Illumio Core 24. 2.10 Administration Guide

Published: 2024

The guides in this category explain how to configure specific Illumio Core features after installing the PCE and VEN software in your environment.

Use the information in this section to keep your Illumio Core deployment running effectively, and to troubleshoot issues that might occur.

Table of Contents

Security Advisories

This category includes announcements of security fixes and updates made in critical patch update advisories, security alerts and bulletins.

September 2024 Security Advisories

Here's a list of the security advisories for 2024.

Ruby SAML gem component authentication bypass vulnerability

The Ruby SAML gem is affected by an authentication bypass vulnerability, which impacts the Illumio PCE in both SaaS and on-premises deployments. An authenticated attacker could potentially leverage this vulnerability to authenticate as another SAML user. For SaaS customers, the target user can be in a different org and on a different cluster.

Severity

Critical: CVSS score is 9.9

CVSS: AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H

Affected Products and Patch Information

Security vulnerabilities addressed by this Security Alert affect the products listed below.

Table 1. Products Affected by the Security Vulnerability

Affected Products	Affected Versions	Fixed Version
Illumio Core PCE	<= 21.5.36	>= 21.5.37
	<= 22.2.42	>= 22.2.43
	<= 22.5.32	>= 22.5.34
	<= 23.2.30	>= 23.2.31
	<= 23.5.21	>= 23.5.22
	<= 24.2.0	>= 24.2.10

Resolution

Upgrade to the latest release for a given major version.

References

- https://nvd.nist.gov/vuln/detail/CVE-2024-45409
- https://github.com/advisories/GHSA-jw9c-mfg7-9rx2

Skipped Critical Patch Updates

Illumio strongly recommends that customers apply security patches as soon as possible. For customers that have skipped one or more Critical Patch Updates and are concerned about products that do not have security patches announced in this Critical Patch Update, please review previous Critical Patch Update advisories to determine appropriate actions.

Discovered By

External Security Firm

Frequently Asked Questions

- What software components are affected?
  Only the Illumio PCE is impacted by this vulnerability.
- What products did this affect?
  This vulnerability impacts the PCE, including Core on-premises deployments, Core SaaS, Endpoint, MSP, and Edge.
- Is Core SaaS affected?
  SaaS PCE clusters were impacted. Those environments have been patched.
- I'm using Cloud. Am I impacted?
  The Cloud platform is not affected.
- Will the patch affect performance?
  The update is not expected to affect performance.
- How can I tell if this vulnerability was used against my on-premises PCE?
  Illumio is creating queries that can be used by customers to detect known vectors for exploitation of this vulnerability. Please contact Illumio Support or your account team for assistance. If you suspect this vulnerability was used within your environment, please reach out to Illumio Support.
- Has Illumio investigated if this vulnerability was used on any SaaS PCEs?
  Illumio is currently investigating all available data from the production SaaS environment and has so far found no indications that the issue has been exploited.
- I can't apply the patch immediately. How can I mitigate the issue in the meantime?
  This vulnerability requires SAML to be enabled on the customer's PCE in order to be exploited. Customers who cannot patch their PCEs immediately, and who wish to mitigate this issue, can choose to disable SAML authentication on the PCE. For details, see the "Authentication" topic in the PCE Administration Guide. Additionally, customers can enable IP restrictions to limit access to only trusted source IPs (for example, for privileged accounts). For details, see the "Configure Access Restrictions and Trusted Proxy IPs" topic in the PCE Administration Guide.
- How long will the upgrade take?
  The fix will be provided in a normal code release so this will take the same amount of time as any PCE upgrade.
- Were any Illumio customers impacted by this vulnerability?

Illumio is not aware of any exploitation of this vulnerability within any customer environments.

Modification History

- September, 2024: Initial Publication of CVE

September 2023 Security Advisories

Here's a list of the security advisories for 2023.

Authenticated RCE due to unsafe JSON deserialization

Unsafe deserialization of untrusted JSON allows execution of arbitrary code on affected releases of the Illumio PCE. Authentication to the API is required to exploit this vulnerability. The flaw exists within the network_traffic API endpoint. An attacker can leverage this vulnerability to execute code in the context of the PCE's operating system user.

Severity

Critical: CVSS score is 9.9

CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H

Affected Products and Patch Information

Security vulnerabilities addressed by this Security Alert affect the products listed below.

Table 2. Products Affected by the Security Vulnerability

Affected Products	Affected Versions	Fixed Version
Illumio Core PCE	<= 19.3.6	>= 19.3.7
	<= 21.2.7	>= 21.2.8
	<= 21.5.35	>= 21.5.36
	<= 22.2.41	>= 22.2.42
	<= 22.5.30	>= 22.5.31
	<= 23.2.10	>= 23.2.11

Resolution

Upgrade to the latest release for a given major version.

References

https://www.cve.org/CVERecord?id=CVE-2023-5183

Skipped Critical Patch Updates

Illumio strongly recommends that customers apply security patches as soon as possible. For customers that have skipped one or more Critical Patch Updates and are concerned about products that do not have security patches announced in this Critical Patch Update, please review previous Critical Patch Update advisories to determine appropriate actions.

Discovered By

External Security Firm

Frequently Asked Questions

- What software components are affected?
  Only the Illumio PCE is impacted by this vulnerability.
- What products did this affect?
  This vulnerability impacts the PCE, including Core on-premises deployments, Core SaaS, Endpoint, MSP, and Edge.
- Is Core SaaS affected?
  SaaS PCE clusters were impacted. Those environments have been patched.
- I'm using Cloud. Am I impacted?
  The Cloud platform is not affected.
- How can I tell if this vulnerability was used against my on-premises PCE?
  Illumio is creating queries that can be used by customers to detect known vectors for exploitation of this vulnerability. Please contact Illumio Support or your account team for assistance. If you suspect this vulnerability was used within your environment, please reach out to Illumio Support.
- Has Illumio investigated if this vulnerability was used on any SaaS PCEs?
  Illumio is currently investigating all available data from the production SaaS environment and has so far found no indications that the issue has been exploited.
- I can't apply the patch immediately. How can I mitigate the issue in the meantime?
  This vulnerability requires SAML to be enabled on the customer's PCE in order to be exploited. Customers who cannot patch their PCEs immediately, and who wish to mitigate this issue, can choose to disable SAML authentication on the PCE.
- Reference
  For details, see the topic Authentication in the PCE Administration Guide.
  Additionally, customers can: Enable IP restrictions to limit access to only trusted source IPs (for example, for privileged accounts). For details, see the topic Configure Access Restrictions and Trusted Proxy IPs in the PCE Administration Guide.

- How long will the upgrade take?
  The fix will be provided in a normal code release so this will take the same amount of time as any PCE upgrade.
- Were any Illumio customers impacted by this vulnerability?
  Illumio is not aware of any exploitation of this vulnerability on any customer environments.

PCE Administration

Overview of PCE Administration

This section describes how to maintain and operate the Policy Compute Engine (PCE). It also includes other tasks required to manage your PCE deployment and help you with ongoing PCE operations and administration.

Before You Begin

Before you begin, become familiar with the following technology:

- Your organization's security goals
- 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

This section gives information about the notational conventions used here.

Review these Notational Conventions

- 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 four main service tiers that are used by both the PCE Web Console UI and the VEN:

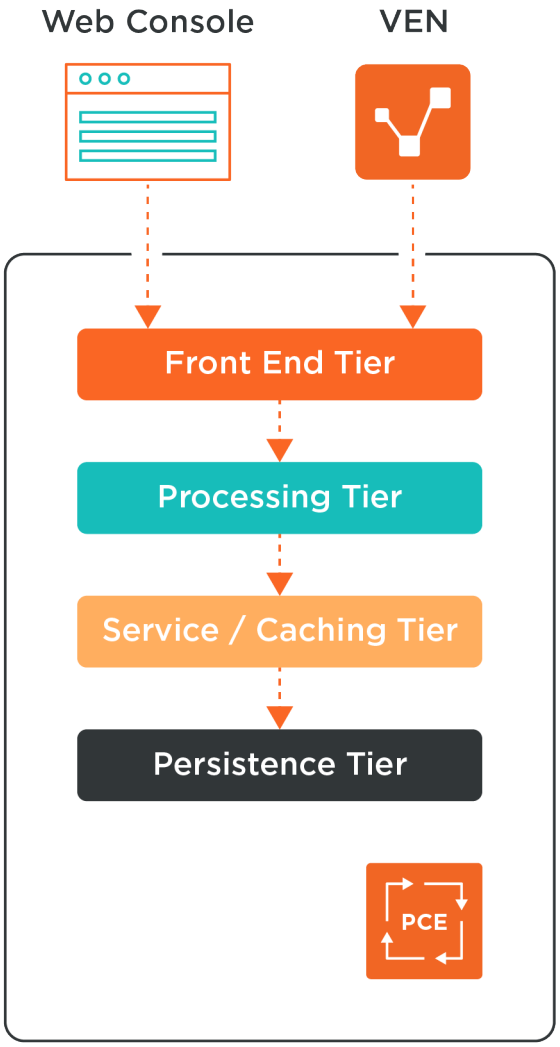

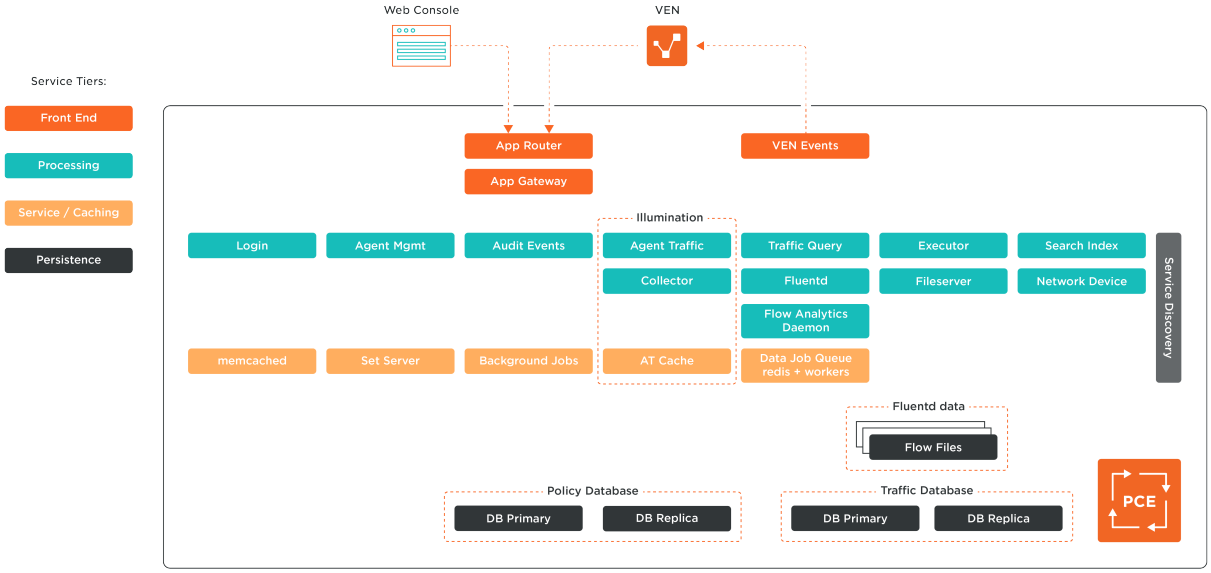
Web Console
VEN
Service Tiers:
Front End
Processing
Service / Caching
Persistence
App Router
App Gateway
VEN Events
Illumination
Login
Agent Mgmt
Audit Events
Agent Traffic
Traffic Query
Executor
Search Index
Collector
Fluentd
Fileserver
Network Device
Flow Analytics Daemon
memcached
Set Server
Background Jobs
AT Cache
Data Job Queue redis + workers
Service Discovery
Fluentd data
Flow Files
Policy Database
DB Primary
DB Replica
Traffic Database
DB Primary
DB Replica
PCE

Description of PCE Components

Tier	PCE component	Description
Front-end	Management interfaces: PCE web console and VEN	Management interfaces include: • PCE web console • REST API • PCE command line • VEN command line
	VEN events	For information, see VEN Administration Guide.
	App Router	Directs requests to the proper service.
	App Gateway	Ensures that all communication between cluster nodes is encrypted and that only cluster nodes can connect to internal services. Most services connect via the application gateway.
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.
	Traffic Query	API for traffic explorer
	Flow Analytics Daemon	Flow analytics daemon
	Network Device	Manages network devices such as switches and server load balancers that are managed by the PCE.
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.”

Tier	PCE component	Description
	Data Job Queue (Redis + workers)	Data job queue
Persistence	Fluentd data	Flow files
	Policy primary database and replica	Postgres database contains all 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 data store. 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.

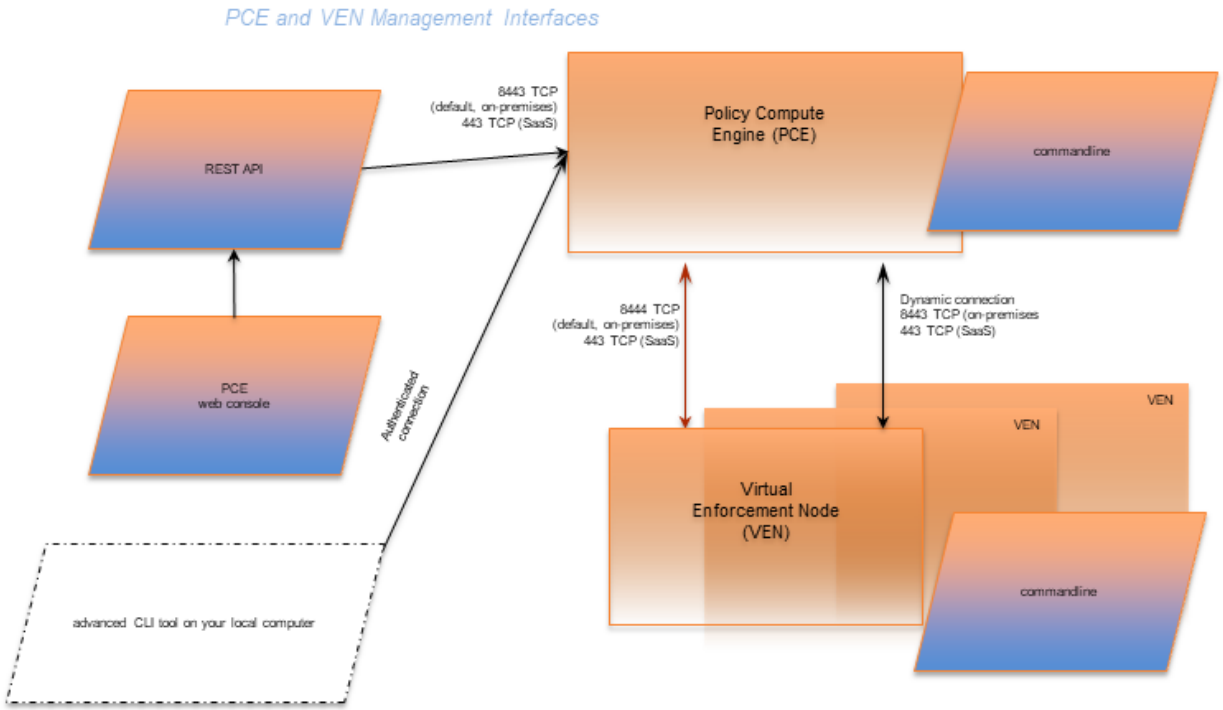

This guide focuses on the use of the `illumio-pce-ctl` control script and related administrative programs on the PCE itself.

Interface	Notes
PCE web console	With the PCE web console, you can perform many common tasks for managing the Illumio Core.
PCE command line	Use of the command line directly on the PCE. The `illumio-pce-ctl` 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.
REST API	With the Illumio Core REST API, you can perform many common management tasks, such as automating 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.
VEN command line	The `illumio-ven-ctl` command-line tool is the primary management tool for the VEN.

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

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 the 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.

Access Configuration for PCE

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

The following topics describe role-based access control features.

Built-in Roles

Illumio Core includes several roles that grant users access to perform operations. Each role is matched with a scope.

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, 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	• 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. **NOTE** You can choose the All Applications, All Environments, and All Locations scope with the Full Ruleset Manager role.
Limited Ruleset Manager	• 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. **NOTE** You cannot choose the All Applications, All Environments, and All Locations scope with the Limited Ruleset Manager role.
Ruleset Viewer	• Read-only access to rules that match the specified scope. • Ruleset Viewers cannot edit rules or rulesets.
Ruleset Provisioner	Provision rulesets within specified scope. **NOTE** You can choose the All Applications, All Environments, and All Locations scope and custom scopes with the Ruleset Provisioner role.
Workload Manager	Manage workloads and pairing profiles within the specified scope. Read-only access provided to all other resources. **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.

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.

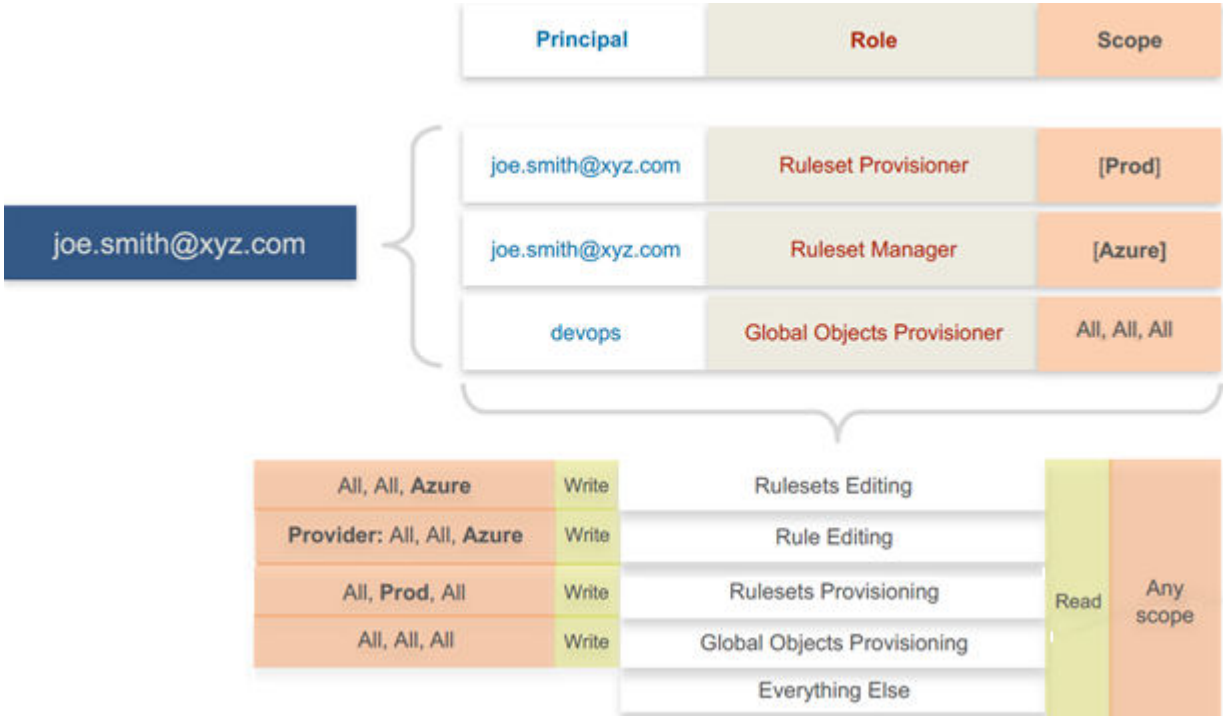

Exercise Caution when Combining Roles

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.

  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.

NOTE

Permission to configure these settings is dependent on your role.

Add a Scoped Role

Add a scoped role to create fine-grained access control to manage security policy for your workloads.

By defining scopes, you can grant different permissions to different users for different resources. For example, you might allow some users to add rulesets for all workloads in your staging environment. You might grant access to all workloads in all environments for other users.

When adding a scoped role:

- use the Access Wizard
- Define the scope of the role by selecting labels or label groups for applications, environment, and location.
- Add a local user, external user, or user group to the role.
- Select roles and confirm your choice.

Manage a Local User

Local users are created in the PCE (an IdP does not manage them). 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 inaccessible.

To add a local user:

- In the Local Users tab, click **Add**.
- Enter a name and an email address. The email address must use the format xxxx@yyyy.zzzz and be 255 characters or less.

  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 to the address you specified an invitation to with a link to create their Illumio user account. The link in the invitation email is valid only for 7 days, after which it expires.
- Select a role for the user: None, Global Organization Owner, Global Administrator, or Global Read Only.

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. The "My Roles" feature allows you to view the list of assigned permissions (roles).

To remove a local user

Select it in the Users and Groups and remove it.

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

In Users and Groups, find the user you want to edit. change the user's name and 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 and Users and Groups pages. The name is not changed in the user's profile or on the RBAC User Activity pages.

