.</p> <p>disk_space_percent_thresholds consists of two values that determ-</p>	<p>Default: The following thresholds trigger the following severity levels:</p> <ul style="list-style-type: none"> NOTICE: disk_space >= 90% or disk_inodes >= 90% WARNING: disk_space >= 95% or disk_inodes >= 95% <p>These default thresholds can be modified using disk_space_percent_thresholds or disk_inode_percent_thresholds.</p> <p>The disk space value is only reported when one of the conditions above is met; otherwise, it reports only disk space. When a node has multiple disk</p>

Category	Key	Description	Severity Levels
		ine the disk space usage percentages that result in NOTICE or WARNING notifications. disk_inode_percent_thresholds consists of two values that determine the disk inode usage percentages that result in NOTICE or WARNING notifications.	mounts, the message might look like: ephemeral_data_root_inodes=n, etc.
Physical Memory	memory, memory_percent_thresholds	Each PCE node reports basic physical memory usage, indicated as: memory=n% memory_percent_thresholds consists of two values that determine the memory usage percentages that result in NOTICE or WARNING notifications.	Default: the following values trigger the following severity levels: <ul style="list-style-type: none"> • NOTICE: memory >= 80% • WARNING: memory >= 95% These default thresholds can be modified using memory_percent_thresholds.
CPU Load	cpu, cpu_max_percent, cpu_tolerance_seconds	Each PCE node reports CPU usage load as cpu=n%. The CPU load is calculated as a percentage between two time slices and represents CPUs of all nodes in the cluster. For example, cpu=100% means all cores are maximized. A notification (NOTICE or WARNING)	Default: the following values trigger the following severity levels: <ul style="list-style-type: none"> • NOTICE: cpu >= 95% for more than 1 minute • WARNING: cpu >= 95% for more than 5 minutes These default

Category	Key	Description	Severity Levels
		<p>is issued when the CPU load exceeds a given percentage for a given amount of time.</p> <p><code>cpu_max_percent</code> is the CPU usage percentage above which the notification timer begins.</p> <p><code>cpu_tolerance_seconds</code> controls the notification timer. It consists of two values that determine how long the CPU is above the maximum usage percentage before a NOTICE or WARNING occurs.</p>	<p>thresholds can be modified using <code>cpu_max_percent</code> and <code>cpu_tolerance_seconds</code>.</p>
Cluster Leader	leader	The IP address of the current leader, or unavailable when no leader exists or it is unreachable.	
Cluster Status	cluster	<p>The overall health of the cluster, reported by the leader only:</p> <ul style="list-style-type: none"> <code>cluster=healthy</code>: Everything is operating properly and all PCE services are running. <code>cluster=degraded</code>: The cluster is running but has unhealthy nodes. <code>cluster=down</code>: The cluster is missing a required service < 5 	<p>These status values trigger the following severity levels:</p> <ul style="list-style-type: none"> NOTICE: <code>cluster=degraded (<2 minutes)</code> WARN: <code>cluster=degraded (>=2 minutes)</code> WARN: <code>cluster=down (<2 minutes)</code> ERROR: <code>cluster=</code>

Category	Key	Description	Severity Levels
		<p>minutes.</p> <ul style="list-style-type: none"> • <code>cluster=failed</code>: The cluster is missing a required service for ≥ 5 minutes. 	<p>=down (≥ 2 minutes)</p> <ul style="list-style-type: none"> • FATAL: <code>cluster=failed</code>
Missing Nodes	<code>missing</code>	The number of nodes that are missing from the cluster. If no nodes are missing, this metric is not reported.	
Replication Lag	<code>database_replication_lag</code>	The number of seconds the database replica is lagging behind the primary database. Output by database replica nodes only.	<p>These thresholds trigger the following severity level:</p> <ul style="list-style-type: none"> • WARNING: ≥ 30 seconds
Disk Latency	<code>policy_disk_latency_milliseconds</code> , <code>traffic_disk_latency_milliseconds</code>	<p>(19.3.2 and later) Average time (in milliseconds) for I/O requests issued to the device to be served. This includes the time spent by the requests in queue and the time spent servicing them. The metric is calculated exactly the same way <code>iostat</code> calculates <code>await</code>.</p> <p>Values: <code>delay</code> (milliseconds), <code>disk</code></p> <p>Usefulness: Indicates Disk I/O, which is especially useful when the DB services are under heavy load.</p>	<ul style="list-style-type: none"> • Normal: ≤ 300 • Warning: $>300 < 800$ • Critical: ≥ 800
Policy Database: Size	<code>policy_data-</code>	(19.3.2 and later) Inform-	

Category	Key	Description	Severity Levels
	base_size_gb	ational. Size of the Policy Database data directory. Provides an indication of disk space requirements of the Policy DB. Depending on size, reported in units of byte, kilobyte, megabyte, gigabyte, terabyte	
Policy Database: Disk Utilization	policy_database_utilization_percentage	(19.3.2 and later) Usage ratio of the disk partition holding the Policy DB. Consequences of the Policy DB running out of disk space can be critical.	<ul style="list-style-type: none"> • Normal: < 90 • Warning: [90 - 95] • Critical: >= 95
Policy Database: Transaction ID Max Age	policy_database_transaction_id_max_age	(19.3.2 and later) Maximum transaction ID (TxID) age of the Policy DB. This does not apply to the Traffic DB. Indicates the risk of the DB running out of TxIDs, which could cause a DB lockdown requiring expensive recovery procedures. The PCE will attempt to automatically detect and recover before this occurs (requires reboot).	<ul style="list-style-type: none"> • Normal: < 1 billion • Warning: [1 billion - 2 billion] • Critical: >= 2 billion
Policy Database: Vacuum Backlog	policy_database_vacuum_backlog_percentage	(19.3.2 and later) Percentage of vacuum-ready rows (a.k.a dead rows) over the total	<ul style="list-style-type: none"> • Normal: < 40 • Warning: 40 - 80 and current number of vacuum-ready rows is

Category	Key	Description	Severity Levels
		number of rows of the Policy database computed over a period of up to 12 hours. This does not apply to the Traffic DB. Indicates how well the auto-vacuum of DB is performing. If the percentage is persistently above Postgres default settings of about 20% of the total number of rows, it is an indication that the auto-vacuum is not working effectively.	above Postgres default minimum to trigger vacuum (20% +50) <ul style="list-style-type: none"> • Critical: ≥ 80 and current number of vacuum-ready rows is above Postgres default minimum to trigger vacuum (20% +50)
VEN Heartbeat Performance: Latency	ven_heartbeat_average_latency_seconds, ven_heartbeat_high_latency_seconds	(19.3.2 and later) (milliseconds) ven_heartbeat_average_latency_seconds is the average over the measurement time period. ven_heartbeat_high_latency_seconds is the average 95% over the measurement time period. Backend processing time of VEN heartbeat requests. Does not include the time spent in the load balancer queues, as the queue time may be influenced by a number of other external factors. The VEN heartbeat uses	<ul style="list-style-type: none"> • Warning: average > 500ms • Critical: average > 5 sec

Category	Key	Description	Severity Levels
		the same PCE services and components as the policy computation and is therefore a good overall indicator for the health of the policy subsystem, including whether system resources are being overwhelmed. Historically, it has reliably indicated I/O and/or policy cache bottleneck (s).	
VEN Heartbeat Performance: Success	ven_heartbeat_success_count_per_hour	(19.3.2 and later) Active VENs send a heartbeat API request to the PCE approximately every 5 minutes. This metric captures the number of VEN heartbeat requests seen on the PCE in approximately the past hour. The count may be transiently inaccurate due to concurrent log rotation or other gaps in the application log files. If the PCE has just started up, this number is expected to ramp up over the first hour. The number of successful VEN heartbeat requests per hour	<ul style="list-style-type: none"> Warning: for any non-2xx code, greater than 1% of total requests for the time window Critical: for any non-2xx code, greater than 20% of total requests for the time window

Category	Key	Description	Severity Levels
		summed across all PCE core nodes should be approximately the number of VENs times 12 (heartbeats happen every 5 minutes per VEN). A low number of successful VEN heartbeats likely indicates issues with VEN connectivity or PCE performance. Depending on the VEN disconnect/offline settings, a low VEN heartbeat success rate may cause traffic to be dropped to/from enforced workloads.	
VEN Heartbeat Performance: Failure	ven_heartbeat_failure_percent, ven_heartbeat_failure_count_per_hour	(19.3.2 and later)	Warning: 5% Critical/Error: 20%
Policy Performance: Latency	ven_policy_average_latency_seconds, ven_policy_high_latency_seconds	(19.3.2 and later) (milliseconds) Average response time for policy. Latency indicates policy complexity and system load/bottlenecks. This metric captures the backend processing time of VEN policy requests. It does not include the time	<ul style="list-style-type: none"> Warning: average > 10 sec Critical: average > 30 sec

Category	Key	Description	Severity Levels
		<p>spent in the load balancer queues, as queue time may be influenced by a number of other external factors.</p> <p>The cost to compute the VEN policy instructions depends on a large number of factors, including but not limited to the rate of change in the environment, the number of rules, the number of actors (workloads, labels, etc.) used in the rules, and the density of desired connectivity between workloads. Abnormally high VEN policy request latency may indicate issues with inadequate system resources, policy changes that result in higher than intended policy complexity, or an abnormally high rate of change to the workload context.</p>	
Policy Performance: Request Count	ven_policy_request_count_per_hour	(19.3.2 and later) (requests/hour) When a new policy is provisioned or the workload context (IP address, label membership, etc.) is changed	<ul style="list-style-type: none"> Warning: > 1M req/hour

Category	Key	Description	Severity Levels
		<p>on the PCE, policy instructions are sent to affected VENs. This metric captures the number of VEN policy requests seen on the PCE in approximately the past hour. The count may be transiently inaccurate due to concurrent log rotation or other gaps in the application log files. When a PCE first starts or is restarted, this number may increase sharply over a short time period as every VEN checks to ensure policy sync.</p> <p>The VEN policy request rate provides an indicator of the rate of policy change across the organization, and therefore, an estimate of the load on the PCE. VEN policy requests are sometimes more expensive to process than other API requests, and frequent policy changes may result in decreased overall performance and longer policy convergence times. Frequent policy changes may also be a symptom of underlying</p>	

Category	Key	Description	Severity Levels
		network or infrastructure issues, such as (but not limited to) frequent IP address changes or improperly cloned VENS.	
Collector: Flow Summaries	collector_summaries_per_second	(19.3.2 and later) Total flow summaries processing rate for a single core PCE node, over the last hour. The sum of these should roughly match the flow summary ingest rate, or the PCE will show an increasing backlog size.	<ul style="list-style-type: none"> Warning: > 12,000 Critical: > 15,000
Collector: Success Rate	collector_post_success_count_per_hour	(19.3.2 and later) Informational. Total flow summary posts accepted by a core machine over the last hour. Posts can be of different sizes, so take longer to process, but you should see roughly the same rates for each core. If counts differ across core machines, ensure intra-PCE latency is within the 10ms limit.	
Collector: Failure Rate	collector_post_failure_count_per_hour	(19.3.2 and later) Informational. Total flow summary failure rate over the last hour. Under normal operational circumstances,	

Category	Key	Description	Severity Levels
		this value should be approximately the same for all core nodes.	
Collector: Failure Percentage	collector_post_failure_percentage	(19.3.2 and later) Failure/total. Failure rate / success ratio over the last hour.	<ul style="list-style-type: none"> Warning: > 10% Critical: > 20%
Traffic Summary: Ingest	traffic_summaries_per_second, total_traffic_summaries_per_second	(19.3.2 and later) The mean rate at which flow summaries are added to the postgresql database over the last hour.	<ul style="list-style-type: none"> Warning: > 12,000 Critical > 15,000
Traffic Summary: Database Size	traffic_database_size_gb, traffic_database_size_days	(19.3.2 and later) Informational. (gigabytes; days)	
Traffic Summary: Database Size: % of Allocated	traffic_database_utilization_percentage	(19.3.2 and later) Informational. The system is behaving normally even if it is near or at configured disk limits. The oldest flows will be dropped to enforce the limit, however, which may not be desirable.	<ul style="list-style-type: none"> Warning: > 10% Critical: > 50%
Traffic Summary: Backlog Size	traffic_backlog_size_gb	(19.3.2 and later) Amount of flows in the backlog that are not in the traffic database, in gigabytes. If the backlog size exceeds a certain limit (default is 10 GB and can be set in runtime environment), flows get dropped.	

Category	Key	Description	Severity Levels
Traffic Summary: Backlog Size: % of Allocated	traffic_backlog_utilization_percentage	(19.3.2 and later) Increasing values indicate that the buffered new flow data is growing, meaning the PCE is unable to keep up with the rate of data posted. The PCE collector flow summary rate and PCE traffic summary ingest rate need to be roughly equal or this buffered backlog will grow.	<ul style="list-style-type: none"> Warning: > 10% Critical: > 50%
Supercluster Replication Lag		<p>(For PCE superclusters only) Number of seconds since a replication event generated by a PCE was processed on another PCE. The supercluster replication engine relies on events to ensure data gets replicated. These are not the same as the PCE audit events.</p> <p>An increasing replication lag usually indicates some issue with the PCE replication engine or network connectivity. The larger the replication lag, the longer it may take a PCE to catch up with other regions once the underlying issue is addressed.</p>	<ul style="list-style-type: none"> Warning: This is an indication that the inter-pce data replication is not working as intended. One or more PCEs may not have the data generated by one or more other PCEs. The supercluster expects that the replication lag will not fall behind by a large margin. If it does, the user may lose some data if the PCE that is ahead fails and is not recoverable.

Support Reports for PCE

To help Illumio troubleshoot issues with your PCE, you can generate support reports to send to Illumio Customer Support. There are two ways to generate support bundles: in the web console or at the command line. The web console is the generally preferable technique.

NOTE:

To generate PCE Support Reports, you must be the Global Organization Owner for your PCE or a member of the Global Administrator role.

To download an already generated support report bundle from the web console, you must be the Global Organization Owner or Global Administrator. See [About Roles, Scopes, and Granted Access](#) for information.

Generate PCE Support Bundle in Web Console

The PCE web console has a Support Bundles page where you can generate PCE support reports. PCE support bundles can also be generated at the command line, but the web console provides a more convenient method which is accessible to more types of users.

To generate a support bundle:

1. Choose **Troubleshooting > PCE Support Bundles** from the main dropdown menu.
2. Click **Generate**.

The support bundle generation dialog box appears.

3. (Optional) Click **Log Collection** and specify the time range.
4. Click **Generate** again in the dialog box.

The dialog disappears. The PCE Support Bundles tab displays the report generation status for each node. When the reports for all nodes are complete, an aggregate support bundle is made available for download.

5. Click **Download**.

Up to five previously generated PCE support bundles remain available for download in a list on the PCE Support Bundles tab.

Generate PCE Support Report at Command Line

Use the PCE `support_report` command-line tool to generate several types of PCE Support Reports:

- **PCE Support Report:** Various diagnostic reports designed to provide Illumio Customer Support with PCE information, such as application logs, process information, and machine statistics.
- **PCE System Inventory Report:** An inventory of the PCE software and all the objects you have created and configured, such as total number of workloads, rules, ruleset scopes, labels, pairing profiles, the number of VENs deployed, OS on deployed VENs, and any modified (non-default) API or object limits.
- **PCE Host Inventory Report:** An inventory of the host, including information such as the number of processors configured on the host and the amount of physical disk space and memory being utilized.
- **PCE Support Report Search Function:** You can search PCE log files by string and by a date range.

The PCE saves the `support_report` command and its arguments in `report_log` so that you can see the command that was used to generate the support report.

Support Report Command-line Syntax

To create a Support Report, follows these general steps:

1. Enter the `support_report` command with options.
2. When you include `support_report` search options (for example, `from=` and `to=`, or combinations), enter the `support_report list` command after entering the search options.

The output is a date-stamped tar file. When the `support_report` command is finished, it displays the path to the file.

Support Report Option	Description
None	Does a system inventory.
system	Generates a node report and inventory report.
inventory	Generates an inventory report only.
list	Runs the report defined by the latest <code>support_report</code> options.
logs (+ optional search arguments)	Includes logs and the optional search criteria described in Search the PCE Log Files .
procs	Includes process details in the Support Report.
stats	Includes statistics in the Support Report.

Run PCE Support or Inventory Report at Command Line

To run the PCE Support Report:

1. To generate the PCE Support Report to collect inventory, logs, statistics, and processes, run this command:

```
$ sudo -u ilo-pce /opt/illumio-pce/illumio/bin/support_report inventory  
system stats procs logs
```

2. To view options for the Support Report, add the help option:

```
$ sudo -u ilo-pce /opt/illumio-pce/illumio/bin/support_report help
```

To run a PCE inventory report:

1. Make sure your shell environment is correctly set up by running this command:

```
$ source /opt/illumio-pce/illumio/bin/illumio/scripts/support
```

2. To run the PCE system inventory report, run this command:

```
$ sudo -u ilo-pce illumio-pce-env inventory system
```

3. To run the PCE host inventory report, run this command:

```
$ sudo -u ilo-pce illumio-pce-env inventory host
```

View Host and System Inventory

You can use the following commands to get a quick source of information for troubleshooting or when working with Illumio Customer Support. Using these commands is a quicker and less detailed alternative to running a PCE support report.

To show host inventory for the "local" node:

```
$ illumio-pce-env show host-inventory
```

To show system inventory for the PCE:

```
$ illumio-pce-env show system-inventory
```

To show host inventory for all PCE nodes and also the PCE system inventory:

```
$ illumio-pce-env show inventory
```


PCE HA and DR

This chapter contains the following topics:

PCE HA and DR Concepts	113
PCE HA and DR Requirements	115
PCE Replication and Failover	117
PCE Failures and Recoveries	126

This section describes how to achieve high availability (HA) for the PCE, and how to handle disaster recovery (DR) if a failure occurs.

PCE HA and DR Concepts

This section describes how the PCE provides high availability (HA) and disaster recovery (DR).

Overview of PCE HA and DR

The PCE provides high availability (HA). In the event of a failure, your PCE cluster's availability and operability can be maintained with zero or minimal data loss and no or limited human intervention, based on the type of failure that occurs.

HA for the PCE depends on the type and severity of failure that occurs. For example, in less severe, non-catastrophic failure cases, such as when a node is powered off, or network connection is lost, the cluster's availability is automatically re-established without human intervention and with no or limited data loss.

In other more severe disaster cases, such as part or all of the PCE is damaged or destroyed, the PCE is designed to be able to recover with minimal data loss and a minimum amount of human intervention.

In all PCE failure cases, the VENs continue to enforce the last known policy until the PCE is recovered.

Design Goals for PCE HA

The PCE is designed to handle system or network failures based on the following design goals:

- **Elimination of single points of failure:** A failure of one component (PCE node or service) does not mean failure of the entire PCE cluster. Recovery from failure is

done with zero or minimal loss of data.

- **Detection of failures as they occur:** The PCE detects failure without human intervention.
- **Reliable recovery:** Recovery from failure is done with zero or minimal loss of data.

Three conditions determine whether the PCE can survive a failure and remain available:

- Quorum
- Service availability
- Capacity

All these conditions must be met for the PCE to be available and provide acceptable performance.

Quorum

A PCE cluster relies on *quorum*, which is a sufficient number of servers to ensure consistent operation. Quorum prevents the so-called “split brain” case where two parts of the cluster are operating autonomously. Any node that becomes disconnected from the quorum is automatically isolated or “fenced” by shutting down most of its services.

All core nodes and the data0 node (an odd number) are voting members of the quorum. The data1 node is not a voting member. A majority of these nodes must be available to maintain quorum and elect a cluster leader.

When a cluster experiences a failure and doesn't have the majority of nodes functioning to maintain quorum, the cluster becomes unavailable until it recovers the minimal number of nodes.

In practice, this means that as long as at least one core node and one data node are available, the PCE remains operational but with restricted functionality.

Service Availability

Another key requirement of PCE high availability is service availability, which means at least one instance of all required PCE services are available.

The Service Discovery Service (SDS) monitors all services running on each node in the cluster. This service must be monitored for failure. See [Monitor PCE Health](#) for information.

For a PCE cluster to provide all its necessary services, even in the event of a partial cluster failure, it must contain at least one functioning data node and at least one core node, with all services fully available on each node.

Node Type	Service Tiers
Core	<ul style="list-style-type: none"> • Front end • Processing • Service and caching
Data	<ul style="list-style-type: none"> • Service and caching • Data persistence (database)

Capacity

Cluster capacity means that at any given time, the PCE is able to provide sufficient compute resources to meet the demands required by the number of workloads deployed.

PCE 2x2 and 4x2 clusters are sized to support the loss of one data node plus half the total number of core nodes and still operate with degraded performance (1+1 redundancy). When more than one data node plus half the total number of core nodes in the cluster is lost, the cluster might not have sufficient capacity to meet demands.

PCE HA and DR Requirements

This section describes how to ensure your underlying systems are sufficient to successfully provide high availability (HA) and disaster recovery (DR) features. Check all of the following system requirements.

PCE Cluster Front End Load Balancing

In order for a PCE cluster to provide high availability, it requires a front-end load balancer to manage traffic distribution and system health checking for the PCE.

The load balancer must be customer-provided and managed, and is not included as part of the PCE software distribution. You have the option of using a traffic load balancer or DNS load balancer.

IMPORTANT:

The load balancer must be able to run application level health checks on each of the core nodes in the PCE cluster, so it can be aware at all times whether each node is available to service requests.

Traffic Load Balancer Requirements

The PCE requires the following traffic load balancer configuration:

- Layer 4 with Secure Network Address Translation (SNAT)
- Least connection (recommended) or round robin load balancing to core nodes
- HTTP health checks from load balancer to core nodes
- High availability capabilities
- A virtual IP (VIP) configured in the `runtime_env.yml` parameter `cluster_public_ips`

For information about setting this parameter, see [Reference: PCE Runtime Parameters](#) in the *PCE Installation and Upgrade Guide*.

NOTE:

Using a traffic load balancer is recommended over DNS, because it provides a quicker failure response, while DNS load balancing typically has a longer failover time.

DNS Load Balancing

Another option for load balancing the PCE cluster is using DNS, where traffic is load balanced to the core nodes based on DNS rather than connection-based load balancing.

When you plan to use DNS for load balancing the PCE software, the PCE requires the following DNS load balancer configuration:

- Round robin load balancing to core nodes
- 30 to 60 second TTL to allow for quick failover
- PCE core node IP addresses configured in the `runtime_env.yml` parameter named `cluster_public_ips`
- HTTP health checks from the load balancer to core nodes

The DNS must be able to run health checks against the PCE `node_available` API, and the DNS load balancer should only serve IP addresses for the cluster FQDN of those nodes that respond to the `node_available` API. See [Node Availability](#) in the *REST API Developer Guide* for more information.

Network Latency Between Nodes

Ensure that network latency between and among the nodes of the clusters does not exceed 10ms. Proper operation of Illumination and Explorer is assured when latency is

10ms or less.

PCE Replication and Failover

(Not supported for Supercluster.)

To increase reliability, you can set up replication and failover for PCEs. Having a PCE on "warm standby," ready to take over if the active PCE fails, contributes to a resilient disaster recovery (DR) plan.

For PCE replication and failover, set up PCEs in pairs. Each pair has an active PCE and a standby PCE. A combination of continuous real-time replication and periodic synchronization is used to keep the standby PCE's data up to date with the active PCE. If the active PCE fails, the standby PCE can take over and become the new active PCE.

The data from the following services are replicated:

- database_service
- citus_coordinator_service
- reporting_database_service
- agent_traffic_redis_server
- data_job_queue_redis_service
- fileserver

Standby PCE Prerequisites

WARNING:Active Standby assumes the same certificate is used for all nodes of the cluster. You cannot use a unique certificate per Core node.

WARNING:The user/secret variable must be set as the ilo-pce user. Alternatively, you need to run it as `sudo -E -u ilo-pce`.

Before designating a standby PCE, perform the following preparation steps.

Set Up Two PCEs

Install PCE software on two machines or find two machines where it is already installed. Be sure the following are true:

- Hardware configuration and capacity are as near identical as possible on the two PCEs.
- PCE software version is the same on both PCEs.

Reset Any Repurposed PCE

If you are repurposing an existing PCE to be the standby, be sure the PCE is completely reset.

1. *On all nodes of the existing PCE*, run the following command to reset the PCE:

```
$ sudo -u ilo-pce illumio-pce-ctl reset
```

2. *On all nodes of the existing PCE*, run the following command to start the PCE and set it to runlevel 1:

```
$ sudo -u ilo-pce illumio-pce-ctl start --runlevel 1
```

3. *On any one data node of the existing PCE*, run the following command to set up the database:

```
$ sudo -u ilo-pce illumio-pce-db-management setup
```

Open Ports Between Active and Standby PCEs

Be sure the required ports are open on both PCEs to allow network traffic between the active PCE and the standby PCE so data replication can occur. Make sure that all the same service ports are opened on the standby PCE and the active PCE. For a list of the required ports, see [Port Ranges for Cluster Communication](#) in the PCE Installation and Upgrade Guide.

Set Up FQDNs

Set up the FQDNs that are required when using active and standby PCEs:

- FQDN of the active PCE.
- FQDN of the standby PCE.
- In the `runtime_env.yml` file, `active_standby_replication:active_pce_fqdn` is always the FQDN of the currently active PCE.

Add `active_standby_replication:active_pce_fqdn` to the `runtime_env.yml` file on both PCEs, active and standby. Example:

```
pce_fqdn: FQDN of the active PCE

active_standby_replication:
  active_pce_fqdn: active-pce-fqdn.com
```

WARNING:
Whether the PCE runs in a standalone or active-standby mode, never remove the setting `active_pce_fqdn` from `runtime_env.yml`. VENS are paired using this FQDN. Removing this entry will break VEN communications.

There are two options for setting up these FQDNs.

Option 1: Use a new FQDN for `active_standby_replication:active_pce_fqdn`.

You can use a FQDN that is not currently assigned to either the active PCE or the standby PCE. Use this option if you do not want to update the FQDN of the currently active PCE. The FQDN assigned to `active_pce_fqdn` should resolve to the currently active PCE. For example:

```
Existing Setup
Active PCE:
  pce_fqdn: active-pce.com

Standby PCE:
  pce_fqdn: standby-pce.com

Before Standby is Set Up

Active PCE:
  pce_fqdn: active-pce.com
  active_standby_replication:
    active_pce_fqdn: active-pce-global.com

Standby PCE:
  pce_fqdn: standby-pce.com
  active_standby_replication:
    active_pce_fqdn: active-pce-global.com
```

The `active_pce_fqdn` always contains the FQDN of the PCE that is currently active in the active-standby pair. When a standby PCE is set up, the VEN master configuration is updated if needed so that it contains the `active_pce_fqdn` FQDN. After the standby PCE is set up, VENS paired to the active PCE contain the `active_pce_fqdn` in their master configuration. If the standby PCE is promoted, reconfigure the load balancer or GTM so that `active_pce_fqdn` resolves to the promoted (new active) PCE.

Option 2: Use the FQDN of the active PCE for `active_standby_replication:active_pce_fqdn`.

You might have scripts that use the `pce_fqdn` of the active PCE. In this case, it is easier to set `active_pce_fqdn` to the same value. Before you set up the standby PCE, change the `pce_fqdn` of the active PCE to something other than the `active_pce_fqdn`. For information about how to do this, see [Update PCE Configuration](#). If necessary, reconfigure your load balancer or global traffic manager (GTM) so that `active_pce_fqdn` and the new `pce_fqdn` of the active PCE resolve to the active PCE. For example:

Existing Setup

Active PCE:

```
pce_fqdn: active-pce.com
```

Standby PCE:

```
pce_fqdn: standby-pce.com
```

Before Standby is Set Up

Active PCE:

```
pce_fqdn: active-pce-updated.com
active_standby_replication:
  active_pce_fqdn: active-pce.com
```

Standby PCE:

```
pce_fqdn: standby-pce.com
active_standby_replication:
  active_pce_fqdn: active-pce.com
```

(Optional) Set DNS TTL Value

The DNS TTL (time to live) setting affects how long it takes for a new active PCE to take over in a failover situation. Consider adjusting the DNS TTL to avoid any delay. A

shorter value, such as 30 minutes, is recommended.

Set Up PCE Certificates

The SSL certificate must include all three FQDNs that are described in [Set Up FQDNs](#).

Set Up VEN Library

The PCE acts as a repository for distributing, installing and upgrading the VEN software. Install or update the VEN library on both the active and standby PCEs. See the VEN Installation and Upgrade Guide.

NOTE:

Be sure the VEN versions in the library are supported by the PCE version that is installed.

Set Up a Standby PCE

To set up a standby PCE and associate it with its active PCE partner, use the following steps.

1. Complete the prerequisite steps in [Standby PCE Prerequisites](#).
2. *On the active PCE*, generate an API key. This API key is used only while setting up the standby PCE.
3. Bring the standby PCE to runlevel 2. *On any node of the standby PCE*, run the following command:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 2
```

The active PCE can remain at runlevel 5.

4. On the standby PCE, run the following commands to set up authentication. In *username*, give the active PCE's API key authentication username. In *secret*, give the API key secret.

```
$ export ILO_ACTIVE_PCE_USER_NAME=username  
$ export ILO_ACTIVE_PCE_USER_PASSWORD=secret
```

5. Link the standby PCE to its active PCE. On the standby PCE, run the following command. For *active_pce_fqdn:front_end_management_https_port*, give the FQDN and port of the current active PCE. The value in `--active-pce` is not the

same as `active_pce_fqdn` in the configuration file `runtime_env.yml`.

```
$ sudo -u ilo-pce --preserve-env illumio-pce-ctl setup-standby-pce --active-pce active_pce_fqdn:front_end_management_https_port
```

WARNING:

Do not bring the standby PCE to runlevel 5.

6. After replication is set up for the first time, the status of some services, such as the `citus_coordinator_service`, might be NOT RUNNING for a long time, and the cluster status is stuck in PARTIAL. This is usually because the service is performing a database backup, which can take time depending on network latency, disk IOPS, traffic flow, and traffic data size. To check whether the backup process is running, use the following command:

```
$ ps -ef | grep pg
```

Example output:

```
pce      84742 73150 18 16:25 ?          00:04:42
          /var/illumio_pce/external/bin/pg_basebackup -d host=10.31.2.172 port=5532 -
D /var/traff_dir/traffic_datastore -v -P -X stream -c fast
pce      84747 84742  7 16:25 ?          00:01:54
          /var/illumio_pce/external/bin/pg_basebackup -d host=10.31.2.172 port=5532 -
D /var/traff_dir/traffic_datastore -v -P -X stream -c fast
```

WARNING:

If the citus coordinator service is busy with a backup, do not restart services yet. Wait until this operation is complete and the service status changes to RUNNING.

7. Restart services on the active PCE. On *any node* of the active PCE, run the following command:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-restart
```

For example:

```
$ export ILO_ACTIVE_PCE_USER_NAME=api_17abrwerwe
$ export ILO_ACTIVE_PCE_USER_PASSWORD=6efefeafe34ewroopp11494934kdf
$ sudo -u ilo-pceillumio-pce-ctl setup-standby-pce --active-pce
active.pce.com:8443
$ sudo -u ilo-pce illumio-pce-ctl cluster-restart
```

Failover to Standby PCE

This section tells how to perform a PCE failover for disaster recovery (DR). The active PCE has failed, and you need to promote the standby PCE so it can take over as the active PCE. Follow these steps.

1. Check to be sure the PCE you are about to promote is actually a standby PCE and that it is at runlevel 2.

```
$ sudo -u ilo-pce illumio-pce-ctl active-standby?
```

The output should say "standby."

2. Check to be sure the active PCE has failed and is offline. There must not be any data replicating to the standby PCE. *On every node of the active PCE*, run the following command:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status
```

The output should contain STOPPED. Be sure to repeat this command on every node of the PCE.

3. *On the standby PCE*, run the following command to promote the standby PCE.