NOTE

Local and external users can change their names when they create their accounts or from their profiles.

To convert a local user

In Users and Groups, select the name of the user and click **Convert**.

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

In Users and Groups, select the name of the user and click **Re-invite**.

You can send a new email to users 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

In Users and Groups, select the name of the user and click **Lock**.

Local users are locked out of their accounts when they fail to log in after five 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. The user's account resets after 15 minutes and does not require an Illumio administrator to unlock it.

Add or Remove an External User

Using RBAC, you can control access to Illumio Core for users who a corporate IdP externally authenticates. 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.

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:

Use the External Users tab to click Add and enter a name, 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 duplicate 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.

Select the role: None, Global Organization Owner, Global Administrator, or Global Read Only.

Users without a role (None) can still log into the PCE to view resources when Read Only User access to the PCE is enabled. 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:

Use the External Users tab to select the user you want to remove and click **Remove**.

Removing an external user removes the user from the External Users tab and all the user's RBAC role memberships. Your corporate IdP still manages the user's authentication.

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 logs the user out for their next action 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 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:

- Use the External Users tab to add an external group
- In the External Group field, enter the group name as it's 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 the PCE, check the **memberOf** attribute in the SAML assertion from your IdP. The **memberOf** attribute is a multiple-value attribute that contains a list of distinguished names for groups that contain the group.

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: Click Edit Group from the group's details page to change an external group's name.

Use the External Users tab to remove an external group, select it, and 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. Your corporate IdP still manages user authentication for the group members.

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.

Change Users and Groups Added to Roles

When you change the membership for a role, the affected users must log out and log in 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 them out of the next action they can take after revoking their access.

- In Global Roles, click the name of the role you want to assign users or groups to
- To remove a user or group from the role, select it and click **Remove**.
- To add a user or group to a role, click **Add**.
- From the first drop-down list, select what (Any Principal Type, Local Users, External Users, or External Groups) you want to add to the role.

  Selecting what you want to add filters the second list to display only those types of users or user groups.
- Select the user or group to add to the role.
- Click **Grant Access**.

Alternatively, you can select users or groups to add to roles from the **Role-Based Access** > **User and Groups** details pages, and select **Add** 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 **Role-Based Access** > **User Activity**
  - Displays session details for each user, including their status, email address, and 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 appear in the User Activity pages can be different from the **Role-Based Access** > **Users and Groups** pages when users edit their profiles or an Organization Owner changes names in the **Role-Based Access** > **Users and Groups** pages.
- **Organization Events:** Go to **Troubleshooting** > **Organization Events**

  The Organization Events page provides an ongoing log of all events 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 has a severity level and are exportable in JSON format. You can narrow the search for many eventsby event type, severity, or time filters.

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.

To change your password

- In My Profile, click on **Change Password**.
- Enter your current password and then your new password twice.
- Click **Change Password**.

Color Vision Deficiency Mode

Users with color vision deficiency (Deuteranopia, Protanopia, or Tritanopia) can select Color Vision Deficiency mode, making it easier for them 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 color vision deficiency mode

- In My Profile, Accessibility section, select the **Color Visioin Deficiency** button.
-

> **NOTE**
> To restore the default setting, select the **Normal Vision** button.

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

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 multiple roles to the same user. The resulting set of permissions is the summation of all permissions included with each stacked. With these updates:

- Existing scoped roles were enhanced to restrict reads by scope.
- The 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 is disabled by default for new PCE installations.
- PCE performance and scale enhanced to support concurrently active users.

Global Roles

Global roles allow the user to view everything and 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 manage most aspects of the PCE, except 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.

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 as 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 scoped objects, provided the objects are inside the user's assigned scope. A user with this role can also 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: This role allows a user to perform workload-specific operations such as pairing, unpairing, label assignment, and changing policy state. A user with this role cannot view policies and traffic and cannot provision changes.

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 upgraded from prior releases, this setting must be manually turned **off** for users to have reads restricted by scope. If this setting is se **On**, users with scoped roles will get global visibility by default.

Figure 1. Global Viewer Setup

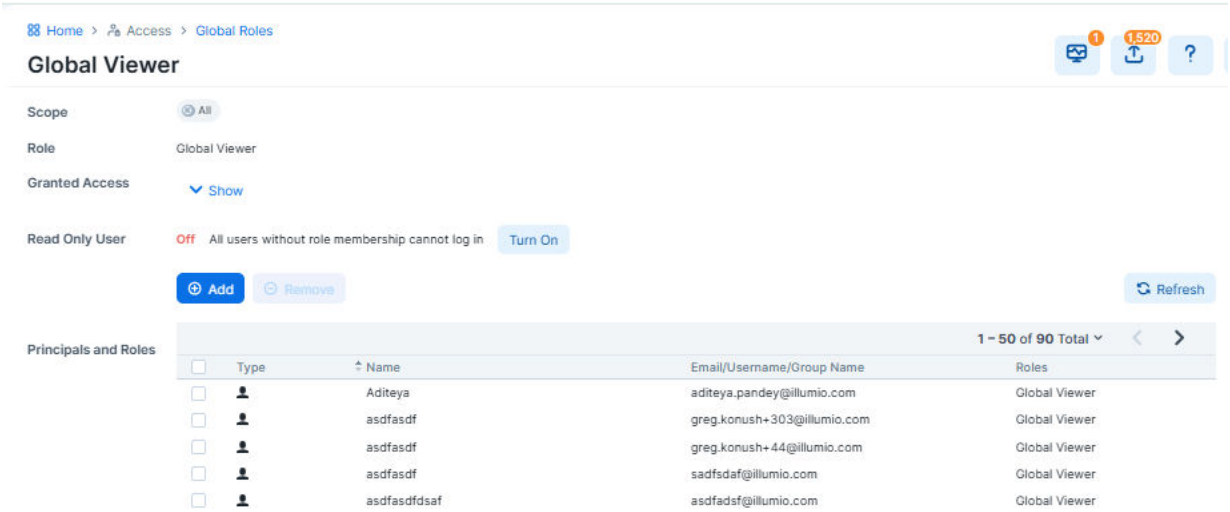

Manage Global Owners

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 Permis-sions)
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 Versions	Read (R)	Read (R)	Read (R)	---	Read (R)
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

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provisioner	Workload Manager	Application Owner (Combined Permissions)
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 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	---	---	---	---	---

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provisioner	Workload Manager	Application Owner (Combined Permis-sions)
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 Re-ports	---	---	---	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
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 Provision-er	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
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

Page	Ruleset Viewer	Ruleset Manager	Ruleset Provision-er	Workload Manager
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. The Illumination page opens when you log in to your account as an Organization Owner. 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.

Labeled Objects

The scope of the user filters labeled objects, such as workloads. 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 logged in as a Ruleset Manager, you cannot manage workloads. So, all the workload-specific operations buttons are disabled. However, you can 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 Rulesets. For example, if you search for Application Labels, 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, etc. The search bar automatically filters the facets to the list of facets in the user's assigned scope.

Global Objects

Scoped users get complete read-only visibility into all global objects. This includes IP Lists, services, labels, label groups, and user groups. However, scoped users cannot create, modify, or provision global objects.

NOTE

Only the Global Organization Owner and Global Administrator can create, modify, and provision global objects.

Rulesets and Rules

Scoped users, except Workload Managers, can see rulesets and rules that apply to their applications. A Ruleset Manager can edit the ruleset, whereas the other scoped roles (Ruleset Viewer and Ruleset Provisioner) can view rulesets. A scoped user can see all the rules within the application ruleset.

When label groups are used within the scope of a ruleset, a Ruleset Manager may not be allowed to edit the ruleset and its 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 that apply to their applications. For example, scoped users can view rules written by other applications that apply to their application. To see those rules, click **Rule Search** from the navigation menu.

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. The application owner provides visibility to all the rules that are applied to your application.

Policy Generator and Explorer

With Policy Generator, scoped users can generate policies only for their applications. Only Ruleset Managers can generate policies with Policy Generator. Ruleset Viewers can preview Policy Generator without the ability to save the 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.

My Roles

"My Roles" is a new feature that allows you to view the list of assigned permissions (roles).

App Group Map

The App Group Map provides visibility into applications and their contents. All scoped users except for Workload Managers can view App Group Maps.

Scoped users get limited visibility for connected App Groups such as Source App Groups and Destination App Groups. Scoped users get limited information on endpoints with traffic flows to their application. For an endpoint in a connected App Group from traffic flow, scoped users can get limited information such as labels, role names, and host names.

Figure 2. App Group Map

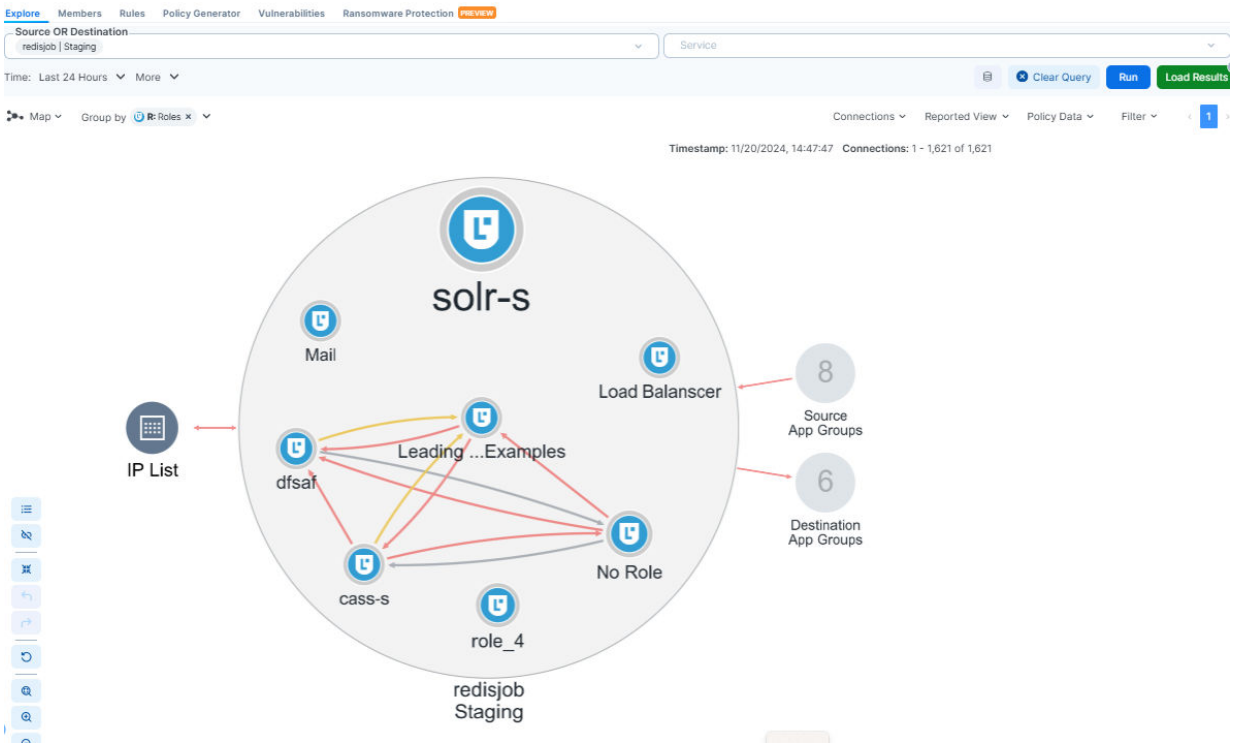

Configure Access Restrictions and Trusted Proxy IPs

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 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 adresses 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 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. Select **Settings > Trusted Proxy**.
3. In the Trusted Proxy IPs page, click **Edit**.
4. A list of trusted proxy IPs is displayed. Proxy configuration can have upto 8 Trusted Proxy IPs.

5. To remove any of the proxies from the list, select the checkbox in front of the proxy address and click **Remove**.
6. To edit Trusted Proxy IPs, click **Edit**.
7. In the Edit Trusted Proxy IPs dialog box, you can add a proxy IP address to the list, or delete any of the existingaddresses by hoiwering oiver the number in front of the address and then clicking the Trash Can icon that shows up.

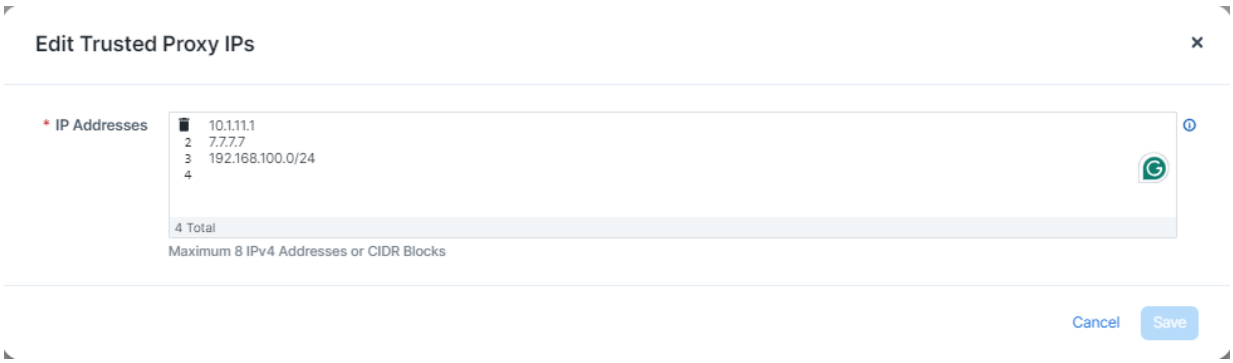

8. Once you have added or deleted the proxy addresses for your needs, click **Save**.

Manage API Keys

You can add and edit API keys using the PCE console.

Creating API Keys

1. In the Web console, type "API keys" in the Search field.
2. In the API Keys page, click **Add**.
3. In the "Create API Key" pop-up dialog, add the
   a. Key Name
   b. Description of the key
   c. Org ID
4. Click **Create**.
5. The confirmation dialog appears to show the data for the created API key.

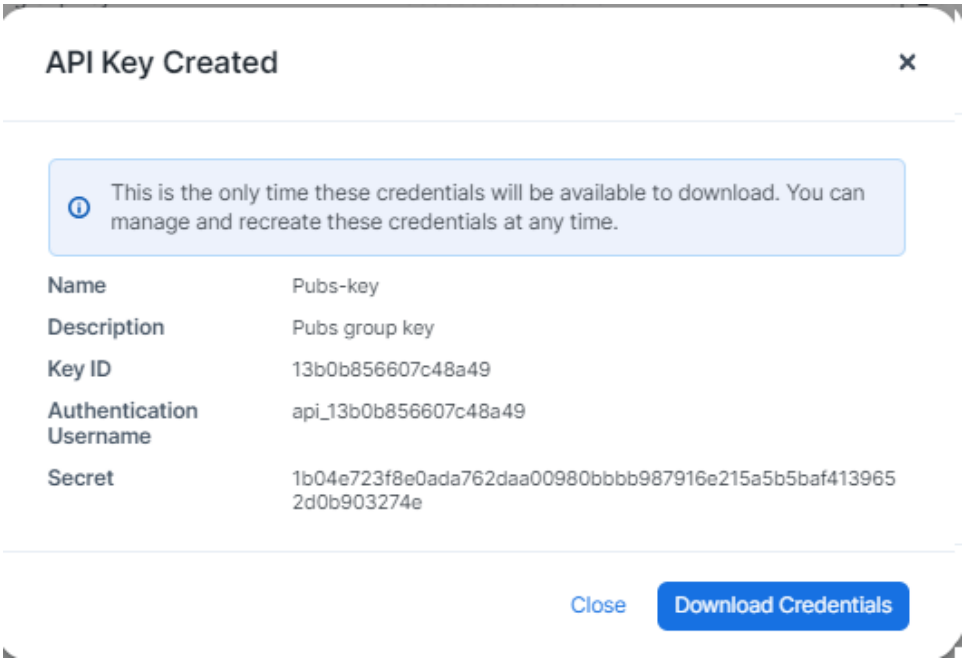

6. To download the credentials, click on **Download Credentials**.

You can download the credentials only after the key is created. You can, however, manage the credentials at any time.

7. The credentials will be downloaded in the default download directory on you hard drive , with the name API-Key-<your-key-name>. The format of the credential is a TXT file.

```
{"key_id":"13b0b856607c48a49","auth_username":"api_13b0b856607c48a49",
"secret":"1b04e723f8e0ada762daa00980bbbb987916e215a5b5baf4139652d0b903
274e"}
```

Editing Expiration of API Keys

To edit expiration of the Service account API keys using the PCE console:

1. Select **Settings > API Keys**.
2. On the API Key Settings page, click **Edit**.
3. By default, API Key for Service Account expires in:
   Select from thr dropdown list: `Never expires`, `1 day`, `30 days`, `60 days`, or `90 days`.
   If you change this setting, expiration of the existing API keys will not be impacted.
4. Keep expired API keys for:
   Select from the dropdown list: `1 day`, `30 days`, `60 days`, `90 days`, or `custom`.

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.

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
Session timeout	The session expiration timeout values must be set accordingly to balance security and usability so that your users can comfortably complete operations within the PCE web console without their session frequently expiring. The timeout value is dependent on how critical the application and its data are. For example, you might set the timeout to 3-5 minutes for high-value applications and 15-30 minutes for low-risk applications. The changed session timeout value applies to new browser sessions. Existing browser sessions are not affected when the session timeout value is changed. The PCE Org owner can go to **Access > Authentication > Local** to configure Session Timeout. This PCE session timeout is applicable to any user belonging to the same organization, regardless whether they are local or external users.
Timeout	30 minutes

SAML (in use)

Information from Identity provider

SAML Identity provider certificate	-----BEGIN CERTIFICATE----- MIICpDC-CAYwCCQD05WZzgx RugDANBgkqhkiG9w0BAQsFADAUMRIwEAYDVQQD-DAlsb2NhbGhvc 3QwHhcNMTgxMTE0MjAyNzM2WhcNMjgxMTExM-jAyNzM2WjAUMRIw EAYDVQQDDAlsb2NhbGhvc3QwggEiMA0GCSqGSIb3DQEBA-QUAA4I BDwAwggEKAoIBAQDXs/OhH90IPQ8qBrUMqzQZb5MI72fu+Ay0s P8gI1v8RiUqSl+WJNo8s9L8GNI9hnQT+OXg99PNmoE41xiAlnx qx8T78Qxb9zX3uc4hec+9bMSF7iieUiFXWQQrIUVM3g8TWI6B5g Uapt0vZcxNok2eNhiFvVTLgPzB06vb2/yU68ilwQ8wz/MGO00Un/ lRw3LORy-nEA1uMeT6terWtX8JQGbvc1qYddnXD86Y5MOP1AXU+ 1w1w1JFxD0uKiuOHJv-NYfJjkisEbDis9bO/EO0SyayVA7ABELaw QTfeWM6xLrNhZCTGeQiKb4XHMBgeliA-loEvNDDofKbLDQrWUyIf7 TAgMBAAEwDQYJKoZIhvcNAQELBQADggEBANLhqsZs-FUnq7kc+B5a vMmOXbCNJmSaASBULsX+akexhyJdMZUxmN6wfLjZ3FOwxvFuhe-Ta Zpkp1UtC+2E9YlxY//FxOX/YyvNT/xfOBzqZ9SCsNxpCBsSRK5X4 DS+2jGQuz3fwbJDxTXP4sKNUZ/E9Z+dC9Npdq7xtcXr7pWhI2qe MO8E9LdvfWLcsqq8Z0VtxyHYYZYNh8KN0Q6ObfK1sPC4QZ/292B 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-authentiucation	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://2x2testllab360.mylabs.io:8443/login/logout/6b5243ef-2305-4ffd-bf81-4fa97fb91a5b
Timeout	30 minutes

5. LDAP authentication is not active. Click **Turn On** to apply on all the LDAP servers.
6. To create an LDAP server, click on **Create Server**.
   To continue with LDAP server configuration, see the "LDAP Authentication" topic.

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.

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 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**.

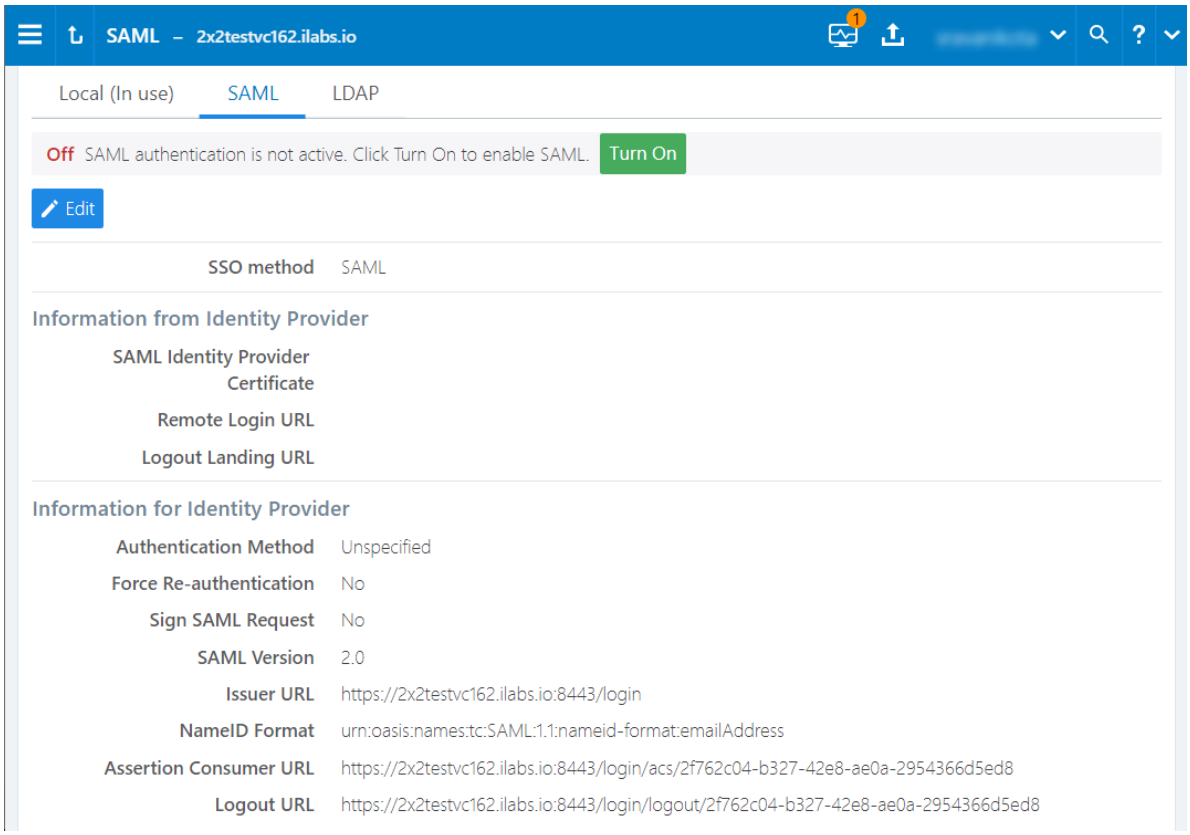

4. In the pop-up screen, click **Confirm**.

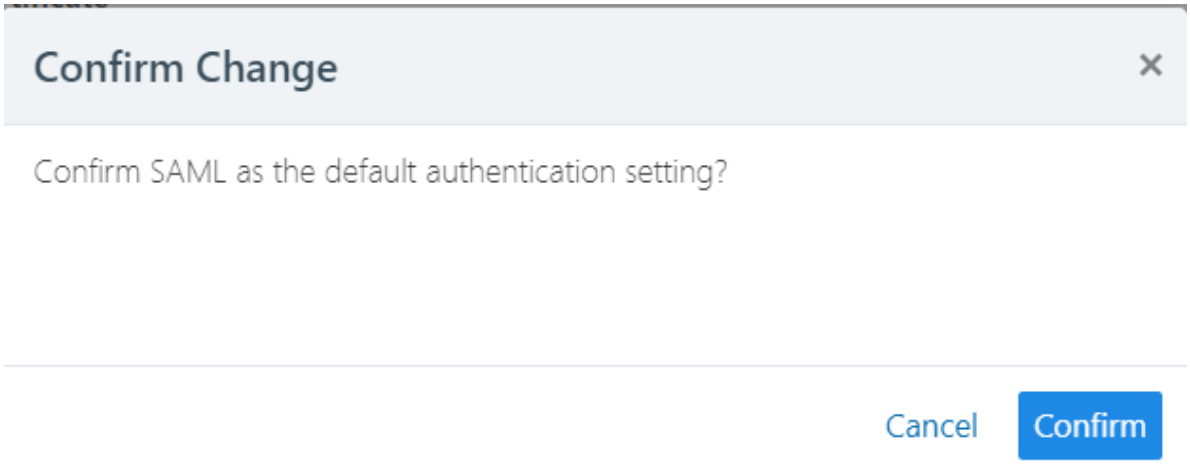

The updated SAML screen shows that SAML authentication is active.

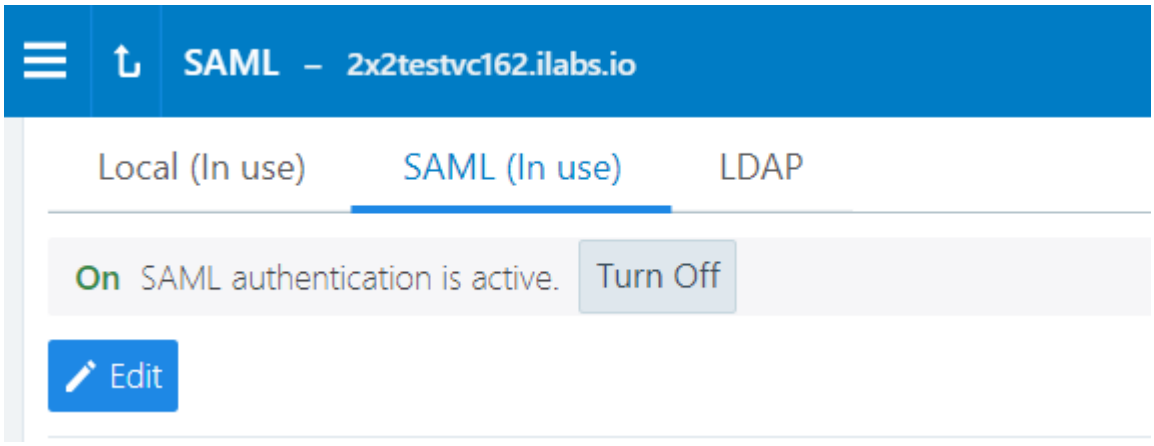

If necessary, you can disable it at any time.

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 [53].
2. Enable the PCE to use LDAP authentication. See Enabling LDAP Authentication [54].
3. Set up an LDAP configuration. See Configuring LDAP Authentication [54].
4. Map your LDAP groups to one or more PCE roles. See Map LDAP Groups to User Roles [55].

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.

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 username. 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.

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 the "Setup for Role-based Access Control" topic.

Modify LDAP Configuration

Follow these steps to update or delete an LDAP configuration in the PCE. It

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).

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.

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 username that's ADUserName@yourADDomanName. 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**.

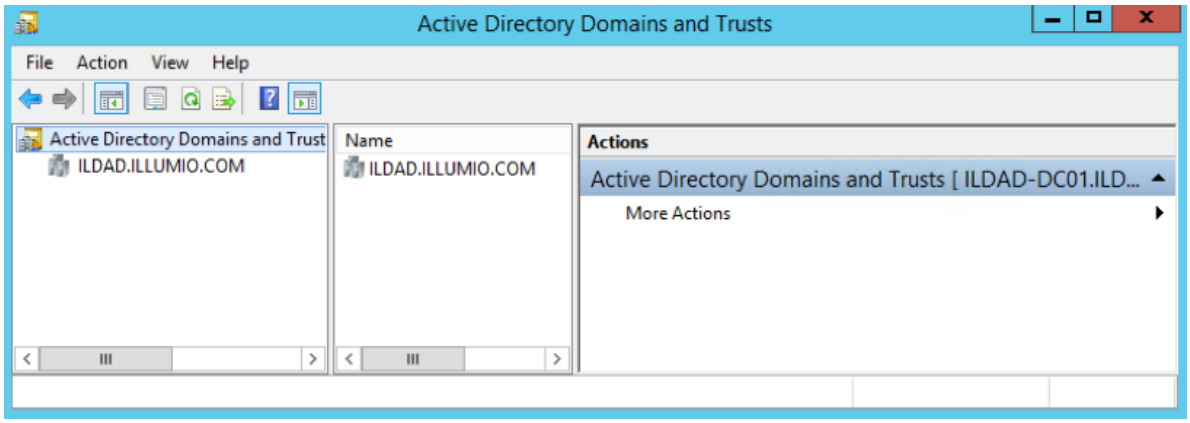

2. From the left side of the window, right-click Active Directory Domains and Trusts, and select **Properties**. In this dialog, you can create new suffixes for Active Directory usernames.

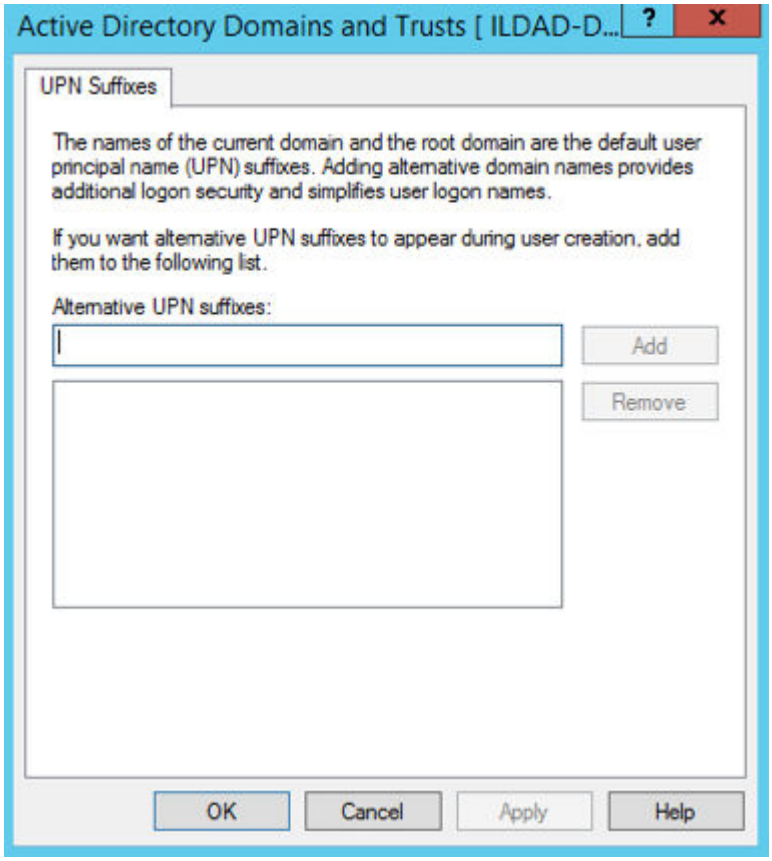

3. Create a suffix that matches the external namespace you'll be using and click **Add**.

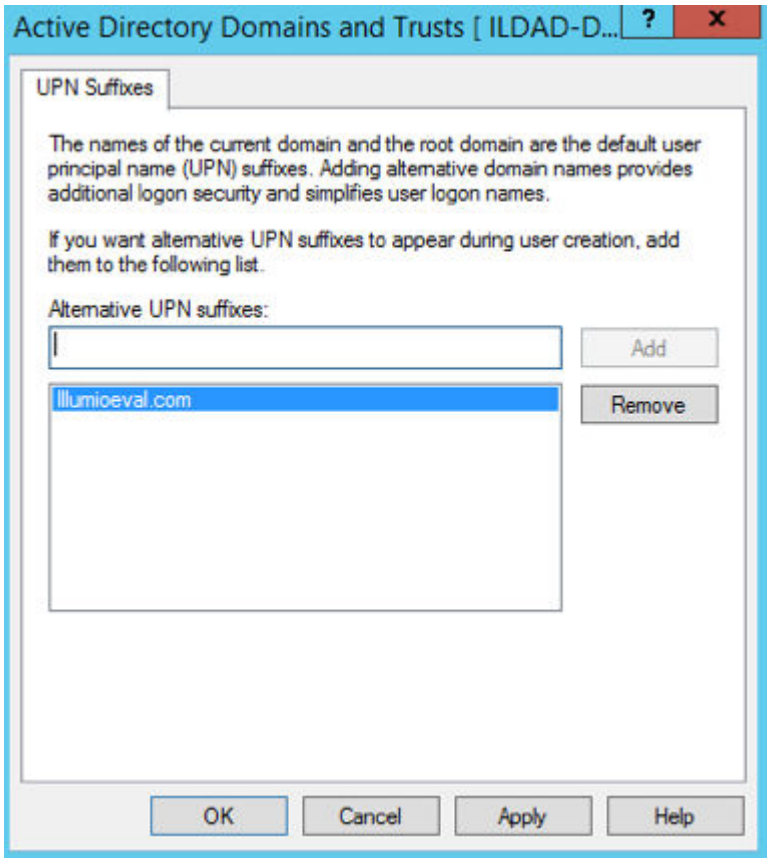

You can now assign an Active Directory user your custom UPN for the SAML response.

4. You can add multiple UPNs if needed. As shown below, you can select the UPN created in the previous steps.

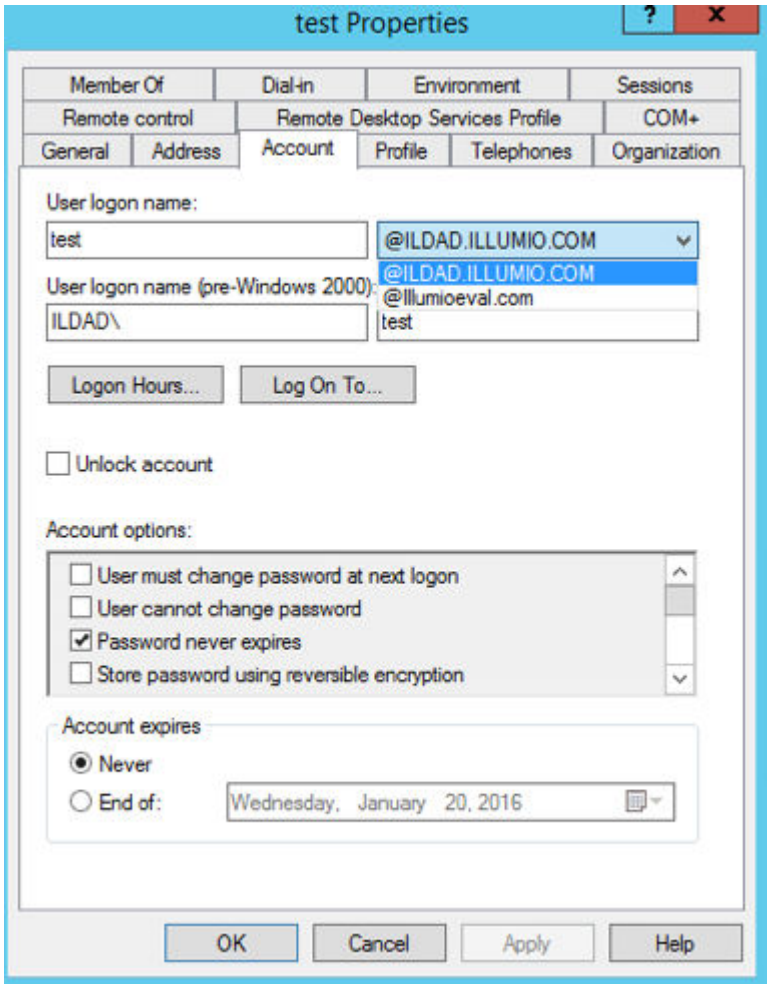

Your UPN configuration is set up and you can begin configuring AD FS for SSO with the PCE.

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.