```
$ sudo -u ilo-pce illumio-pce-ctl promote-standby-pce
```

When the active PCE is down, this command promotes this PCE to be the new primary. If the active PCE is not down, the standby PCE will not be promoted, and a message like "Active PCE is still reachable" is generated.

4. Make sure that DNS recognizes this as the new active PCE FQDN so devices in your network can find the PCE. Make sure that the values for both `active_standby_replication` and `active_pce_fqdn` in the configuration file `runtime_env.yml` are the PCE FQDN of the former standby (new active) PCE. For example,

reconfigure the PCE FQDN on load balancers. The steps depend on your devices and configuration. For more information about the PCE FQDN, see [Standby PCE Prerequisites](#).

5. Check the VEN synchronization status by running the following command:

```
$ sudo -u ilo-pce illumio-pce-ctl promote-standby-check
```

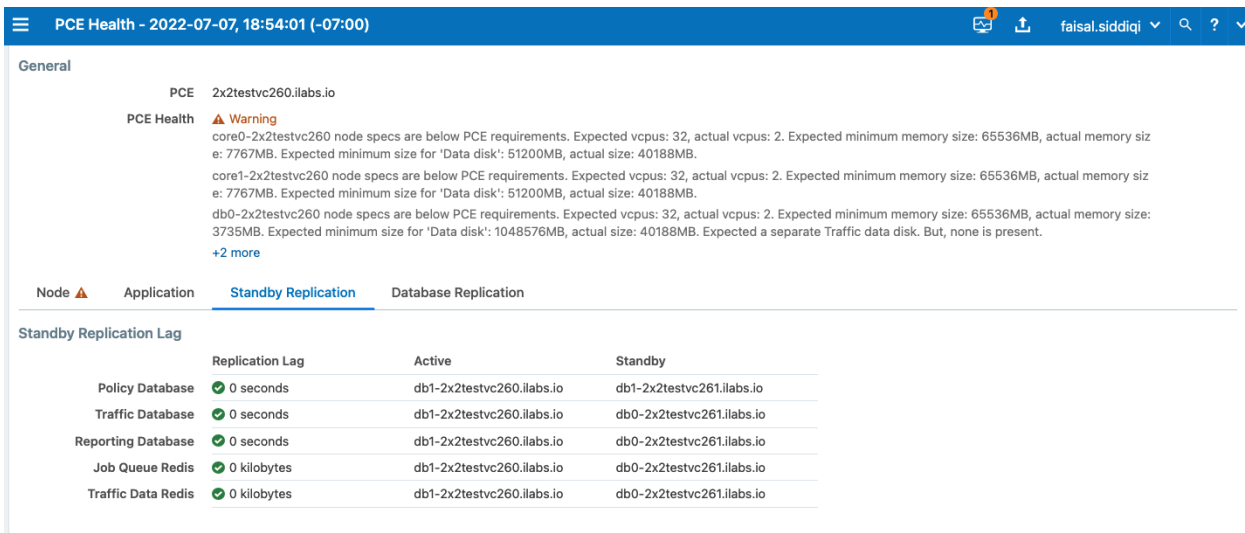
Run the command repeatedly and watch the output to make sure the VEN sync count increases. This indicates that the DNS change is in effect and the new active PCE has been promoted successfully.

The DNS update for the new PCE FQDN can take some time, depending on the DNS TTL value.

6. When you are ready, connect a new standby PCE to the new active PCE. Repeat the steps in [Standby PCE Prerequisites](#) and [Set Up a Standby PCE](#).

Monitoring Replication

In the Health page of the PCE web console, use the Standby Replication tab to monitor replication between the active PCE and standby PCE. The Standby Replication tab shows the replication lag on the active and standby PCEs for the traffic database, the policy database, the reporting database, the job queue redis, and traffic data redis. (The fileserver is not shown.)



PCE Health - 2022-07-07, 18:54:01 (-07:00)

General

PCE 2x2testvc260.ilabs.io

PCE Health ▲ Warning

core0-2x2testvc260 node specs are below PCE requirements. Expected vcpus: 32, actual vcpus: 2. Expected minimum memory size: 65536MB, actual memory size: 7767MB. Expected minimum size for 'Data disk': 51200MB, actual size: 40188MB.

core1-2x2testvc260 node specs are below PCE requirements. Expected vcpus: 32, actual vcpus: 2. Expected minimum memory size: 65536MB, actual memory size: 7767MB. Expected minimum size for 'Data disk': 51200MB, actual size: 40188MB.

db0-2x2testvc260 node specs are below PCE requirements. Expected vcpus: 32, actual vcpus: 2. Expected minimum memory size: 65536MB, actual memory size: 3735MB. Expected minimum size for 'Data disk': 1048576MB, actual size: 40188MB. Expected a separate Traffic data disk. But, none is present.

[+2 more](#)

Node ▲ Application **Standby Replication** Database Replication

Standby Replication Lag

	Replication Lag	Active	Standby
Policy Database	✔ 0 seconds	db1-2x2testvc260.ilabs.io	db1-2x2testvc261.ilabs.io
Traffic Database	✔ 0 seconds	db1-2x2testvc260.ilabs.io	db0-2x2testvc261.ilabs.io
Reporting Database	✔ 0 seconds	db1-2x2testvc260.ilabs.io	db0-2x2testvc261.ilabs.io
Job Queue Redis	✔ 0 kilobytes	db1-2x2testvc260.ilabs.io	db0-2x2testvc261.ilabs.io
Traffic Data Redis	✔ 0 kilobytes	db1-2x2testvc260.ilabs.io	db0-2x2testvc261.ilabs.io

Another way that the PCE administrator can monitor replication is by watching the service discovery log for WAL segment missing errors. This error may appear when

the standby traffic database service could not keep up with synchronization from the active traffic database service. When this error occurs, the log looks like the following:

```
2022-06-30T15:43:19.556560+00:00 level=warning host=db0-4x2systest50 ip=127.0.0.1
program=illumio_pce/service_discovery| sec=603799.555 sev=ERROR pid=12416 tid=2440
rid=0 [citus_coordinator_service] Health Check: WAL segment 105/2B95FD98 is
missing. Full base backup marker file set.
```

When this situation arises, the `citus_coordinator_service` causes the service to restart and perform the full database backup again. The network latency, disk IOPS, traffic flow, and traffic data size affect the replication latency. If you experience this issue, make any improvements you can to these factors.

For example, you can increase the value of the `wal_keep_segments` setting in the `traffic_datastore` section of the `runtime_env.yml` configuration file. Increasing this value comes at the expense of disk space cost. Each WAL segment is 16 MB, so 5120 WAL segments would use about 82 GB of extra space.

```
traffic_datastore:
  wal_keep_segments: 5120
```

Limitations and Constraints

When using active and standby PCEs for replication, be aware of the following limitations and constraints:

- Fileserver replication lag is not shown in the Standby Replication tab of the Health page.
- Support reports are replicated, but support bundles are not replicated.
- In an active-standby PCE pair, it is not necessary to perform database backups in the same way you would with a standalone PCE. However, if you wish to do so, take the backups from the active PCE. It is also not normally necessary to restore a database backup on the active PCE or the standby PCE. If one of the PCEs fails, the other takes over as active PCE, and it already has an up-to-date copy of the data because of the ongoing replication between the two PCEs.

WARNING:

If it becomes necessary to restore data from a backup (for example, if both PCEs fail), you must restore the same backup to both the active PCE and the standby PCE.

PCE Failures and Recoveries

This section describes how the PCE handles various types of possible failures. It tells whether the failure can be handled automatically by the PCE and, if not, what manual intervention you need to perform to remedy the situation.

Types of PCE Failures

These are the general kinds of failures that can occur with a PCE deployment:

- **PCE-VEN network partition:** A network partition occurs that cuts off communication between the PCE and VENS.
- **PCE service failure:** One or more of the PCE's services fail on a node.
- **PCE node failure:** One of the PCE's core or data nodes fails.
- **PCE split cluster failure (site failure):** One data plus half the total number of core nodes fail.
- **PCE cluster network partition:** A network partition occurs between two halves of a PCE cluster but all nodes are still functioning.
- **Multi-node traffic database failure:** If the traffic database uses the optional multi-node configuration, the coordinator and worker nodes can fail.
- **Complete PCE failure:** The entire PCE cluster fails or is destroyed and must be rebuilt.

Failure-to-Recovery Stages

For each failure case, this document provides the following information (when applicable):

Stage	Details
Preconditions	Any required or recommended pre-conditions that you are responsible for to be able to recovery from the failure. For example, in some failure cases, Illumio assumes you regularly exported a copy of the primary database to an external system in case you needed to recover the database.
Failure beha-	The behavior of the PCE and VENS from the time the failure occur to

Stage	Details
vior	recovery. Can be caused by the failure itself, or by execution of recovery procedures.
Recovery	<p>A description of how the system recovers from the failure incident to resume operations, which might be automatic or require manual intervention on the PCE or VEN. When intervention is required, the steps are provided.</p> <p>Includes the following items:</p> <ul style="list-style-type: none"> • Recovery type: Whether the PCE and VENs can automatically recover from the failure or human intervention is required to resume operations. • Recovery procedure (when required): When human intervention is required on the PCE or VENs, the recovery procedures are provided. • Recovery Time Objective (RTO): The average time it takes to detect and recover from a failure. • Recovery Point Objective (RPO): The amount of data loss due to the failure.
Full Recovery (not always applicable)	In some cases, additional steps might be required to revert the PCE back to its normal, pre-failure operating state. This situation is usually a planned activity that can be scheduled.

Legend for PCE Failure Diagrams


The following diagram symbols illustrate the affected parts of the PCE in a failure:

- **Dotted red line:** Loss of network connectivity, but all nodes are still functioning
- **Dotted red X:** Failure or loss of one or more nodes, such as when a node is shut down or stops functioning

PCE-VEN Network Partition

In this failure case, a network partition occurs between the PCE and VENs, cutting off communication between the PCE and all or some of its VENs. However, the PCE and VENs are still functioning.

Disaster Recovery: Failure Scenario 1




Stage	Details
Preconditions	None
Failure Behavior	<p>PCE</p> <ul style="list-style-type: none"> • Users cannot provision any changes to the VENs until the connection is re-established. • The information displayed in the Illumination map in the PCE web console is only as current as the last time the VENs reported to the PCE. • The PCE ignores any disconnected VENs until at least one hour has passed. • When the outage persists longer than one hour, the PCE marks unreachable VENs as offline. When any existing policy allows the offline VENs to communicate with other VENS, the PCE recalculate its current policy and exclude those workloads marked as offline. <p>VENS</p> <ul style="list-style-type: none"> • VENs continue to enforce their last known good policy. • All VEN state and flow updates are cached locally on the workload where the VEN is installed. The VEN stores up to 24 hours of flow data then purges the oldest data first during an extended event. • After missing 3 heartbeats (approximately 15 minutes), the VEN

Stage	Details
	enters a degraded state, during which the VEN ignores all asynchronous commands received as lightning bolts from the PCE, except commands that initiate software upgrade and Support Reports.
Recovery	<ul style="list-style-type: none">• Recovery type: Automatic. The VEN tries to connect to the PCE every 5 minutes. After PCE-VEN network connectivity is restored, the VENs automatically reconnect to the PCE and resume normal operations:<ul style="list-style-type: none">◦ Policy for the VEN is automatically synchronized (when new policy from PCE was provisioned during failure).◦ Cached state and flow data from the VEN is sent to the PCE.◦ After three successful heartbeats (approximately 15 minutes), the VEN comes out of the degraded state.• Recovery procedure: None required.• RTO: Customer dependent based on the time it takes for PCE-VEN network connectivity to be restored, plus approximately 15 minutes for three successful heartbeats.• RPO: Zero.

Service Failure

In this failure case, one of the PCE's services fails on a node.

Disaster Recovery: Failure Scenario 2



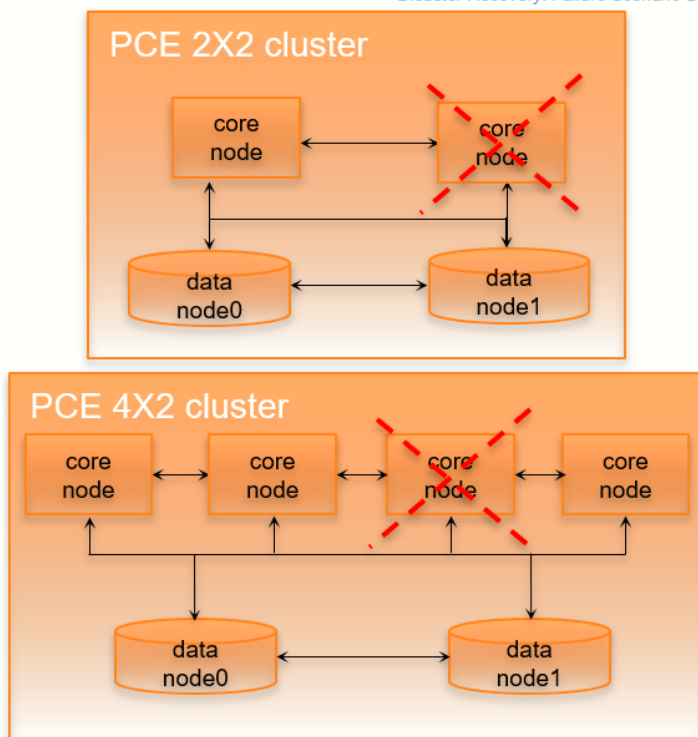
Stage	Details
Preconditions	None.
Failure Behavior	<p>PCE</p> <ul style="list-style-type: none"> The PCE might be temporarily unavailable. Users might be unable to log into the PCE web console. The PCE might return an HTTP 502 response and the <code>/node_available</code> API request might return an HTTP 404 error. Other services that are dependent on the failed services might be restarted within the cluster. <p>VENs</p> <ul style="list-style-type: none"> VENs are not affected. VENs continue to enforce the current policy. When a VEN misses a heartbeat to the PCE, it retries in 5 minutes.
Recovery	<ul style="list-style-type: none"> Recovery type: Automatic. The PCE's SDS ensures that all PCE services are running, including itself. When any service fails, SDS restarts it. Recovery procedure: None required. RTO: Variable depending on which service failed and how many

Stage	Details
	dependent services must be restarted. Typically 30 seconds to 2 minutes. • RPO: Zero.

Core Node Failure

In this failure case, one of the core nodes completely fails. This situation occurs anytime a node is not communicating with any of the other nodes in the cluster; for example, a node is destroyed, the node's SDS fails, or the node is powered off or disconnected from the cluster.

Disaster Recovery: Failure Scenario 2




Stage	Details
Preconditions	The load balancer must be able to run application level health checks on each of the core nodes in the PCE cluster, so that it can be aware at all times whether a node is available. IMPORTANT: When you use a DNS load balancer and need to provision a new core node to recover from this failure, the <code>runtime_env.yml</code> file parameter named <code>cluster_public_ips</code> must include the IP address of your existing core nodes and the IP addresses of the replacement nodes.

Stage	Details
	<p>When this is not configured correctly, VENs will not have outbound rules programmed to allow them to connect to the IP address of the replacement node. Illumio recommends that you preallocate these IP addresses so that, in the event of a failure, you can restore the cluster and the VENs can communicate with the replacement node.</p>
<p>Failure Behavior</p>	<p>PCE</p> <ul style="list-style-type: none"> • The PCE is temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return an HTTP 502 response and the <code>/node_available</code> API call might return an HTTP 404 error. • Other services that are dependent on the failed services might be restarted within the cluster. <p>VENs</p> <ul style="list-style-type: none"> • VENs are not affected. • VENs continue to enforce the current policy. • When a VEN misses a heartbeat to the PCE, it retries in 5 minutes.
<p>Recovery</p>	<ul style="list-style-type: none"> • Recovery type: Automatic. The cluster has multiple active core nodes for redundancy. • Recovery procedure: None required. • RTO: 5 minutes. • RPO: Zero. No data loss occurs because the core nodes are stateless.
<p>Full Recovery</p>	<p>Either recover the failed node or provision a new node and join it to the cluster.</p> <p>For information, see Replace a Failed Node.</p>

Data Node Failure

In this failure case, one of the data nodes completely fails.

Disaster Recovery: Failure Scenario 4




Stage	Details
Preconditions	<p>You should continually monitor the replication lag of the replica database to make sure it is in sync with the primary database.</p> <p>You can accomplish this precondition by monitoring the <code>illumio_pce/system_health</code> syslog messages or by running the following command on one of the <i>data nodes</i>:</p> <pre data-bbox="451 1291 1421 1375">\$ sudo -u ilo-pce illumio-pce-db-management show-replication-info</pre>
Failure Behavior	<p>PCE</p> <ul data-bbox="440 1455 1421 1862" style="list-style-type: none"> • The PCE is temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return a HTTP 502 response and the <code>/node_available</code> API call might return an HTTP 404 error. • Other services that are dependent on the failed services might be restarted within the cluster. • When the <code>set_server_redis_server</code> service is running on the failed data node, the VENs go into the syncing state and policy is recomputed for each VEN, even when no new policy has been pro-

Stage	Details
	<p>visioned. The CPU usage on the PCE core nodes might spike and stay at very high levels until policy computation is completed.</p> <p>VENs</p> <ul style="list-style-type: none"> • VENs are not affected and continue to enforce the current policy. • When a VEN misses a heartbeat to the PCE, it retries in 5 minutes.
Recovery	<ul style="list-style-type: none"> • Recovery type: Automatic. The PCE detects this failure and automatically migrates any required data services to the surviving data node. When the failed node is the primary database, the PCE automatically promotes the replica database to be the new primary database. • Recovery procedure: None required. • RTO: 5 minutes, with the following caveats for specific PCE services: <ul style="list-style-type: none"> ◦ <code>set_server_redis_server</code>: Additional time is required for all VENs to synchronize. This time is variable based on the number of VENs and complexity of the policy. • RPO: Service-specific based on the data services that were running on the failed data node. <ul style="list-style-type: none"> ◦ <code>database_service</code>: Implies the failed data node was the primary database. All data committed to the primary database, and not replicated to the replica, is lost. Typically under one second. ◦ <code>database_slave_service</code>: Implies the failed data node is the replica database. No data is lost. ◦ <code>agent_traffic_redis_server</code>: All traffic data is lost. ◦ <code>fileserver_service</code>: All asynchronous query requests and Support Reports are lost.
Full Recovery	<p>When the failed data node is recovered or a new node is provisioned, it registers with PCE and is added as an active member of the cluster. This node is designated as the replica database and will replicate all the data from the primary database.</p> <p>For recovery information, see Replace a Failed Node.</p>

Primary Database Doesn't Start

In this failure case, the database node fails to start.

Disaster Recovery: Failure Scenario 4



Stage	Details
Preconditions	The primary database node does not start.
Failure Behavior	The database cannot be started. Therefore, the entire PCE cluster cannot be started.
Full Recovery	<p>Recovery type: Manual. You have two recovery options:</p> <ul style="list-style-type: none"> Find the root cause of the primary database failure and correct it. Contact Illumio Customer Support for assistance if needed. Promote the replica data node to be the primary data node. <div style="border: 1px solid red; padding: 10px; margin: 10px 0;"> <p>WARNING: Promoting a replica to primary risks data loss</p> <p>Illumio strongly recommends that this option be a last resort because of the potential for data loss.</p> </div> <p>When you decide on the second option, see Configure Data1 and Core Nodes as Standalone Cluster.</p> <p>When the PCE Supercluster is affected by this problem, you must also restore data on the promoted primary database.</p>

Primary Database Doesn't Start When PCE Starts

In this failure case, the database node fails to start when the PCE starts or restarts.

The following recovery information applies only when the PCE starts or restarts. When the PCE is already running and the primary database node fails, database failover will occur normally and automatically, and the replica database node will become the primary node.

Stage	Details
Preconditions	The primary database node does not start during PCE startup. This issue could occur because of an error on the primary node. Even when no error occurred, you might start the replica node first and then be interrupted, causing a delay in starting the primary node that exceeds the timeout.
Failure Behavior	The database cannot be started. Therefore, the entire PCE cluster cannot be started.
Full Recovery	<p>Recovery type: Manual. You have two recovery options:</p> <ul style="list-style-type: none"> • Find and correct the root cause of the primary database failure. Contact Illumio Customer Support for help if needed. • Promote the replica data node to primary data node. <div style="border: 1px solid red; padding: 10px; margin: 10px 0;"> <p>WARNING: Promoting replica to primary risks data loss</p> <p>Consider this option as a last resort because of the potential for data loss, depending on the replication lag.</p> </div> <p>When you decide on the second option, on the <i>replica database node</i>, run the following command:</p> <pre style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;">\$ sudo ilo-pce illumio-pce-ctl promote-data-node <core-node-ip-address></pre> <p>This command promotes the node to be the primary database for the cluster whose leader is at the specified IP address.</p>

Site Failure (Split Clusters)

In this failure type, one of the data nodes plus half the total number of core nodes fail, while the surviving data and remaining core nodes are still functioning.

For example:

In a 2x2 deployment, a split cluster failure means the loss of one of these node combinations:

- Data0 and one core node
- Data1 and one core node

In a 4x2 deployment, a split cluster failure means the loss of one of these node combinations::

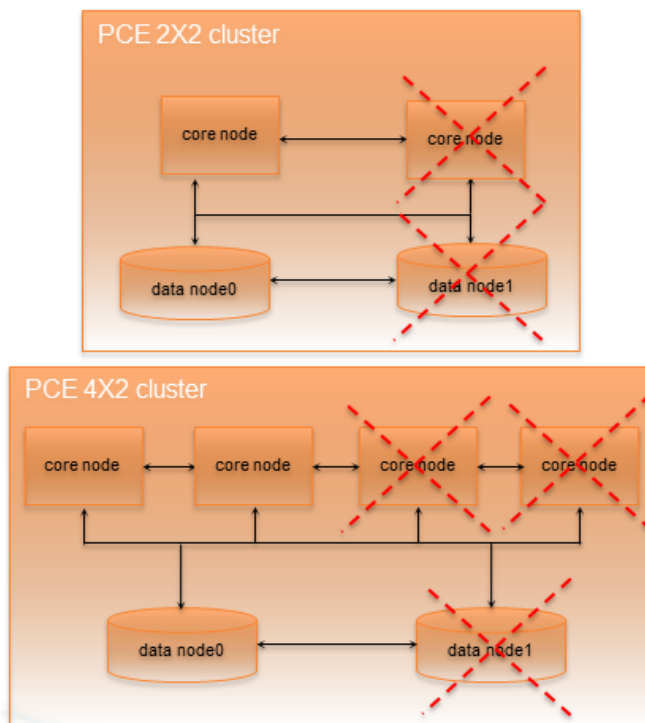
- Data0 and two core nodes
- Data1 and two core nodes

This type of failure can occur when the PCE cluster is split across two separate physical sites or availability zones with network latency greater than 10ms, and a site failure causes half the nodes in the cluster to fail. A site failure is one case that can cause this type of failure; however, split cluster failures can also occur in a single site deployment when multiple nodes fail simultaneously for any reason.

Split Cluster Failure Involving Data1

In this failure case, data1 and half the core nodes completely fail.

Disaster Recovery: Failure Scenario 6




Stage	Details
Preconditions	None.
Failure Behavior	<p>PCE</p> <ul style="list-style-type: none"> The PCE is temporarily unavailable. Users might be unable to log into the PCE web console. The PCE might return a HTTP 502 response and the <code>/node_available</code> API request might return an HTTP 404 error. Other services that are dependent on the failed services might be restarted within the cluster. <p>VENs</p> <ul style="list-style-type: none"> VENs are not affected. VENs continue to enforce the current policy. When a VEN misses a heartbeat to the PCE, it retries in 5 minutes.
Recovery	<ul style="list-style-type: none"> Recovery type: Automatic. Because quorum is maintained, the data0 half of the cluster can operate as a standalone cluster. When data1 is the primary database, the PCE automatically promotes data0 to be the new primary database. Recovery procedure: None. RTO: 5 minutes. RPO: Service specific based on which data services were running on data1 at the time of the failure: <ul style="list-style-type: none"> <code>database_service</code>: Data1 node was the primary database. All database data committed on data1 and not replicated to data0 is lost. Typically under one second. <code>database_slave_service</code>: Data1 node was the replica database. No database data is lost. <code>agent_traffic_redis_server</code>: All traffic data is lost. <code>fileserver_service</code>: All asynchronous query requests and Support Reports are lost.
Full Recovery	<p>Either recover the failed nodes or provision new nodes and join them to the cluster.</p> <p>For recovery information, see Replace a Failed Node.</p>

Split Cluster Failure Involving Data0

In this failure case, data0 and half of the total number of core nodes completely fail.

Disaster Recovery: Failure Scenario 7



Stage	Details
Preconditions	<p>CAUTION: When reverting the standalone cluster back to a full cluster, you must be able to control the recovery process so that each recovered node is powered on and re-joined to the cluster one node at a time (while the other recovered nodes are powered off). Otherwise, the cluster could become corrupted and need to be fully rebuilt.</p>
Failure Behavior	<p>PCE</p> <ul style="list-style-type: none"> The PCE is unavailable because it does not have the minimum number of nodes to maintain quorum. <p>VENs</p> <ul style="list-style-type: none"> The VEN continues to enforce its last known good policy. The VEN's state and flow updates are cached locally on the workload where the VEN is installed. The VEN stores up to 24 hours of flow data, then purges the oldest data first during an extended event. After missing 3 heartbeats (approximately 15 minutes), the VEN

Stage	Details
	enters a degraded state. While it is in the degraded state, the VEN ignores all asynchronous commands received as lightning bolts from the PCE, except the commands that initiate software upgrade and Support Reports.
Recovery	<ul style="list-style-type: none"> • Recovery type: Manual intervention is required to recover from this failure case. • Recovery procedure: See Configure Data1 and Core Nodes as Standalone Cluster for information. • RTO: Customer dependent based on how long it takes you to detect this failure and perform the manual recovery procedures. • RPO: Service specific based on which data services were running on data0 at the time of the failure: <ul style="list-style-type: none"> ◦ database_service: Data0 node was the primary database. All database data committed on data0 and not replicated to data1 is lost. Typically under one second. ◦ database_slave_service: Data0 node was the replica database. No database data is lost. ◦ agent_traffic_redis_server: All traffic data is lost. ◦ fileserver_service: All asynchronous query requests and Support Reports are lost.
Full Recovery	See Revert Standalone Cluster Back to a Full Cluster for information.

Configure Data1 and Core Nodes as Standalone Cluster

To enable the surviving data1 and core nodes to operate as a standalone 2x2 or 4x2 cluster, follow these steps in this exact order.

1. On the *surviving data1 node and all surviving core nodes*, stop the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

2. On *any surviving core node*, promote the core node to be a standalone cluster leader:

```
$ sudo -u ilo-pce illumio-pce-ctl promote-cluster-leader
```

3. On the *surviving data1 node*, promote the data1 node to be the primary database for the new standalone cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl promote-data-node <promoted-core-node-ip-address>
```

For the IP address, enter the IP address of the promoted core node from step 2.

4. **(4x2 clusters only)** On the *other surviving core node*, join the surviving core node to the new standalone cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-join <promoted-core-node-ip-address> --split-cluster
```

For the IP address, enter the IP address of the promoted core node from step 2.

5. Back up the surviving data1 node. For information, see [Back Up the Policy Database](#).

Revert Standalone Cluster Back to a Full Cluster

To revert back to a 2x2 or 4x2 cluster, follow these steps in this exact order:

IMPORTANT:

When you plan to recover the failed nodes and the PCE software is configured to auto-start when powered on (the default behavior for a PCE RPM installation), you *must* power on every node and re-join them to the cluster *one node at a time*, while the other nodes are powered off and the PCE is *not* running on the other nodes. Otherwise, your cluster might become corrupted and need to be fully rebuilt.

1. Recover one of the failed core nodes or provision a new core node.
2. If you provisioned a new core node, run the following command on any existing node in the cluster (not the new node you are about to add). For *ip_address*, substitute the IP address of the new node.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-nodes allow ip_address
```

3. On the *recovered or new core node*, start the PCE software and enable the node to join the cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-join <promoted-core-node-ip-address>
```

For the IP address, enter the IP address of the promoted core node.

4. **(4x2 clusters only)** For the *other recovered or new core nodes*, repeat steps 1-3.
5. Recover the failed data0 nodes or provision a new data0 node.
6. If you provisioned a new data node, run the following command on any existing node in the cluster (not the new node you are about to add). For *ip_address*, substitute the IP address of the new node.

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-nodes allow ip_address
```

7. On the *recovered data0 or new data0 node*, start the PCE software and enable the node to join the cluster:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-join <promoted-core-node-ip-address>
```

For the IP address, enter the IP address of the promoted core node.

8. On the *surviving data1 node and all core nodes*, remove the standalone configuration for the nodes that you previously promoted during failure:

```
$ sudo -u ilo-pce illumio-pce-ctl revert-node-config
```

NOTE:

Run this command so that the nodes that you previously promoted during the failure no longer operate as a standalone cluster.

9. Verify that the cluster is in the RUNNING state:

```
$ sudo -u ilo-pce illumio-pce-ctl cluster-status --wait
```

10. Verify that you can log into the PCE web console.

NOTE:

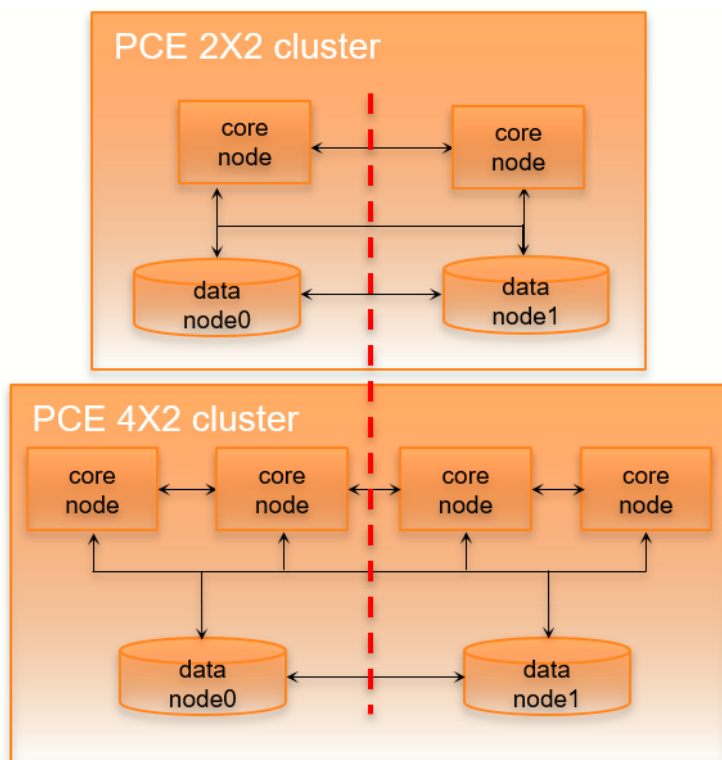
In rare cases, you might receive an error when attempting to log into the PCE web console. When this happens, restart all nodes and try logging in again:

```
$ sudo -u ilo-pce illumio-pce-ctl restart
```

Cluster Network Partition

In this failure case, the network connection between half your PCE cluster is severed, cutting off all communication between the each half of the cluster. However, all nodes in the cluster are still functioning.

Illumio defines “half a cluster” as one data node plus half the total number of core nodes in the cluster.



Stage	Details
Preconditions	None.
Failure behavior	<p>PCE</p> <ul style="list-style-type: none"> The PCE is temporarily unavailable. Users might be unable to log into the PCE web console.

Stage	Details
	<ul style="list-style-type: none"> The PCE might return an HTTP 502 response and the /node_available API request might return an HTTP 404 error. Other services that are dependent on the failed services might be restarted within the cluster. <p>VENs</p> <ul style="list-style-type: none"> VENs are not affected. VENs continue to enforce the current policy. When a VEN misses a heartbeat to the PCE, it retries in 5 minutes.
Recovery	<ul style="list-style-type: none"> Recovery type: Automatic: Having two sides of the PCE cluster operate independently of each other (“split brain”) could cause data corruption. To prevent this situation, the PCE stops services on the nodes that are not part of the quorum (namely, nodes in the data1 half of the cluster). Additionally, the PCE automatically migrates any required data services to the data0 node. When data1 was the primary database, the PCE automatically promotes data0 to be the new primary database. Recovery procedure: None required. RTO: 5 minutes. RPO: Service specific based on which data services were running on data1 at the time of the partition: <ul style="list-style-type: none"> database_service: Data1 node was the primary database. All database data committed on data1 and not replicated to data0 is lost. Typically under one second. database_slave_service: Data1 node was the replica database. No database data is lost. agent_traffic_redis_server: All traffic data is lost. fileserver_service: All asynchronous query requests and Support Reports are lost.
Full Recovery	<p>No additional steps are required to revert the PCE back to its normal, pre-failure operating state. When network connectivity is restored, the data1 half of the cluster automatically reconnects to the data0 half of the cluster. The PCE then restarts all services on the data1 half of the cluster.</p>

Multi-Node Traffic Database Failure

If the traffic database uses the optional multi-node configuration, the coordinator and worker nodes can fail.

For information about multi-node traffic database configuration, see [Scale Traffic Database to Multiple Nodes](#) in the PCE Installation and Upgrade Guide.

Coordinator Primary Node Failure

If the coordinator master completely fails, all the data-related PCE applications might be unavailable for a brief period of time. All other PCE services should be operational.

Recovery is automatic after the failover timeout. The coordinator replica will be promoted to the primary, and all data-related applications should work as usual when the recovery is done.

WARNING:

Any unprocessed traffic flow data which remained on the coordinator primary will be lost until the coordinator primary is back to normal.

Coordinator Primary Does Not Start

If the coordinator primary does not start, the PCE will not function as usual.

There are two options for recovery:

- Find the root cause of the failure and fix it. Contact Illumio Support if needed.
- Promote a replica coordinator node to primary.

WARNING:

Promoting a replica coordinator to primary can result in data loss. Use this recovery procedure only as a last resort.

To promote a replica coordinator node to primary:

```
sudo -u ilo-pce illumio-pce-ctl promote-coordinator-node cluster-leader-address
```

Worker Primary Node Failure

If the worker primary node completely fails, all data-related applications might be unavailable for a brief period of time. All other PCE services should be operational.

Recovery is automatic after the failover timeout. The worker replica will be promoted to the primary. All data-related applications should work as usual once the recovery is done.

WARNING:
Any data which was not replicated to the replica worker node before the failure will be lost.

Worker Primary Does Not Start

If the worker primary does not start, the PCE will not function as usual.

There are two options for recovery:

- Find the root cause of the failure and fix it. Contact Illumio Support if needed.
- Promote the corresponding replica worker node to primary.

WARNING:
Promoting a replica worker to primary can result in data loss. Use this recovery procedure only as a last resort.

To promote a replica worker node to primary, find out the corresponding replica worker for the failed primary node. Run the following command to list the metadata information for all the workers. Get the IP address of the replica for the failed primary:

```
sudo -u ilo-pce illumio-pce-db-management traffic citus-worker-metadata
```


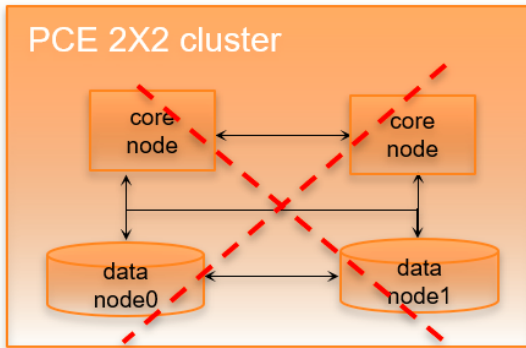
Promote the replica worker node to primary:

```
sudo -u ilo-pce illumio-pce-ctl promote-worker-node core-node-ip
```

Complete Cluster Failure

In this rare failure case, the entire PCE cluster has failed.

Disaster Recovery: Failure Scenario 8



Stage	Details
Preconditions	<p>For this failure case, Illumio assumes that you have met the following conditions before the failure occurs.</p> <p>IMPORTANT:</p> <p>You must consistently and frequently back up the PCE primary database to an external storage system that can be used for restoring the primary database after this type of failure. To recover from this failure case, you need access to this backup database file.</p> <p>The <code>runtime_env.yml</code> file parameter named <code>cluster_public_ips</code> must include the front end IP addresses of the primary and secondary cluster. When this is not configured properly, VENS will not have outbound rules programmed to allow them to connect to the secondary cluster in a failure case. Illumio recommends that you preallocate these IP addresses so that, in the event of a failure, you can restore the cluster and the VENS can communicate with the newly restored PCE.</p> <ul style="list-style-type: none"> • Regularly back up of the PCE <code>runtime_env.yml</code> file for each node in the functioning cluster before failure. • Have a secondary PCE cluster deployed in a different datacenter

Stage	Details
	<p>than the primary cluster. The secondary PCE cluster can have different IP addresses and hostnames than the primary cluster.</p>
<p>Failure behavior</p>	<p>PCE</p> <ul style="list-style-type: none"> The PCE is unavailable. <p>VENs</p> <ul style="list-style-type: none"> The VEN continues to enforce its last known good policy. The VEN's state and flow updates are cached locally on the workload where the VEN is installed. The VEN stores up to 24 hours of flow data then purges the oldest data first during an extended event. After missing 3 heartbeats (approximately 15 minutes), the VEN enters a degraded state. While it is in the degraded state, the VEN ignores all asynchronous commands received as lightning bolts from the PCE, except the commands that initiate software upgrade and Support Reports.
<p>Recovery</p>	<ul style="list-style-type: none"> Recovery type: Manual intervention is required to achieve full recovery from this failure case. Recovery procedure: See Complete Cluster Recovery for information. RTO: Customer dependent based on how long it takes to detect this failure and perform the manual recovery procedures. RPO: Customer dependent based on your backup frequency and time of the last backup.
<p>Full Recovery</p>	<p>See Complete Cluster Recovery for full recovery information; perform all the listed steps on the restored primary cluster.</p>

Complete Cluster Recovery

Recovering from this failure case requires performing the following tasks:

1. Power on *all nodes* in the secondary PCE cluster.
2. Use the database backup file saved from your most recent backup and restore the backup on the *primary database node*.

To restore the PCE database from backup, perform the following steps:

1. On *all nodes* in the PCE cluster, stop the PCE software:

```
$ sudo -u ilo-pce illumio-pce-ctl stop
```

2. On *all nodes* in the PCE cluster, start the PCE software at runlevel 1:

```
$ sudo -u ilo-pce illumio-pce-ctl start --runlevel 1
```

3. Determine the primary database node:

```
$ sudo -u ilo-pce illumio-pce-db-management show-master
```

4. On the *primary database node*, restore the database:

```
$ sudo -u ilo-pce illumio-pce-db-management restore --file <location of prior  
db dump file>
```

5. Migrate the database by running this command:

```
$ sudo -u ilo-pce illumio-pce-db-management migrate
```

6. Copy the Illumination data file from the primary database to the other data node. The file is located in the following directory on both nodes:

```
<persistent_data_root>/redis/redis_traffic_0_master.rdb
```

7. Bring the PCE cluster to runlevel 5:

```
$ sudo -u ilo-pce illumio-pce-ctl set-runlevel 5
```

8. Verify that you can log into the PCE web console.

PCE-Based VEN Distribution Recovery

When you rely on the PCE-based distribution of VEN software, after you have recovered from a PCE cluster failure, you need to reload or redeploy PCE VEN Library.

- When you have at least one PCE core node that is not affected by the failure, you can redeploy the VEN library to the other nodes.

- When the failure is catastrophic and you have to replace the entire PCE cluster, you need to reload the PCE's VEN library. See the *VEN Administration Guide* for information.

Restore VENs Paired to Failed PCE

A failed PCE does not receive information from VENs paired with it. This lack of connectivity can result in stale IP addresses and other information recorded for the VENs. Additionally, other PCEs might also have this stale information only. When the PCE regains connectivity, the PCE eventually marks those uncommunicative VENs “offline” and removes them from policy.

To resolve this situation, you must delete the “offline” workloads from the PCE by using the PCE web console or the REST API. After deleting the VENs, you can re-install and re-activate the affected VENs on the affect workloads. See the *VEN Installation and Upgrade Guide* for information.

Connectivity Configuration for PCE

This chapter contains the following topics:

Connectivity Settings	151
SecureConnect Setup	159
AdminConnect Setup	165

This section describes how to configure connectivity to control access to network resources and communication between workloads.

Connectivity Settings

This section describes how to modify PCE settings that affect connectivity.

NOTE:

Permission to edit these settings is dependent on your role. See [About Roles, Scopes, and Granted Access](#) for information.

Private Data Centers

The PCE uses connectivity settings to decide whether workloads are allowed to communicate with each other in private datacenters, private clouds, and shared network environments (private datacenter and public cloud).

By default, the Private Data Center connectivity setting is set and intended for workloads that are hosted in private datacenters, which do not have duplicate IP addresses in the network. When your network environment hosts workloads in your own private datacenter and in a public cloud, and you want to change this setting, contact Illumio Support.

Offline Timers

You can configure Offline Timers in the PCE web console and choose appropriate settings for your workloads.

NOTE:

To configure Offline Timers, you must be the Global Organization Owner for your PCE or a member of the Global Administrator role. See [About Roles, Scopes, and Granted Access](#) for information.

WARNING:

Disabling the Offline Timer setting degrades your security posture because the PCE will not remove IP addresses that belonged to workloads that have been disconnected from those that were allowed to communicate with the disconnected workloads. You need to remove the disconnected workloads from the PCE to ensure that its IP addresses are removed from the policy.

The PCE isolates a workload from the other workloads when the workload goes offline. The VEN sends a heartbeat message every 5 minutes and a goodbye message when it is gracefully shutdown. The PCE marks a workload offline when these conditions occur:

- The PCE hasn't received a heartbeat message from the VEN for 3600 seconds (1 hour).
- The PCE receives a goodbye message from the VEN.

You can change the default Offline Timer settings before putting your workloads in enforcement under the following conditions:

- The default setting might potentially disrupt your critical applications.
- Application availability is more important than security.

NOTE:

How you configure this setting is a tradeoff between benefiting from an increased zero-churn outage time window versus increasing the window of time where IP addresses could be reused. You should weigh the operational and security benefits and find a balance suitable for your applications.

Decommission and IP Cleanup Timer

Sets the time period to wait after a managed workload sends a goodbye message to mark it offline. By default, the *High Security* setting is *Wait 15 minutes before IP Cleanup*. This default setting has the following affect on the PCE:

1. Listens for Goodbye messages from the VEN.

NOTE:

The default VEN goodbye timeout was increased from zero to 15 minutes. When required, you can reset it to 0.

2. Pushes an updated policy to the peer workloads that were previously allowed to communicate with the removed workloads.
3. Immediately cleans up those workloads IP addresses from its active policy.

WARNING:For VENs installed on endpoints: Offline Timers are hardcoded for 24 hrs and can't be modified.

Disconnect and Quarantine Timer

WARNING:Do not apply on endpoints!

Sets the time period to wait with no heartbeat before a managed workload is marked offline.

By default, the *High Security* setting is *Wait One Hour before Timeout*. This default setting has the following affect on the PCE:

1. Waits for an hour for the disconnected workloads to heartbeat and then quarantine those workloads that do not respond at the end of the hour.
2. Removes the quarantined workloads IP addresses from its active policy.
3. Pushes an updated policy to the peer workloads that were previously allowed to communicate with the quarantined workloads.

Edit Offline Timers Settings

Edit the Offline Timers setting to change the values from the default settings.

1. From the PCE web console menu, choose **Settings > Offline Timers**.

The Settings page for Offline Timers appears, which displays the current settings for the timers.

2. Click **Edit** to change the settings from the default values.
3. **Disconnect and Quarantine Timer:** Select a setting from the drop-down list to change the value from the High Security (Default) setting:

- *Never Timeout or Quarantine - Highest Availability*

This setting has the following affect on the PCE:

- Never disconnects or quarantines workloads that fail to heartbeat.
- Keeps all IP addresses in policy and never automatically removes unused IP addresses.
- Requires a removal of those unused IP addresses.

- *Custom Timeout - Wait a Specified Time before Quarantine*

Enter a time period; the minimum wait time is 300 seconds.

The PCE performs the following actions:

- a. Waits for the specified time period for the disconnected workloads to heartbeat.
- b. Quarantines those workloads that do not respond at the end of that time period.
- c. Removes the quarantined workloads IP addresses from its active policy.
- d. Pushes an updated policy to the peer workloads that were previously allowed to communicate with the quarantined workloads.

4. **Decommission and IP Cleanup Timer:** Select a setting from the drop-down list to change the value from the Highest Security (Default) setting:

- *Never clean up - Highest Availability*

This setting has the following affect on the PCE:

- Ignores Goodbye messages from workloads.
- Keeps all IP addresses in policy and never automatically remove unused IP addresses.
- Requires a removal of those unused IP addresses.

- *Custom Timeout - Wait a Specified Time before IP Cleanup*

Enter a time period; the minimum wait time is 0 seconds.

The PCE performs the following actions:

- a. Listens for Goodbye messages from the VEN.
- b. Waits for the specified time period before cleanup of those workloads IP addresses from its active policy.
- c. Pushes an updated policy to the peer workloads that were previously allowed to communicate with the removed workloads.

5. Click **Save**.

A message appears displaying your current and new settings.

Confirm Timer Setting Changes

Disconnect and Quarantine Timer ~~Wait One Hour before Timeout—High Security (Default)~~
~~Never Timeout or Quarantine - Highest Availability~~

1. Never disconnect or quarantine workloads that fail to heartbeat,
2. Keep all IP addresses in policy and never automatically remove unused IP addresses, and
3. Require a removal of those unused IP addresses.

Cancel **OK**

6. Click **OK** to save the new settings.

Set the IP Version for Workloads

This section describes how to enforce a preference for IPv4 over IPv6 addresses.

Change Linux Workloads to Prefer IPv4

To ensure that your paired Linux VEN workloads prefer IPv4 over IPv6 addresses in your PCE organization, edit the `/etc/gai.conf` file on the VEN by adding the following line:

```
$ precedence ::ffff:0:0/96 100
```

This change will cause `getaddrinfo` system calls to return the IPv4 addresses before IPv6 addresses.

This method works when you assign IPv4 addresses to your workloads. However, it doesn't work when your workloads only have IPv6 addresses (meaning, no IPv4 addresses for the hosts) or the software installed is hard coded to look for IPv6 addresses.

Change Windows Workloads to Prefer IPv4

When you choose to allow only IPv4 traffic for your PCE organization, the VENS on your workloads drop IPv6 traffic when they are in Enforced mode. This decision can lead to delays and communication failures in applications because applications will wait for IPv6 connection attempts to time out before attempting to connect over IPv4.

The problem occurs because, by default, the Windows OS prefers IPv6 over IPv4 and will attempt to connect over IPv6 before IPv4. As a workaround, you can change the order of connection attempts so that IPv4 is preferred over IPv6. With this change, applications will connect over IPv4 first and succeed or fail as governed by the workload's firewall policies.

For information about changing the connection order to prefer IPv4 over IPv6, see the Microsoft KB article [Guidance for configuring IPv6 in Windows for advanced users](#).

As explained in the KB article, run the following command and reboot the Windows workload:

```
reg add hklm\system\currentcontrolset\services\tcpip6\parameters /v
DisabledComponents /t REG_DWORD /d 0x20
```

To avoid rebooting the Windows workload, run the following commands:

```
netsh interface ipv6 delete prefixpolicy ::ffff:0:0/96
netsh interface ipv6 add prefixpolicy ::ffff:0:0/96 60 4
```

Manage Security Settings

You can manage security settings by accessing the page **Settings -> Security**:

Security for		Options	Description
VENS (Versions 20.2.0.and higher)	IPv6 traffic	Allow IPv6 traffic	Allowed based on policy

Security for		Options	Description
		Block IPv6 traffic	Blocked only in Enforcement state. Always allowed on AIX and Solaris workloads
VENS (Versions lower than 20.2.0)	IPv6 traffic	Allow IPv6 traffic	All IPv6 traffic allowed
		Block IPv6 traffic	Blocked only in Enforcement state. Always allowed on AIX and Solaris workloads
IKE Authentication	Authentication type	PSK	Use Pre-shared Keys for authentication
		Certificate	Use certificates for authentication
Public cloud configuration	NAT Detection	Private Data Center or Public Cloud with 1:1 NAT (default)	For workloads in a known public cloud (such as AWS or Azure) the public IP address of the workload as seen by the PCE is distributed along with the IP addresses of the interfaces on the workload. Use this setting only if there are no shared SNAT IP addresses for egress traffic from the public cloud workloads.
		Public Cloud with SNAT/NAT Gateway (recommended setting if using a NAT gateway in AWS or Azure or the	The PCE will ignore the public IP address of the workload in policy computation. This setting is used in environments where

Security for		Options	Description
		default outbound access in Azure	workloads in a known public cloud (e.g, AWS or Azure) that connect to other workloads or the PCE outside the VPC or cloud via the SNAT IP address or SNAT pool (e.g, NAT Gateway in AWS) as the public IP seen by the PCE is nit specific to any workloads. Only the IP address of the network interfaces on the workload (usually the private IP addresses) is distributed in the policy.

Enable IP Forwarding

(For Linux VEnS only)

In PCE versions earlier than 21.5.10, IP forwarding is automatically enabled for hosts in a container cluster that is reported by Kubelink to the PCE or hosts explicitly set to use the Containers Inherit Host Policy feature.

Starting in PCE version 21.5.10, you can enable IP forwarding on hosts without using any container segmentation features. To enable this feature, contact Illumio Support.

1. In the PCE web console, choose **Security > IP Forwarding**. The IP Forwarding tab appears if the feature is enabled.
2. In this tab, you can use labels and label groups to enable IP forwarding for the workloads that match the label combination. Use combinations of Role, Application, Environment, and Location labels and label groups in the same way that you would to specify workloads for any other purpose; for example, in a Rule or any of the tabs under the Security Settings page.

Workloads with IP forwarding enabled will configure the host firewall to allow all forwarded traffic without visibility, including traffic forwarded through the host.

SecureConnect Setup

Enterprises have requirements to encrypt in transit data in many environments, particularly in PCI and other regulated environments. Encrypting in transit data is straightforward for an enterprise when the data is moving between datacenters. An enterprise can deploy dedicated security appliances (such as VPN concentrators) to implement IPsec-based communication across open untrusted networks.

However, what if an enterprise needs to encrypt in transit data within a VLAN, data-center, or PCI environment, or from a cloud location to an enterprise datacenter? Deploying a dedicated security appliance to protect every workload is no longer feasible, especially in public cloud environments. Additionally, configuring and managing IPsec connections becomes more difficult as the number of hosts increases.

Features of SecureConnect

SecureConnect has the following key features. Platforms Supported by SecureConnect

SecureConnect works for connections between Linux workloads, between Windows workloads, and between Linux and Windows workloads.

IPsec Implementation

SecureConnect implements a subset of the IPsec protocol called Encapsulating Security Payload (ESP), which provides confidentiality, data-origin authentication, connectionless integrity, an anti-replay service, and limited traffic-flow confidentiality.

In its implementation of ESP, SecureConnect uses IPsec transport mode. Using transport mode, only the original payload is encrypted between the workloads. The original IP header information is unchanged so all network routing remains the same.

However, the protocol being used will be changed to reflect the transport mode (ESP).

Making this change causes no underlying interfaces to change or be created or any other underlying networking infrastructure changes. Using this approach simply obfuscates the data between endpoint workloads by encrypting the data between them.

If SecureConnect is unable to secure traffic between two workloads with IPsec, it will block unencrypted traffic when the policy was configured to encrypt that traffic.

IKE Versions Used for SecureConnect

SecureConnect connections between workloads use the following versions of Internet Key Exchange (IKE) based on workload operating system:

- Linux ↔ Linux: IKEv2
- Windows ↔ Windows: IKEv1
- Windows ↔ Linux: IKEv1

For a list of supported operating systems for managed workloads, see [VEN OS Support and Package Dependencies](#) on the Illumio Support portal.

Existing IPsec Configuration on Windows Systems

Installing a VEN on a Windows system does not change the existing Windows IPsec configuration, even though SecureConnect is not enabled. The VEN still captures all logging events (`event.log`, `p1atform.log`) from the Windows system that relate to IPsec thereby tracking all IPsec activity.

Performance

The CPU processing power that a workload uses determines the capacity of the encryption. The packet size and throughput determine the amount of power that is required to process the encrypted traffic using this feature.

In practice, enabling SecureConnect for a workload is unlikely to cause a big spike in CPU processing or a decrease in network throughput. However, Illumio recommends benchmarking performance before enabling SecureConnect and comparing results after enabling it.

Use Pre-Shared Keys with SecureConnect

SecureConnect includes the option of using pre-shared keys (generated by the PCE) or client-side PKI certificates for IKE authentication.

You can configure SecureConnect to use pre-shared keys (PSKs) to build IPsec tunnels that are automatically generated by the PCE. SecureConnect uses one key per organization. All the workloads in that organization share the one PSK. SecureConnect uses a randomly generated 64-character alpha-numeric string, for example:

```
c4aeb6230c508063db3e3e1fac185bea9c4d17b4642a87e091d11c9564fbd075
```

When SecureConnect is enabled for a workload, you can extract the PSK from a file in the `/opt/illumio` directory, where the VEN stores it. You cannot force the PCE to regen-

erate and apply a new PSK. If you feel the PSK has been compromised, contact [Technical Support](#).

NOTE:

Illumio customers accessing the PCE from the Illumio cloud can have multiple Organizations. However, the Illumio PCE does not support multiple Organizations when you have installed the PCE in your datacenter.

Use PKI Certificates with SecureConnect

SecureConnect allows you to use client-side PKI certificates for IKE authentication and IPsec communication between managed workloads. If you have a certificate management infrastructure in place, you can leverage it for IKE authentication between workloads because it provides higher security compared to using pre-shared keys (PSKs).

Certificate-based SecureConnect works for connections between Linux workloads, between Windows workloads, and between Linux and Windows workloads.

The IPsec configuration uses the certificate with the distinguished name from the issuer field that you specify during PCE configuration for IKE peer authentication.

Prerequisites, Limitations, and Caveats

Before configuring your workloads to use SecureConnect, review the following prerequisites and limitations, and consider the following caveats.

PKI Certificates with SecureConnect

The following prerequisites and limitations apply when configuring SecureConnect to use certificates:

- You must have a PKI infrastructure to distribute, manage, and revoke certificates for your workloads. The PCE does not manage certificates or deliver them to your workloads.
- The PCE supports configuring only one global CA ID for your organization.
- The VEN on a workload uses a Certificate Authority ID (CA ID) to authenticate and establish a secure connection with a peer workload.

Connected workloads must have CA identity certificates signed by the same root certificate authority. When workloads on either end of a connection use different CA IDs, the IKE negotiation between the workloads will fail and the workloads will not be able to communicate with each other.

VEN Versions

To use PKI certificates with SecureConnect, your workloads must be running VEN version 17.2 or later.

Maximum Transmission Unit (MTU) Size

IPsec connections cannot assemble fragmented packets. Therefore, a high MTU size can disrupt SecureConnect for the workloads running on that host.

Illumio recommends setting the MTU size at 1400 or lower when enabling SecureConnect for a workload.

Ports

Enabling SecureConnect for a workload routes all traffic for that workload through the SecureConnect connection using ports 500/UDP and 4500/UDP for NAT traversal and for environments where ESP traffic is not allowed on the network (for example, when using Amazon Web Services). You must allow 500/UDP and 4500/UDP to traverse your network for SecureConnect.

Unsupported SecureConnect Usage

SecureConnect is not supported in the following situations:

- SecureConnect cannot be used between a workload and unmanaged entities, such as the label “Any (0.0.0.0/0 and ::/0)” (such as, the internet).
- SecureConnect is not supported on virtual services.
- SecureConnect is not supported on workloads in the Idle policy state. If you enable it for a rule that applies to workloads that are in both Idle and non-Idle policy states, you can impact the traffic between these workloads.
- SecureConnect is not supported on AIX and Solaris platforms.

SecureConnect

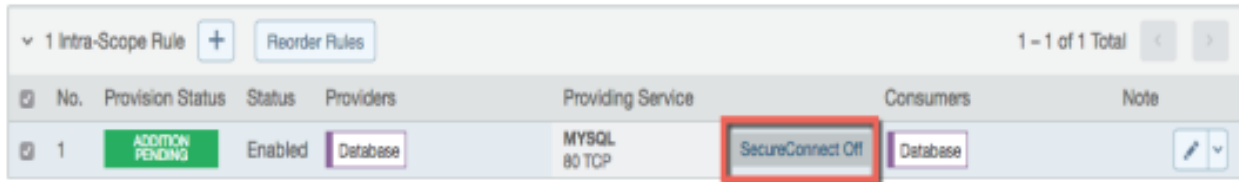
When you configure workloads to use SecureConnect be aware of the following caveat.

SecureConnect encrypts traffic for workloads running in all policy states except Idle.

SecureConnect Host-to-Host Encryption

When you configure workloads to use SecureConnect be aware of the following caveat.

SecureConnect encrypts traffic between workloads on a host-to-host basis. Consider the following example.



No.	Provision Status	Status	Providers	Providing Service	Consumers	Note
1	ADDITION PENDING	Enabled	Database	MYSQL 80 TCP	SecureConnect Off Database	

In this example, it appears that enabling SecureConnect will only affect MySQL traffic. However, when you enable SecureConnect for a rule to encrypt traffic between a database workload and a web workload over port 3306, the traffic on all ports between the database and web workloads is protected by IPsec encryption.

Configure SecureConnect to Use Pre-Shared Keys

You can configure SecureConnect to use pre-shared keys (PSKs) for IKE authentication and IPsec communication between managed Workloads. SecureConnect uses one key per Organization. All the Workloads in that organization share the one PSK. SecureConnect generates a random 64-character alpha-numeric string for this key.

1. From the PCE navigation menu, choose **Settings > Security Settings**.
2. Choose **Edit > Configure SecureConnect**.
The page refreshes with the settings for SecureConnect.
3. In the Default IPsec Authority field, select the **PSK** option.
4. Click **Save**.

Configure SecureConnect to Use Certificates

SecureConnect allows you to use client-side PKI certificates for IKE authentication and IPsec communication between managed Workloads. The PCE supports configuring only one global CA ID for your organization. Configuring SecureConnect to use certificates applies the setting to All Roles, All Applications, All Environments, and All Locations.

Configuring SecureConnect to use PKI certificates in the global Security Settings page does not manage certificates for your organization or deliver them to your Workloads.

NOTE:

You must independently set up certificates on your Windows and Linux Workloads. For information, see [Requirements for Certificate Setup on Workloads](#).

1. From the PCE web console menu, choose **Settings > Security Settings**.
2. Choose **Edit > Configure SecureConnect**.
The page refreshes with the settings for SecureConnect.
3. In the *Default IPsec Authority* field, select the **Certificate Authority** option.
4. In the *Global Certificate ID* field, enter the distinguished name from the Issuer field of your trusted root certificate. (This certificate is used globally for all workloads in your organization enabled with SecureConnect.)
5. Click **Save**.

Requirements for Certificate Setup on Workloads

To use PKI certificates with SecureConnect, you must independently set up certificates on your Windows and Linux workloads.

Generate or obtain certificates from a trusted source in your organization. You should only use certificates obtained from trusted sources.

File Requirements

File	Requirements
Issuer's certificate	<p>The global CA certificate, either root or intermediate, in PEM or DER format</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>NOTE: On Linux, the issuer's certificate must be readable by the Illumio user.</p> </div>
pkcs12 container	<p>Archive containing the public key, private key, and identity certificate generated for the workload host.</p> <p>Sign the identity certificate using the global root certificate.</p> <p>You can password protect the container and private key but do not password protect the public key.</p>

Certificate Requirements

There are certain requirements regarding certificate use that you have to follow for the installation:

- x509 certificate must contain fields `SubjectName` and `SubjectAltName`.
- `SubjectNameDN` should contain `CN`, which has to match to DNS name of `SubjectAltName`.

```
X509v3 Subject Alternative Name:  
          DNS:centos6, email:centos6@ilabs.io  
Subject: OU=VEN, CN=centos6.ilabs.io
```

- x509v3 extension with the key usage must have Digital Signature, Key agreement
- x509v3 extension with extended key usage must contain either "Any Extended Key Usage" or "IPSec End System, IPSec User, TLS Web Server Authentication"
- x509v3 extension with the authority Key Identifier field is required as well

Installation Locations

Windows Store

Use the Windows OS, for example Microsoft Management Console (MMC), to import the files into these locations of the local machine store (not into your user store).

- Root certificate: Trusted Root Certificate Store
- pkcs12 container: Personal ("My") certificate store

WARNING:If the Windows machine cert storage /My/Personal contains more than one certificate issued by the same issuer, Windows doesn't know what to pick to program Cert for SecureConnect.

WARNING:Make sure to have one single certificate for IPSec (not for SSL or other purposes) that is signed by a distinguished Issuer.

Linux Directories

Copy the files into the following Linux directories. (You cannot change these directories.)

- Root certificate: /opt/illumio_ven/etc/ipsed.d/cacert
- pkcs12 container: /opt/illumio_ven/etc/ipsed.d/private

AdminConnect Setup

Relationship-based access control rules often use IP addresses to convey identity. This authentication method can be effective. However, in certain environments, using IP addresses to establish identity is not advisable.

When you enforce policy on servers for clients that change their IP addresses frequently, the policy enforcement points (PEPs) continuously need to update security rules for IP address changes. These frequent changes can cause performance and scale challenges, and the ipsets of protected workloads to churn.

Additionally, using IP addresses for authentication is vulnerable to IP address spoofing. For example, server A can connect to server B because the PEP uses IP addresses in packets to determine when connections originate from server A. However, in some environments, bad actors can spoof IP addresses and impact the PEP at server B so that it mistakes a connection as coming from server A.

Illumio designed its AdminConnect (Machine Authentication) feature with these types of environments in mind. Using AdminConnect, you can control access to network resources based on Public Key Infrastructure (PKI) certificates. Because the feature bases identity on cryptographic identity associated with the certificates and not IP addresses, mapping users to IP addresses (common for firewall configuration) is not required.

With AdminConnect, a workload can use the certificates-based identity of a client to verify its authenticity before allowing it to connect.

Features of AdminConnect

Cross Platform

Microsoft Windows provides strong support for access control based on PKI certificates assigned to Windows machines. Modern datacenters, however, must support heterogeneous environments. Consequently, Illumio designed AdminConnect to support Windows and Linux servers and Windows laptop clients.

AdminConnect and Data Encryption

When only AdminConnect is enabled, data traffic does not use ESP encryption. This ensures that data is in cleartext even though it is encapsulated in an ESP packet.

When AdminConnect and SecureConnect are enabled for a rule, the ESP packets are encrypted.

Ease of Deployment

Enabling AdminConnect for identity-based authentication is easy because it is a software solution and it does not require deploying any network choke points such as firewalls. It also does not require you to deploy expensive solutions such as Virtual Desktop Infrastructure (VDI) or bastion hosts to control access to critical systems in your datacenters.

Prerequisites and Limitations

Prerequisites

You must meet the following prerequisites to use AdminConnect:

- You must configure SecureConnect to use certificate-based authentication because both features rely on the same PKI certificate infrastructure. See the following topics for more information:
 - [Configure SecureConnect to Use Certificates](#)
 - [Requirements for Certificate Setup on Workloads](#)
 - [Certificates for AdminConnect](#)
- - AdminConnect must be used with VEN version 17.3 and later.
 - AdminConnect supports Linux/Windows IKE v1 (client only) with unmanaged workloads.

Limitations

You cannot enable AdminConnect for the following types of rules:

- Rules that use All services
- Rules with virtual services in providers or consumers
- Rules with IP lists as providers or consumers
- Stateless rules

AdminConnect is not supported in these situations:

- AdminConnect does not support “TCP -1” (TCP all ports) and “UDP -1” (UDP all ports) services.
- You cannot use Windows Server 2008 R2 or earlier versions as an AdminConnect server.
- Windows Server does not support more than four IKE/IPsec security associations (SAs) concurrently from the same Linux peer (IP addresses).

Certificates for AdminConnect

AdminConnect relies on PKI certificates for relationship-based access control of workloads.

The feature uses the same certificate infrastructure enabled for SecureConnect. If you have not set up certificate for SecureConnect, see [Configure SecureConnect to Use Certificates](#) and [Requirements for Certificate Setup on Workloads](#) for information.

The same prerequisites and limitations for certificate set up apply for AdminConnect. Additionally, because you can use AdminConnect to control access for laptops, certificates on laptops must meet these additional requirements:

- The certificate must have a unique Subject Name and Subject Alt Name.
- The certificate must be enabled with all extended key usage to check trust validation.

Secure Laptops with AdminConnect

You can use Illumio to authenticate laptops and grant them access to managed workloads. To manage a laptop with AdminConnect, complete the following tasks:

1. Deploy a PKI certificate on the laptop. See [Certificates for AdminConnect](#).
2. Add the laptop to the PCE by creating an unmanaged workload and assign the appropriate labels to it to be used for rule writing
3. Create rules using those labels to grant access to the managed workloads. For information, see [Enable AdminConnect for a Rule](#) in the *Security Policy Guide*.
4. Configure IPsec on a laptop.

To add a laptop to the PCE by creating an unmanaged workload:

To manage a laptop with AdminConnect, add the laptop to the PCE as an unmanaged workload.

1. From the PCE web console menu, choose **Workloads > Add > Add Unmanaged Workload**.

The Workloads - Add Unmanaged Workload page appears.

2. Complete the fields in the *General*, *Labels*, *Attributes*, and *Processes* sections. See [Add an Unmanaged Workload](#) in the *Security Policy Guide* for information.
3. In the *Machine Authentication ID* field, enter all or part of the DN string from the *Issuer* field of the end entity certificate (CA Subject Name). For example:

CN=win2k12, O=Illumio, OU=Portal, ST=CA, C=US, L=Sunnyvale

TIP:

Enter the exact string that you get from the `openssl` command output.

4. Click **Save**.

To configure IPsec on a laptop:

To use the AdminConnect feature with laptops in your organization, you must configure IPsec for these clients.

See the Microsoft Technet article [Netsh Commands for Internet Protocol Security \(IPsec\)](#) for information about using netsh to configure IPsec.

See also the following examples for information about the IPsec settings required to manage laptops with the AdminConnect feature.