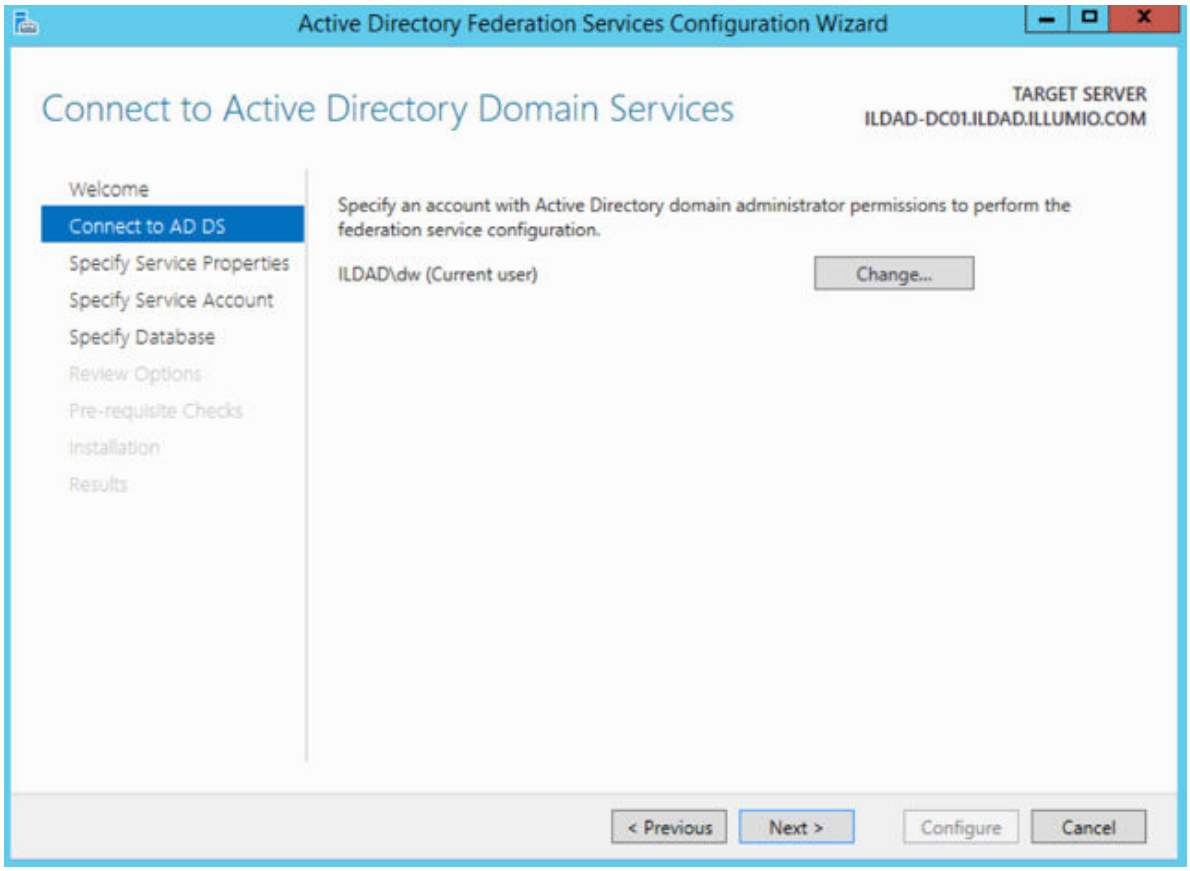

5. Select or import a certificate. This certificate can be a self-signed certificate.

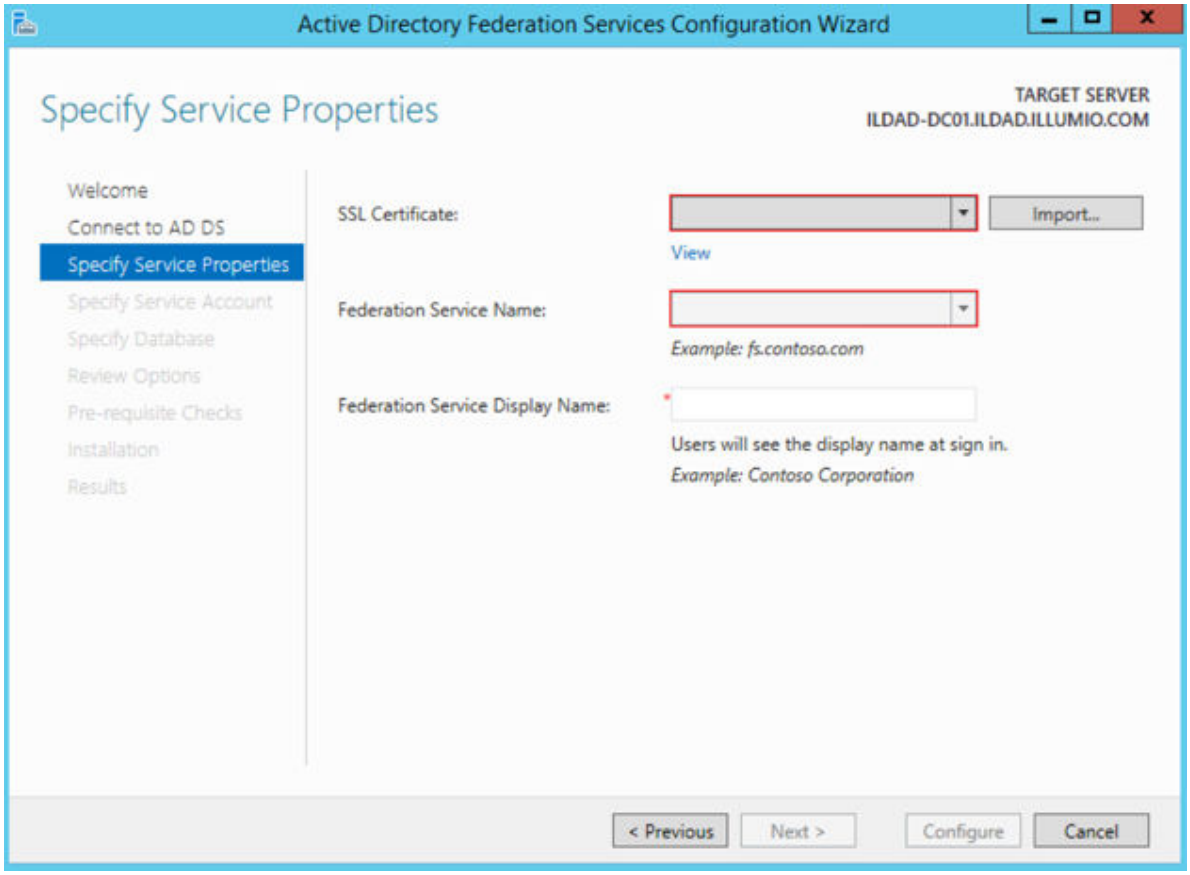

6. Specify your Federation Service Name, enter a display name for this instance of AD FS, and click **Next**

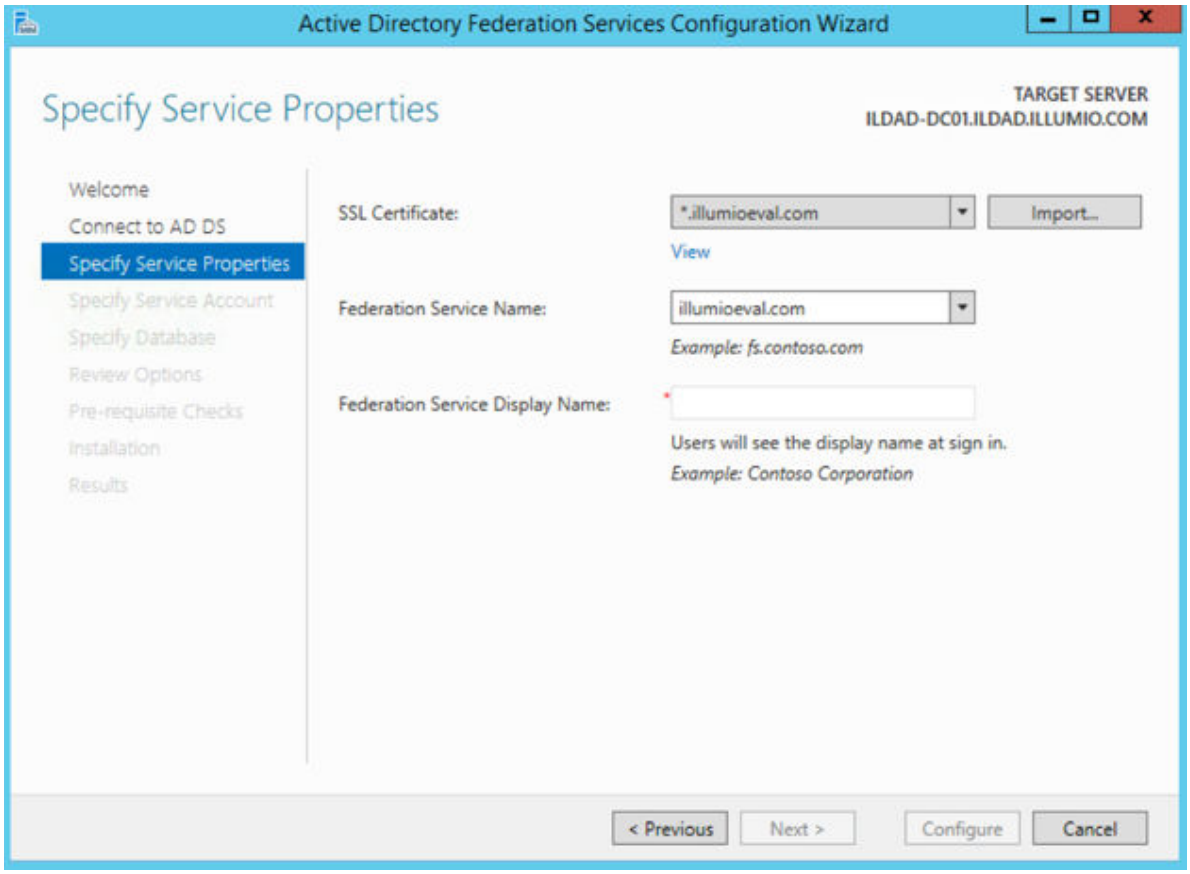

7. Specify your service account and click **Next**.

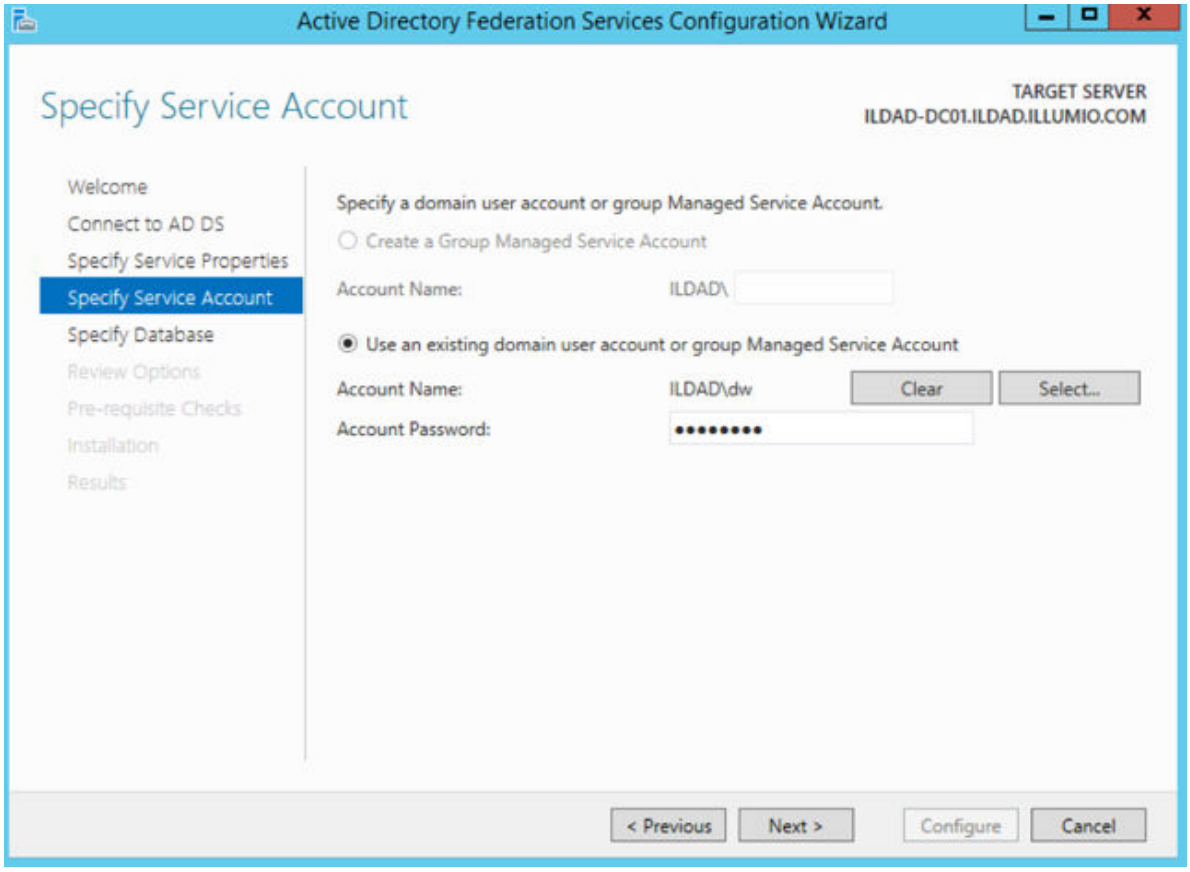

8. Select “Create a database on this server using Windows Internal Database” or choose the SQL server option, and click **Next**.

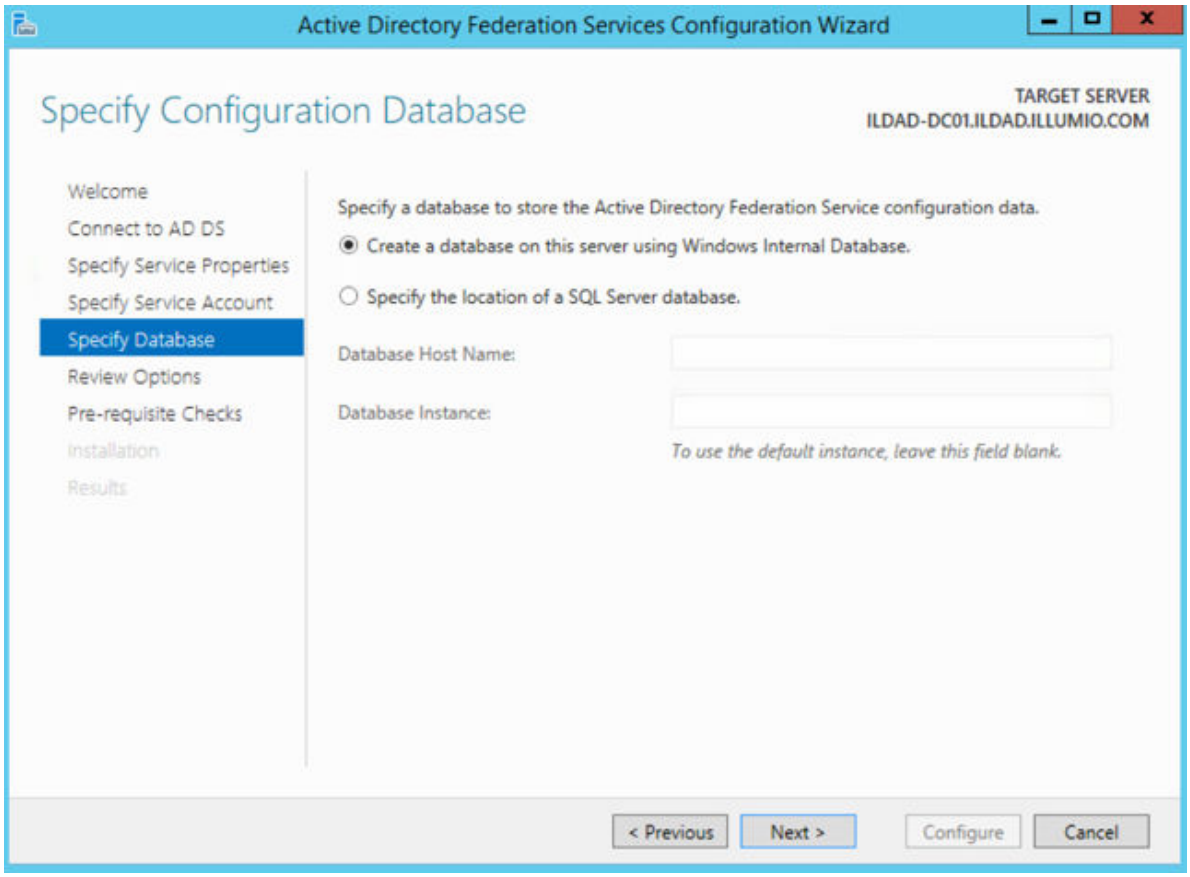

9. Review your selected options and click **Next**.

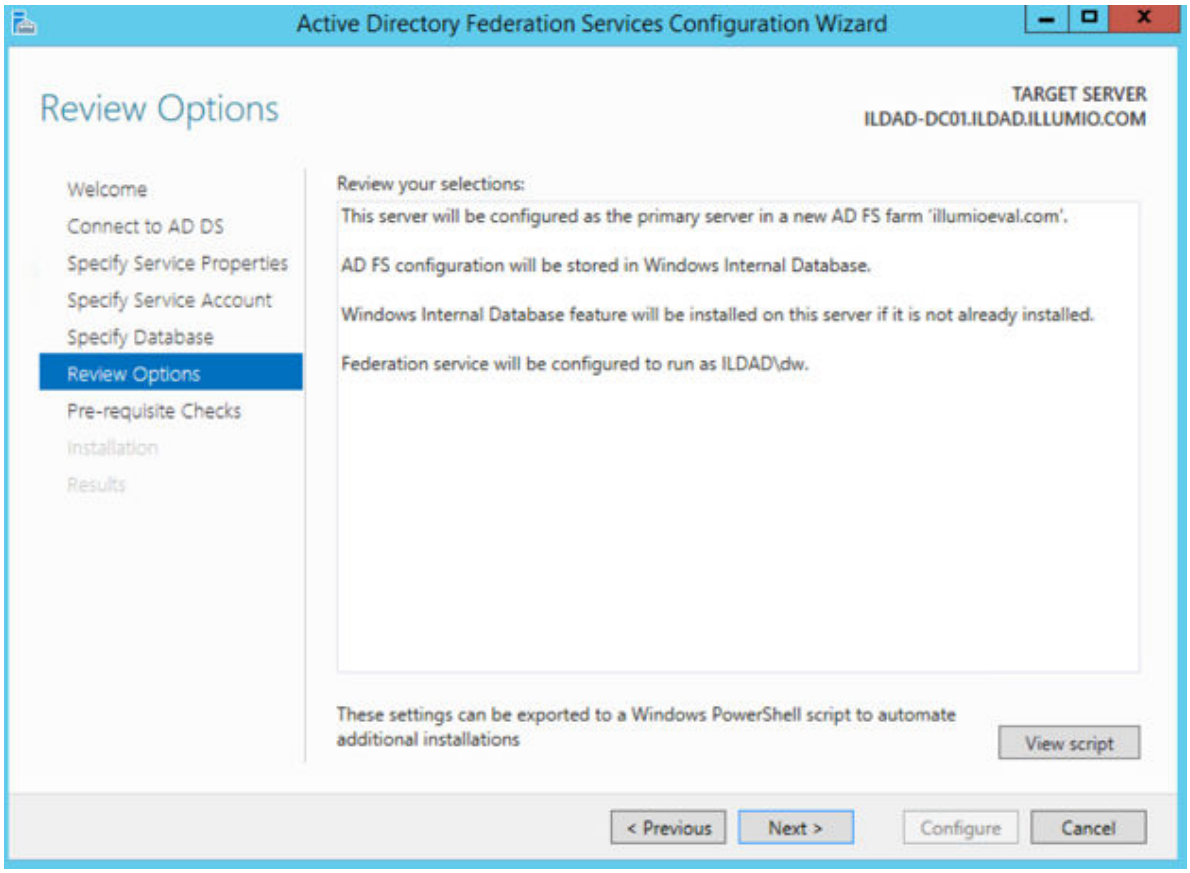

10 Click **Configure** to finish the basic configuration of AD FS.

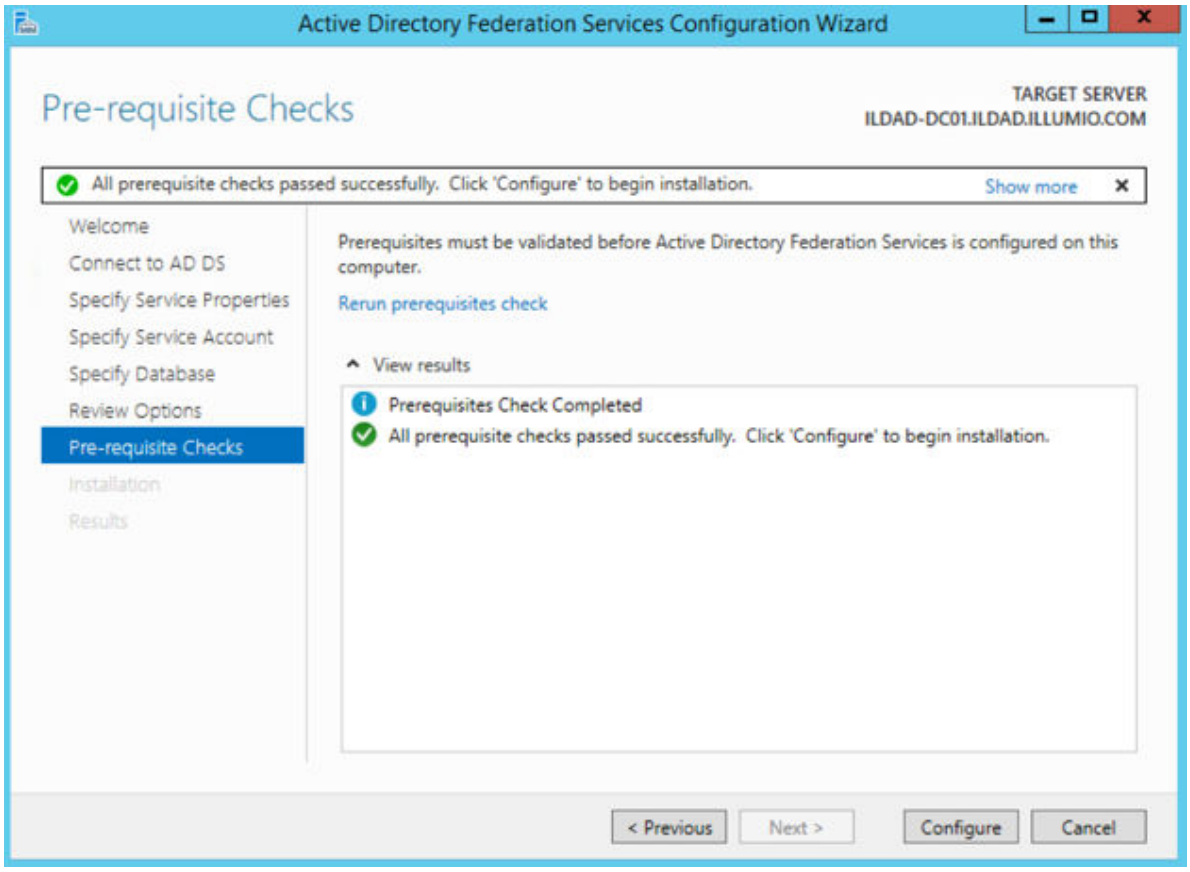

11. In the results screen, click **Close**.
AD FS is now installed with the basic configuration on this host.

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**.

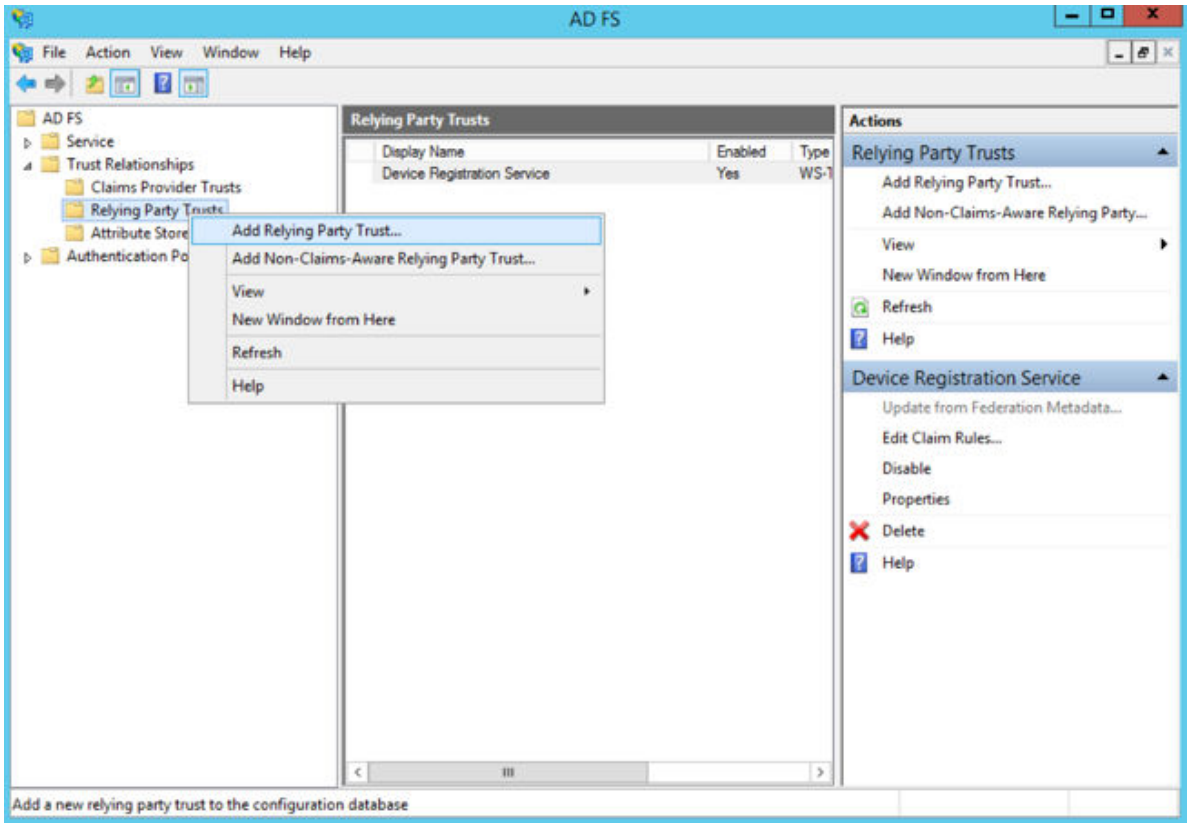

The Add Relying Party Trust Wizard appears.

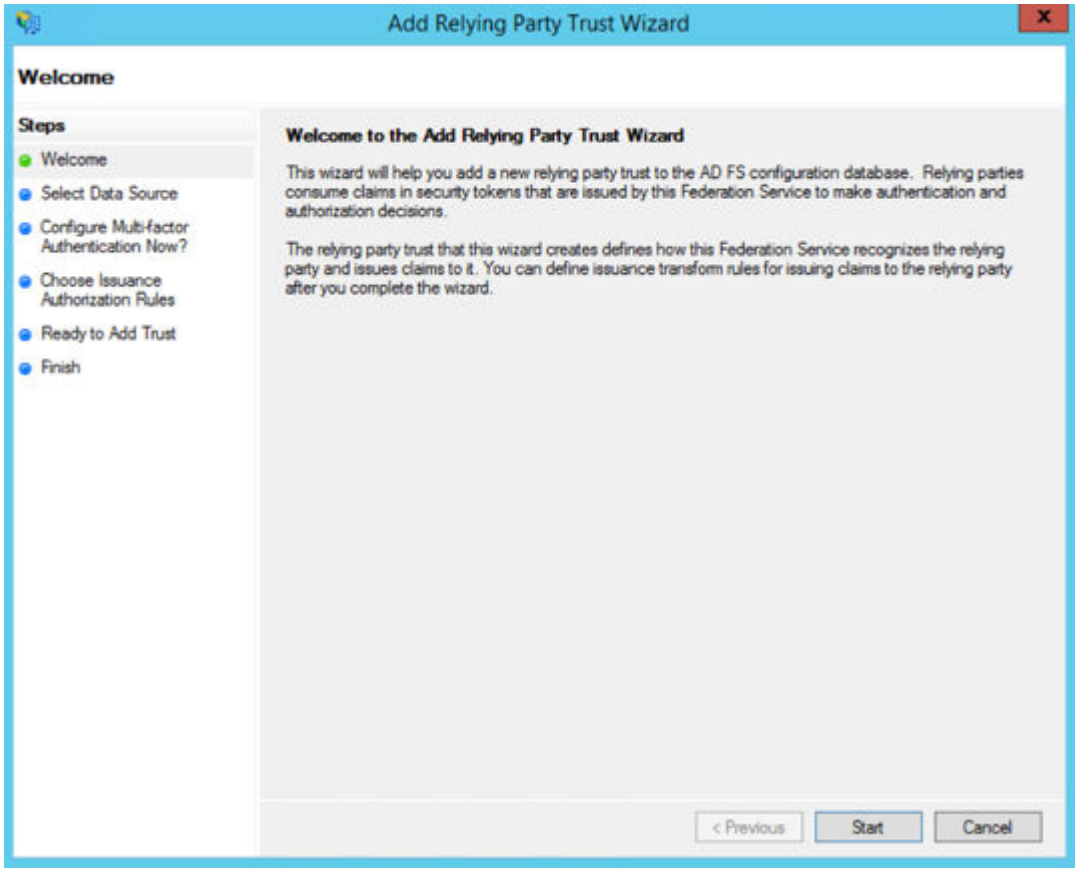

3. Click **Start**.
4. Select the "Enter data about the relying party manually" option and click **Next**.

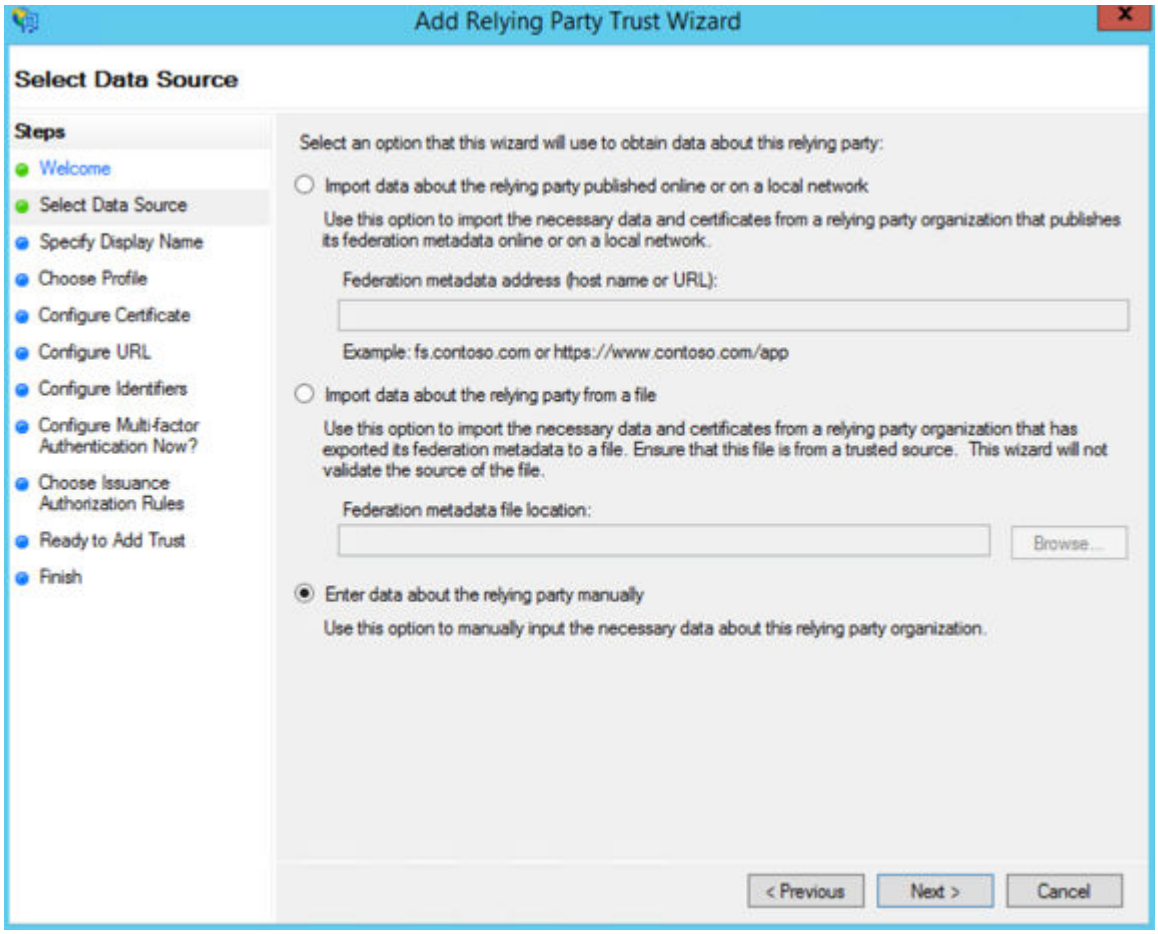

5. Name your Relying Party Trust and click **Next**.

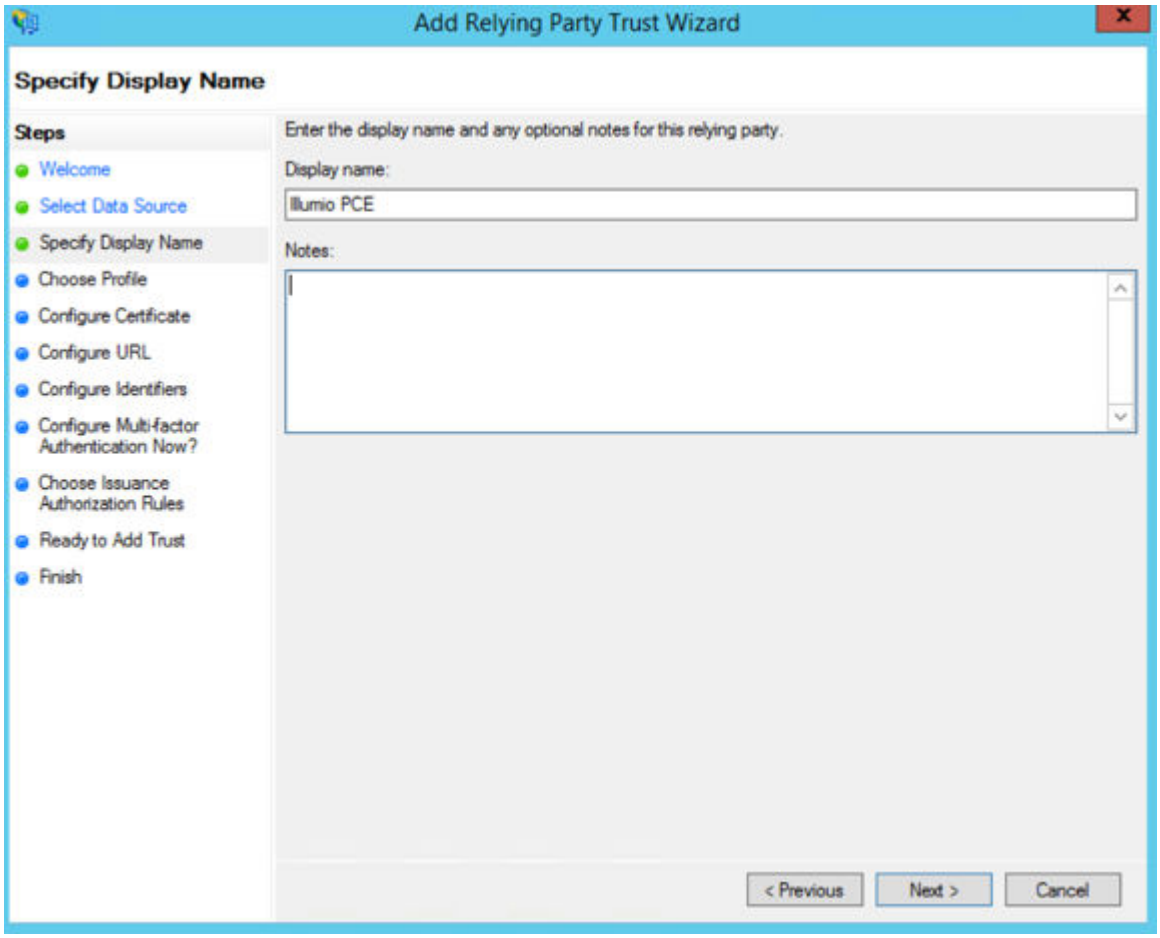

6. Select “ADFS profile” and click **Next**.

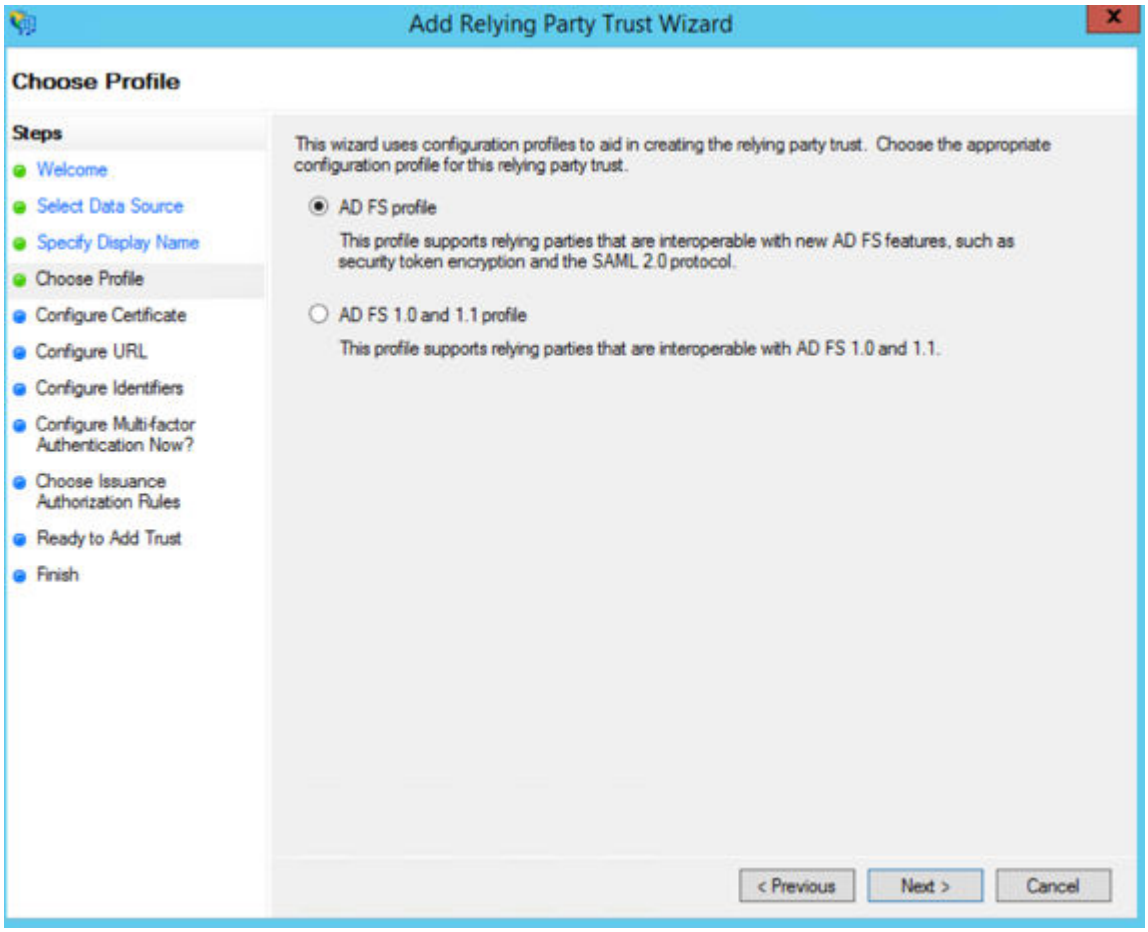

7. When you have a separate certificate for token encryption, browse to, select it, and click **Next**.

> **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**.

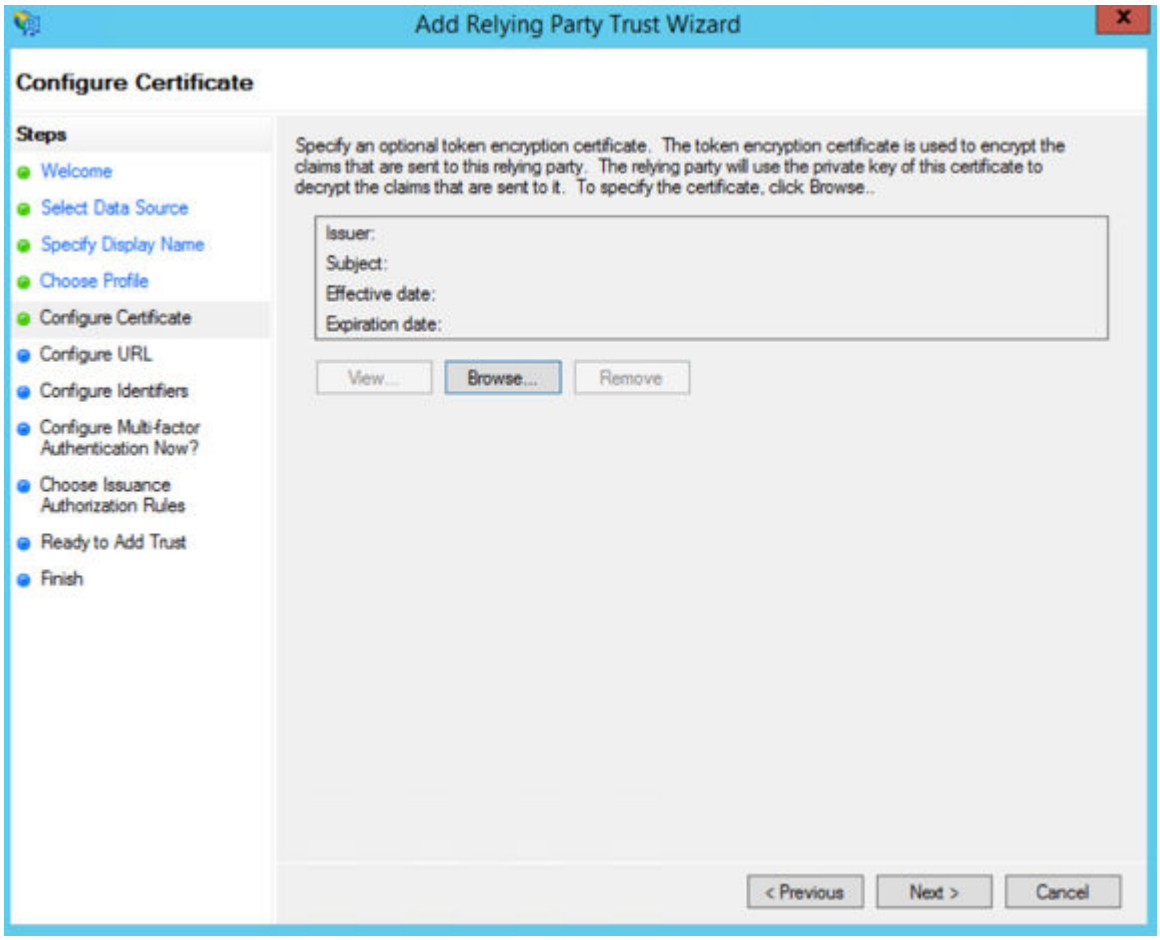

8. Select “Enable support for the SAML 2.0 WebSSO protocol.” In the Relying party SAML 2.0 SSO service URL field, add your “Assertion Consumer URL” (obtained from the PCE web console).

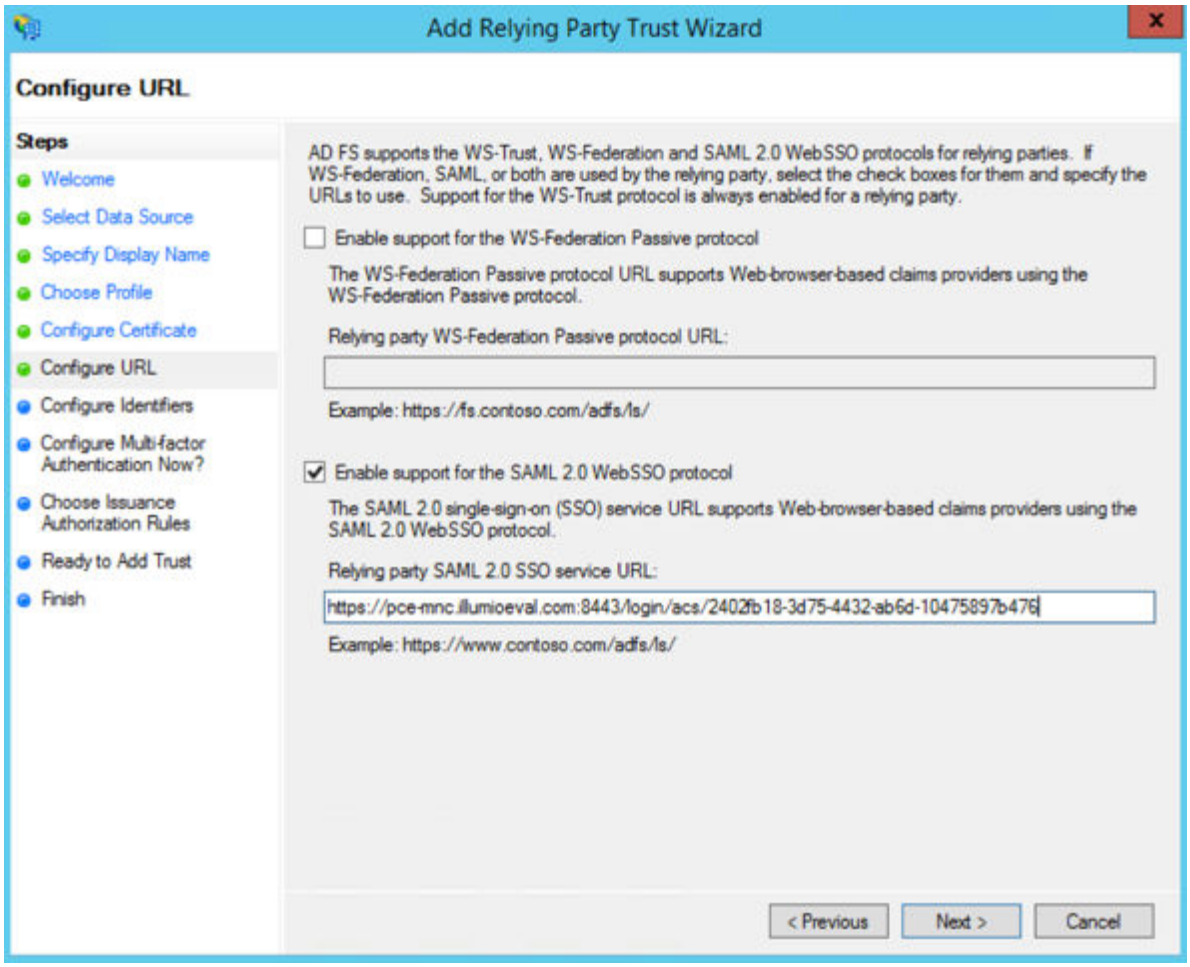

To locate the "Assertion Consumer URL," go to **Settings** > **Authentication** > **Information for Identity Provider** in the PCE web console:

9. 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**.