```
PS C:\WINDOWS\system32> netsh advfirewall show global

Global Settings:
-----
IPsec:
StrongCRLCheck                0:Disabled
SAIdleTimeMin                 5min
DefaultExemptions             NeighborDiscovery,DHCP
IPsecThroughNAT               Server and client behind NAT
AuthzUserGrp                  None
AuthzComputerGrp              None
AuthzUserGrpTransport         None
AuthzComputerGrpTransport     None

StatefulFTP                   Enable
StatefulPPTP                  Enable

Main Mode:
KeyLifetime                   60min,0sess
SecMethods                    ECDHP384-AES256-SHA384
ForceDH                       Yes

Categories:
BootTimeRuleCategory          Windows Firewall
FirewallRuleCategory          Windows Firewall
StealthRuleCategory           Windows Firewall
ConSecRuleCategory            Windows Firewall

Ok.
```

```
PS C:\WINDOWS\system32> netsh advfirewall consec show rule name=all

Rule Name:                                telnet
-----
Enabled:                                   Yes
Profiles:                                 Domain,Private,Public
Type:                                      Static
Mode:                                      Transport
Endpoint1:                                Any
Endpoint2:                                10.6.3.189/32,10.6.4.35/32,192.168.41.163/32
Port1:                                     Any
Port2:                                     23
Protocol:                                  TCP
Action:                                    RequireInRequireOut
Auth1:                                     ComputerKerb,ComputerCert
Auth1CAName:                               CN=MACA, O=Company, OU=engineering, S=CA,
C=US, L=Sunnyvale, E=user@sample.com
Auth1CertMapping:                          No
Auth1ExcludeCAName:                        No
Auth1CertType:                              Intermediate
Auth1HealthCert:                           No
MainModeSecMethods:                         ECDHP384-AES256-SHA384
QuickModeSecMethods:                        ESP:SHA1-AES256+60min+100256kb
ApplyAuthorization:                         No
Ok.
```

Access Configuration for PCE

This chapter contains the following topics:

Role-based Access Control	171
Setup for Role-based Access Control	182
Role-based Access for Application Owners	192
Configure Access Restrictions and Trusted Proxy IPs	208
Password Policy Configuration	211
Authentication	215
Active Directory Single Sign-on	225
Azure AD Single Sign-on	258
Okta Single Sign-on	270
OneLogin Single Sign-on	272
Ping Identity Single Sign-on	274

This section describes how to configure the PCE to control access.

Role-based Access Control

This section describes the concepts of role-based access control (RBAC) and how it works with the PCE.

Overview of Role-based Access Control

Security-oriented companies should grant employees the exact permissions they need based on their role. Illumio Core uses role-based access control (RBAC) to

deliver security at an enterprise scale in the following ways:

- Assign your users the least required privilege they need to perform their jobs. Limit access for your users to the smallest operation-set they need to perform their jobs; for example, monitor for security events.
- Implement separation of duties. Delegate the responsibility to manage a zone to a specific team or delegate authority to application teams; for example, delegate a team to manage security for the US-West Dev zone, or assign the DevOps team to set security policy for the HRM application they manage.
- Grant access to users based on two dimensions: roles and scopes. Each role grants access to a set of capabilities in Illumio Core. Scopes define the workloads in your organization that users can access, and are based on labels. A common set of label types include Application, Environment, and Location, but you may define additional label types and values using Flexible Labels. The scopes specify the boundaries of the sphere of influence granted to a user. For example, a user can be added to the Ruleset Provisioner role with the scope Application CRM, Environment Staging, and Location US. With that access, the user could provision rulesets for workloads that are part of your CRM application in the Staging environment located in the US.
- Centrally manage user authentication and authorization for Illumio Core. Configure single sign-on with your corporate Identity Provider (IdP) and designate which external IdP groups should have access roles. Group membership is managed by your IdP while resource authorization is configured in Illumio Core.

Use Cases

Illumio designed our RBAC feature around a set of use cases based on the way that enterprises manage the security of the computing assets in their environment. These use cases encompass common security workflows for the modern, security-conscious enterprise. The personas include different levels of security professionals.

Support the Security Workflow

Customers can configure the RBAC feature to support any type of responsibility bifurcation that they have in their workflow models. For example, the following workflows are supported:

- Architect-level professionals define all security policy for an enterprise by adding rulesets and rules in the PCE.
- Junior-level professionals provision rulesets and rules to workloads during maintenance windows. Junior personnel cannot edit any policy items in the Illumio PCE.
- Some users only view the infrastructure and alert senior team members when security issues occur.

Manage Security for Specific Workloads

When you combine Illumio Core RBAC roles with scopes, you can secure access for IT teams who support specific applications or different geographic locations. For example, customers could delegate authority for workloads in the following ways:

- To manage security for workloads around silos; for example, a particular cloud provider like AWS.
- To decentralize their security policy to specific application teams allowing them to act quickly when managing application security without waiting for the central security team.
- To bifurcate the security of their infrastructure in such a way that one user is responsible only for the West coast assets and another user is responsible for the East coast assets.

Features of Role-based Access Control

Built-in Roles

Illumio Core includes several roles that grant users access to perform operations. Each role is matched with a scope. See [About Roles, Scopes, and Granted Access](#) for information.

Granular Permissions

You can assign multiple roles to one user and by mixing and matching the different roles, you can achieve different levels of granularity of permissions.

You can grant different permissions to different users for different resources by defining scopes. For example, you might allow some users complete access to add rulesets for all workloads in your staging environment. For other users, you might grant access to all workloads in all environments. Users can be assigned exactly one role, representing their singular job function while other users can be assigned multiple roles, representing multiple job functions.

Identity Federation Using External Users and Groups

You can connect to external LDAP directories to manage users and user groups by configuring single sign-on (SSO) for the PCE.

Using this feature, you can create and manage users locally in PCE, or use an IdP to manage users and user groups from an existing directory. External user and user groups authenticate with the external IdPs.

Custom Role Assignments

You can customize access to suit your organization by specifying specific scopes for the Ruleset Manager and Ruleset Provisioner roles.

Audit Information

You can access an audit trail of user activity through the following reports:

- The User Activity page, which displays the authentication details for each user, when they logged in, and whether they are online.
- The Organization Events page, which displays when Organization Owners granted users access, when users logged in and out, and the actions they performed.

About Roles, Scopes, and Granted Access

Illumio Core includes several roles that grant users access to perform operations. Each role is matched with a scope. You can add users (local and external) and groups to all the roles.

Roles with Global Scopes

These Global Roles use the scope All Applications, All Environments, and All Locations. You cannot change the scope for these roles. The roles have the following capabilities in Illumio Core.

Role	Granted Access
Global Organization Owner	Perform all actions: add, edit, or delete any resource, security settings, or user account
Global Administrator	Perform all actions except user management: add, edit, or delete any resource or organization setting
Global Viewer	View any resource or organization setting. They cannot perform any operations. This role was previously called "Global Read Only."
Global Policy Object Provisioner	Provision rules containing IP lists, services, and label groups. They cannot provision rulesets, virtual services, or virtual servers,

Role	Granted Access
	or add, modify, or delete existing policy items.

NOTE:
You can add, modify, and delete your API keys because you own them.

About Read Only Users in the Global Viewer User Role

The Read Only User role applies to all users in your organization—local, external, and users who are members of external groups managed by your IdP. This role allows users to view resources in Illumio Core when they are not explicitly assigned to roles and scopes in the PCE.

For example, you configure single sign-on for your corporate Microsoft Active Directory Federation Services (AD FS) so that users managed by AD FS can log into the PCE by using their corporate usernames and passwords. However, you haven't added all your external users to the PCE or assigned them to roles. These users can still log into the PCE by authenticating with the corporate IdP and view resources in the PCE.

The Read Only User role is not listed in the **Access Management > Global Roles** or **Scopes** pages because it is considered a default, catchall type of role. Users have access to this role on an organization-wide basis because you either enable or disable it for your entire organization. Additionally, you do not see it in the list of a user's role assignments when you view the user's details page (**Access Management > External Users** or **Local Users**). However, when the role is enabled for your organization, you see it listed in the **Access Management > User Activity** details for each user.

NOTE:
You can enable and disable the Read Only User role from the **Access Management > Global Roles** page, by clicking the **Global Viewer** role.

When the Read Only User role is disabled for your organization, users who are not assigned to roles cannot access Illumio managed resources. When attempting to log into the PCE, they are still authenticated by their corporate IdP but the PCE immediately logs them out because they do not have access (even read-only access) to any Illumio managed assets.

Roles with Custom Scopes

You can apply the following roles to specific scopes. These roles are called "Scoped Roles."

Role	Granted Access
Full Ruleset Manager	<ul style="list-style-type: none"> • Add, edit, and delete all rulesets within the specified scope. • Add, edit, and delete rules when the provider matches the specified scope. The rule consumer can match any scope. <div data-bbox="435 401 1419 604" style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>NOTE: You can choose the All Applications, All Environments, and All Locations scope with the Full Ruleset Manager role.</p> </div>
Limited Ruleset Manager	<ul style="list-style-type: none"> • Add, edit, and delete all rulesets within the specified scope. • Add, edit, and delete rules when the provider and consumer match the specified scope. • Ruleset Managers with limited privileges cannot manage rules that use IP lists, custom iptables rules, user groups, label groups, iptables rules as consumers, or have internet connectivity. <div data-bbox="435 909 1419 1115" style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>NOTE: You cannot choose the All Applications, All Environments, and All Locations scope with the Limited Ruleset Manager role.</p> </div>
Ruleset Viewer	<ul style="list-style-type: none"> • Read-only access to rules that match the specified scope. • Ruleset Viewers cannot edit rules or rulesets.
Ruleset Provisioner	<p>Provision rulesets within specified scope.</p> <div data-bbox="435 1283 1419 1486" style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>NOTE: You can choose the All Applications, All Environments, and All Locations scope and custom scopes with the Ruleset Provisioner role.</p> </div>
Workload Manager	<p>Manage workloads and pairing profiles within the specified scope. Read-only access provided to all other resources.</p>

Role	Granted Access
	<p>NOTE: The 19.1.0 PCE does not support unpairing multiple managed workloads via the REST API when you are logged in as a Workload Manager. You can unpair workloads using the PCE web console because it restricts selection of workloads by the user's scope. However, via the REST API, the bulk unpair operation fails when multiple workloads are selected and one or more of the workloads are out of the user's scope.</p>

Workload Manager Role

Use Case 1

You want to use scripts in your development environment to programmatically spin up and bring down workloads; your scripts create pairing profiles and generate pairing keys without you granting elevated Admin privileges to the scripts.

Use Case 2

Your application teams are in charge of changing the security posture of workloads, such as changing the policy enforcement states. You want to allow your application teams to manage workload security without granting them broad privileges, such as All access (for the standard Application, Environment, and Location label types, or for any customer label types you have defined).

Use Case 3

You want to prevent your PCE users from accidentally changing workload labels by moving the workloads in Illumination or Illumination Plus.

Solution

Users with the Workload Manager role can create, update, and delete workloads and pairing profiles. This role is a scoped role; when you assign a user to a scope, they can only manage workloads within the allocated scope. The Workload Manager can pair, unpair, and suspend VENs and change the policy state. It is an additive role; you can assign the Workload Manager role to a user and combine it with any other PCE role to provide additional privileges for that user.

Configuration

1. Create a local user with “None” or the Global Viewer role (with Read Only User turned on).
2. Assign the Workload Manager role to the user.
3. (Optional) Provide the invitation link to the new workload manager user.
4. The workload manager can then log into the PCE and manage workloads and pairing profiles per the allocated scope.

The Workload Manager role is available under **Scopes**. Users assigned this role can view applications that are outside their scopes but can only modify those applications that are within their scopes.

NOTE:

A workload manager user cannot clear traffic counters from workloads within their scope.

Example: Limited Ruleset Manager Role

A user has the role Full Ruleset Manager role and access to the following scope:

All Applications | Production Environment | All Locations

The user can create and manage:

- Any ruleset that matches the Production environment
- Intra- or extra-scope rules that match this scope:

All Applications | Production Environment | All Locations

Where the provider and consumer of the rule are both within the Production environment scope.

For intra-scope rules, all workloads can communicate within their group (as defined by the scope), so the rule consumer is not restricted. However, in extra-scope rules, the Environment label of the resource selected as the consumer must match the label in the scope exactly.

The user cannot create a rule with the scope “All | All | All” because that scope is broader than the user’s access, which is only for the Production environment.

Because the user is a member of the Limited Ruleset Manager role, the user cannot manage custom iptables rules and the following resources cannot be selected as consumers in extra-scope rules:

- IP lists
- Label groups
- User groups
- Workloads

Combine Roles to Support Security Workflows

Illumio includes fine-grained roles to manage security policy. The roles control different aspects of the security workflow. By mixing and matching them, you can effectively control the access needed by your company.

Ruleset Only Roles

You can add users to the Full Ruleset Manager and Ruleset Provisioner roles so that they can edit the security policies on the workloads within their assigned scopes without affecting other entities, such as services, virtual services, or virtual servers.

- The full Ruleset Manager can add, edit, and delete rules when the provider matches a specified scope.
- The limited Ruleset Manager can add, edit, and delete rules when the provider and consumer match the specified scope. Ruleset managers with limited privileges cannot manage rules that use IP lists, user groups, label groups, iptables rules as consumers, or rules that allow internet connectivity.
- The Ruleset Provisioners can provision rulesets within a specified scope. They cannot provision virtual servers, virtual services, SecureConnect gateways, security settings, IP lists, services, or label groups. Provision rulesets within a specified scope.

If you are granting a user or group the Ruleset Manager or the Ruleset Provisioner role, you can also associate a scope to the role so you can control which rulesets they can add and provision.

Ruleset Plus Global Policy Object Provisioner Roles

You can add users to the Ruleset Manager (Full or Limited) role and the Global Policy Object Provisioner role so that they can control the security policy for workloads.

The rule consumer can match any scope.

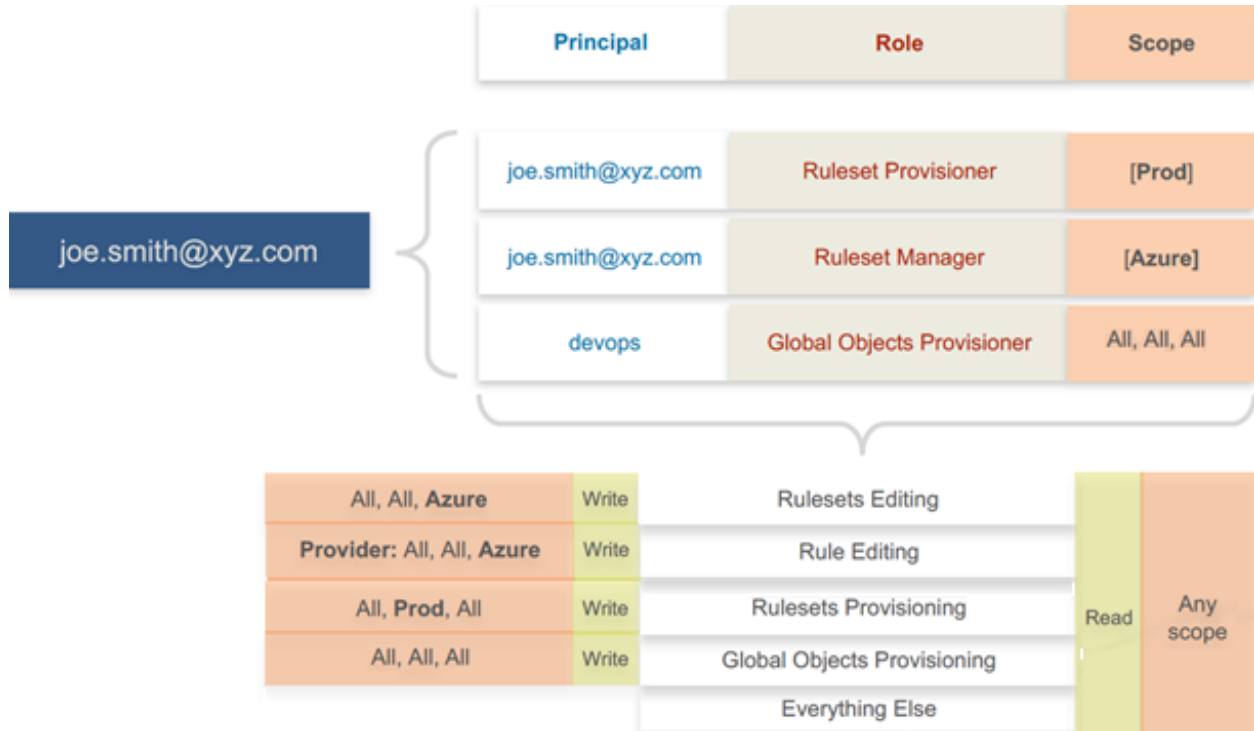
Global Organization Owner or Administrator Roles

You can add architect-level professionals to the Global Organization Owner or Global Administrator role so that they can define all security policy for an enterprise.

They have the capability to modify global objects, such as services and labels, add workloads, pair workloads, and change workload modes to function as a security policy administrator.

Role Access is Additive

In the following example, Joe Smith is added to two user roles and one external group and each is assigned a specific role and scope. Joe’s ability to manage security for his company is a union of the roles and scopes he is assigned to.



Exercise Caution when Combining Roles

Because role access is additive, some caution is advisable when assigning more than one role to a user. Be sure you do not grant permissions beyond what is intended. For example, suppose you are assigning a scoped role to a user. The user’s access will be restricted to workloads within the defined scope. If you then assign the Global Read Only role to the same user, the user will be able to view all workloads, including those outside the scope that was defined in the first role.

Example Role Workflows

The following example shows the hand offs between a user who is a member of the Global Organization Owner role and a member of a Ruleset Manager role.

1. An Organization Owner grants access to one or more scopes for a Ruleset Manager by selecting specific labels, which define the permitted scopes for the Ruleset Manager.
2. The Ruleset Manager logs in and creates rules that conform to the specified scopes, as defined by the labels that are accessible to that user.
3. The Ruleset Manager has read-only access to all other PCE resources, such as services or rulesets with different scopes from the scopes that the Ruleset Manager can access.
4. The Organization Owner reviews the rules created by the Ruleset Manager and provisions them as needed.

Prerequisites and Limitations

- You must be a member of the Global Organization Owner role to manage users, roles, and scopes in the PCE.
- Configuring SSO for an Illumio supported IdP is required for using RBAC with external users and groups. See [Authentication](#) for information.

If you have not configured SSO, you can still add external users and external groups to the PCE; however, these users will not be able to log into the PCE because they will not be able to reach the IdP or SAML server to authenticate.

- Illumio resources that are not labeled are not access restricted and are accessible by all users.
- External users who are designated by username and not an email address in your IdP will not receive an automatic invitation to access the PCE. You must send them the PCE URL so they can log in.
- You cannot change the primary designation for users and groups in the PCE; specifically, the email address for a local user, the username or email address for an external user, or the contents of the External Group field for an external group. To change these values, you must delete the users or groups and re-add them to the PCE.
- An App Owner who is in charge of the application in both production and development environments does not have permissions to write extra-scope rules between production and development.

Local users are not locked out of their accounts when they fail to log in. After 5 consecutive failures, the PCE emails the user that their account might be compromised.

Locked users retain all their granted access to scopes in the PCE; however, they cannot log into the PCE.

Setup for Role-based Access Control

This section describes how to configure role-based access control (RBAC) for the PCE. Before doing these tasks, be sure to understand the concepts in [Role-Based Access Control](#).

NOTE:

Permission to configure these settings is dependent on your role. See [About Roles, Scopes, and Granted Access](#) for information.

Add a Scoped Role

Add a scoped role to create fine-grained access control to manage security policy for your workloads.

You can grant different permissions to different users for different resources by defining scopes. For example, you might allow some users complete access to add rulesets for all workloads in your staging environment. For other users, you might grant access to all workloads in all environments.

1. From the PCE web console menu, choose **Access Management** > **Scopes**.
2. Click **Add**.
The Access Wizard appears.
3. Define the scope for the role by selecting labels or label groups for Applications, Environment, and Location, or for any other custom label types you have defined using Flexible Labels.
4. Add a principal -- a local user, external user, or user group -- to the role.
5. Select roles. For a description of these role, see [About Roles, Scopes, and Granted Access](#).
6. Click **Grant Access** > **Confirm**.

The newly-added role is displayed on the Scopes page and you can select it to edit or remove access.

Manage a Local User

Local users are created in the PCE (they are not managed by an IdP). When they log into the PCE, they must enter their email addresses and passwords. The Illumio PCE encrypts and stores their passwords.

When you install the PCE, the first user account it creates is a local user. You can create additional local users as a backup in case your external IdP goes offline or the SAML server is not accessible.

To add a local user:

1. From the PCE web console menu, choose **Access Management > Local Users** tab.
2. Click **Add**.
3. Enter a name and an email address.
The email address must use the format `xxxx@yyyy.zzzz` and be 255 characters or less. From the 20.1.0 release onwards, you can add email addresses with an apostrophe (') in them.
In the PCE, you can have duplicate names for local users but you cannot have duplicate email addresses.
The PCE emails the user at the address you specify an invitation with a link to create their Illumio user account. The link in invitation email is valid only for 7 days after which it expires.
4. Select a role for the user:
 - None
 - Global Organization Owner
 - Global Administrator
 - Global ViewerFor a description of these roles, see [About Roles, Scopes, and Granted Access](#).
5. Choose an access restriction for the user (or None for no restrictions).

You can change a user's role membership after adding them by going to the user's details page or from a role details page. From the 20.1.0 release onwards, the "My Roles" feature allows you to view the list of assigned permissions (roles).

To remove a local user:

1. From the PCE web console menu, choose **Access Management > Local Users**.
2. Select the user you want to remove.
3. Click **Remove User**.

When you remove a local user while the user is online, the PCE logs the user out as soon as the user is removed.

The user is removed from the Local Users tab; however, the user remains in the User Activity page and is designated as offline. The user's actions remain in the Organization Events page.

You can re-add the user to the PCE as a local or external user with the same name and email address or username.

To edit a local user:

1. From the PCE web console menu, choose **Access Management > Local Users**.
2. Click the name of the user you want to edit.
3. Click **Edit User**.
4. Change the user's name or access restriction and click **Save**.

You cannot edit a user's email address. You must remove and re-add the user with the new email address.

Changing a local user's name only changes it in the RBAC Roles pages and the Users and Groups page. The name is not changed in the user's personal profile or in the RBAC User Activity pages.

NOTE:

Local and external users can change their name when they create their accounts or from their profiles.

To convert a local user:

1. From the PCE web console menu, choose **Access Management > Local Users**.
2. Click the name of the user.
3. Click **Convert User**.

You can convert a local user to an external user so that your corporate IdP manages the user authentication credentials. When you convert a user to an external user, the user retains all their role memberships.

To invite a local user:

1. From the PCE web console menu, choose **Access Management > Local Users**.
2. Click the name of the user.
3. Click **Re-Invite**.

You can send a new email to a user to create their account when they haven't responded to the original email. An invitation remains valid for 7 days.

To lock or unlock a local user:

1. From the PCE web console menu, choose **Access Management > Local Users**.
2. Click the name of the user.
3. Click **Lock** or **Unlock**.

Manually locking out users with the above steps locks the user out indefinitely until you manually unlock their access.

Local users are automatically locked out of their accounts when they fail to log in after 5 consecutive failures.

Locked users retain all their granted access to scopes in the PCE; however, they cannot log into the PCE. When an account is locked, the PCE web console reports that the username or password is invalid even when a user enters valid credentials. If automatically locked out for consecutive login failures, the user's account resets after 15 minutes and does not require an Illumio administrator to unlock it.

Manage a Service Account

An API key can be created by the Global Organization Owner without creating a new user account to be associated with the API key. The API key can instead be associated with a service account. The service account is a security principal, just as a user is.

- A service account can perform any API operation using its API key.
- Permissions for service accounts are specified with a combination of one or more PCE roles (Global Owner, Global Admin, etc.). You can include multiple roles for a single service account, just as you can for a user account.
- Access restrictions are supported. You can limit the use of service account API keys by IP addresses, just as you can for user API keys.
- Audit events are supported with service accounts. All audit events triggered by a service account indicate the name of the service account and ID of its API key

To create a service account:

1. Choose **Access Management > Service Accounts**.
2. Click **Add**, and give the account a unique name.
3. Enter an optional description for the account.
4. Specify the access restrictions for the account.

5. Optionally change the API key expiration duration from the default value. This duration cannot exceed your organization's setting.
6. Set the roles and scopes that determine the permissions granted to the service account. Click **Add** to assign an additional role or scope in the Roles and Scopes table.
7. To create an API key for the service account, click **Save**, then click **Download Credentials**.

The new credentials are saved in the API Key section of the **Service Accounts** page.

Add or Remove an External User

Using RBAC, you can control access to Illumio Core for users who are externally authenticated by a corporate IdP. Your corporate IdP manages authentication so that when these users log into the PCE, they are redirected to the IdP to authenticate. The PCE does not validate their usernames or passwords. See [Authentication](#) for more information.

Using RBAC, you control the access external users have to Illumio Core features and functionality. When you add an external user to the PCE, you specify that user's access by assigning the user to Illumio roles and scopes.

To add an external user:

1. From the PCE web console menu, choose **Access Management > External Users** tab.
2. Click **Add**.
3. Enter a name and an email address or username.

Whether you enter an email address or username for the user depends on how you have configured your IdP to identify corporate users.

The username can contain up to 225 alphanumeric and special characters (. @ / _ % + -).

In the PCE, you can have duplicate names for external users but you cannot have duplicates email addresses or usernames.

When your IdP is configured to identify users by using email addresses, the PCE emails the user at the address you specify an invitation with a link to create their Illumio user account.

If your IdP is configured to use usernames, you must provide the user your Illumio PCE web console URL.

4. Select a role for the user:
 - None
 - Global Organization Owner
 - Global Administrator
 - Global Viewer

For a description of these roles, see [About Roles, Scopes, and Granted Access](#).

5. Specify an access restriction for the user, or leave it as None.

In general, a user must have a role in order to access the PCE. However, when default read only is ON, any user accessing the org inherits the Global Read Only role and is able to access the PCE. You can enable and disable Read Only User access in the Global Read Only role.

You can change a user's role membership after adding them by going to the user's details page or from a role details page.

To change an external user's name, click **Edit User** from the user's details page. You cannot edit the email address or username for an external user. You must remove and re-add the user with the new information.

To remove an external user:

1. From the PCE web console menu, choose **Access Management > External Users** tab.
2. Select the user you want to remove.
3. Click **Remove**.

Removing an external user removes the user from the **External Users** tab and all the user's RBAC role memberships. The user's authentication is still managed by your corporate IdP.

If Read Only User access to the PCE is enabled for your organization, the user can still log into the PCE and view resources after you remove the user.

When you remove an external user while the user is online, the PCE log the user out the next action they make after being removed.

Add or Remove an External Group

The RBAC feature in Illumio Core integrates with the user groups maintained in your corporate IdP so that you can manage user authentication centrally for the Illumio Core. In the PCE, you assign roles and scopes to the groups managed by your IdP to control the access that Illumio users have to their Illumio managed resources.

With user groups, you can authorize your teams to manage the security for the applications they manage without waiting for a centralized security team to delegate authority.

When a user who is a member of an external group logs into the PCE, the corporate IdP authenticates the user and returns the list of groups the user belongs to. For each of those groups, the PCE determines what roles and scopes are assigned to the group. The user is granted access to the resources associated with the roles and scopes.

A user can belong to multiple external groups. When a user belongs to multiple groups, the user is granted access to Illumio resources based on the most permissive role and scopes defined for each group.

To add an external group:

1. From the PCE web console menu, choose **Access Management > External Groups** tab.
2. Click **Add**.
3. In the *Name* field, enter up to 225 alphanumeric or special characters.
4. In the *External Group* field, enter the group name as it is configured in your IdP.

In your IdP, the group is designated by a simple group name (for example “Sales”) or by a group name in distinguished name (DN) format (for example “CN=Sales, OU=West”). To verify the correct format to enter in the PCE, check the `memberOf` attribute in the SAML assertion from your IdP.

The `memberOf` attribute is a multiple-value attribute that contains the list of distinguished names for groups that contain the group as a member.

5. Click **Add**.

To change an external group’s name, click **Edit Group** from the group’s details page. You cannot edit the *External Group* field. You must remove and re-add the group with the new information.

To remove an external group:

1. From the PCE web console menu, choose **Access Management > External Groups** tab.
2. Select the external group you want to remove.
3. Click **Remove**.

Removing an external group from the PCE removes all the group’s RBAC role memberships and, therefore, removes access for all the group members. User authentication for the group members is still managed by your corporate IdP.

If Read Only User access to the PCE is enabled, the external group members can still log into the PCE and view resources after you remove the group. See [About Roles, Scopes, and Granted Access](#) for more information.

Change Users and Groups Added to Roles

When you change the membership for a role, the affected users must log out and log in again to access the new capabilities.

When you revoke a user's access to scopes or global objects while the user is online, the PCE logs the user out the next action they make after having their access revoked.

1. From the PCE web console menu, choose **Access Management > Global Roles**.
2. Click the name of the role you want to change users or groups in.
3. To remove one or more users or groups from the role, select them, and click **Remove**.
4. To add a user or group to a role, click **Add**.
5. Click in the *Add Principals* field to select a user or group to add to the role. Continue to click in the field to select additional principals to add.
6. Click **Grant Access**.

Alternatively, you can select users or groups to add to roles from the various group and users details pages under **Access Management**, by clicking the desired user or group name, click **Add Role**, choose **Add Global Role** or **Add Scoped Role**, and follow the steps in the Access Wizard.

View User Activity

You can access a historical audit trail of user activity through the following reports:

- **User Activity:** Go to **Access Management > User Activity**
 - Displays session details for each user, including their status, email address, when they were last logged in.
 - Click a user, to view all the roles and scopes that are assigned to that user.

The User Activity page also displays users who were removed and are designated as offline.

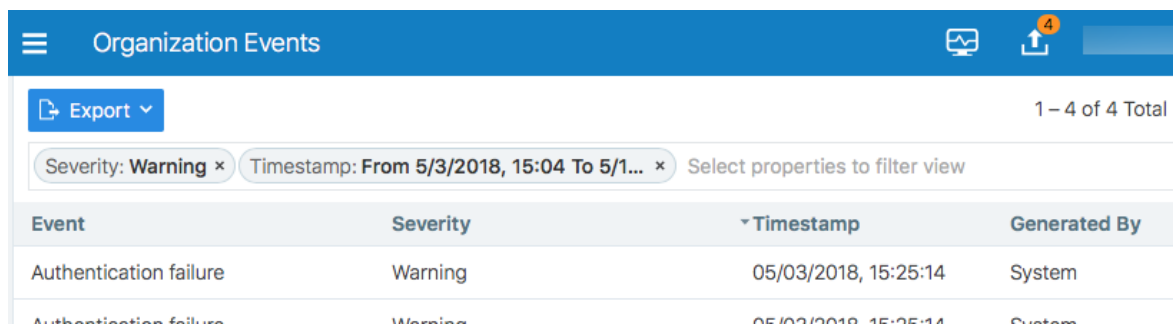
NOTE:

The names that appears in the User Activity pages can be different from the various user and groups pages under **Access Management > Users and Groups** when users edit their profiles or an Organization Owner changes names at those pages.

- **Events:** Go to **Troubleshooting > Events**

The Events page provides an ongoing log of all Organization events that occur in the PCE. For example, it captures actions, such as users logging in and logging out, and failed log in attempts; when a system object is created, modified, deleted, or provisioned; and when a workload is paired or unpaired.

Each of these events have a severity level and they are exportable in JSON format. For a large number of events, you can narrow the search by event type, severity, status, timestamp, user-generated, or agent-generated filters.



Event	Severity	Timestamp	Generated By
Authentication failure	Warning	05/03/2018, 15:25:14	System
Authentication failure	Warning	05/03/2018, 15:25:14	System

Change Your Profile Settings

If you want to change the password you use to access the PCE web console, you can do so from your User menu located at the top right corner of the PCE web console.

My Profile

Save Cancel

Personal

Email Address/Username @illumio.com

Name

Time Zone America/Los_Angeles

Accessibility

Color Mode

Normal vision
Optimize the color palette for normal vision

Color vision deficiency
Optimize the color palette for Deuteranopia, Protanopia, and Tritanopia vision

Change Password

Click to change your user account password

Change Password

To change your password:

1. From the User menu in the PCE web console, select **My Profile**.
2. Click **Change Password**.
3. On the change password screen, enter your current password, and then your new password twice.
4. Click **Change Password**.

Color Vision Deficiency Mode

Users with color vision deficiency (Deuteranopia, Protanopia, or Tritanopia) can select Color Vision Deficiency mode, which makes it easier for color vision deficiency users to distinguish between blocked and allowed traffic lines in the Illumination map. This mode can be enabled on a per-user basis.

The color vision deficiency mode is disabled by default. To enable it:

1. From the User menu in the PCE web console, select **My Profile**.
2. In the *Accessibility* section, select the **Color vision deficiency** radio button.

NOTE:

To restore the default setting, select the **Normal vision** radio button.

3. Click **Save**.

Role-based Access for Application Owners

The enhancements made to the Role-based Access Control (RBAC) framework in the Illumio Core 20.1.0 release enable organizations to address several use cases related to application owners.

Overview

These enhancements include:

- Delegation of policy writing to downstream application teams.
- Assigning read-only privileges to application owners. Those users get read access based on the assigned scopes.
- Flexibility to assign read/write or read-only privileges to the same user for different applications. For example, the same user can have read/write privileges in a staging environment but has read-only privileges in a production environment.

Although the RBAC controls in releases prior to Illumio Core 20.1.0 restricted "writes" based on user role and scope, users had visibility into all aspects of the PCE irrespective of the role. With these new RBAC controls, application owners get visibility into the applications within their assigned scopes, specifically the PCE information relevant to their applications. Depending on the user's role, application owners can:

- Read/write policies to manage application segmentation.
- View inbound and outbound traffic flows as well as use Explorer.
- View labeled objects used in policies.
- View details of global objects such as, IP Lists and Services used by their applications.

Benefits

The key benefits of the RBAC framework in the PCE are as follows:

- Provides a label based approach to define user permissions.
- Provides roles based on application owner personas to manage application segmentation.
- Provides a building block based approach to stack permissions for users.
- Offers flexibility to delegate read/write and read-only privileges to same user for different sets of applications.
- Enables enforcement of least privilege by hiding information outside of an application scope.

- Allows application owners to effectively manage segmentation for their applications.

Updates to Roles

As described in [About Roles, Scopes, and Granted Access](#), Illumio Core provides two types of user roles - Global and Scoped. It also provides the ability to stack multiple roles for the same user. A PCE owner can assign a combination of multiple roles to the same user. The resulting set of permissions is the summation of all permissions included with each of the stacked roles. With these updates:

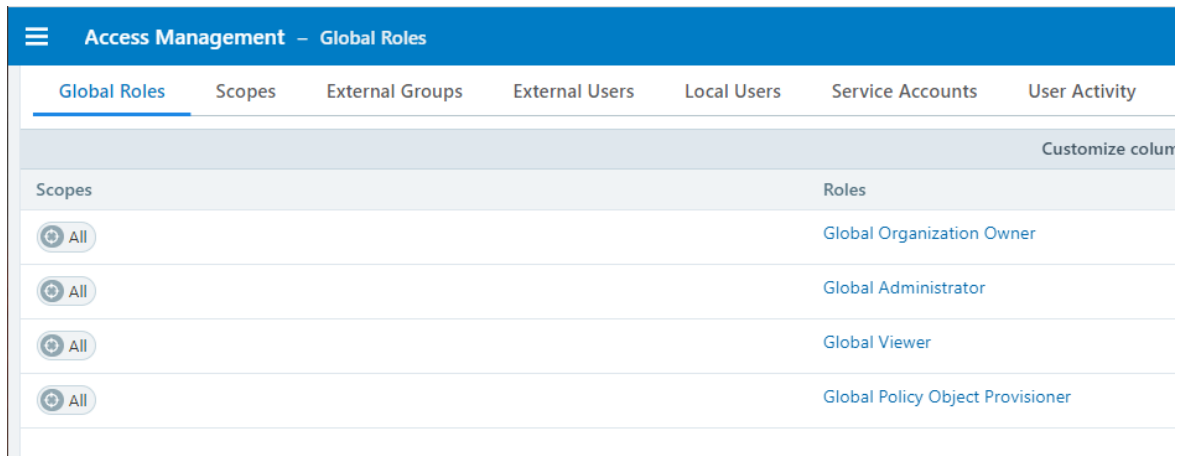
- Existing scoped roles enhanced to restrict reads by scope.
- New scope based *read-only* role limits read access by labels.
- Scoped users get limited visibility into objects 1-hop away (this applies to Explorer, App Group Maps, Rule Search, and Traffic).
- Global read-only disabled by default for new PCE installations.
- PCE performance and scale enhanced to support concurrently active users.

Global Roles

Global roles provide the user with permissions to view everything and to perform operations globally. The four Global roles are :

- Global Organization Owner: Allowed to manage all aspects of the PCE, including user management.
- Global Administrator: Allowed to managed most aspects of the PCE, with the exception of user management.
- Global Viewer: Allowed to view everything within the PCE in a read-only capacity. This role was previously called "Global Read-only".
- Global Policy Object Provisioner: Allowed to provision global objects that

require provisioning such as, Services and Label Groups.



Access Management – Global Roles	
Global Roles	Scopes
	Roles
All	Global Organization Owner
All	Global Administrator
All	Global Viewer
All	Global Policy Object Provisioner

Scoped Roles

The Scoped roles are defined using labels. The permissions included with the assigned role apply only to the assigned scope where the scope is defined using a combination of as many label types you have defined (and with only one label value per type). To provide permissions to different applications for a user, each of the application scopes has to be added to the same user.

All the Scoped roles have been enhanced to restrict reads and writes by Scope. The Scoped roles are :

- **Ruleset Viewer:** A new scope-based read-only role. A user with this role has read-only permissions within the assigned scope. The user can view policy, application groups, incoming and outgoing traffic, and labeled objects such as, workloads, within the assigned scope.
- **Ruleset Manager (Limited or Full):** An existing scope-based read/write role. A user with this role can read/write policy within the assigned scope. The user can also view application groups, incoming and outgoing traffic, and labeled objects, within the assigned scope.
- **Ruleset Provisioner:** This role allows a user to provision changes to the scoped objects, provided the objects are inside the user's assigned scope. A user with this role can provision changes to policies within the assigned scope. The user can also view application groups, incoming and outgoing traffic, and labeled objects, within the assigned scope.
- **Workload Manager:** Allows a user to perform workload-specific operations such as pairing, unpairing, assignment of labels, and changing of policy state. A user

with this role cannot view policies and traffic, and cannot provision changes.

1 Choose a Scope

America-KK x Dollar x Marketing x env group x testsw x Select Labels And Label Groups

2 Add Principals

Select a principal

Type	Name	Email/Username/Group Name	Roles
	TestTest	testtest@test.com	Ruleset Viewer x

3 Select Roles

Ruleset Management

- Ruleset Viewer**
Read-only access to Rules that match the scope. Does not permit editing of Rulesets and Rules.
- Limited Ruleset Manager**
Manage Rulesets that match the scope and Rules where the Provider and Consumer match the scope.
- Full Ruleset Manager**
Manage Rulesets that match the scope and Rules where the Provider matches the scope.
- Ruleset Provisioner**
Provision Rulesets that match the scope.
- Workload Manager**
Manage Workloads and Pairing Profiles that match scope.

Rulesets and Rules	View Scope
Workloads and VENS	View Scope
Illumination Map	None
App Group Map	View Scope
App Groups List	View Scope
Illumination Plus	View Scope
Scopes and Roles	None
Users and Groups	None
Services	View
IP Lists	View
User Groups	View
Label Groups	View
Virtual Services	View Scope
Virtual Servers	View
Labels	View
Pairing Profiles	None
Infrastructure	None
Blocked Traffic	View Scope
Security Settings	None
App Group Configuration	None
My Profile	View, Modify
My API Keys	View, Add, Modify, Delete
SSO Config	None

Summary Scope: America-KK, Dollar, Marketing, env group, testsw
Principals: TestTest
Role: Ruleset Viewer

[Cancel](#) [+ Grant Access](#)

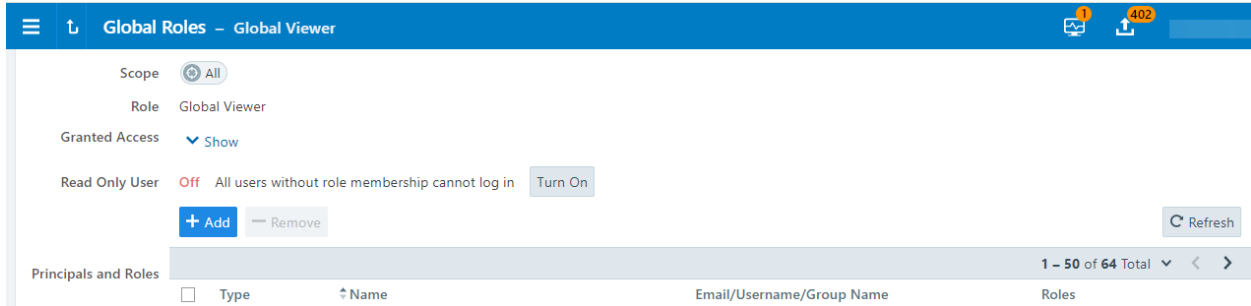
Configuration

The Global Read-only user setting should be disabled to enforce scoped reads for users with scoped roles. To disable this setting, make sure that the *Read Only User* setting under **Access Management > Global Roles > Global Viewer** is set to **Off**.

NOTE:

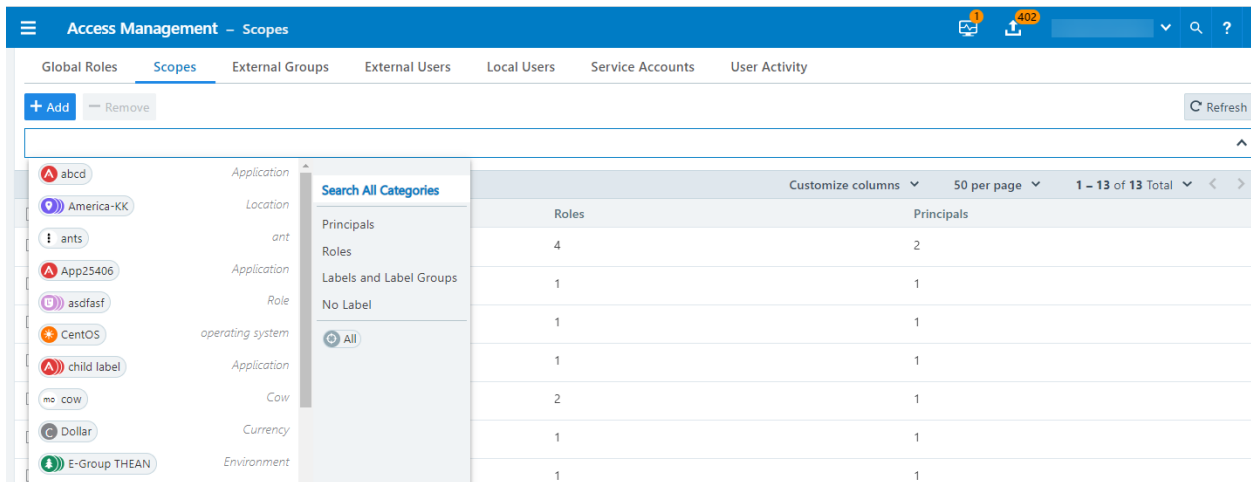
In PCE versions 20.1.0 and higher, the Global Read-only user setting is disabled by default.

On PCE versions that are upgraded from prior releases, this setting must be manually turned off for users to have reads restricted by scope. If this setting is set to **On**, users with scoped roles will get global visibility by default.



Facet Searches for Scoped Roles

The Scopes page now features a search bar with auto-complete and facets. This is restricted to users with a Global Organization Owner role. To use this feature, navigate to **Access Management > Scopes**. The search bar allows Organization Owners to query a list of users by a user's role. They can search by labels and label groups to get a list of users with the selected label(s) in their assigned scope(s), or for users with no labels assigned. They can also select Principals to search for a specific user.



Ruleset Viewer

Ruleset Viewer is a new scope-based read-only role. When assigned, a user get read-only visibility into the assigned application scope. As a Ruleset Viewer, you can view all the Rulesets and Rules within the assigned scope. However, you cannot edit any of the rules or create new rules. You can use Policy Generator to preview the policies that will be generated. However, you are not allowed to save policy after previewing it using Policy Generator.

A Ruleset Viewer is allowed to view everything that a Ruleset Manager with the same scope is allowed to view. This includes traffic flows, labeled objects, application groups, global objects, and so on. The only difference between a Ruleset Manager and

a Ruleset Viewer is the absence of write privileges for a Ruleset Viewer. A Ruleset Manager is allowed to create and update policy within the application scope.

Scoped Roles and Permissions

The following table provides a summary of the different permissions provided with each of the scoped roles.

- (R) = Restricted based on scope
- (T) = Restricted based on resource type
- --- = Not applicable

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provisioner	Workload Manager	Application Owner (Combined Permissions)
Traffic - Illumination, App Group, Explorer					
Illumination Location Map	---	---	---	---	---
App Group Policy Map	Read (R)	Read (R)	Read (R)	---	Read (R)
App Group Vulnerability Map	Read (R)	Read (R)	Read (R)	---	Read (R)
App Group List	Read (R)	Read (R)	Read (R)		Read (R)
Explorer	Read (R)	Read (R)	Read (R)	---	Read (R)
Blocked Traffic	Read (R)	Read (R)	Read (R)	---	Read (R)
Policy					
Policy Generator	Read (R)	Read+Write (R)	Read (R)	---	Read+Write (R)
Rulesets and Rules	Read (R)	Read+Write (R)	Read (R)	---	Read+Write (R)
Rule Search	Read (R)	Read (R)	Read (R)	---	Read (R)
Policy Check	Read (R)	Read (R)	Read (R)	---	Read (R)