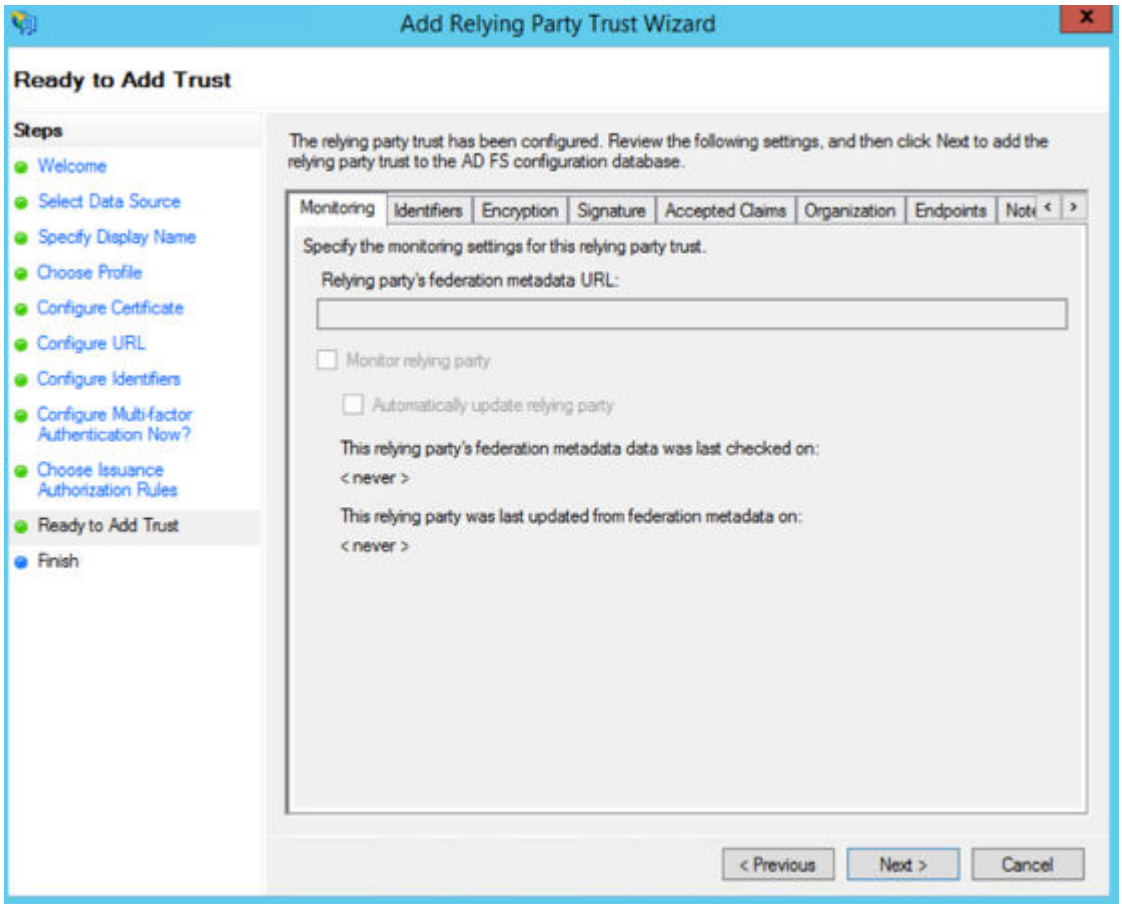

13. Leave the Open the Edit Claim Rules checkbox selected and click **Close**.

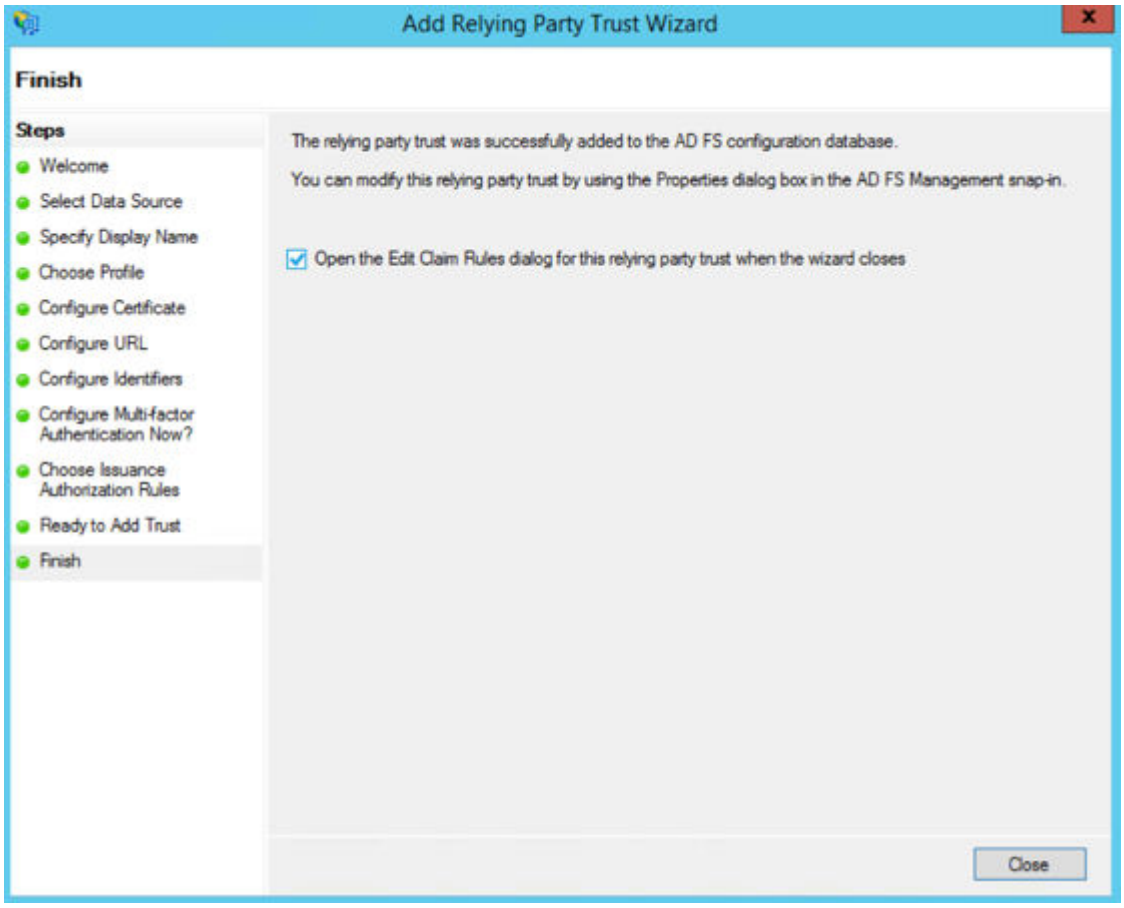

Create Claim Rules

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**.

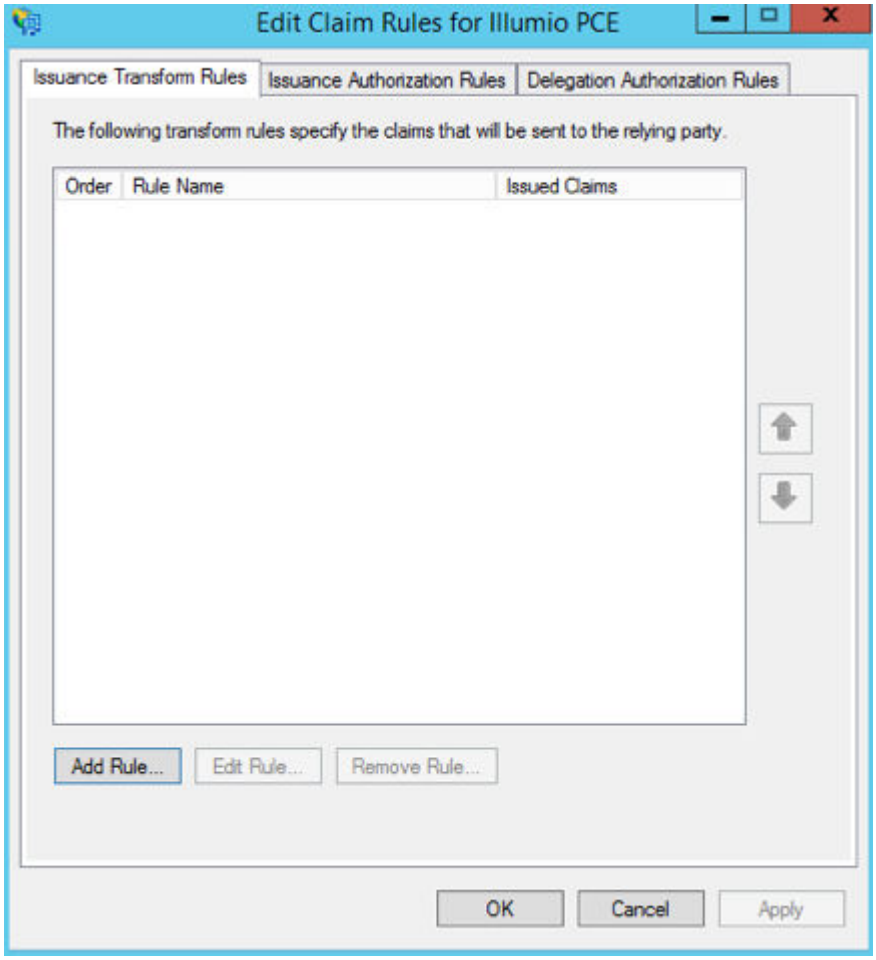

2. Under Select Rule Template, select “Send LDAP Attributes as Claims” and click **Next**.

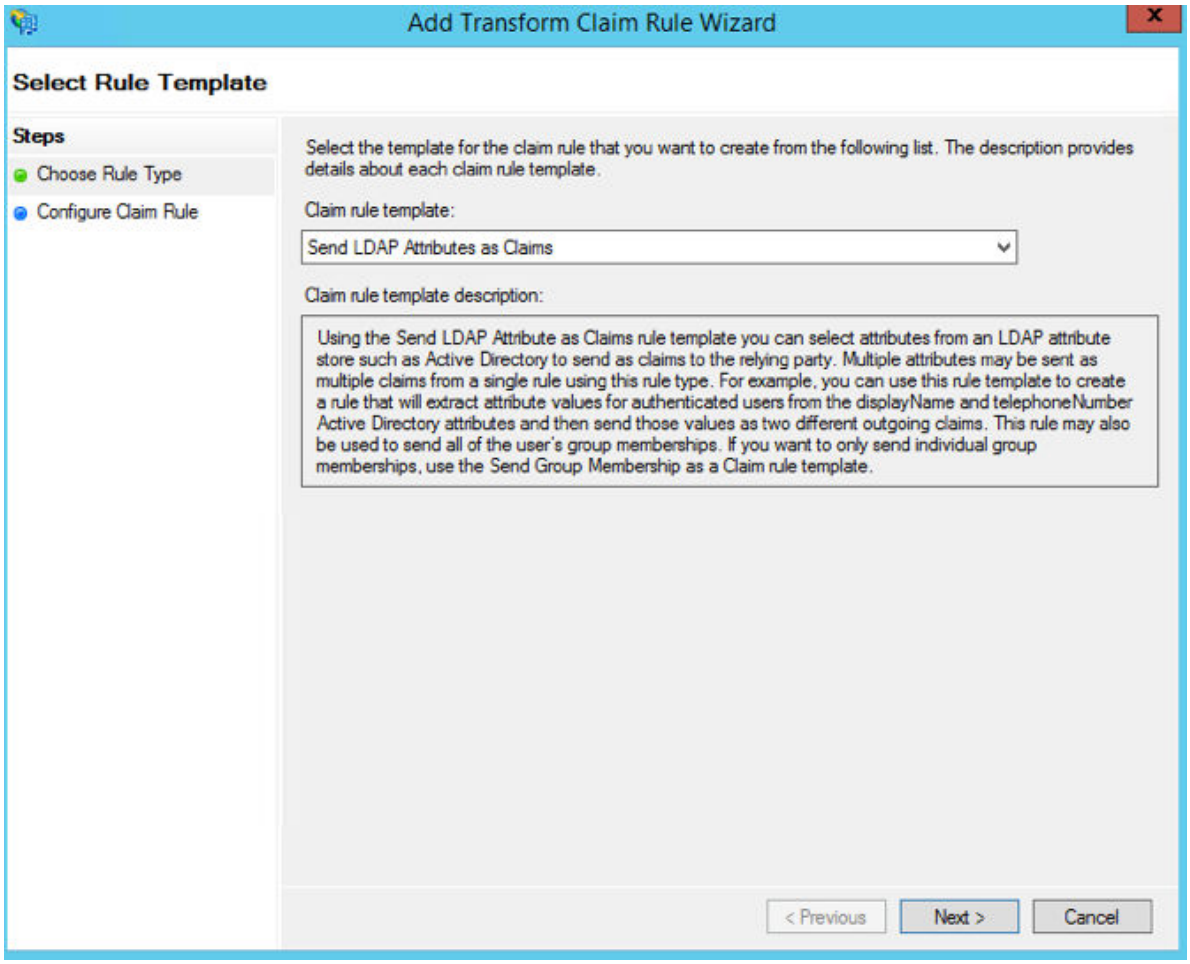

3. Name the Claim rule "Illumio Attributes" and select **Active Directory** as the Attribute store. Under the first attribute, select "User-Principal-Name" and "E-Mail Address" as the outgoing. Select "Surname" and type the custom field name of "User.LastName" in the outgoing field. Repeat the values for "Given-Name" and "User.FirstName" and click **Finish**.