Provisioning Draft Changes	Read (R)	Read (R)	Read+Write (R)	---	Read+Write (R)
Policy Ver-	Read (R)	Read (R)	Read (R)	---	Read (R)

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provisioner	Workload Manager	Application Owner (Combined Permissions)
sions					
Provisioning Status	Read (R)	Read (R)	Read (R)	---	Read (R)
Labeled Objects					
Workloads	Read (R)	Read (R)	Read (R)	Read+Write (R)	Read+Write (R)
Container Workloads	Read (R)	Read (R)	Read (R)	Read (R)	Read (R)
Virtual Enforcement Nodes	Read (R)	Read (R)	Read (R)	Read+Write (R)	Read+Write (R)
Pairing Profiles	---	---	---	Read+Write (R)	Read+Write (R)
Virtual Services	Read (R)	Read (R)	Read (R)	Read (R)	Read (R)
Virtual Servers	Read	Read	Read	Read	Read
Global Policy Objects					
Services	Read	Read	Read	Read	Read
IP Lists	Read	Read	Read	Read	Read
User Groups	Read	Read	Read	Read	Read
Labels	Read	Read	Read	Read	Read
Label Groups	Read	Read	Read	Read	Read
Settings					
Segmentation Templates	---	---	---	---	---
Role-Based Access Global Roles	---	---	---	---	---
Role-Based Access Scoped Roles	---	---	---	---	---
Role-Based	---	---	---	---	---

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provisioner	Workload Manager	Application Owner (Combined Permissions)
Access Users and Groups					
Role-Based Access User Activity	---	---	---	---	---
Load Balancers	---	---	---	---	---
Container Clusters	---	---	---	---	---
Bi-directional Routing Networks	---	---	---	---	---
Event Settings	---	---	---	---	---
Setting Security	---	---	---	---	---
Setting Single Sign-On	---	---	---	---	---
Setting Password Policy	---	---	---	---	---
Setting Offline Timers	---	---	---	---	---
VEN Library	---	---	---	Read	Read
My Profile	Read+Write	Read+Write	Read+Write	Read+Write	Read+Write
My API Keys	Read+Write	Read+Write	Read+Write	Read+Write	Read+Write
Other					
Support Reports	---	---	---	Read+Write (R)	Read+Write (R)
Events	---	---	---	---	---
Reports	Read (R, T)	Read (R, T)	Read (R, T)	Read (R, T)	Read (R)
Support	Read	Read	Read	Read	Read
PCE Health	---	---	---	---	---
Product Version	Read	Read	Read	Read	Read


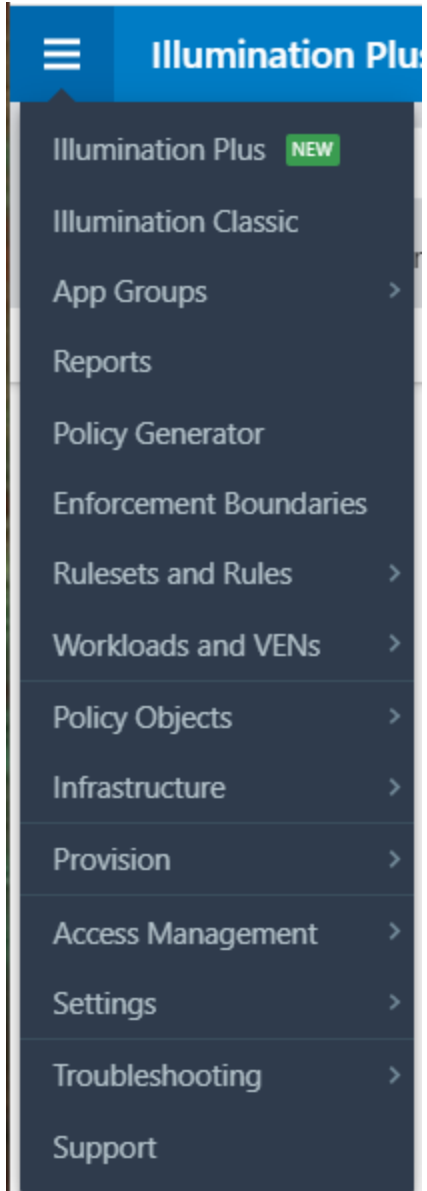
Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provisioner	Workload Manager	Application Owner (Combined Permissions)
Help	Read	Read	Read	Read	Read
Terms	Read	Read	Read	Read	Read
Privacy	Read	Read	Read	Read	Read
Patents	Read	Read	Read	Read	Read
About Illumio	Read	Read	Read	Read	Read

Scoped Users and PCE

Each scoped role has different permissions that impact an application owner's visibility into various aspects of the PCE. Application owners can be assigned scoped roles that come with different permissions.

Navigation Menus

The PCE navigation menu options vary based on the user's role. The navigation menu options available for Application Owner are limited. For example, a user is logged in as a Global Organization Owner has more (complete) menu options displayed than when a user logs in as a scoped user (Application Owner).



The following table provides the menu options available for different scoped users.

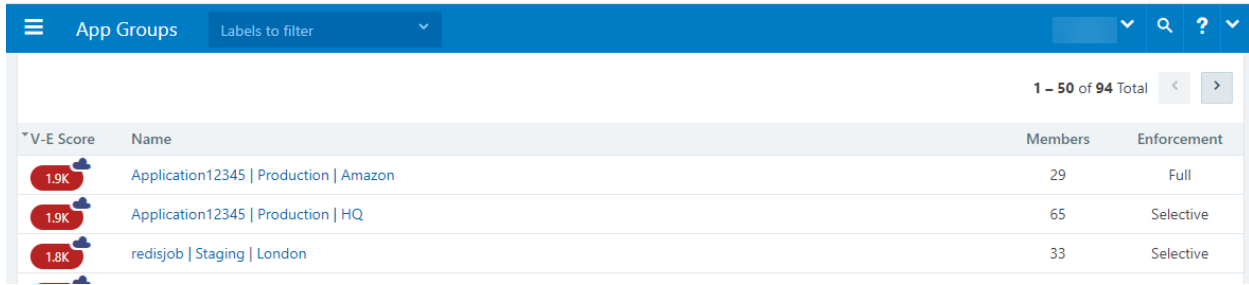
- Y = Yes (menu option is displayed for the user)
- N/A = Not applicable (menu option is hidden from the user)

Page	Ruleset Viewer	Ruleset Manager	Ruleset Provisioner	Workload Manager
Illumination Map	N/A	N/A	N/A	N/A
Role-based Access	N/A	N/A	N/A	N/A
Policy Objects > Segmentation Templates	N/A	N/A	N/A	N/A
Policy Objects > Pairing Profiles	N/A	N/A	N/A	Y
Infrastructure	N/A	N/A	N/A	N/A
Troubleshooting > Events	N/A	N/A	N/A	N/A
Troubleshooting > Support Reports	N/A	N/A	N/A	Y
Settings	N/A	N/A	N/A	See row below
Settings > VEN Library	N/A	N/A	N/A	Y
PCE Health	N/A	N/A	N/A	N/A
App Groups > Map	Y	Y	Y	N/A (App Group Members are visible)
App Groups > List	Y	Y	Y	Y
App Groups > Vulnerability Map	Y	Y	Y	N/A
Explorer	Y	Y	Y	N/A
Policy Generator	Y	Y	Y	N/A
Rulesets and Rules	Y	Y	Y	N/A
Rule Search	Y	Y	Y	N/A
Workload Management > Workloads	Y	Y	Y	Y
Workload Management > Container Workloads	Y	Y	Y	Y
Workload Management > Virtual Enforcement Nodes (Agents)	Y	Y	Y	Y

Page	Ruleset Viewer	Ruleset Manager	Ruleset Provisioner	Workload Manager
Provision > Draft Changes	Y	Y	Y	N/A
Provision > Policy Versions	Y	Y	Y	N/A
Policy Objects > IP Lists	Y	Y	Y	Y
Policy Objects > Services	Y	Y	Y	Y
Policy Objects > Labels	Y	Y	Y	Y
Policy Objects > User Groups	Y	Y	Y	Y
Policy Objects > Label Groups	Y	Y	Y	Y
Policy Objects > Virtual Services	Y	Y	Y	Y
Policy Objects > Virtual Servers	Y	Y	Y	Y
Troubleshooting > Blocked Traffic	Y	Y	Y	N/A
Troubleshooting > Export Reports	Y	Y	Y	Y
Troubleshooting > Policy Check	Y	Y	Y	N/A
Troubleshooting > Product Version	Y	Y	Y	Y
Support	Y	Y	Y	Y
My Profile	Y	Y	Y	Y
My Roles	Y	Y	Y	Y
My API Keys	Y	Y	Y	Y
Help	Y	Y	Y	Y
Terms	Y	Y	Y	Y
Patents	Y	Y	Y	Y
Privacy	Y	Y	Y	Y
About Illumio	Y	Y	Y	Y

Landing Page

The PCE landing page changes dynamically based on the user's role. When you log in to your account as an Organization Owner, the Illumination page opens. However, when you log in as a Scoped user, the landing page changes to the App Groups List page where you can see the list of App Groups assigned.



V-E Score	Name	Members	Enforcement
1.9K	Application12345 Production Amazon	29	Full
1.9K	Application12345 Production HQ	65	Selective
1.8K	redisjob Staging London	33	Selective

Labeled Objects

Labeled objects, such as workloads are filtered by the scope of the user. On the Workloads page, you will only see the list of the workloads within the application scope. You cannot see any workloads that are outside the application scope. This applies to any labeled object, such as workloads, containers, Virtual Services, and Virtual Enforcement Nodes (VENs).

The menu functions and buttons change dynamically to reflect a user’s permissions. If you are logged in as a Ruleset Manager, you are not allowed to manage workloads. So, all the workload-specific operations buttons are disabled. However, you are allowed to view the list of workloads within the scope and get details for individual workloads, except for Virtual Servers.

NOTE:

While Virtual Servers are considered labeled objects, they are visible to all scoped users regardless of object scope.

Facet Searches and Auto-complete

The search bar with auto-complete and facets is scoped for labeled objects and Rule-sets. For example, you search for Application Labels, then you can only select the Application Labels under the assigned scope. This applies to other label types such as Environment labels and Location labels. However, Role labels are excluded since Role labels are not part of the user scope. The restriction of visibility by scope applies to facets such as hostname, IP address, and others. The search bar automatically filters the facets to the list of facets in the user’s assigned scope.

Global Objects

Scoped users get full read-only visibility into all global objects. This includes IP Lists, services, labels, label groups, and user groups. However, scoped users are not allowed to create, modify, or provision global objects.

NOTE:
Only Global Organization Owner and Global Administrator can create, modify, and provision global objects.

Rulesets and Rules

Scoped users, with the exception of Workload Managers, are allowed to see rulesets and rules which apply to their application. A Ruleset Manager is allowed to edit the ruleset whereas the other scoped roles (Ruleset Viewer and Ruleset Provisioner) are allowed to view rulesets. A scoped user can see all the rules within the application rule-set.

When label groups are used within the scope of a ruleset, a Ruleset Manager may not be allowed to edit the ruleset and it's rules even if there is a scope match between the user's assigned scope and the underlying scope of the ruleset. The user will however be able to view the rules within such a ruleset.

In addition, scoped users can also see rules which apply to their application. For example, scoped users are allowed to view rules written by other applications that apply to their application. To see those rules, click **Rule Search** from the navigation menu.

Provision Status	Status	Name	Scope (App Env Loc)	Last Modified On
<input type="checkbox"/>	Enabled	AO_App1 AO_Env1 AO_Loc1	AO_App1 AO_Env1 AO_Loc1	02/03/2020, 10:08:55
<input type="checkbox"/>	Enabled	erTests	CRM Production Amazon	02/10/2020, 11:12:57

On the Rule Search page, a scoped user can see all the rules that apply to their application. This includes rules for incoming and outgoing traffic flows. The rules highlighted in the screenshot below are the outbound rules which are for your application. Application Owner provides the visibility to see all the rules that are applied to your application.

Rule Search PG_AppOwnerUser1_RV... 🔍 ?

Draft Rules ▾ Basic ▾ Exact Results ▾ 1 – 5 of 5 Total 🔄

Columns ▾ Download 📄 Reset Filters 🔄

Filter by Labels and Rule attributes Go


Providers	Providing Service	Consumers	Ruleset	Access	Note
Web	33 UDP	All Workloads	ingA...	Editable	
unmanaged-dhcp	67 TCP	nfs	AO_App1 AO_Env1 AO_Loc1	Read-only	
All Workloads	468 UDP	All Workloads	app1Exp env1Exp loc1Exp	None	
Mail	33 UDP	Web Production Amazon CRM	CRM Test Rackspace	None	
All Workloads	Service - ICMP ICMP	Production Amazon ex-w1 CRM	Sales Staging Rackspace	None	

1 – 5 of 5 Total 🔄

App Group Map

The App Group Map provides complete visibility into applications and everything inside the application. Scoped users, with the exception of Workload Managers, can view App Group Maps. Scoped users get complete visibility into everything inside their application group. Scoped users can see workload objects, labels, traffic flows and every other detail within their application group.

For connected App Groups such as Providing App Groups and Consuming App Groups, scoped users get limited visibility. Scoped users get limited information on endpoints with traffic flows to their application. For an endpoint in a connected App Group from which there is traffic flow, scoped users can get limited information such as labels, role name, and hostname. The scoped user is not allowed to view any other endpoints in the connected App Group from which there are no traffic flows. To view the Illumination Map, the user should be assigned a Global role such as Global Organization Owner, Global Administrator, or Global Viewer.



NOTE:
For Scoped Roles, only the App-Group Map is available and the Illumination Map is not available.

Policy Generator and Explorer

With Policy Generator, scoped users can generate policies only for their applications. Only Ruleset Managers are allowed to generate policy with Policy Generator. Ruleset Viewers are allowed to preview Policy Generator without the ability to save policy.

Explorer views are also filtered for scoped users. To use Explorer, one of the endpoints has to be within the scoped user's application. The same applies to Blocked Traffic.

Reported Policy Decision	Connection State	Consumer	Consumer Labels	Provider	Provider Labels	Port/Process [User]	Flows	First Detected
Potentially Blocked by Provider	Closed	192.168.125.37		Ubuntu-Linux-7 45 Unicast	ex-wl1 CRM Production Amazon	22 TCP sshd [root]	2	02/13/2020 16:41:07
Potentially Blocked by Provider	Active	192.168.125.37		Ubuntu-Linux-7 1.45 Unicast	ex-wl1 CRM Production	22 TCP	2	02/13/2020 19:42:30

My Roles

"My Roles" is a new feature that allows you to view the list of assigned permissions (roles).

Type	Scope (App Env Loc)	Roles
Scoped	CRM Production Amazon	Ruleset Manager, Ruleset Viewer
Scoped	AO_App1 AO_Env1 AO_Loc1	Ruleset Viewer

Configure Access Restrictions and Trusted Proxy IPs

To employ automation for managing the PCE environment, you can use API Keys created by an admin user and automate PCE management tasks. This section tells how you can restrict the use of API keys and the PCE web interface by IP address. In this way, you can block API requests and users coming in from non-allowed IP addresses.

Configure Access Restrictions

This section tells how to use the Illumio web console UI to configure access restrictions. You can also configure access restrictions programmatically using the REST API calls described in [Access Restrictions and Trusted Proxy IPs](#) in the *REST API Developer Guide*.

- You must have the global Org Owner role to view or change access restrictions.
- A maximum of 50 access restrictions can be defined.

To configure access restrictions:

1. Log in to the PCE web console as a user with the Global Org Owner role.
2. Open the menu and choose **Access Management - Access Restrictions**.

The Access Restriction page opens with a list that shows which IP addresses are allowed and where the restrictions have been applied.

3. To add a new restriction, click **Add**.

The Add Access Restriction page opens.

Provide the required attributes:

- Provide a name.
 - In **Restriction Applies To**, choose User Session, API Key, or Both. Access restrictions can be applied to these different types of user authentication.
 - List a maximum of eight IPv4 addresses or CIDR blocks.
4. Click **Edit** to edit the restriction.
 5. View the access restrictions applied to local users. The default is blank, no restrictions.
 6. You can assign access restrictions to local and external users or user groups. To add a local user:
 - a. Click **Add**.
 - b. In **Access Restriction**, choose the type of access restriction.
 - c. Click **Add**.
 7. View the local user's detail page. To modify the user settings, click **Edit User**.
 8. Use the Edit User dialog to apply restrictions.

If an Org Owner assigns an access restriction to any Org Owner, a warning is shown, because this can result in the Org Owner user losing access to the PCE.


9. View the list of API keys in the API Keys page and the Event page.

Configure Trusted Proxy IPs


This section tells how to use the Illumio web console UI to configure trusted proxy IPs. You can also configure trusted proxy IPs programmatically using the REST API calls as described in [Access Restrictions and Trusted Proxy IPs](#) in the *REST API Developer Guide*.

When a client is connected to the PCE's haproxy server, this connection can traverse one or more load balancers or proxies. Therefore, the source IP address of a client connection to haproxy might not be the actual public IP address of the client.

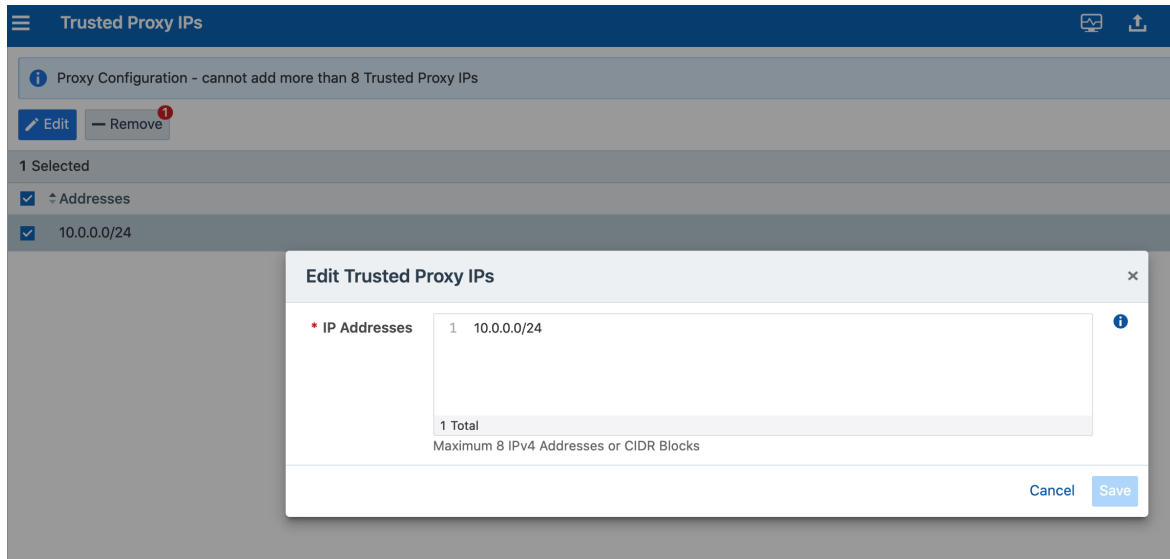
1. Log in to the PCE web console as a user with the Global Org Owner role.
2. Open the menu and choose **Settings - Trusted Proxy IPs**.



3. Click **Edit**.



4. In **IP Addresses**, enter up to eight IPv4 addresses or CIDR blocks.



5. Click **Save**.

Password Policy Configuration

The PCE enforces password policies that only a Global Organization Owner can configure. In the PCE web console, you set password policies that the PCE enforces, such as password length, composition (required number and types of characters), and password expiration, re-use, and history.

About Password Policy for the PCE

You need to be a Global Organization Owner to view the Password Policy feature under the Settings > Authentication menu options.

Prior to Illumio Core 18.2.0, a Global Organization Owner set the password in the PCE by using the PCE runtime script. The settings in the PCE runtime script are the same as before Illumio Core 18.2.0, except that the password length can now be set to a maximum of 64 characters.

NOTE:

The Password Policy feature is not applicable for organizations using SAML authentication.

NOTE:

Permission to edit this setting is dependent on your role. See [About Roles, Scopes, and Granted Access](#) for information.

Password Requirements

The password requirements you set are displayed to users when they are required to change their passwords. You can set the minimum character length, ranging from a minimum of 8 characters to a maximum of 64 characters. The default length is 8 characters.

A Global Organization Owner should configure passwords based on the following categories:

- Uppercase English letters
- Lowercase English letters
- Numbers 0 through 9 inclusive
- Any of the following special characters: ! @ # \$ % ^ & * < > ? .

WARNING: Any other special characters are neither tested nor supported.

You have to select at least three of the above categories. The default password requirement is one number, one uppercase character, and one lowercase character. You can set the password to use either one or two characters from each category.

Password Expiration and Reuse

You can set the password expiration range from 1 day to 999 days. The default setting for password expiration is “Never.”

You can set the password reuse history from 1 to 24 passwords before a user can reuse the old password. The default setting is five password changes before reuse of the password is allowed.

NOTE:

The number of password changes before password reuse is allowed is the value you enter + 1 (the current password). For example, when you specify 3, the number of passwords before reuse is allowed is 4.

You can also set the similarity of a password by not allowing a user to change their password unless it changes from a minimum of 1 to a maximum of 4 characters and positions from their current password.

Allowable password reuse and password history can be set to from 1 to 24 passwords before reuse is allowed. The default setting for password reuse is five password changes before reuse is permitted.

Caveats

- When a Global Organization Owner increases the required minimum password length policy or increases the password complexity requirements and enables the password expiration (1-999 days), all the existing users must reset their passwords based on the new policy.
- When a Global Organization Owner configures the password to never expire, all users who were migrated from an older release to 18.2.0 must reset their passwords when they next log in.

Change Password Policy Settings

1. From the PCE web console menu, choose **Access -> Authentication**.
2. In the Authentication Settings screen, choose the Authentication Method to authenticate users for accessing the PCE:
 LOCAL (IN USE) : User will sign in to the PCE only with a local credential provided by the user’s organization password policy.
 SAML (IN USE) : SAML users can also authenticate to the PCE using local credentials.
 LDAP: LDAP user can also authenticate to the PCE using local credentials>
3. Once you decide which option to take, click on the **Configure** button.
4. Depending on the authentication method, these are the available options:
 Choose option LOCAL, SAML, or LDAP:

LOCAL (in use)	
Password requirements	
Min lengths	8 characters
Character categories	A-Z (required), a-z (required), 0-9 (required)
Min characters per category	1
Password expiration and reuse	
Expiration	Never
Reuse history	1 password changes
Similarity	1 character and position from the current password

SAML (in use)	
	xm2ckxsWDTyONV8ytLQKwp93exxqmzzpbz6qi23y0B4u4af+/SW9 ukjzD/atP34bY1YjeLBCsKEgy1nDTVgypAZSEy46kJ9mAu6t3r4/gEg XTkMYQDtrPA= -----END CERTIFICATE-----
Remote login URL	https://hohoho.illumio.com
Logout landing URL	https://hohoho.illumios.com/1logout
Information for identity provider	
Authentication method	unspecified
Force re-authentication	no
Sign SAML request	no
SAML version	2.0
Issuer URL	https://2x2testlab360.ilabs.io:8443/login
NameID format	urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress
Assertion consumer URL	https://2x2testlab360.mylabs.io:8443/login/acs/6b5243ef-2305-4ffd-bf81-4fa97fb91a5b
Logout URL	https://2x2testlab360.mylabs.io:8443/login/logout/6b5243ef-2305-4ffd-bf81-4fa97fb91a5b
Timeout	30 minutes

- LDAP authentication is not active. Click **Turn On** to apply on all the LDAP servers.
- To create an LDAP server, click on **Create Server**.
To continue with LDAP server configuration, see [Enable LDAP Authentication](#).

Authentication

The Illumio PCE supports the use of either SAML SSO or LDAP as an external authentication method. Both SAML SSO and LDAP cannot be used at the same time. When

LDAP is turned on, the use of SAML SSO, if already configured, is disabled. Similarly, enabling SAML SSO after LDAP is enabled will disable LDAP authentication.

SAML SSO Authentication

When you use a third-party SAML-based Identity provider (IdP) to manage user authentication in your organization, you can configure that IdP to work with the PCE. By configuring a single sign-on (SSO) IdP in the PCE, you can validate usernames and passwords against your own user management system, rather than having to create additional user passwords managed by the Illumio Core.

Illumio Core currently supports the following SAML-based IdPs:

- Azure AD
- Microsoft Active Directory Federation Services (AD FS)
- Okta
- OneLogin
- Ping Identity

NOTE:

You can use other SAML-based IdPs; however, configuring those IdPs is your responsibility as an Illumio customer.

Before you configure SSO in the PCE, you need to configure SSO on your chosen IdP and obtain the required SSO information. After obtaining the IdP SSO information, log into the PCE web console and complete the configuration.

PCE Information Needed to Configure SSO

Before you configure SSO in the PCE, obtain the following information from your IdP:

- x.509 certificate
- Remote Login URL
- Logout Landing URL

The PCE supports the following optional attributes in the SAML response from the IdP:

- User.FirstName - First Name
- User.LastName - Last Name
- User.MemberOf - Member of

Details

User email address is the primary attribute used by the PCE to uniquely identify users.

IMPORTANT:

The client browser must have access to both the PCE and the IdP service.
The Illumio PCE uses HTTP-redirect binding to transmit SAML messages.

To obtain the SSO information from the PCE:

1. From the PCE web console menu, choose **Access Management > Authentication**.
2. On the Authentication Settings screen, locate the SAML configuration panel and click **Configure**.
3. Use the displayed information (as shown in the example below) while configuring your specific IdP.

Information for Identity Provider

Authentication Method	Unspecified
Force Re-authentication	No
SAML Version	2.0
Issuer	https://c...../3/login
NameID Format	urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress
Assertion Consumer URL	https://...../3/login/acs/a63e.....49598e
Logout URL	https://.....43/login/logout/a63e.....49598e

NOTE:

Even though the SAML NameID format specifies an emailAddress, the PCE can support any unique identifier such as, userPrincipalName (UPN), common name (CN), or samAccountName as long as the IdP is configured to map to the corresponding unique user identifier.

Signing for SAML Requests

There are four new APIs you can use to sign SAML requests:

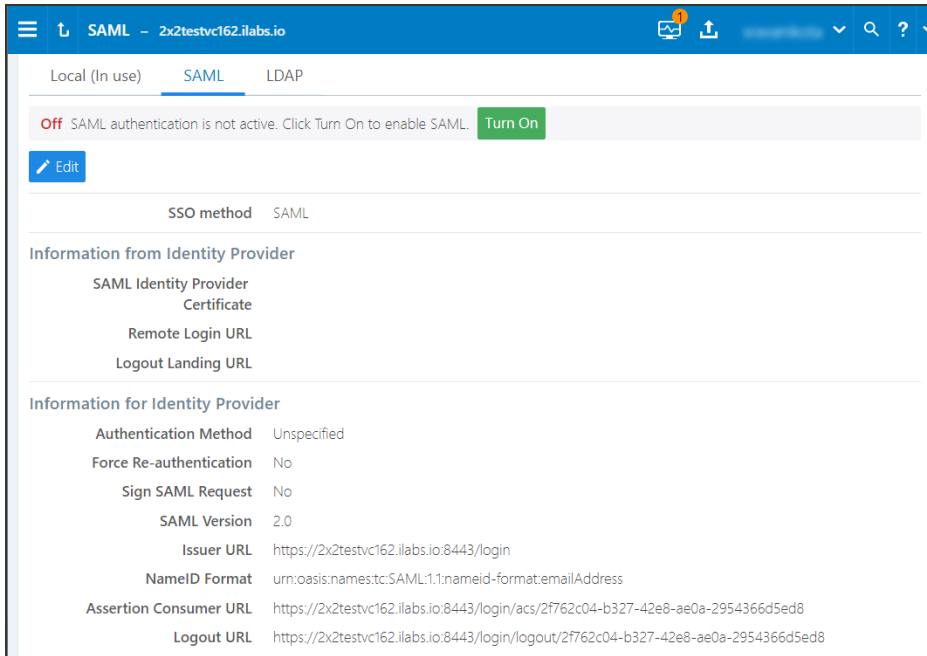
- GET /authentication_settings/saml_configs
- GET /authentication_settings/saml_configs/:uuid
- PUT /authentication_settings/saml_configs/:uuid
- POST /authentication_settings/saml_configs/:uuid/pce_signing_cert

These APIs are covered in detail in the *REST API Developer Guide*.

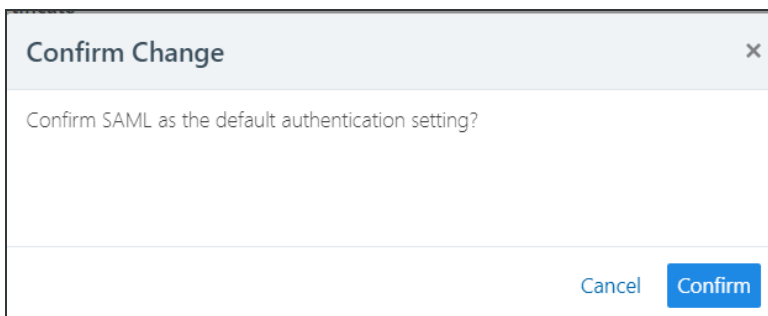
Signing of SAML requests is, however, disabled by default.

To enable SAML request signing:


1. Using the Web Console, go to **Access Management > Authentication**.
2. In the *Authentication Setting* screen, select **Configure** button for SAML.
3. In the SAML screen, click **Turn On**.



4. In the pop-up screen, click **Confirm**.



The updated SAML screen shows that SAML authentication is active.



If necessary, you can disable it at any time.

Once configured using these steps, the lifetime of the SAML certificate is ten years.

LDAP Authentication

The PCE supports LDAP authentication for users with OpenLDAP and Active Directory. The PCE supports user and role configuration for LDAP users and groups. You can configure up to three LDAP servers and map users and user groups from your LDAP servers to PCE roles. Core Cloud does not support LDAP authentication.

To use LDAP authentication:

1. Review the [Prerequisites and Limitations](#).
2. Enable the PCE to use LDAP authentication. See [Enable LDAP Authentication](#).
3. Set up an LDAP configuration. See [Configure LDAP Authentication](#).
4. Map your LDAP groups to one or more PCE roles. See [Map LDAP Groups to User Roles](#).

Prerequisites and Limitations

Before configuring LDAP for authentication with the PCE, complete the following prerequisites, and review the limitations.

Determine Your User Base DN (Distinguished Name)

Before you map your LDAP settings to PCE settings, determine your user base distinguished name ("DN"). The DN is the location in the directory where authentication information is stored.

If you are unable to get this information, contact your LDAP administrator for assistance.

Additional Considerations

When configuring the PCE to work with LDAP, be aware of the following support:

- PCE uses LDAP protocol version 3 ("v3").
- Supported LDAP distributions include OpenLDAP 2.4 and Active Directory.
- Supported LDAP protocols include LDAP, LDAPS, or LDAP with STARTTLS.

Limitations

- Any user that is created locally will have precedence over an LDAP user of the same name. For example, if the LDAP server has a user with a username attribute (such as, cn or uid) of johndoe and the default PCE user of the same name is present, the PCE user takes precedence. Only the local password will be accepted and on login, the roles mapped to the local user will be in effect. To work around this limitation, you must delete the specific local user.
- LDAP and SAML single sign-on cannot be used together. An organization can either use LDAP or SAML single sign-on for authenticating external users.

Enable LDAP Authentication

To enable LDAP authentication:

1. Log in to the PCE web console as a Global Organization Owner.
2. Choose **Access** > **Authentication**.
3. In the Authentication Settings screen, locate the LDAP configuration panel and select **Configure**.
4. In the LDAP Authentication screen, select **Create Server**.

Configure LDAP Authentication

Follow these steps to configure LDAP authentication on the PCE. Make sure you have first followed the steps in [Enable LDAP Authentication](#).

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access** > **Authentication**.
3. On the Authentication Settings screen, locate the LDAP configuration panel and click **Configure**.
4. In the LDAP Authentication screen, make sure LDAP is enabled.
5. Click **+ Create Server**.
6. In the LDAP Server Create Screen, enter information to configure LDAP as follows:
 - Name: Enter a friendly name for the LDAP server.
 - IP Address or Hostname: The IP address or hostname of the LDAP server.

- Protocol: Select one from LDAP, LDAPS (Secure LDAP) or LDAP with STARTTLS.
 - Port: Enter a port number if you are not using a default port. Default ports are 389 for standard LDAP, 636 for LDAPS, and 389 for LDAP with STARTTLS.
 - Anonymous Bind: When using an Open LDAP server, you can use anonymous bind. Choose **Allow** if you want to use anonymous bind. When using Active Directory, the use of Anonymous Bind is not recommended. Choose **Do not Allow** and specify values for Bind DN and Bind Password.
 - Bind DN: Distinguished name (DN) used to bind to the LDAP server. The bind DN is required only when Anonymous Bind is set to **Do not Allow**.
 - Bind Password: Required only when Bind DN is required. When using Anonymous Bind, no bind password is used.
 - Request Timeout Period: This is the number of seconds to wait for a response from the LDAP server. The default is 5 seconds. It can be configured to any value from 1-60 seconds.
 - Trusted CA Bundle: The bundle of certificates including the chain of trust to use when the LDAP server uses either LDAPS or LDAP with STARTTLS.
 - Verify TLS: Enabled by default. This flag specifies whether to verify the server certificate when establishing an SSL connection to the LDAP server. Disabling this is not recommended.
 - User Base DN: Base DN of the LDAP directory to search for users.
 - User Search Filter: Search filter used to query the LDAP tree for users.
 - User Name Attribute: Attribute on a user object that contains the user-name. For example, uid, sAMAccountName, userPrincipalName.
 - Full Name Attribute: Attribute of a user object that contains the full name. For example, cn, commonName, displayName.
 - Group Membership Attribute: Attribute of a user object containing group membership information. For example, memberOf, isMemberOf.
7. Click **Test Connection** to verify that the PCE is able to successfully connect to the LDAP server. If Test Connection fails, check your LDAP configuration and retry.

You can enter up to three LDAP server configurations for a PCE. For more information about using multiple LDAP servers, see [How the PCE Works with Multiple LDAP Servers](#).

Map LDAP Groups to User Roles

After you configure the PCE to use LDAP authentication, map PCE user roles to the LDAP server's groups. When a user attempts to log in, the PCE queries the server(s) to find the user. It grants the user permissions based on any PCE user roles associated with the LDAP groups in which the user is a member.

To change user permissions, use one of the following options:

- To change the permissions for a group of users, you can remap the LDAP group to a different PCE role.
- To change the permissions for an individual user, you can move the user to an LDAP group mapped to a different PCE role. You do this action on the LDAP server.

You can also perform these user management activities:

- Add a user to a PCE role: On the PCE, map the PCE role to an LDAP group. Then, on your LDAP server, add the user to that LDAP group.
- Remove a user from a PCE role: Remove the user from the corresponding LDAP group on your LDAP server.

A user can have membership in several roles. In that case, the user has access to all the capabilities available for any of those roles. For example, if a user is a member of both the docs and eng LDAP server groups, and the docs group is mapped to the PCE user role "Ruleset Manager" and the eng group is mapped to "Ruleset Provisioner," the user obtains all permissions assigned to both the "Ruleset Manager" and "Ruleset Provisioner" roles.

NOTE:

The PCE checks LDAP membership information when a user attempts to log in. You do not need to reload the authentication configuration when adding or removing users.

For details about how to map external groups to PCE user roles, see [Setup for Role-based Access Control](#).

Modify LDAP Configuration

Follow these steps to update or delete an LDAP configuration in the PCE. It is assumed you have already followed the steps in [Enable LDAP Authentication](#) and [Configure LDAP Authentication](#).

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access Management > Authentication**.
3. On the Authentication Settings screen, locate the LDAP configuration panel and click **Configure**.
4. In the LDAP Authentication screen, make sure LDAP is enabled.
5. Choose the desired action:
 - To delete a configuration, click the **Remove** icon.
 - To modify a configuration, click the **Edit** icon.

Verify LDAP Connectivity

Follow these steps to test the PCE's connection to the LDAP server(s). It is assumed you have already followed the steps in [Enable LDAP Authentication](#) and [Configure LDAP Authentication](#).

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access Management > Authentication**.
3. On the Authentication Settings screen, locate the LDAP configuration panel and click **Configure**.
4. In the LDAP Authentication screen, make sure LDAP is enabled.
5. The LDAP Authentication screen displays a list of configured LDAP server entries. Click **Test Connection** next to each entry to check whether the configuration is working.

Secure LDAP with SSL/TLS Certificates

The PCE supports LDAPS and LDAP with STARTTLS. To use the PCE with secure LDAP, add the certificate chain to the local certificate store on the PCE. Follow these steps to configure secure LDAP. It is assumed you have already followed the steps in [Enable LDAP Authentication](#) and [Configure LDAP Authentication](#).

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access Management > Authentication**.
3. On the Authentication Settings screen, locate the LDAP configuration panel and click **Configure**.
4. In the LDAP Authentication screen, make sure LDAP is enabled.
5. Select your LDAP server from the list of configured server entries and click the **Edit** icon.
6. Make sure **Protocol selected** is set to either LDAPS or LDAP with StartTLS.

7. For the Trusted CA bundle, click **Choose File** and upload the chain of certificate authority (CA) certificates for the LDAP server.
8. If your LDAP server uses self-signed certificates, uncheck the **Verify TLS** option.

NOTE:

The use of self-signed certificates for an LDAP server is not recommended. Illumio recommends the use of certificates signed by a valid CA.

Authentication Precedence

PCE local authentication takes precedence over any external systems. When the PCE authenticates a user, it follows this order:

1. The PCE attempts local authentication first. If the account is expired or otherwise fails, the PCE does not attempt to log in by using LDAP authentication.
2. If the local user does not exist, the PCE attempts LDAP login (if enabled).

How the PCE Works with Multiple LDAP Servers

You can configure up to three LDAP servers for each PCE. In a PCE supercluster deployment, the Illumio Core platform can support up to three LDAP servers per region.

When attempting to connect to an LDAP server, the PCE follows the order in which the servers were configured. When the request timeout expires, the PCE attempts to connect to the next server in the configuration. The PCE request timeout is configurable. By default, the timeout is 5 seconds.

For example, assume that you configure three LDAP servers in this order: A, B, C. The PCE attempts to connect to the servers in that order: A, B, C. If the PCE fails to connect to A, it attempts to connect to the remaining servers: first B, then C, after the expiration of the connection timeout.

When the PCE successfully connects to an LDAP server, it searches for the user on that server. If the user is found, the PCE stops looking. If the user is found on server A, even if the user also exists on B and C, the PCE will only use A's credentials for that user.

If the PCE successfully connects to an LDAP server but the user is not found, the PCE attempts to connect to the next server in the configured order, and searches for the user again.

You can not dynamically change the order in which the LDAP servers are contacted. To change this priority order, delete the configured entries and add them back in the desired order.

Active Directory Single Sign-on

This section describes how to configure Microsoft Active Directory Federation Services (AD FS) 3.0 for Single Sign-on (SSO) 2.0 authentication with the PCE.

Overview of AD FS SSO Configuration

To enable AD FS for the PCE, the PCE needs three fields returned as claims from:

- NameID
- Surname
- Given Name


There are two ways for AD FS to produce the NameID claim for an SSO user. The first uses the email field in an Active Directory user account for the NameID.

The second way to return a NameID of an Active Directory user is to use the User Principal Name (UPN). Each user created in Active Directory has an extension to their user name that's ADUserName@yourADDomainName. For example, a user named "test" in an Active Directory domain called "testing.com" would have a UPN of test@testing.com.


Configure AD Users to Use Different UPN Suffixes

To configure different UPN suffix as the source for NameID:


1. Add a UPN suffix. On your system under Server Manager Tools, click **Active Directory Domains and Trusts**.



- From the left side of the window, right-click *Active Directory Domains and Trusts*, and select **Properties**. In this dialog, you can create new suffixes for Active Directory usernames.




- Create a suffix that matches the external namespace you'll be using and click **Add**.



You can now assign an Active Directory user your custom UPN for the SAML response.

4. You can add multiple UPNs if needed. As shown below, you can select the UPN created in the previous steps.



Your UPN configuration is set up and you can begin configuring AD FS for SSO with the PCE.

Initial AD FS SSO Configuration

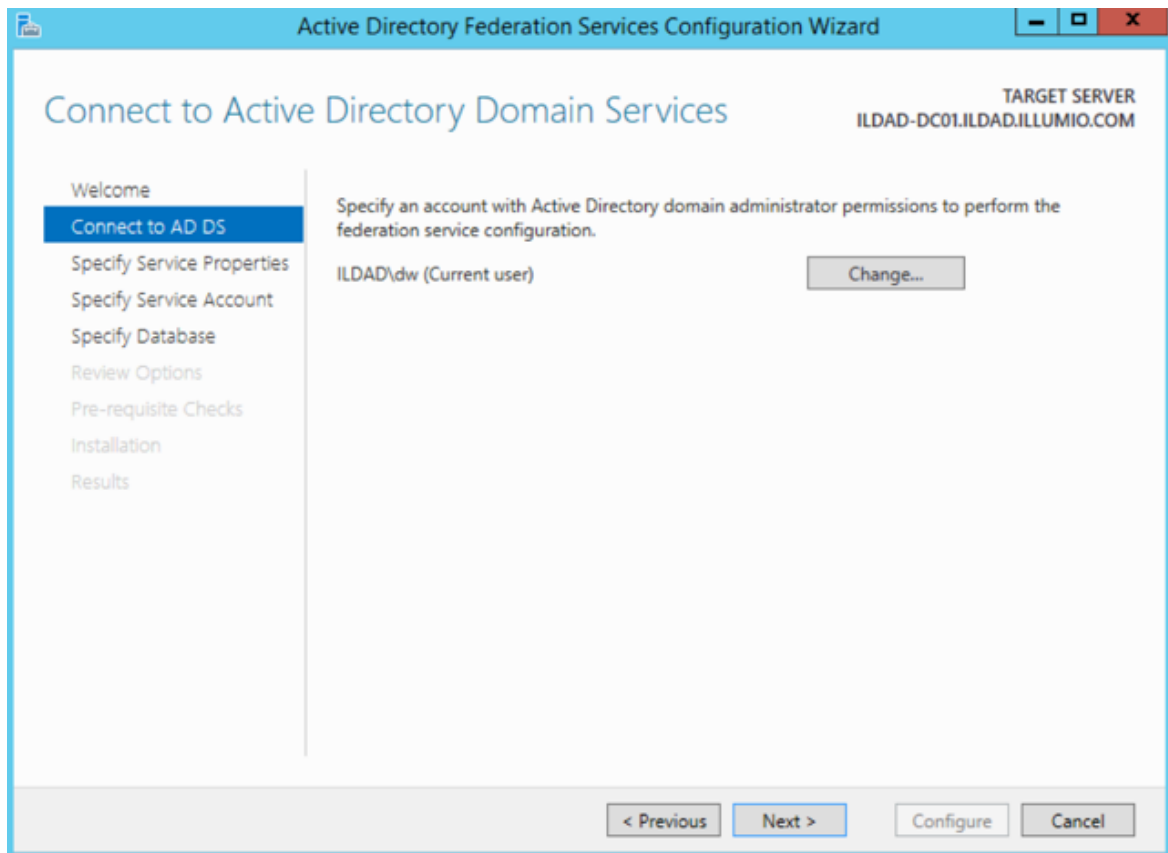
This task explains how to perform the initial configuration of AD FS to be your SSO IdP for Illumio Core.

To configure AD FS:


1. Open Microsoft Server Manager and click the notification icon.




2. Click the “Configure the federation service on this server” link.
3. Select “Create the first federation server in a federation server farm” option and click **Next**.
4. Specify a domain admin account for AD FS configuration.




5. Select or import a certificate. This certificate can be a self-signed certificate.




6. Specify your Federation Service Name, enter a display name for this instance of AD FS, and click **Next**




7. Specify your service account and click **Next**.




8. Select "Create a database on this server using Windows Internal Database" or choose the SQL server option, and click **Next**.



9. Review your selected options and click **Next**.



10. Click **Configure** to finish the basic configuration of AD FS.




11. In the results screen, click **Close**.

AD FS is now installed with the basic configuration on this host.


Create a Relying Party Trust

To start configuring AD FS for SSO with the PCE, you need to create a Relying Party Trust for your Illumio PCE.

1. From Server Manager/Tools, open the AD FS Manager.
2. From the left panel, choose **Relying Party Trusts > Add Relying Party Trust**.



The Add Relying Party Trust Wizard appears.



3. Click **Start**.
4. Select the “Enter data about the relying party manually” option and click **Next**.

The screenshot shows the 'Add Relying Party Trust Wizard' dialog box, specifically the 'Select Data Source' step. The window title is 'Add Relying Party Trust Wizard'. On the left, a 'Steps' pane lists the following steps: Welcome, Select Data Source (highlighted), Specify Display Name, Choose Profile, Configure Certificate, Configure URL, Configure Identifiers, Configure Multi-factor Authentication Now?, Choose Issuance Authorization Rules, Ready to Add Trust, and Finish. The main area contains the following text and options:

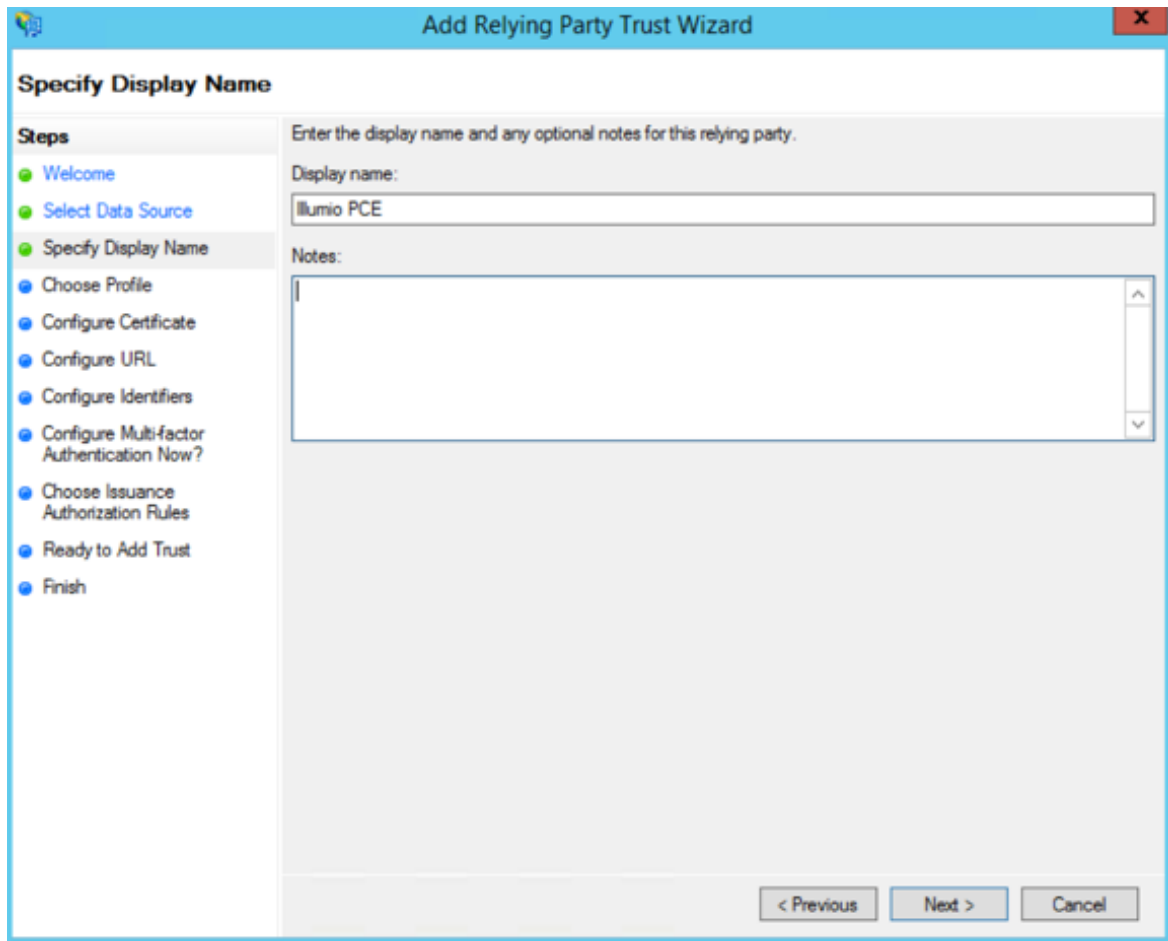
Select an option that this wizard will use to obtain data about this relying party:

- Import data about the relying party published online or on a local network
Use this option to import the necessary data and certificates from a relying party organization that publishes its federation metadata online or on a local network.
Federation metadata address (host name or URL):

Example: fs.cortoso.com or https://www.cortoso.com/app
- Import data about the relying party from a file
Use this option to import the necessary data and certificates from a relying party organization that has exported its federation metadata to a file. Ensure that this file is from a trusted source. This wizard will not validate the source of the file.
Federation metadata file location:
 - Enter data about the relying party manually
Use this option to manually input the necessary data about this relying party organization.


At the bottom right, there are three buttons: '< Previous', 'Next >', and 'Cancel'.

5. Name your Relying Party Trust and click **Next**.



The screenshot shows the 'Add Relying Party Trust Wizard' window. The title bar reads 'Add Relying Party Trust Wizard'. The main heading is 'Specify Display Name'. On the left, a 'Steps' pane lists the following steps: Welcome, Select Data Source, Specify Display Name (highlighted), Choose Profile, Configure Certificate, Configure URL, Configure Identifiers, Configure Multi-factor Authentication Now?, Choose Issuance Authorization Rules, Ready to Add Trust, and Finish. The main area contains the instruction 'Enter the display name and any optional notes for this relying party.' Below this, there is a 'Display name:' label and a text box containing 'illumio PCE'. Underneath is a 'Notes:' label and a large text area. At the bottom right, there are three buttons: '< Previous', 'Next >', and 'Cancel'.


6. Select "ADFS profile" and click **Next**.




7. When you have a separate certificate for token encryption, browse to, select it, and click **Next**.

NOTE:

To use the standard AD FS certificate (created during AD FS installation) for token signing, don't select anything in this step and click **Next**.



8. Select “Enable support for the SAML 2.0 WebSSO protocol.” In the *Relying party SAML 2.0 SSO service URL* field, add your “Assertion Consumer URL” (obtained from the PCE web console).



To locate the “Assertion Consumer URL,” go to **Settings > Authentication > Information for Identity Provider** in the PCE web console:


Information for Identity Provider	
Default User Role	Read Only
SAML Version	2.0
Issuer	https://pce-mnc.illumioeval.com:8443/login
NameID Format	urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress
Assertion Consumer URL	https://pce-mnc.illumioeval.com:8443/login/acs/2402fb18-3d75-4432-ab6d-10475897b476
Logout URL	https://pce-mnc.illumioeval.com:8443/login/logout/2402fb18-3d75-4432-ab6d-10475897b476

- On the Configure Identifiers page, use the same URL for the Relying party trust identifier, without the /acs/<randomNumbers>.


For example: https://pce.domain.com:8443/login.

Click **Next**.

10. Select the radio button “I do not want to configure multi-factor authentication settings for this relying party at this time” and click **Next**.
11. Select “Permit all users to access this relying party” and click **Next**.
12. On the Ready to Add Trust page, click **Next**.



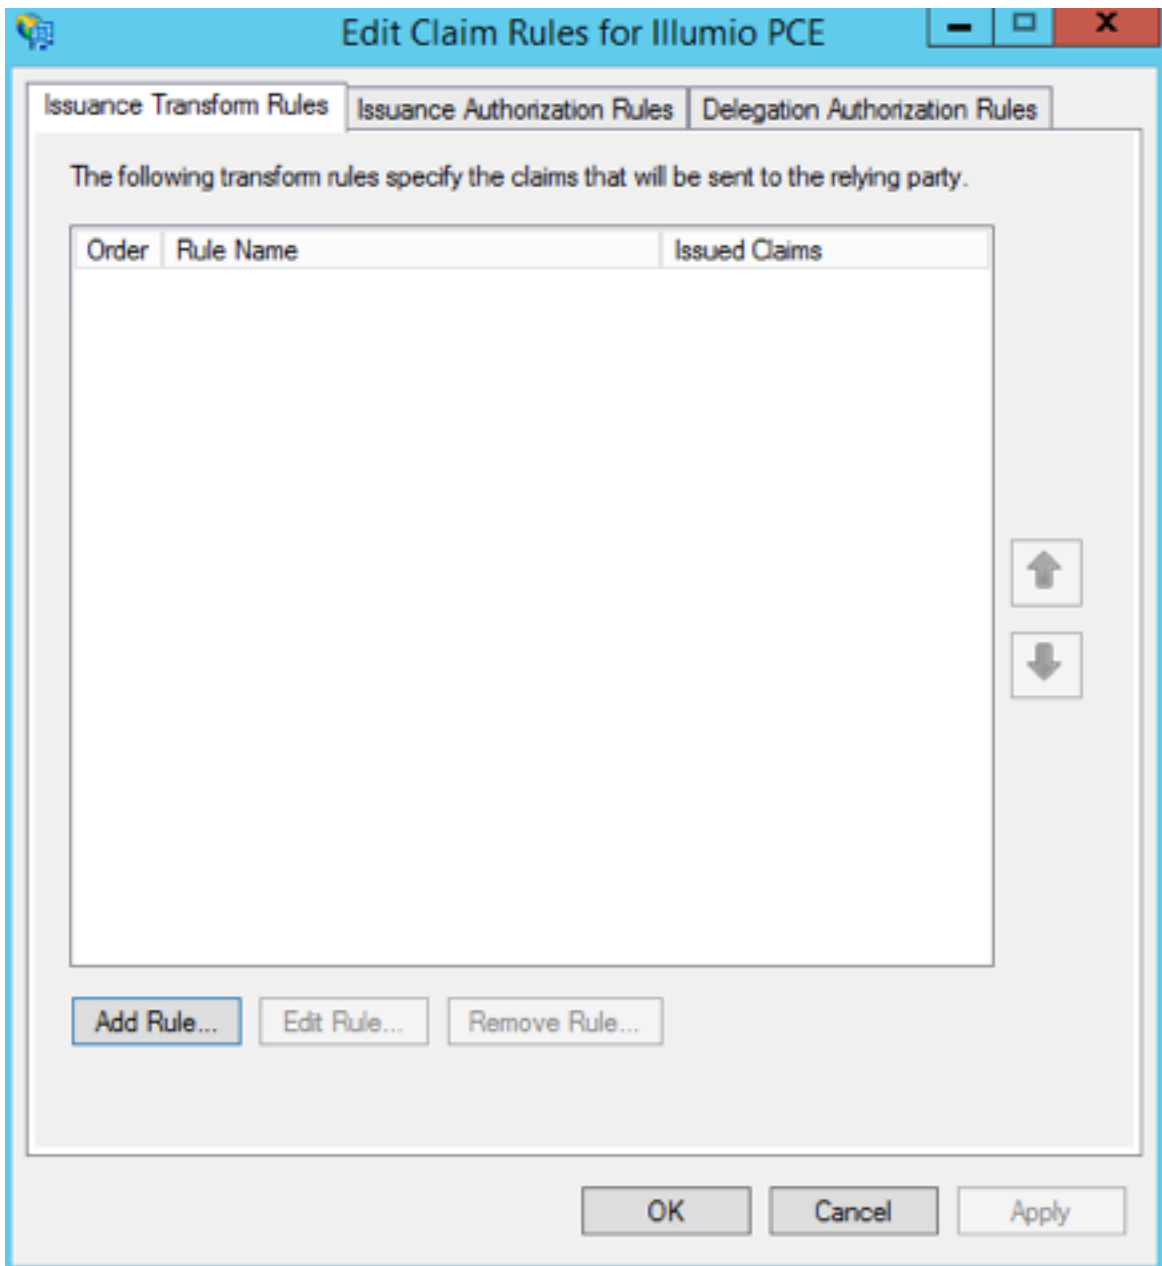
13. Leave the *Open the Edit Claim Rules* checkbox selected and click **Close**.




Create Claim Rules

You need to create claim rules to enable proper communication between AD FS and the PCE.


1. In the Edit Claim Rules dialog, click **Add Rule**.




2. Under Select Rule Template, select “Send LDAP Attributes as Claims” and click **Next**.



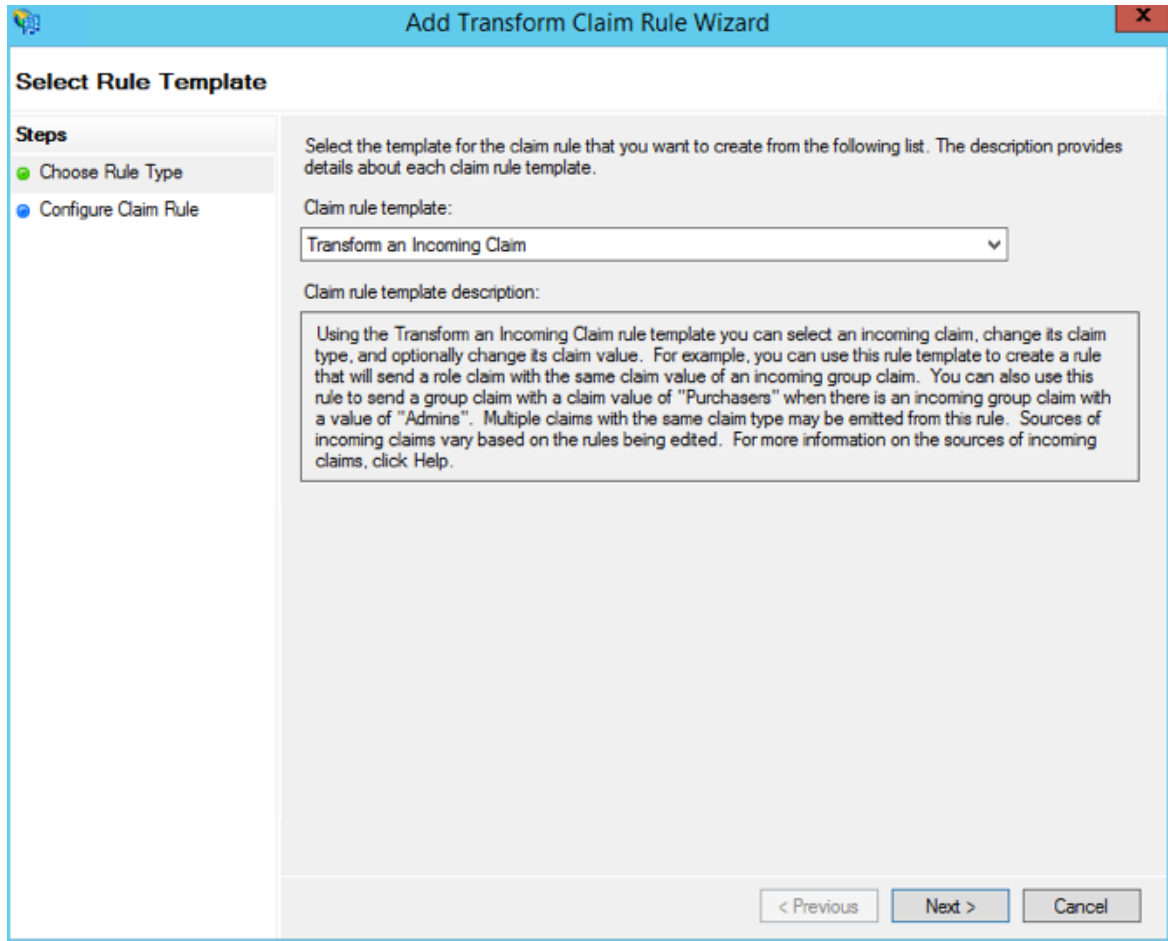
3. Name the Claim rule "Illumio Attributes" and select **Active Directory** as the Attribute store. Under the first attribute, select "User-Principal-Name" and "E-Mail Address" as the outgoing. Select "Surname" and type the custom field name of "User.LastName" in the outgoing field. Repeat the values for "Given-Name" and "User.FirstName" and click **Finish**.



- In the Edit Claim Rules dialog with your new rule added, click **Add Rule** to add the final rule.



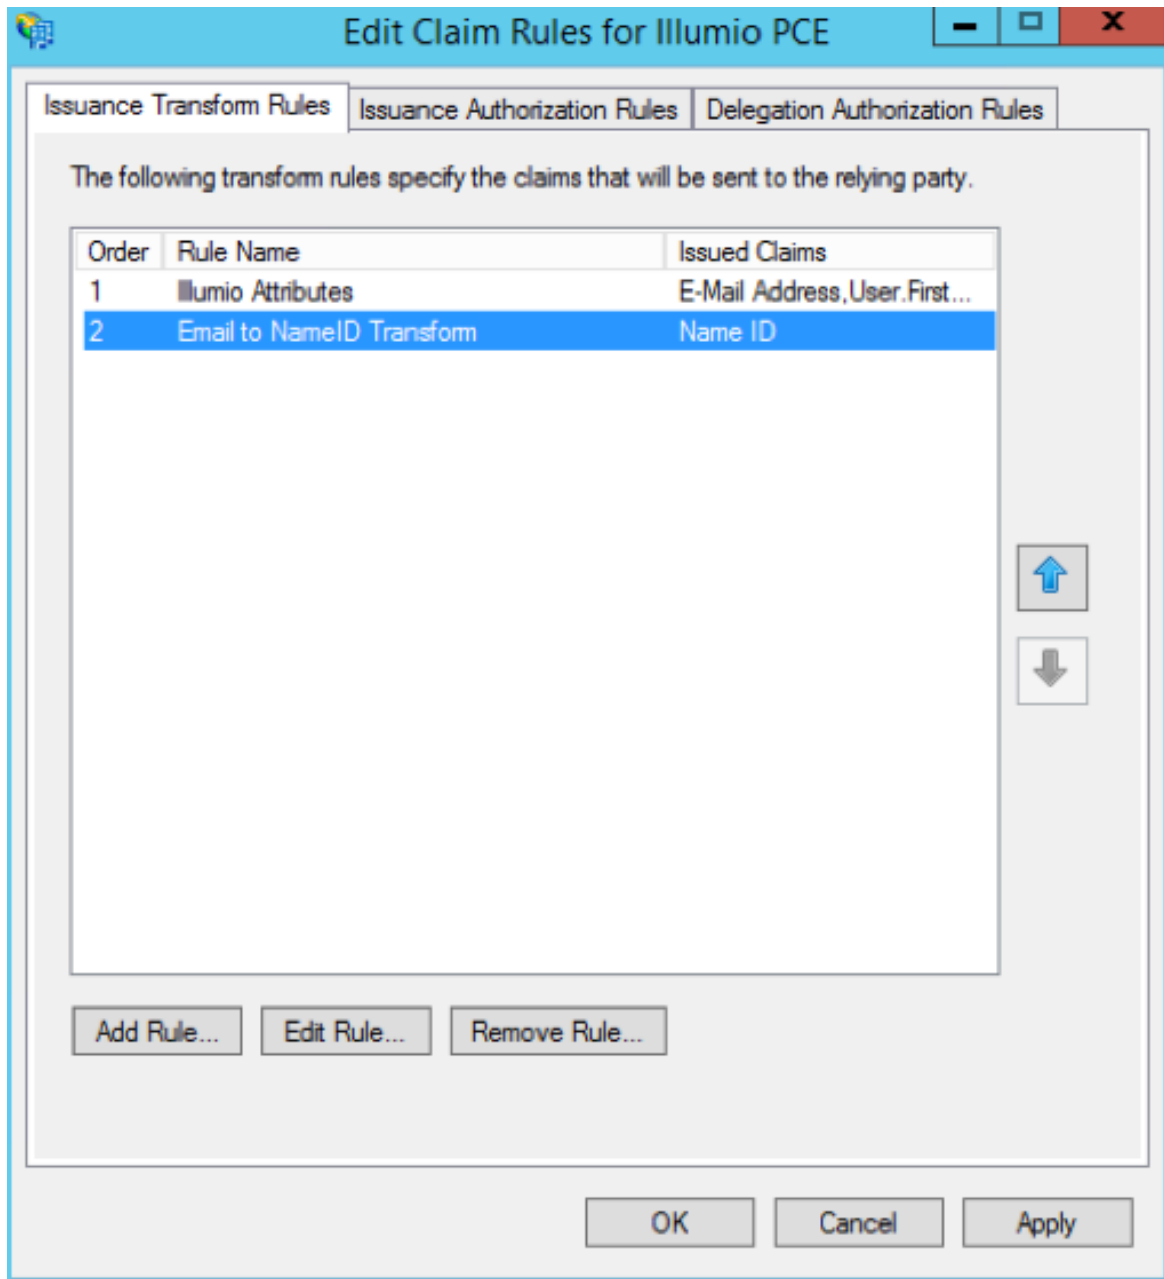
5. Under the Claim Rule Template, select “Transform and Incoming Claim” and click Next.



6. Name the rule "Email to NameID Transform" and change the incoming claim type to "E-Mail Address." Set the Outgoing claim type to "Name ID" and the Outgoing name ID format to "Email" and click **Finish**.

The screenshot shows a window titled "Add Transform Claim Rule Wizard" with a "Configure Rule" tab. On the left, a "Steps" pane shows "Choose Rule Type" and "Configure Claim Rule". The main area contains a text box for "Claim rule name" with the value "Email to NameID Transform". Below it, "Rule template" is set to "Transform an Incoming Claim". There are four dropdown menus: "Incoming claim type" (E-Mail Address), "Incoming name ID format" (Unspecified), "Outgoing claim type" (Name ID), and "Outgoing name ID format" (Email). Three radio buttons are present: "Pass through all claim values" (selected), "Replace an incoming claim value with a different outgoing claim value", and "Replace incoming e-mail suffix claims with a new e-mail suffix". The second option has input fields for "Incoming claim value" and "Outgoing claim value" with a "Browse..." button. The third option has a "New e-mail suffix" field with the example "fabrikam.com". At the bottom are buttons for "< Previous", "Finish", and "Cancel".

The Edit Claim Rules window opens.

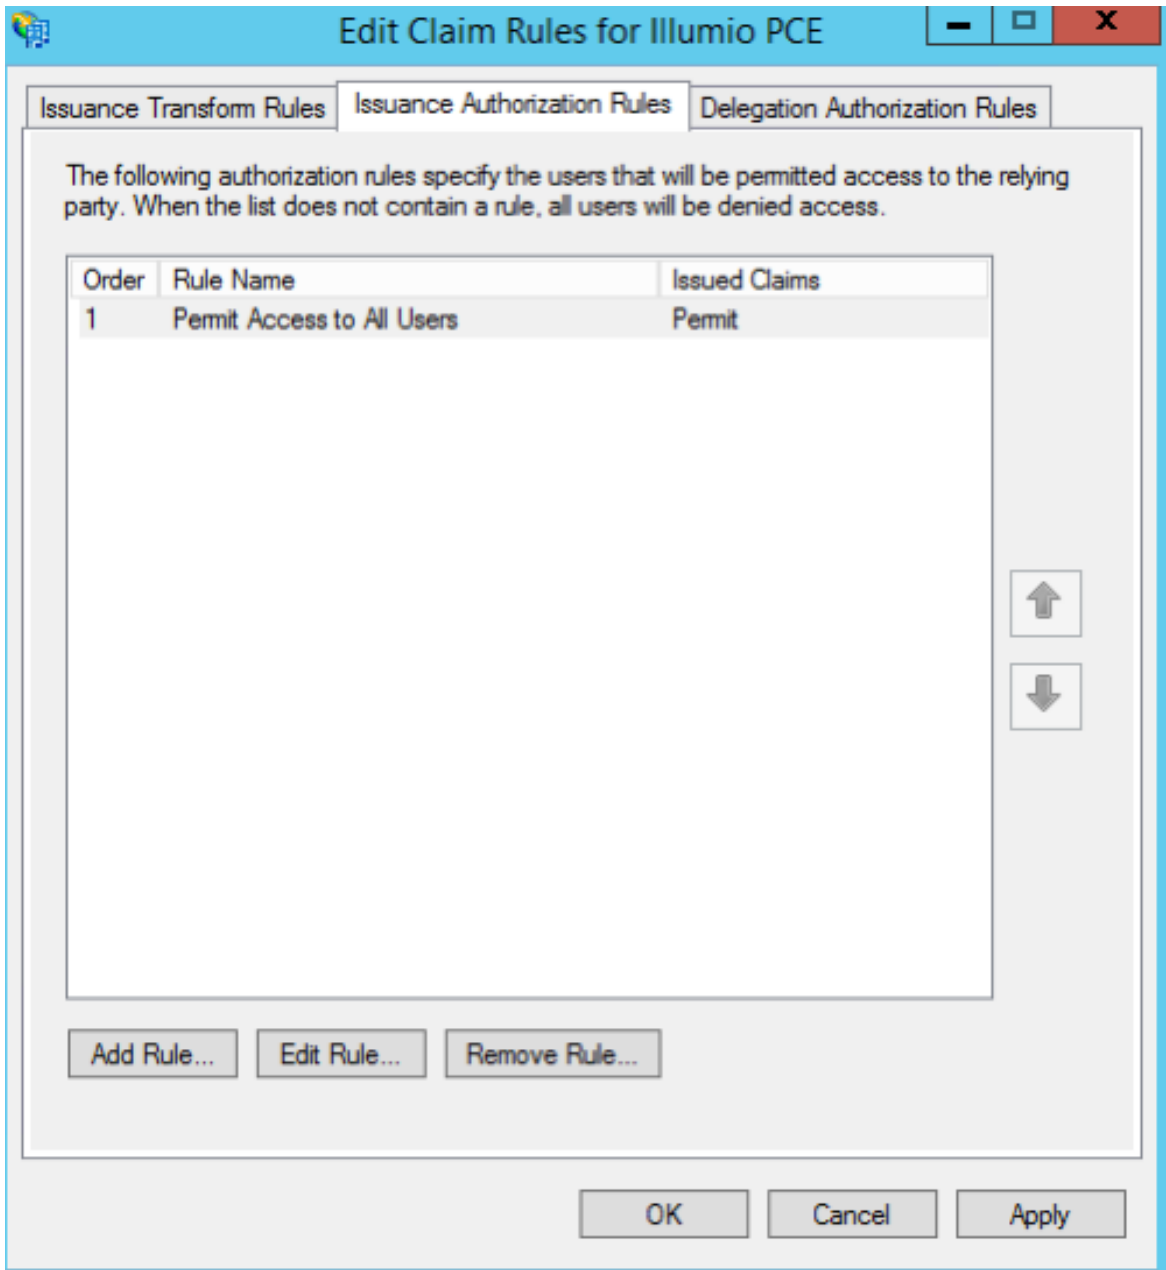


7. (Windows 2016 and Windows 2019) Skip to step 12.


The Edit Claim Rules window has three tabs. You have already filled out the first tab. The other two tabs are not available in Windows 2016 or Windows 2019. Therefore, skip steps 8 - 11.

8. Select the Issuance Authorization Rules tab.
9. To allow all your Active Directory Users to access the PCE, leave the “Permit Access to All Users” as is. Otherwise, you should restrict access to a single group

or groups of users.



10. Select “Permit or Deny Users Based on an Incoming Claim” and click **Next**.



11. Name the rule "AD FS Users" and change the Incoming claim type to "Group SID" (you might have to scroll to find it). In Incoming claim value, browse to the group of users you want to give access. Make sure "Permit access" is selected and click **Finish**.

The screenshot shows a Windows-style dialog box titled "Add Issuance Authorization Claim Rule Wizard" with a close button (X) in the top right corner. The main area is titled "Configure Rule". On the left, a "Steps" pane shows two steps: "Choose Rule Type" (completed) and "Configure Claim Rule" (current step). The main content area contains the following fields and options:

- Instructional text: "You can configure this rule to permit or deny users based on an incoming claim. Specify the incoming claim type, claim value, and whether the users should be permitted or denied access to the relying party."
- Text field: "Claim rule name:" with the value "AD FS Users".
- Text label: "Rule template: Authorize Users Based on an Incoming Claim".
- Dropdown menu: "Incoming claim type:" with "Group SID" selected.
- Text field: "Incoming claim value:" with "ILDAD\ADFS Users" and a "Browse..." button to its right.
- Radio button options: "Select one of the following options to indicate whether users with this claim will be permitted or denied access to the relying party."
 - Permit access to users with this incoming claim
 - Deny access to users with this incoming claim
- Navigation buttons at the bottom: "< Previous", "Finish", and "Cancel".

12. If you are using RBAC with groups, you need to create a Group Claim Rule.

To add groups to AD FS claim rule configuration, click **Edit Rule**. Add the requirement for "LDAP Attribute: memberOf" by selecting the Outgoing Claim Type as "User.MemberOf." Click **OK**.

Edit Rule - Groups
X

You can configure this rule to send the values of LDAP attributes as claims. Select an attribute store from which to extract LDAP attributes. Specify how the attributes will map to the outgoing claim types that will be issued from the rule.

Claim rule name:

Rule template: Send LDAP Attributes as Claims

Attribute store:

Mapping of LDAP attributes to outgoing claim types:

	LDAP Attribute (Select or type to add more)	Outgoing Claim Type (Select or type to add more)
▶	Token-Groups - Unqualified Names ▼	User.MemberOf ▼
*	▼	▼


Obtain ADFS SSO Information for the PCE

Before you can configure the PCE to use AD FS for SSO, obtain the following information from your AD FS configuration:

- x.509 certificate supplied by ADFS
- Remote Login URL
- Logout Landing URL

To obtain the AD FS SSO information for the PCE:

1. To find the certificate in your AD FS configuration, log into the AD FS server and open the management console.
2. Browse to the certificates and export the Token-Signing certificate.
3. Right-click the certificate and select **View Certificate**.
4. Select the **Details** tab.
5. Click **Copy to File**.



6. When the Certificate Export Wizard launches, click **Next**.
7. Verify that the “No - do not export the private key” option is selected and click **Next**.
8. Select Base 64 encoded binary X.509 (.cer) and click **Next**.
9. Select where you want to save the file, name the file, and click **Next**.
10. Click **Finish**.
11. After exporting the certificate to a file, open the file with a text editor. Copy and paste the contents of the exported x.509 certificate, including the BEGIN CERTIFICATE and END CERTIFICATE delimiters in to the SAML Identity Provider Certificate field.

- To find the **Remote Login URL** (which AD FS calls “Sign-On URL”), download and open the following metadata file from your AD FS server by navigating to `https://server.mydomain/FederationMetadata/2007-06/FederationMetadata.xml` and search for `SingleSignOnService`.