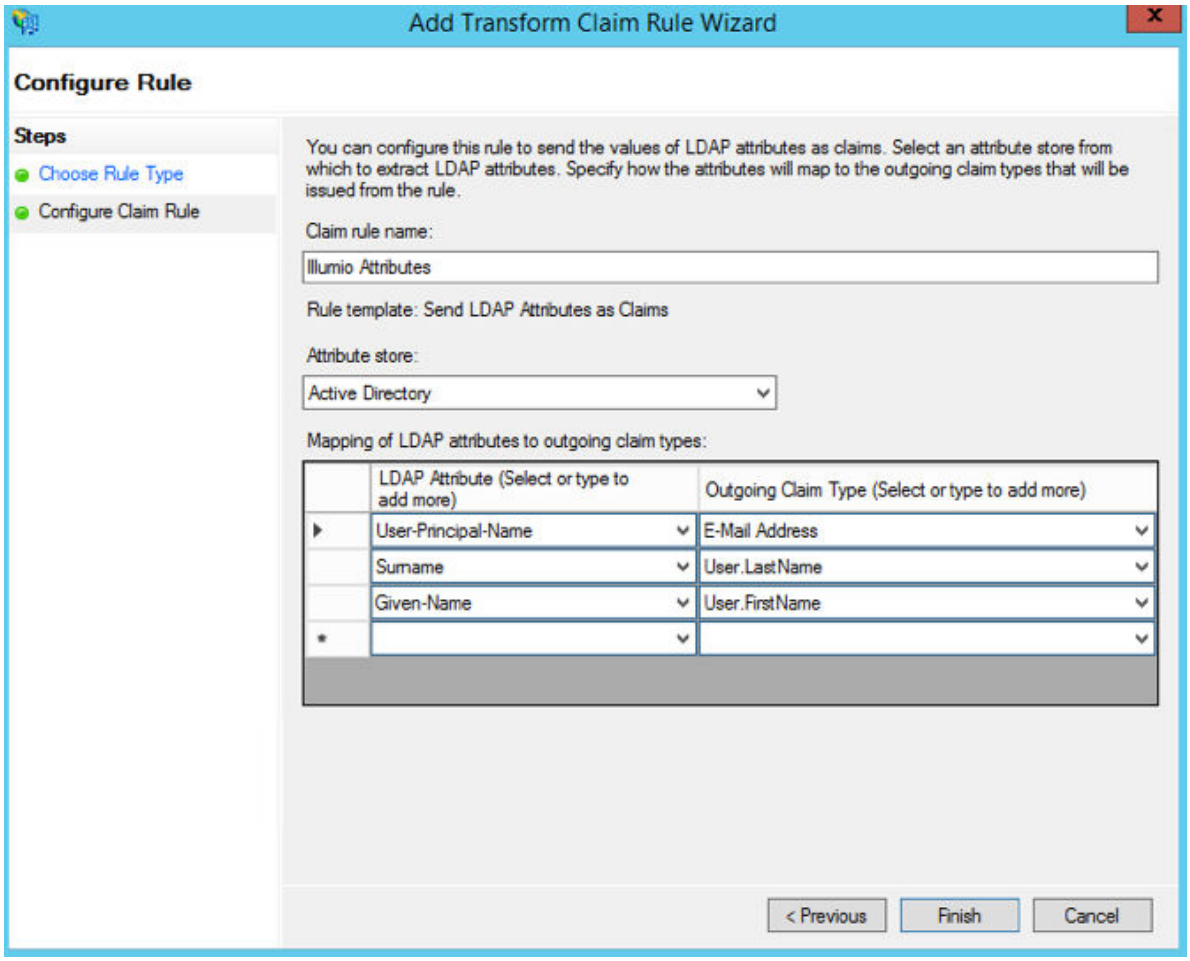

4. In the Edit Claim Rules dialog with your new rule added, click **Add Rule** to add the final rule.

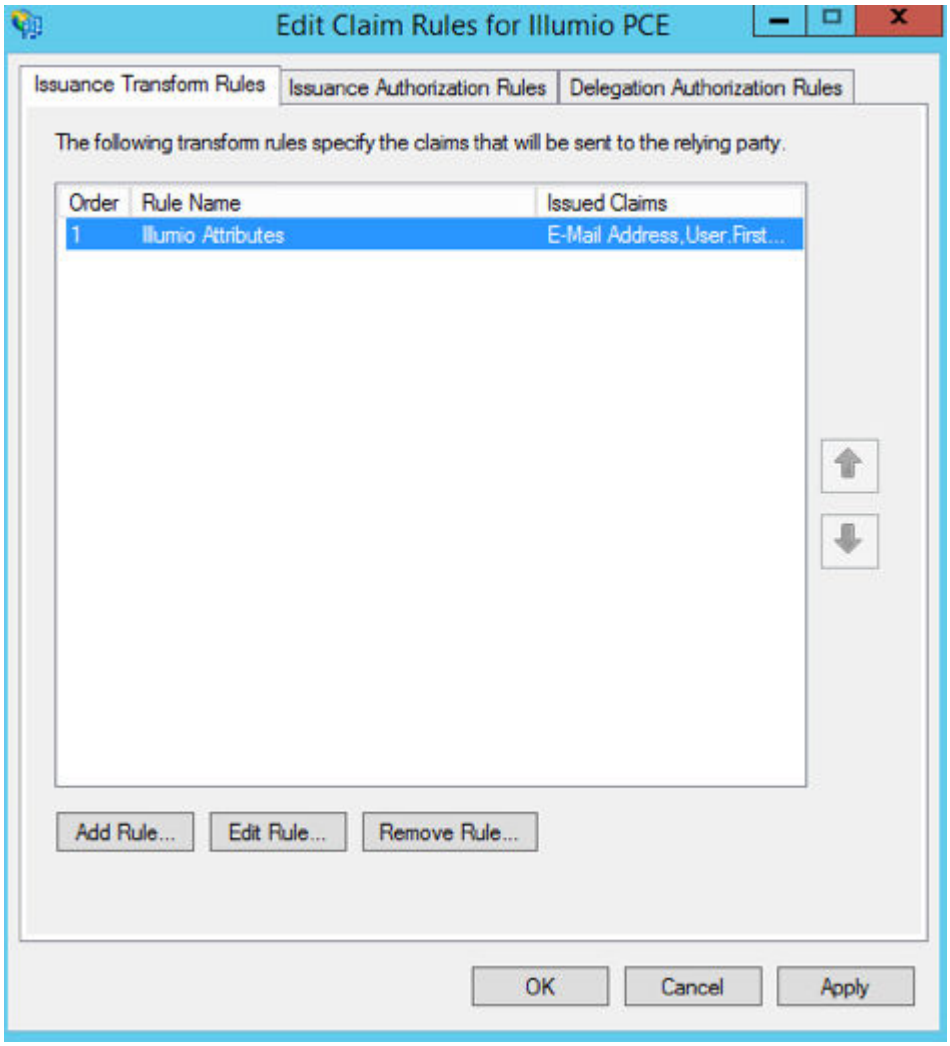

5. Under the Claim Rule Template, select “Transform and Incoming Claim” and click **Next**.

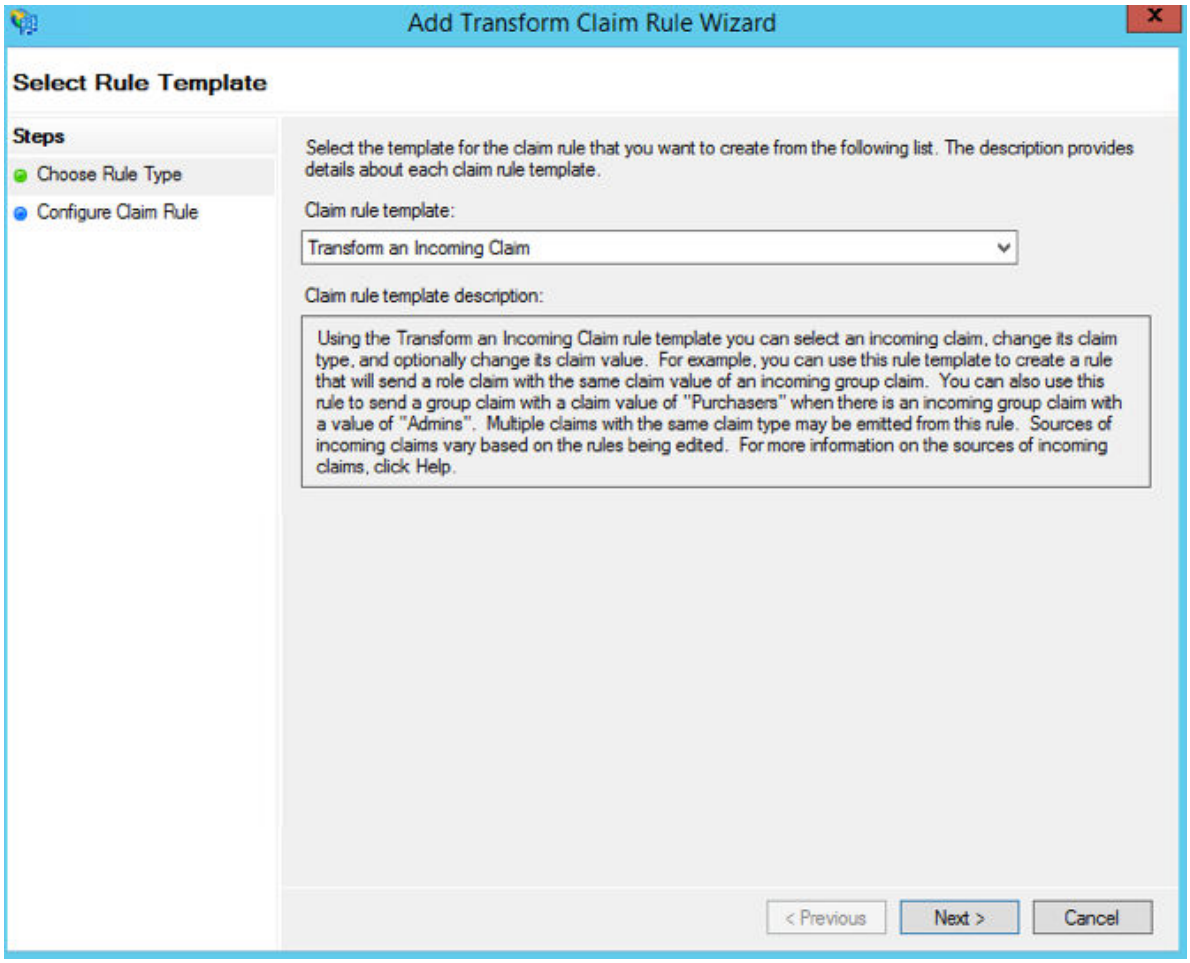

6. Name the rule “Email to NameID Transform” and change the incoming claim type to “E-Mail Address.” Set the Outgoing claim type to “Name ID” and the Outgoing name ID format to “Email” and click **Finish**.

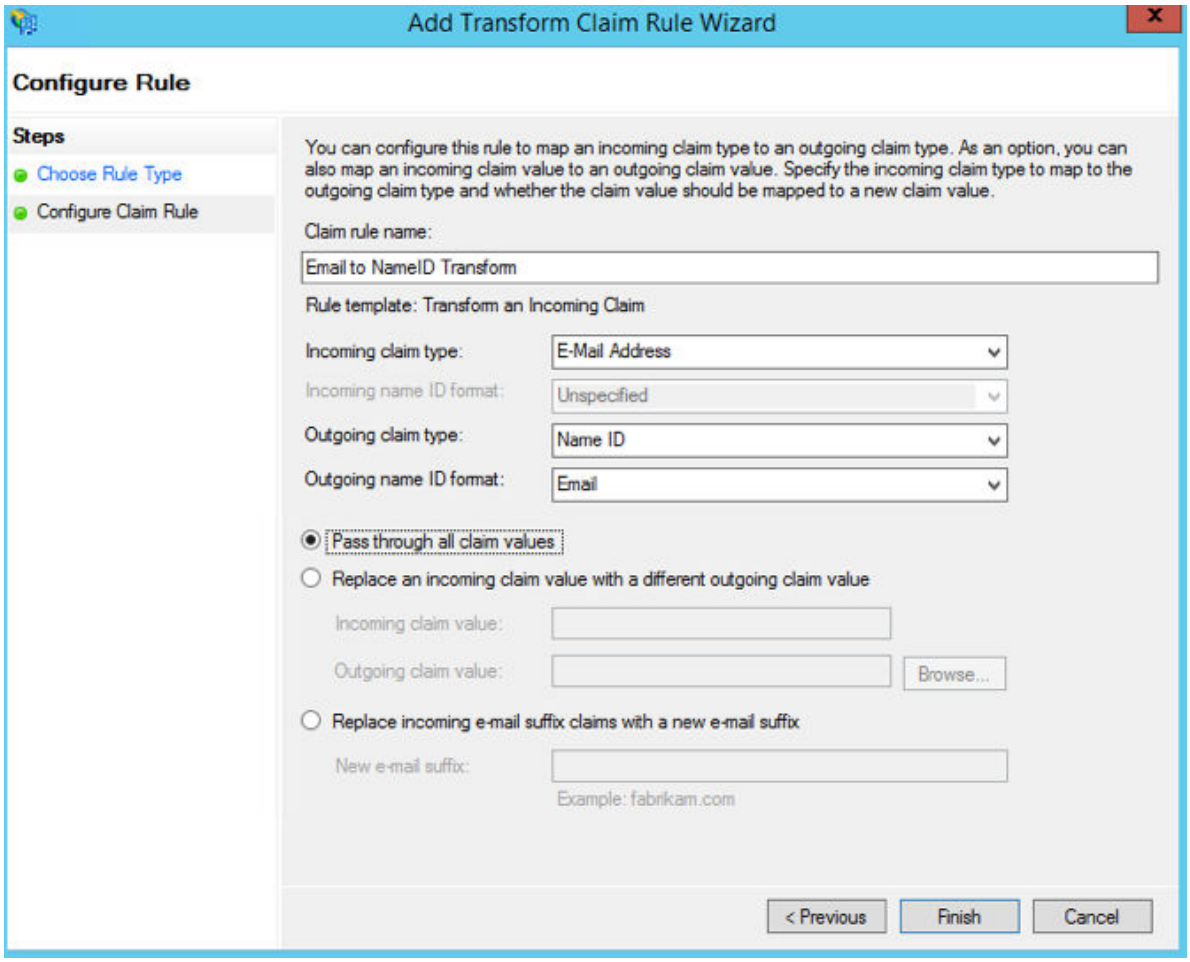

The Edit Claim Rules window opens.

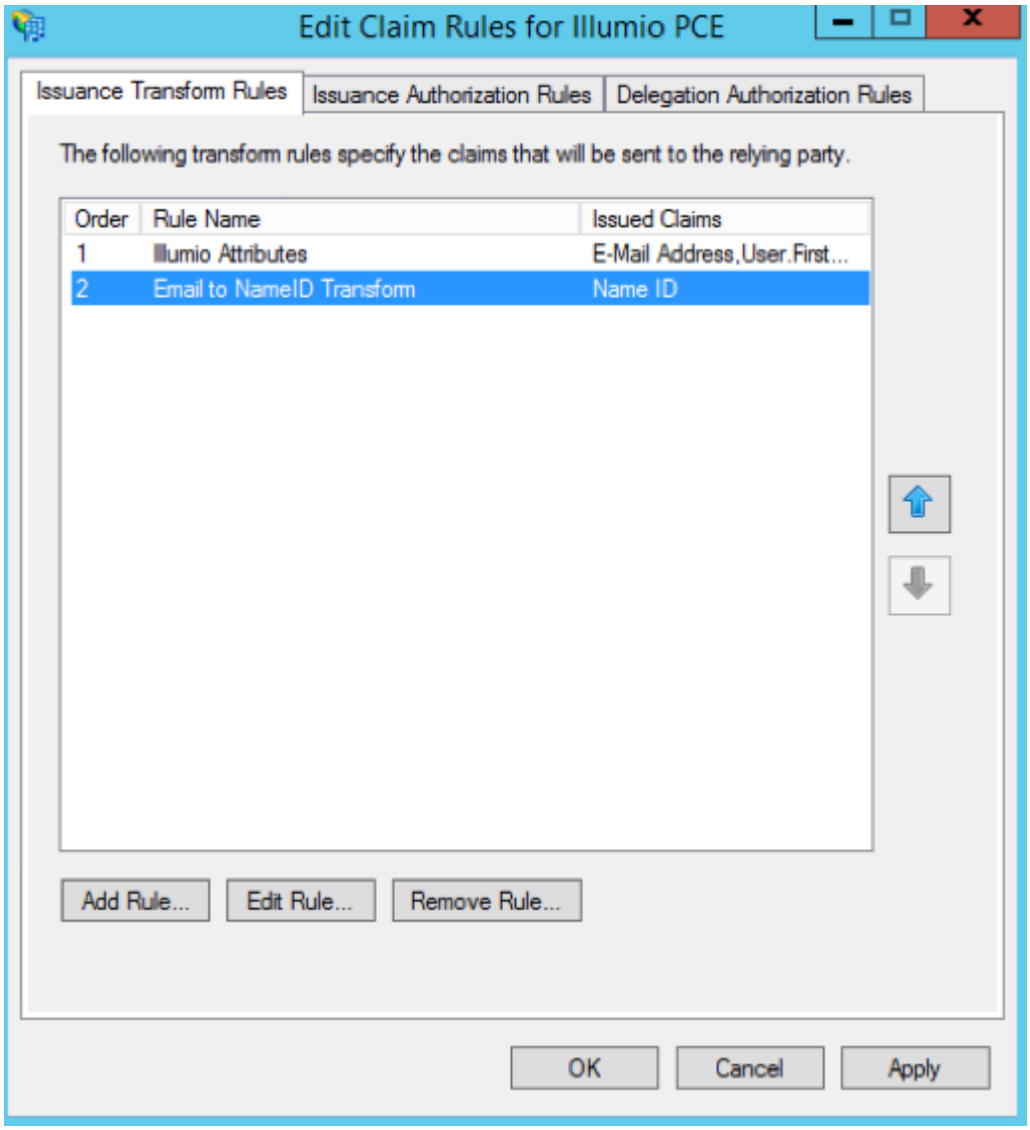

7. (Windows 2016 and Windows 2019) Skip to step 12.
   The Edit Claim Rules window has three tabs. You have already filled out the first tab. The other two tabs are not available in Windows 2016 or Windows 2019. Therefore, skip steps 8 - 11.
8. Select the Issuance Authorization Rules tab.
9. To allow 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.

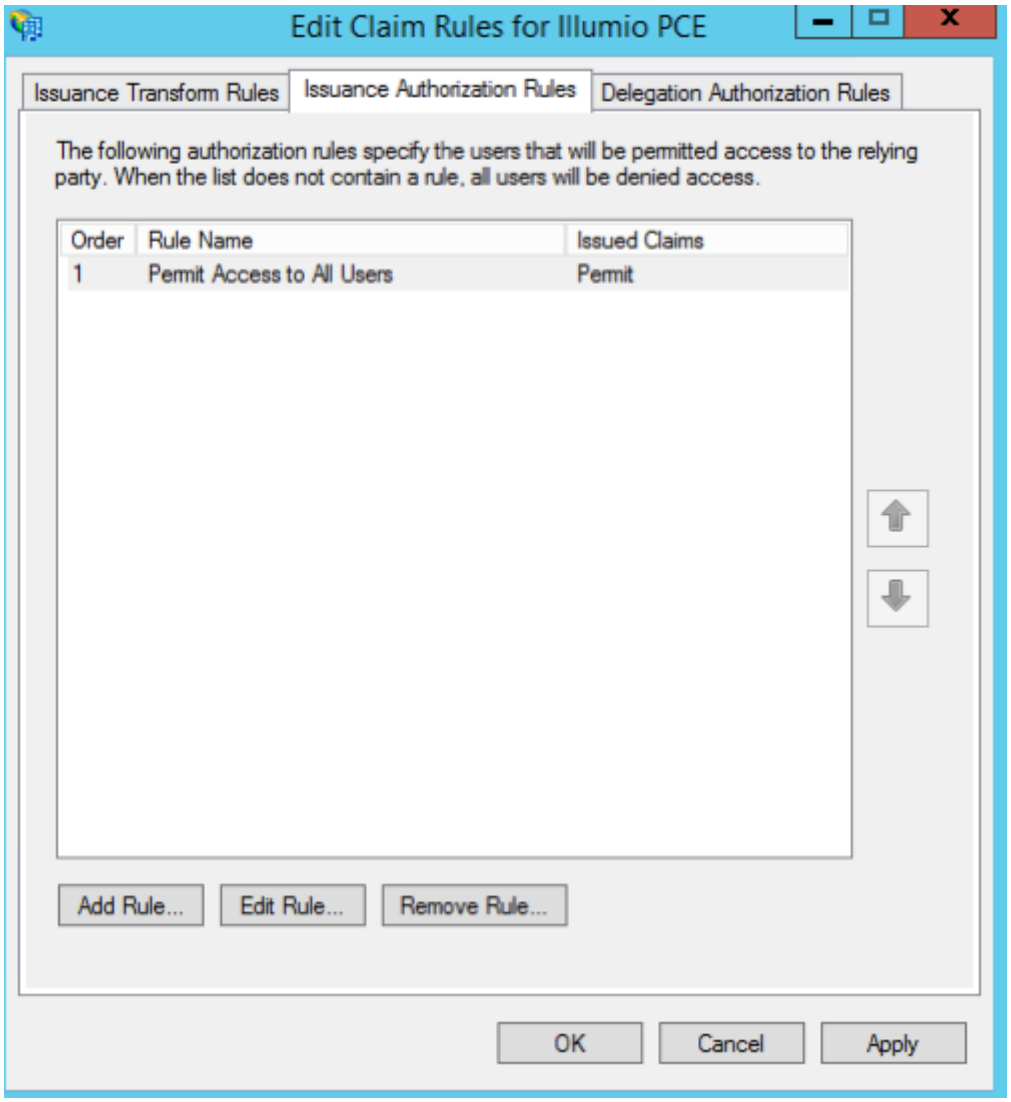

10 Select “Permit or Deny Users Based on an Incoming Claim” and click **Next**.

.

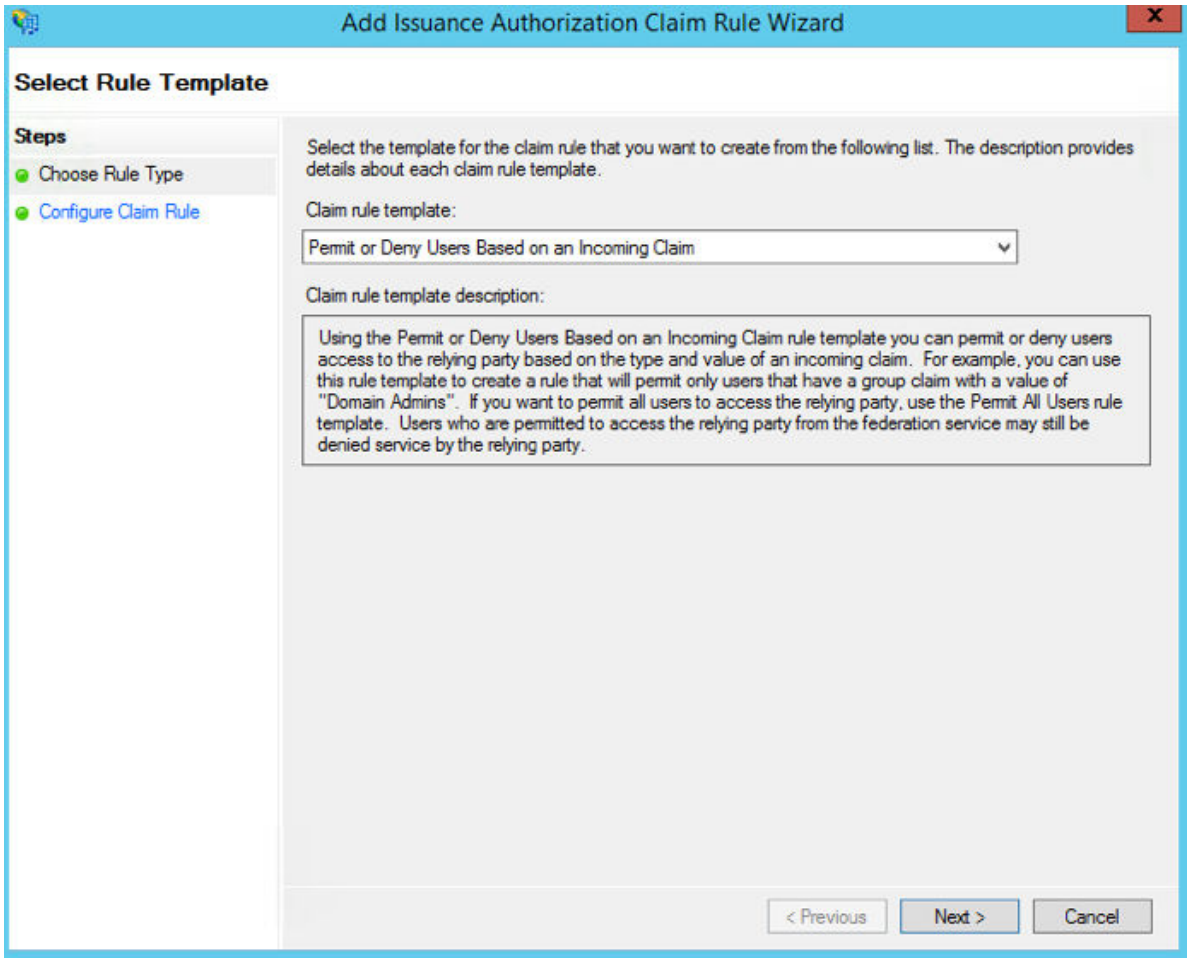

11. Name the rule “AD FS Users” and change the Incoming claim type to “Group SID” (you might have to scroll to find it). In Incoming claim value, browse to the group of users you want to give access. Make sure “Permit access” is selected and click **Finish**.