```
format:persistent</NameIDFormat><NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid
-format:transient</NameIDFormat><SingleSignOnService

Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
Location="https://.illumio.com/adfs/ls/"><SingleSignOnService
Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
Location="https://anoop.illumioeval.com/adfs/ls/"><Attribute
Name="http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress"
```

- To find the **Logout Landing URL** for the PCE, you can use the login URL of the PCE (preferred):

```
https://<myPCENAMEAndPort>/login
```

Or, a generic logout URL of AD FS:

```
https://<URLToMyADFSserver>/adfs/ls/?wa=wsignout1.0
```

You are now ready to configure the PCE to use AD FS for SSO.

Configure the PCE for AD FS SSO

Before you configure the PCE to use Microsoft AD FS for SSO, make sure you have the following information provided by your AD FS, which you configure in the PCE web console:

- x.509 certificate supplied by ADFS
- Remote Login URL
- Logout Landing URL

For more information, see [Obtain ADFS SSO Information for the PCE](#).

NOTE:

When SSO is configured in Illumio Core and for the IdP, the preferences in Illumio Core are used. When SSO is not configured in Illumio Core, the default IdP settings are used.

To configure the PCE for AD FS:

1. From the PCE web console menu, choose **Settings >SSO Config**.
2. Click **Edit**.
3. Select the *Enabled* checkbox next to SAML Status.
4. In the *Information From Identity Provider* section, enter the following information:
 - SAML Identity Provider Certificate
 - Remote Login URL
 - Logout Landing URL
5. Select the authentication method from the drop-down list:
 - **Unspecified:** Uses the IdP default authentication mechanism.
 - **Password Protected Transport:** Requires the user to log in with a password using a protected session; select this option and check the Force Re-authorization checkbox to force user re-authorization.
6. To require users to re-enter their login information to access Illumio (even if the session is still valid), check the Force Re-authentication checkbox. This allows users to log into the PCE using a different login than their default computer login and is disabled by default.

NOTE:

You must select "Password Protected Transport" as the authentication method and check the Force Re-authentication checkbox to force users to re-authenticate.

7. Click **Save**.

Your PCE is now configured to use AD FS for SSO authentication.

Azure AD Single Sign-on

This topic describes how to configure Azure Active Directory (AD) to provide SSO authentication to the Illumio PCE.

TIP:

Because you'll configure settings in both the Illumio PCE Web Console and in Azure AD, have both applications open in adjacent browser tabs.

Prerequisites

To perform this configuration, you need the following:

- An Azure AD subscription. If you don't have a subscription, you can get a [free account](#).
- An Illumio single sign-on (SSO) enabled subscription.

STEP 1: Obtain URLs from the Illumio PCE Web Console

In this step you'll copy and preserve URLs from the Illumio PCE for use in [STEP 2: Configure SSO settings in Azure AD](#).

1. Log in to the PCE as a Global Organization Owner.
2. Go to **Access Management > Authentication**.
3. On the **SAML** tile, click **Configure**.
4. Copy and preserve the following URLs needed to complete the Azure configuration in a later step:

TIP:

Make sure to replace the x's in the URLs below with the actual values from your implementation.


- **Issuer:** https://PCE.xxxx:8443/login
- **NameID Format:** urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress
- **Assertion Source URL:** https://PCE.xxxx:8443/login/acs/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx
- **Logout URL:** https://pce.xxxx:8443/login/logout/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx

STEP 2: Configure SSO settings in Azure AD


NOTE:

Only an Azure Application Administrator can configure Azure AD.

1. In a different browser tab, log in to Azure AD as an Application Administrator.
2. Go to **Enterprise applications > All applications**.
3. Search for the **Illumio SSO** app and then click the app.



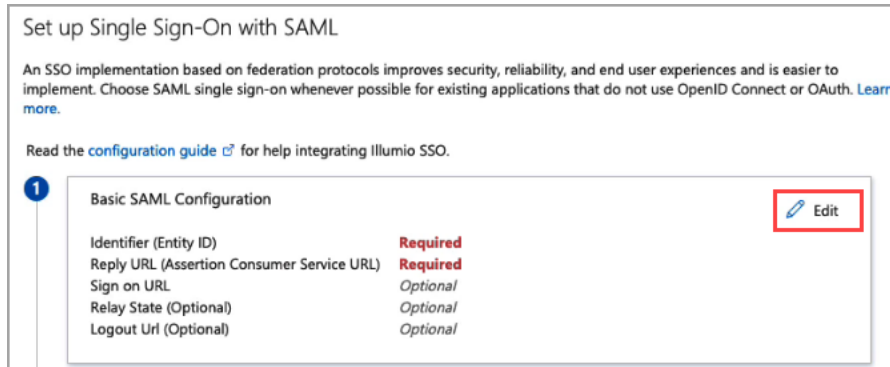
4. In the center of the page under **Getting Started**, click **Get started** on the **Set up single sign on** tile.



5. If prompted to select a single sign-on method, click **SAML**.


6. Configure Basic SAML:

- a. On the **Set up Single Sign-On with SAML** page **Basic SAML Configuration** tile, click **Edit**.




- b. On the **Basic SAML Configuration** panel that opens, populate the following fields with the values you copied and preserved in [STEP 1: Obtain URLs from the Illumio PCE Web Console](#):
 - In the **Identifier (Entity ID)** field, paste the **Issuer URL** you copied from the Illumio PCE.
 - In the **Reply URL (Assertion Consumer Service URL)** field, click **Add reply URL** and then paste the **Assertion Source URL** you copied from the Illumio PCE. **Note:** Your Reply URL must have a subdomain such as `www`, `wd2`, `wd3`, `wd3-impl`, `wd5`, `wd5-impl`. For example, `http://www.myIllumio.com` will work but `http://myIllumio.com` won't.
- c. Click **Save** and close the **Basic SAML Configuration** panel.


7. Click **Edit** on the **Attributes & Claims** tile.




8. Under **Required claim**, update the **Claim name**:



- a. Click the three dots.
 - b. On the **Manage claim** page, click in the **Source attribute** field and select **user.mail** from the dropdown.
 - c. Click **Save**.
9. Back on the **Attributes & Claims** page, delete **all** of the existing claims in the **Additional claims** section by clicking the three dots for each one and then clicking **Delete**.



10. Click **Add new claim** and add three new claims:



Given Name

- **Name:** Enter **Given Name**.
- **Source attribute:** Enter **user.givenname**

Surname

- **Name:** Enter **Surname**
- **Source attribute:** Enter **user.surname**

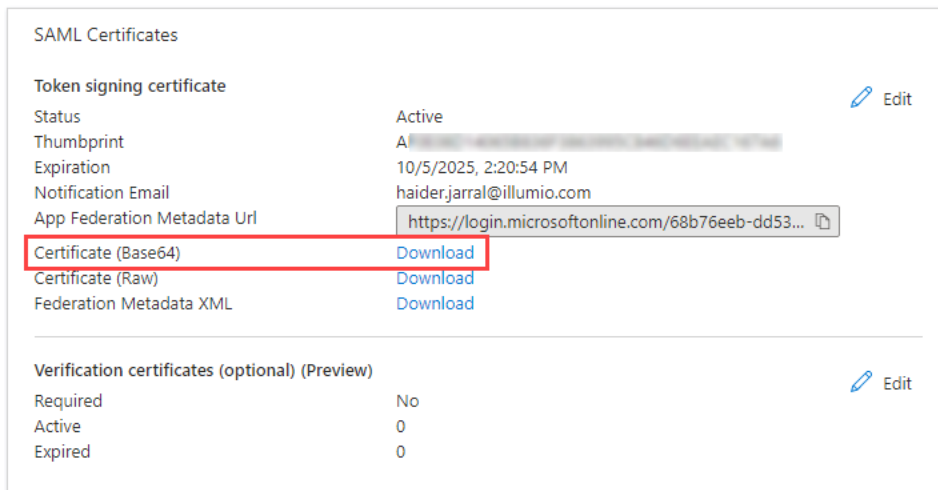
User.MemberOf

- **Name:** Enter **User.MemberOf**
- **Source attribute:** Enter **user.assignedroles**

STEP 3: Obtain SAML certificate and URLs from Azure AD

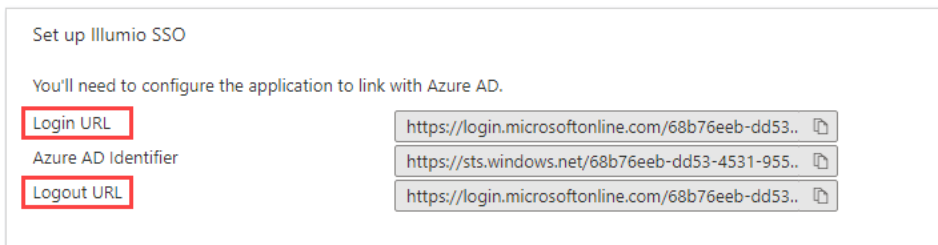
In this step you'll download a certificate and copy two URLs that you'll later paste into the Illumio PCE SAML setup in [STEP 4: Configure SAML SSO settings in the Illumio PCE](#).

1. On the **SAML Certificates** tile, click **Download** for the **Certificate (Base64)** certificate and save the certificate to your computer.



SAML Certificates		Edit
Token signing certificate		
Status	Active	
Thumbprint	A[REDACTED]	
Expiration	10/5/2025, 2:20:54 PM	
Notification Email	haider.jarral@illumio.com	
App Federation Metadata Url	https://login.microsoftonline.com/68b76eeb-dd53...	
Certificate (Base64)	Download	
Certificate (Raw)	Download	
Federation Metadata XML	Download	
Verification certificates (optional) (Preview)		
Required	No	
Active	0	
Expired	0	

2. On the **Set up Illumio SSO** tile, copy and preserve the following URLs that you'll later paste into the Illumio PCE SAML setup in [STEP 4: Configure SAML SSO settings in the Illumio PCE](#).



Set up Illumio SSO	
You'll need to configure the application to link with Azure AD.	
Login URL	https://login.microsoftonline.com/68b76eeb-dd53..
Azure AD Identifier	https://sts.windows.net/68b76eeb-dd53-4531-955..
Logout URL	https://login.microsoftonline.com/68b76eeb-dd53..

- **Login URL.** You'll paste this in the **Remote Login URL** field in the PCE Web Console.
- **Logout URL.** You'll paste this in the **Logout Landing URL** field in the PCE Web Console.

STEP 4: Configure SAML SSO settings in the Illumio PCE

In this procedure you'll paste the following information that you copied and preserved from Azure in [STEP 3: Obtain SAML certificate and URLs from Azure AD](#):

- Certificate (Base64)
 - Azure Login URL
 - Logout URL
1. In the Illumio PCE Web Console, go to **Access Management > Authentication**.
 2. On the **SAML** tile, click **Configure**.
 3. Click **Edit**.
 4. In the **Information from Identity Destination** section, enter the following information that you obtained from Azure AD:
 - **SAML Identity Destination Certificate:** Open the certificate that you downloaded in [STEP 3: Obtain SAML certificate and URLs from Azure AD](#), and then copy and paste the contents.
 - **Remote Login URL:** Paste the Login URL you copied from Azure AD.
 - **Logout Landing URL:** Paste the Logout URL you copied from Azure AD.
 5. In the **Information for Identity Destination** section:
 - a. Choose an authentication method:
 - **Unspecified** uses the IdP default authentication mechanism.
 - **Password Protected Transport** requires the user to log in with a password in a protected session.
 - b. If you want to require users to re-enter login credentials to access Illumio (even if the session is still valid), select **Force Re-authentication**. This allows users to log in to the PCE using login credentials different from their default computer login credentials.
 6. Click **Save**.


STEP 5: Create App Roles in Azure AD

In this step you'll create app roles in Azure AD that you'll map to roles in the Illumio PCE Web Console in [STEP 7: Add External Groups and assign roles in the PCE Web Console](#).

For reference in this step, here's a list of the Global Roles available in the PCE Web Console:

- Global Organization Owner
- Global Administrator
- Global Viewer
- Globally Policy Object Provisioner

1. In Azure AD, go to **Users and Groups** and then click **application registration**.



2. Create the roles you want by clicking **+ Create app role** and entering the required information for each role:
 - **Display name:** For example, enter one of the Global Roles that appear in the PCE Web Console.
 - **Value:** This must match the name you'll enter in the **Add External Groups** dialog box in [STEP 7: Add External Groups and assign roles in the PCE Web Console](#).
 - **Description:** The description will appear as help text in the app assignment and consent experiences.
3. Click **Apply** for each role that you create.
4. Delete the default app role **msiam_access**.

Note: You first need to disable the default app role before you can delete it.

- a. Click **msiam_access** to open the **Edit app role** panel.
- b. Deselect **Do you want to enable the app role?**

- c. Click **Apply**. The side panel closes.
- d. Click `msiam_access` again to to open the **Edit app role** panel again.
- e. Click **Delete**.

When you're done creating roles in Azure AD, the **App roles** section should look similar to this:

Display name	Description	Allowed member types	Value	ID
Global Organization O...	Global Organization Owner	Users/Groups	GOO	309c156d
Global Administrator	Global Administrator	Users/Groups	GA	f6473e65-
Global Viewer	Global Viewer	Users/Groups	GV	cb677852
Global Policy Object P...	Global Policy Object Provisioner	Users/Groups	GPOP	d07b17b1

STEP 6: Assign users and groups to app roles in Azure AD

In this step, you'll assign users and groups to the app roles you created in [STEP 5: Create App Roles in Azure AD](#).

1. In Azure AD, go to **Users and groups**.
2. Select the Illumio SSO app.
3. Click **Remove** to remove the current app assignments.
4. Click **Yes** to confirm removal.
5. Click **Add user/group**.
6. On the **Add Assignment** page, assign desired role(s) to users or groups:
 - a. Under **User and groups**, click **None Selected**.
 - b. In the **Users and groups** panel that opens, search for your desired user-/group, click to select it, and then click **Select** at the bottom of the panel.
 - c. Back on the **Add Assignment** page, under **Select a role***, click **None Selected**.
 - d. In the **Select a role** panel that opens, find and click the role you want to assign, and then click **Select** at the bottom of the panel.
 - e. Back on the **Add Assignment** page, click **Assign** at the bottom of the page.

- f. Repeat these sub-steps for each user and/or to which you want to assign app roles.

STEP 7: Add External Groups and assign roles in the PCE Web Console

In this step, you'll add external groups in the PCE Web Console and assign them the relevant global or scoped roles in Illumio RBAC.

TIP:

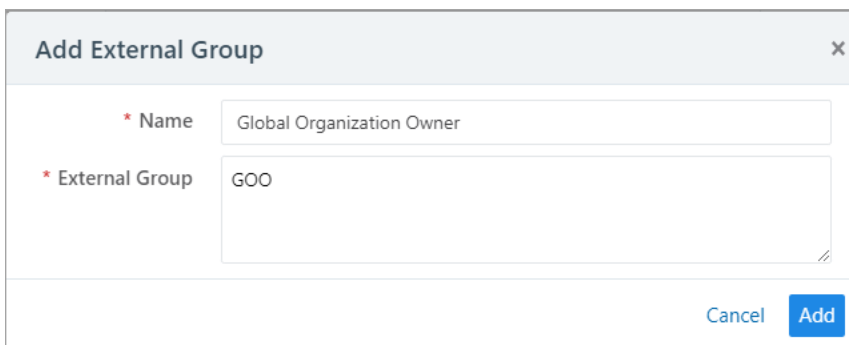
Alternatively, you can add individual users by going to the **External Users** tab and following the onscreen prompts.

1. On the PCE Web Console, go to **Access Management > External Groups**.
2. Click **Add**.
3. In the **Add External Group** dialog box:
 - Enter a **Name**.
 - Enter an **External Group**.


IMPORTANT:

This must match the **Value** that you specified for the app role in [STEP 5: Create App Roles in Azure AD](#)


- Click **Add**.
4. Repeat for additional groups.



The screenshot shows a dialog box titled "Add External Group" with a close button (X) in the top right corner. The dialog contains two input fields, both marked with an asterisk (*). The first field is labeled "Name" and contains the text "Global Organization Owner". The second field is labeled "External Group" and contains the text "GOO". At the bottom right of the dialog, there are two buttons: "Cancel" and "Add".




5. Click to open a group you created in the above step.
6. Click **Add Role > Add Global Role** or **Add Scoped Role**.
7. In the **Access Wizard**, select the appropriate **Role** and then click **Grant Access**.
8. Repeat for additional groups.



STEP 8: Turn on SAML authentication in the PCE Web Console

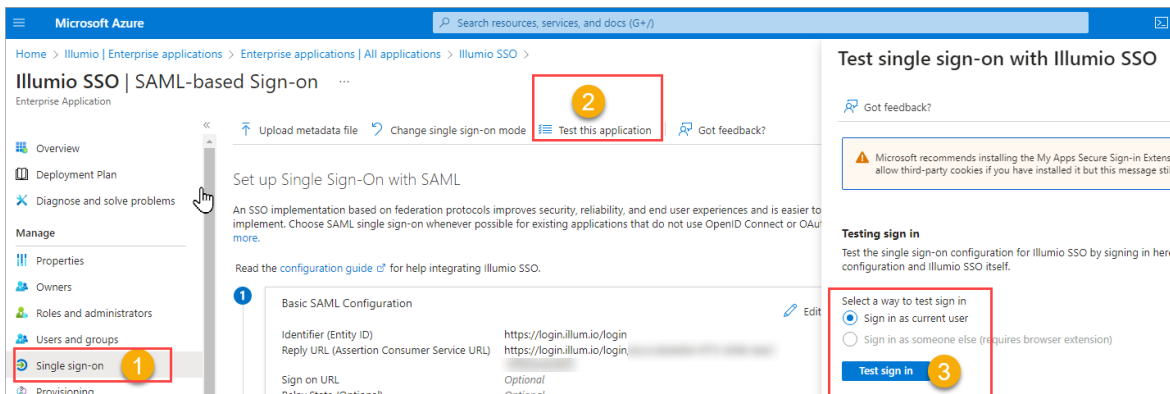
1. In the PCE Web Console, go to **Access Management > Authentication**.
2. On the SAML tile, click **Configure**.
3. On the SAML page, click **Turn On** and then click **Confirm**.



STEP 9: Test SSO

Perform this procedure to test the SSO authentication you configured in the previous steps.

1. In Azure AD, go to **Single sign-on**.
2. Click **Test this application**.
3. In the panel that opens, select a way to sign in and then click **Test sign in**.



4. If the test is successful, the PCE will log you in to the **Welcome to Illumio** screen.

Okta Single Sign-on

This section explains how to configure SSO for user authentication with the PCE using Okta as your IdP.

Prerequisite for Okta SSO

Before you begin, make sure you have the following information from your Okta account:

- x.509 certificate
- Remote Login URL
- Logout Landing URL

NOTE:

Your PCE user account must have Owner or Admin privileges to perform this task.

Configure the PCE for Okta SSO

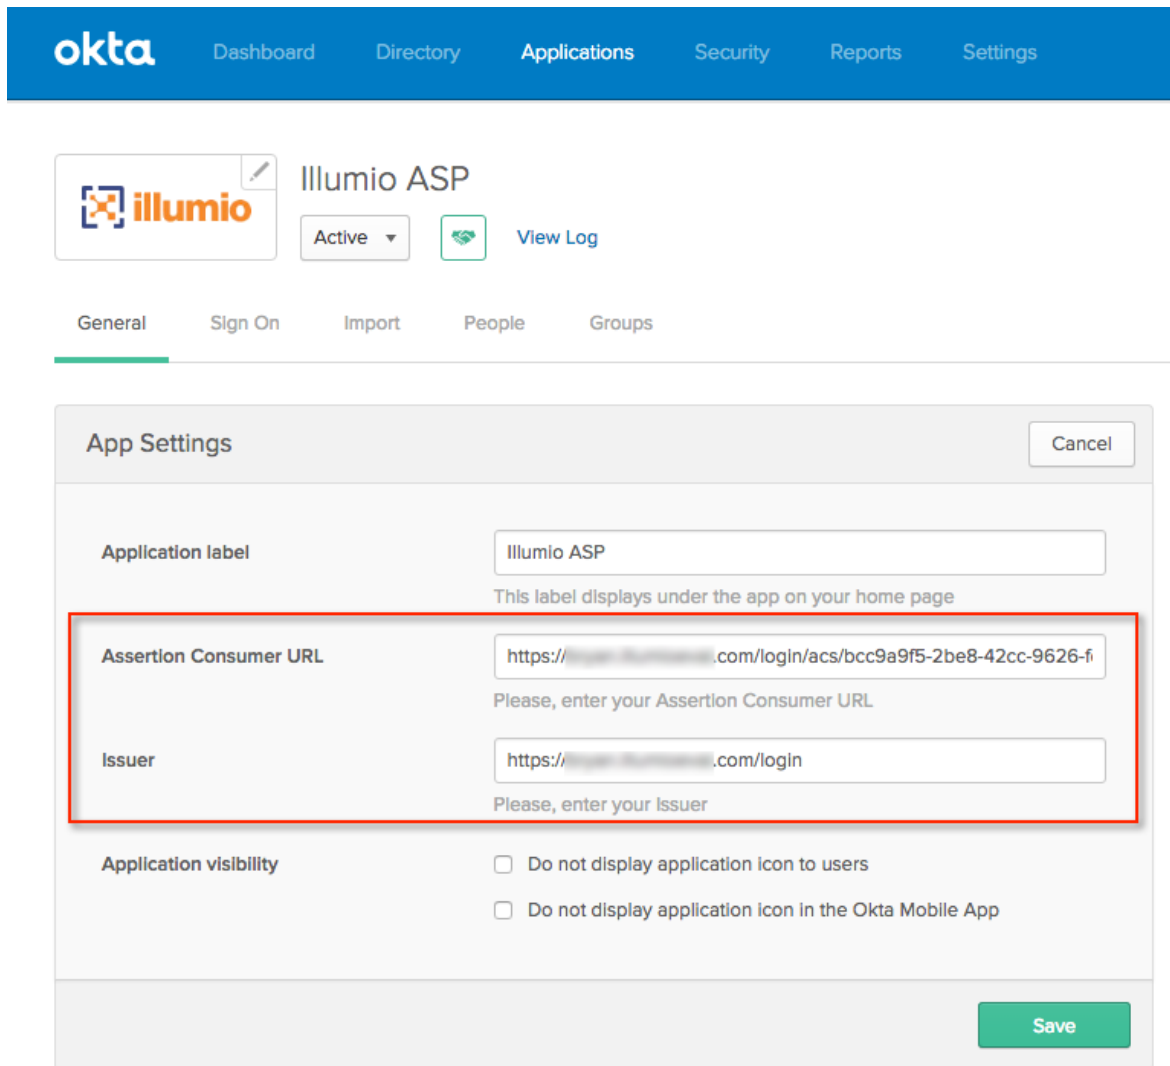
1. From the PCE web console menu, choose **Access Management** > **Authentication**.
2. On the Authentication Settings screen, locate the SAML configuration panel and click **Configure**.
3. Enter the following information:
 - **SAML Identity Provider Certificate:** Paste your Okta x.509 certificate (in PEM text format):
 - **Remote Login URL:** Enter the Okta Remote Login URL.
 - **Logout Landing URL:** Enter the Okta Logout Landing URL.
4. In the Information for Identity Provider section, choose the Access Level for the users who will use Okta to authenticate with the PCE. When you select No Access, SSO users from your Okta account will have to be added manually before they can log into the PCE. (For more information on PCE user permissions, see [Role-based Access Control](#).)
5. In the Information for Identity Provider section, make note of the following fields:
 - Issuer
 - Assertion Consumer URL
6. Select the authentication method from the drop-down list:

- **Unspecified:** Uses the IdP default authentication mechanism.
 - **Password Protected Transport:** Requires the user to log in with a password using a protected session.
7. To require users to re-enter their login information to access Illumio (even if the session is still valid), check the Force Re-authentication checkbox. This allows users to log into the PCE using a different login than their default computer login and is disabled by default.

NOTE:

When SSO is configured both in Illumio Core and for the IdP, the preferences in Illumio Core are used. When SSO is not configured in Illumio Core, the default IdP settings are used.

8. Click **Save**.
9. Log into your Okta account.
10. Select the Illumio Core app, select the General tab, and click **Edit**.
11. Enter the values you copied from the Information for Identity Provider section of the PCE SSO Configuration page.



The screenshot shows the Okta Admin Console interface. At the top, there is a navigation bar with the Okta logo and menu items: Dashboard, Directory, Applications, Security, Reports, and Settings. Below this, the 'illumio ASP' application is shown with a status of 'Active' and a 'View Log' button. A tabbed interface below the application name includes 'General', 'Sign On', 'Import', 'People', and 'Groups'. The 'App Settings' modal is open, displaying the following configuration options:

- Application label:** illumio ASP (with a note: "This label displays under the app on your home page")
- Assertion Consumer URL:** https://[redacted].com/login/acs/bcc9a9f5-2be8-42cc-9626-f (with a note: "Please, enter your Assertion Consumer URL")
- Issuer:** https://[redacted].com/login (with a note: "Please, enter your Issuer")
- Application visibility:**
 - Do not display application icon to users
 - Do not display application icon in the Okta Mobile App

Buttons for 'Cancel' and 'Save' are visible in the modal.

12. Click **Save**.

Your PCE is now configured to use Okta SSO for authenticating users with the PCE.

OneLogin Single Sign-on

This section describes how to configure SSO for OneLogin.

Configure SSO for OneLogin

This task shows you how to configure SSO for authenticating users with the PCE using OneLogin as your Identity Provider (IdP).

Before you begin, make sure you have the following information from your OneLogin account:

- x.509 certificate
- SAML 2.0 Endpoint (HTTP)
- SLO Endpoint (HTTP)

NOTE:

Your PCE user account must have Owner or Admin privileges to perform this task

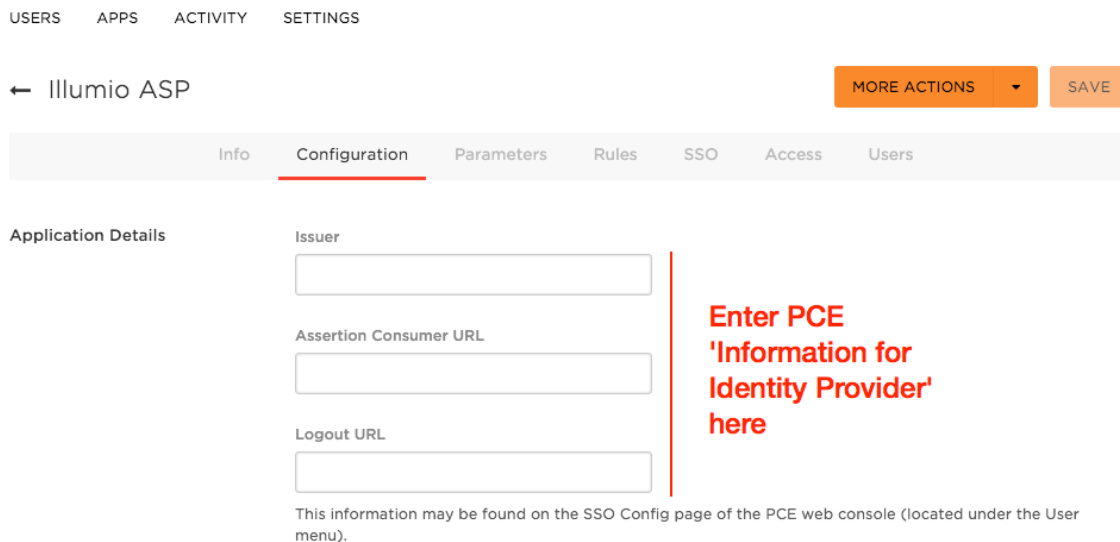
To configure the PCE for OneLogin SSO:

1. From the PCE web console menu, choose **Settings > SSO Config**.
2. Click **Edit**.
3. Select the Enabled checkbox for SAML Status.
4. Enter the following information:
 - **SAML Identity Provider Certificate:** Paste your OneLogin x.509 certificate (in PEM text format).
 - **Remote Login URL:** Enter the OneLogin SAML 2.0 Endpoint (HTTP) URL.
 - **Logout Landing URL:** Enter the OneLogin SLO Endpoint (HTTP) URL.
5. In the Information for Identity Provider section, choose the Access Level for the users who use OneLogin to authenticate with the PCE. When you select No Access, SSO users from your OneLogin account will have to be added manually before they can log in to the PCE. (For more information on PCE user permissions, see [Role-based Access Control](#).)
6. In the Information for Identity Provider section, make note of the following fields:
 - Issuer
 - Assertion Consumer URL
 - Logout URLYou will enter this information into your OneLogin SSO configuration.
7. Select the authentication method from the drop-down list:
 - **Unspecified:** Uses the IdP default authentication mechanism.
 - **Password Protected Transport:** Requires the user to log in with a password using a protected session.
8. To require users to re-enter their login information to access Illumio (even if the session is still valid), check the Force Re-authentication checkbox. This allows users to log in to the PCE using a different login than their default computer

login and is disabled by default.

NOTE:
When SSO is configured both in Illumio Core and for the IdP, the preferences in Illumio Core are used. When SSO is not configured in Illumio Core, the default IdP settings are used.

9. Click **Save**.
10. Log in to your OneLogin account.
11. Select the Illumio Core app, and then click the Configuration tab.
12. Enter the values copied from the Information for Identity Provider section of the PCE SSO configuration page.



The screenshot shows the 'Illumio ASP' configuration page. At the top, there are navigation tabs: 'USERS', 'APPS', 'ACTIVITY', and 'SETTINGS'. Below these is a breadcrumb '← Illumio ASP' and two buttons: 'MORE ACTIONS' and 'SAVE'. A secondary set of tabs includes 'Info', 'Configuration' (which is selected and underlined), 'Parameters', 'Rules', 'SSO', 'Access', and 'Users'. Under the 'Configuration' tab, there is a section titled 'Application Details' with three input fields: 'Issuer', 'Assertion Consumer URL', and 'Logout URL'. A red vertical line and text callout on the right side of these fields reads: 'Enter PCE Information for Identity Provider here'. Below the input fields, a small note states: 'This information may be found on the SSO Config page of the PCE web console (located under the User menu)'.

13. Click **Save**.

Your PCE is now configured to use OneLogin SSO for authenticating users with the PCE.

Ping Identity Single Sign-on

This section explains how to configure SSO for authentication users with the PCE using Ping Identity as your Identity Provider (IdP).

Configure SSO for Ping Identity

Before you begin, make sure you have this information from your Ping Identity SSO account:

- x.509 certificate
- Remote Login URL
- Logout Landing URL

NOTE:

Your PCE user account must have Owner or Admin privileges to perform this task.

To configure the PCE for Ping Identity SSO:

1. From the PCE web console menu, choose **Settings > SSO Config**.
2. Click **Edit**.
3. Select SAML from the Select SSO method drop-down list and click **Configure**.
4. Enter the following information:
 - **SAML Identity Provider Certificate:** Paste your Ping Identity x.509 certificate (in PEM text format).
 - **Remote Login URL:** Enter the Ping Identity Remote Login URL.
 - **Logout Landing URL:** Enter the Ping Identity Logout Landing URL.
5. In the Information for Identity Provider section, make note of the following fields:
 - Issuer
 - NameID Format
 - Assertion Consumer URL
 - Logout URL
6. Select the authentication method from the drop-down list:
 - **Unspecified:** Uses the IdP default authentication mechanism.
 - **Password Protected Transport:** Requires the user to log in with a password using a protected session.
7. To require users to re-enter their login information to access Illumio (even if the session is still valid), check the Force Re-authentication checkbox. This allows users to log in to the PCE using a different login than their default computer login and is disabled by default.

NOTE:

When SSO is configured both in Illumio Core and for the IdP, the preferences in Illumio Core are used. When SSO is not configured in Illumio Core, the default IdP settings are used.

8. Click **Save**.
9. Log in to your Ping Identity account.
10. Select the Applications tab and add the Illumio app.
11. Click **Edit** and enter the following values you just noted from Illumio:
 - **ACS URL:** Enter the value from the Assertion Consumer URL field in the PCE web console.
 - **Entity ID:** Enter the value from the Issuer field in the PCE web console.
 - **Single Logout Endpoint:** Enter the value from the Logout URL field in the PCE web console.
 - **Single Logout Response Endpoint:** Enter the value from the Logout URL field in the PCE web console.


The screenshot shows the 'My Applications' configuration page in the Ping Identity Admin console. The application 'Illumio ASP' is listed with a status of 'Incomplete' and is not yet enabled. The configuration step '1. Configure your connection' is active, requiring the user to assign attribute values for single sign-on (SSO). The form includes fields for ACS URL, Entity ID, Single Logout Endpoint, and Single Logout Response Endpoint, all of which are highlighted with red boxes in the image. The ACS URL and Entity ID fields contain placeholder text: 'https://\${Enter Assertion Consumer U}' and '\${Enter Issuer from the SSO Config p}' respectively. The Single Logout Endpoint and Single Logout Response Endpoint fields contain placeholder text: 'https://\${Enter Logout URL from the S'.

12. Click **Continue to Next Step**.
13. You will now configure the SAML_SUBJECT attribute mapping. Under Advanced Attribute Mapping, next to the Name ID Format to send to SP, select urn:oid:is:names:tc:SAML:1.1:nameid-format:emailAddress.

Advanced Attribute Options

Advanced Attribute Options for SAML_SUBJECT

Advanced Attribute Options

NameIDFormat 

Name ID Format to send to SP:

Attribute Mapping

You can build an attribute mapping using the following syntax:

An example of a possible SAML_SUBJECT value is:

```
firstName + "." + lastName + "@domain.com"
```

IDP Attribute Name or Literal Value	As Literal	Function
1 SAML_SUBJECT	<input type="checkbox"/> As Literal	<input type="text" value=""/>

14. Click **Save**.

Your PCE is now configured to use Ping Identity SSO for authenticating users with the PCE.

PCE Administration Troubleshooting

This chapter contains the following topics:

PCE Administration Troubleshooting Scenarios	279
--	-----

This section describes issues that can arise during ongoing PCE operation and how to resolve them.

PCE Administration Troubleshooting Scenarios

This section describes issues that can arise during ongoing PCE operations and how to resolve them.

Transaction ID Wraparound in PostgreSQL Database

Symptom:

The PCE uses PostgreSQL databases to store data. Under certain conditions, PostgreSQL may issue warnings about transaction ID wraparound.

WARNING:
These messages indicate a very serious condition. The database is not functional, and the PCE will not work as expected. Immediate remediation from Illumio Support is required.

In `illumio-pce.log` and `postgresq1.log`, look for messages like the following:

```
ERROR: database is not accepting commands to avoid wraparound data loss in
database "<database_name>"
```

```
Stop the postmaster and vacuum that database in single-user mode.
```

Cause:

In a PostgreSQL database, transaction wraparound (also known as transaction ID exhaustion) can occur if a very large number of transactions have occurred and the transaction ID reaches its maximum possible value and is forced to begin again at zero. As a result, transactions from the past suddenly have a higher transaction ID than the current ID, and therefore appear to be in the future – and therefore inaccessible. The result is extreme loss of data. The database stops accepting requests.

The only way to recover from transaction ID wraparound is to manually execute commands.

To avoid this situation, PostgreSQL provides an autovacuum feature which recovers disk space, by doing things like removing dead row versions, before transaction ID wraparound can occur. The PCE databases use the PostgreSQL autovacuum feature to prevent transaction wraparound. However, in the following situations, autovacuum might not succeed:

- Vacuum did not run on the tables.
- Temporary tables remained in the database, rather than being dropped as they should be. Temporary tables are not vacuumed.

For details about autovacuum and transaction ID wraparound, see the PostgreSQL documentation page [Preventing Transaction ID Wraparound Failures](#).

Monitoring and Diagnosis:

Use the `dbcheck` tool to periodically monitor the system for early detection of any potential transaction ID wraparound condition. It is vital to act before the situation develops into transaction ID wraparound failure. See [Monitor Database Replication](#).

WARNING:

If you find messages that indicate a risk of transaction wraparound, immediately contact Illumio Support for assistance.