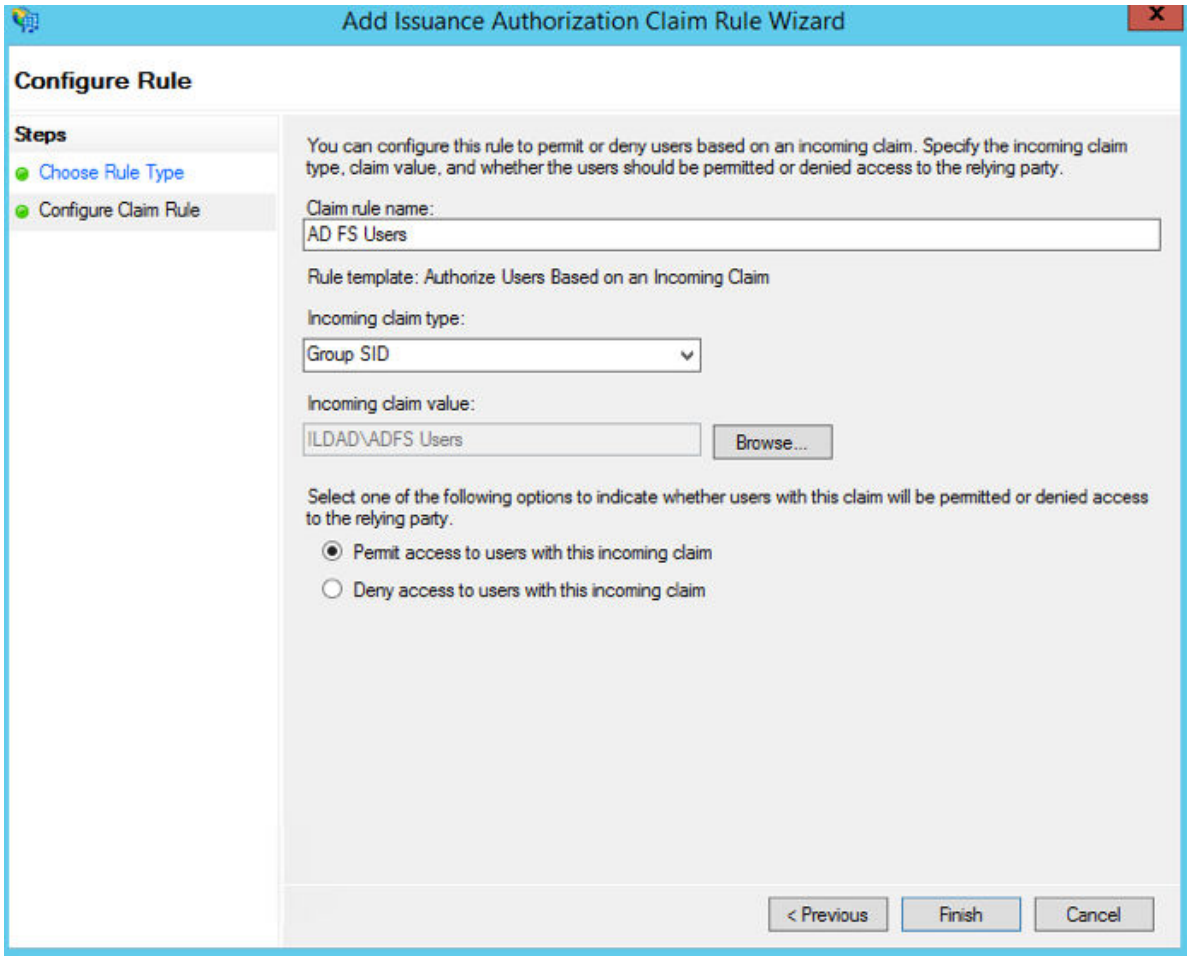

12. If you are using RBAC with groups, you need to create a Goup Claim Rule.

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**.

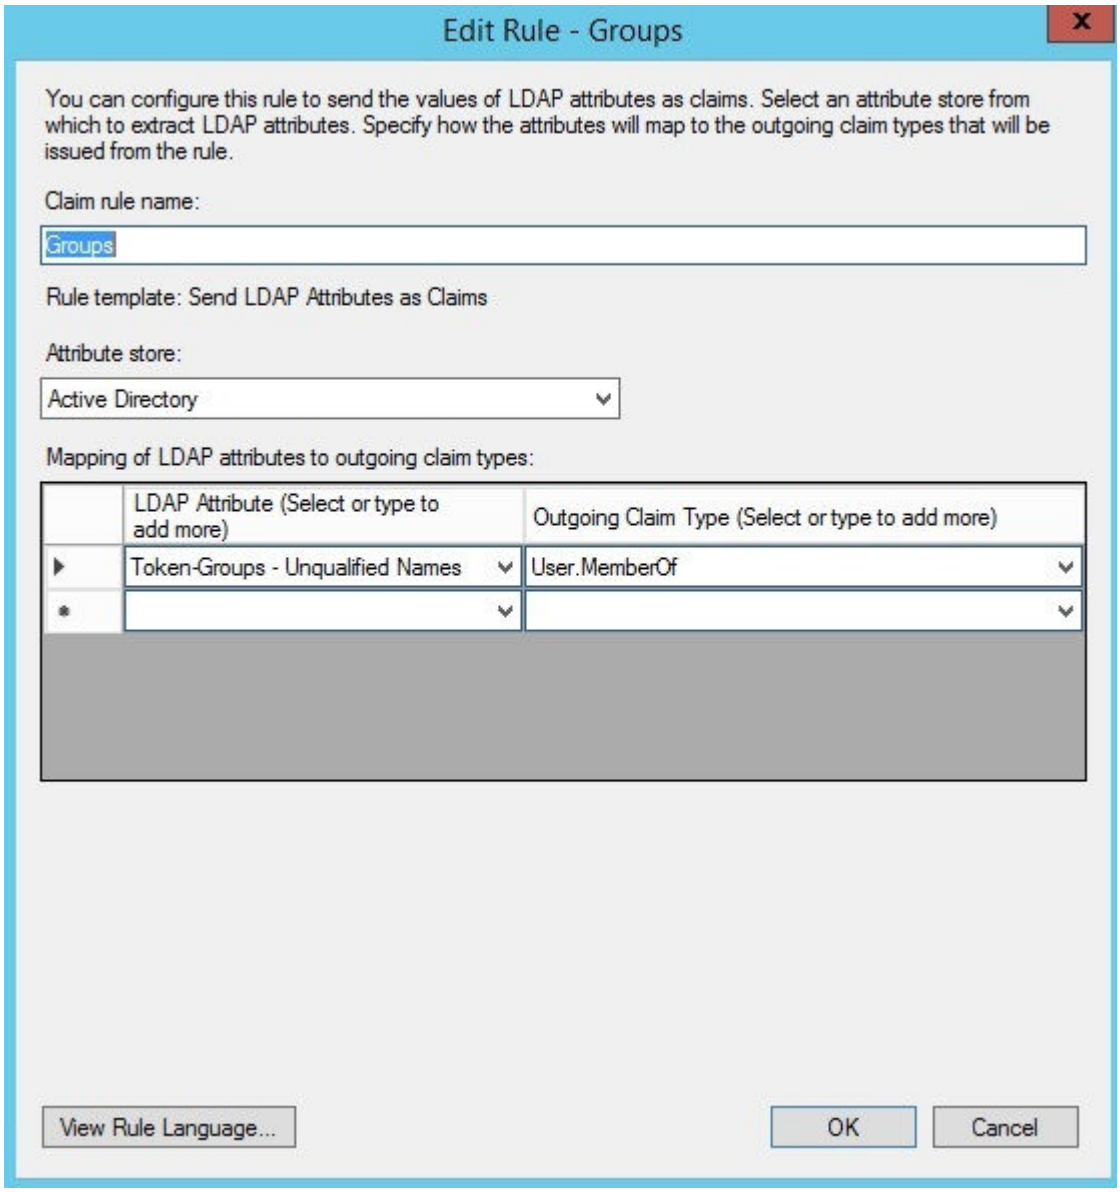

Obtain ADFS SSO Information for the PCE

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**.

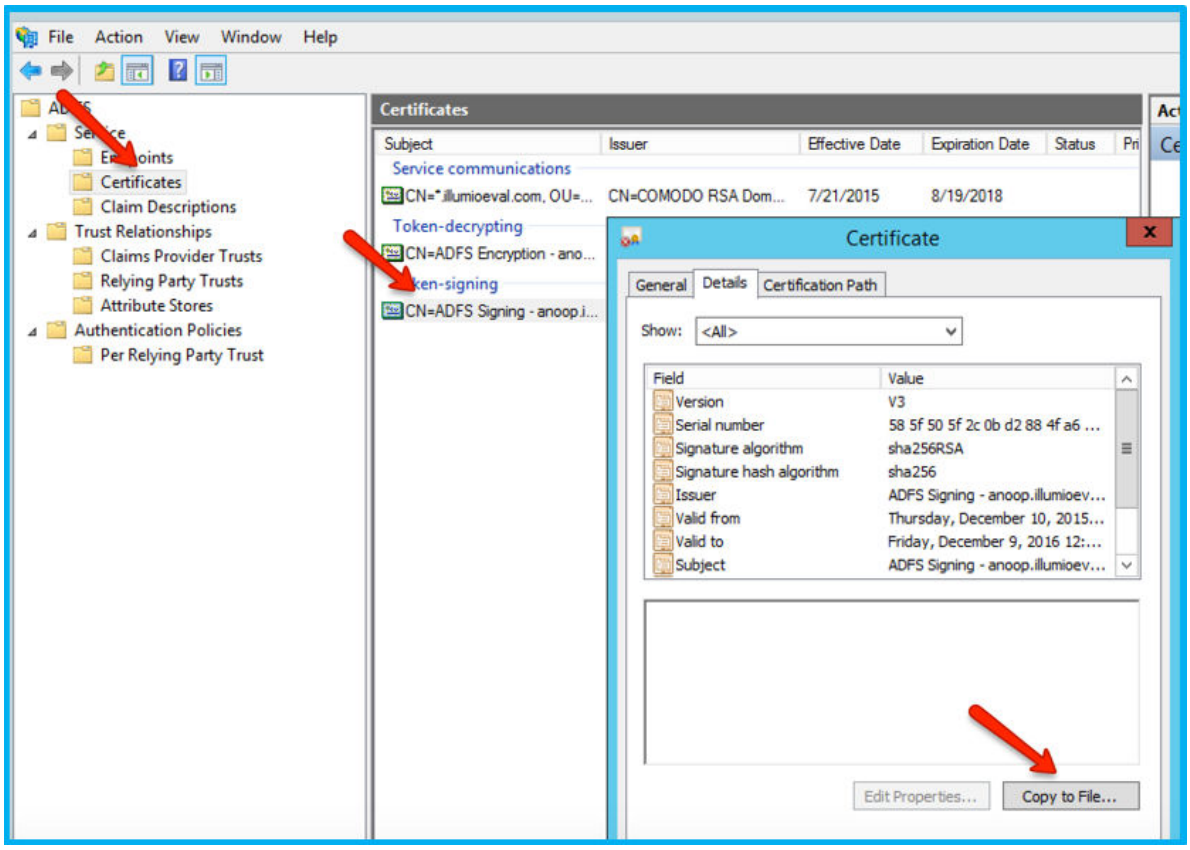

6. When the Certificate Export Wizard launches, click **Next**.
7. Verify that the "No - do not export the private key" option is selected and click **Next**.
8. Select Base 64 encoded binary X.509 (.cer) and click **Next**.
9. Select where you want to save the file, name the file, and click **Next**.
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.
12. 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`.

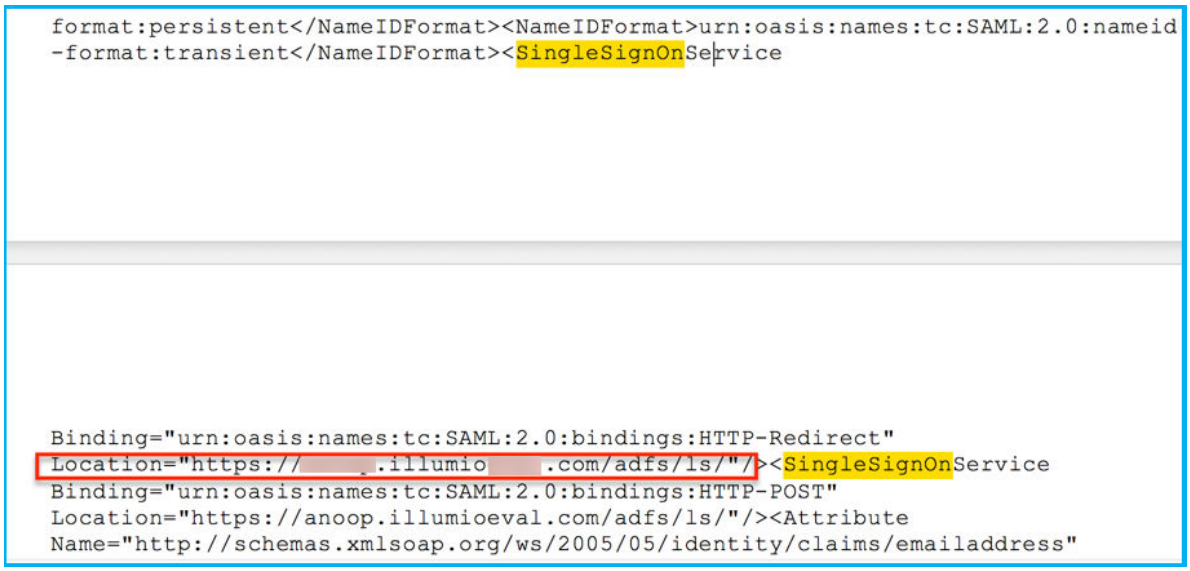

13. To find the **Logout Landing URL** for the PCE, you can use the login URL of the PCE (preferred):

    `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 [82].

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 Step2.

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.

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.

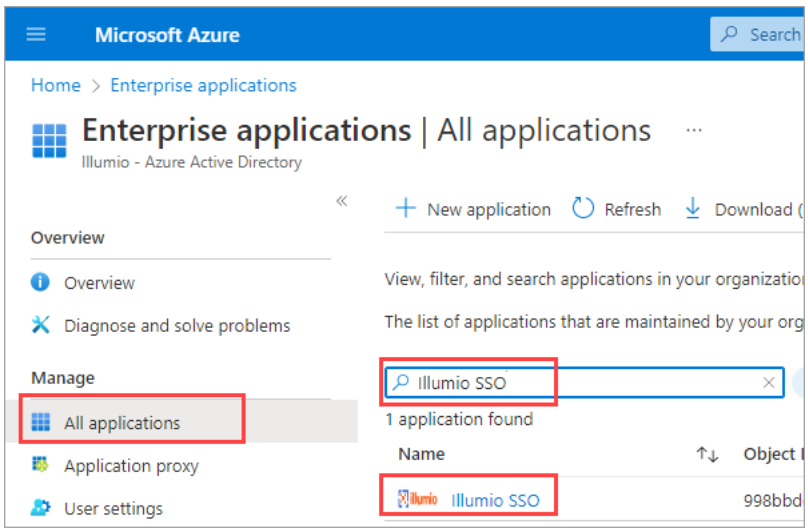

4. In the center of the page under **Getting Started**, click **Get started** on the **Set up single sign on** tile.

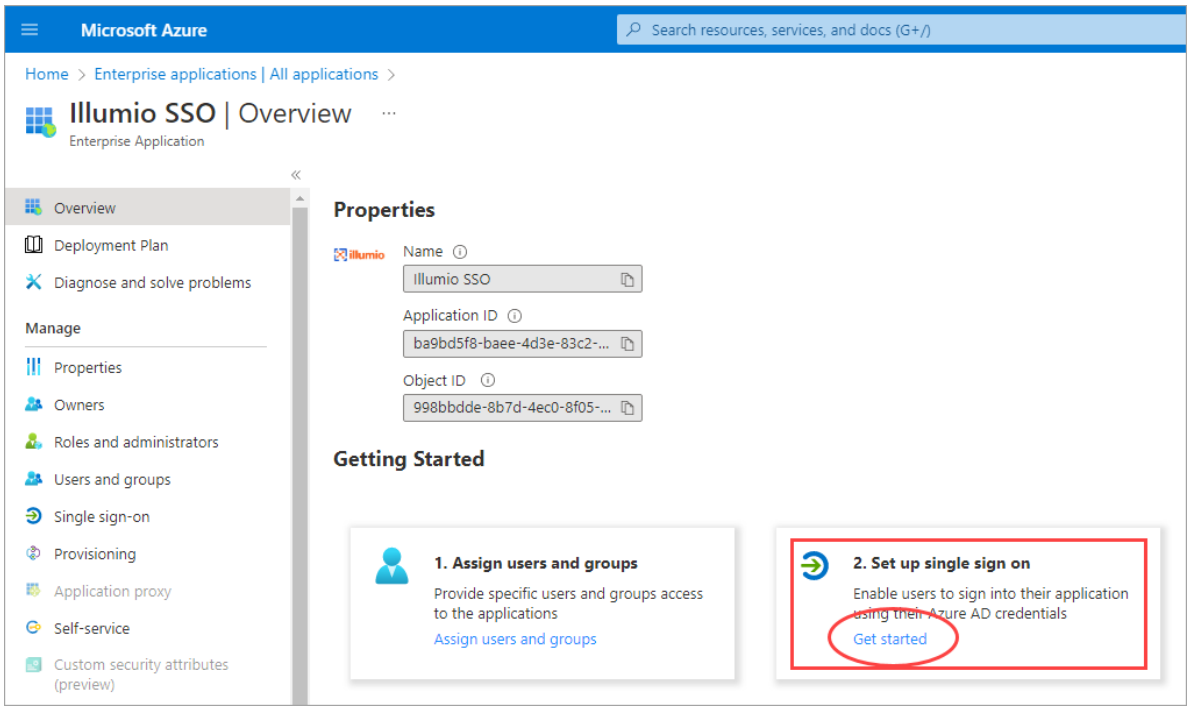

5. If prompted to select a single sign-on method, click **SAML**.
6. Configure Basic SAML:
   a. On the **Set up Single-Sign On with SAML** page **Basic SAML Configuration** tile, click **Edit**.

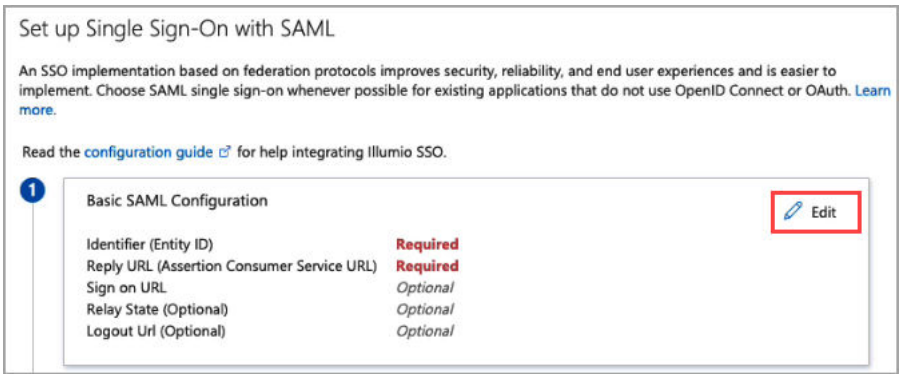

   b. On the **Basic SAML Configuration** panel that opens, populate the fields with the values you copied and preserved.
      - In the **Identifier (Entity ID)** field, paste the **Issuer URL** you copied from the Illumio PCE.
      - In the **Reply URL (Assertion Consumer Service URL** field, click **Add 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.

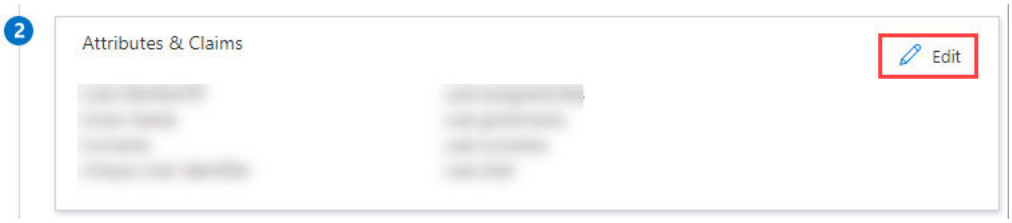

8. Under **Required claim**, update the **Claim name**:

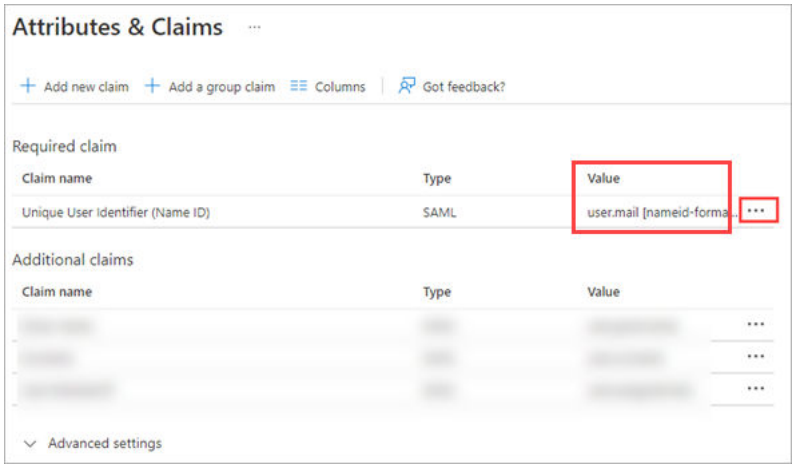

   a. Click the three dots.
   b. On the **Manage claim** page, click in the **Source attribute** field and select **user.mail** from the dropdown.
   c. Click **Save**.

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**.

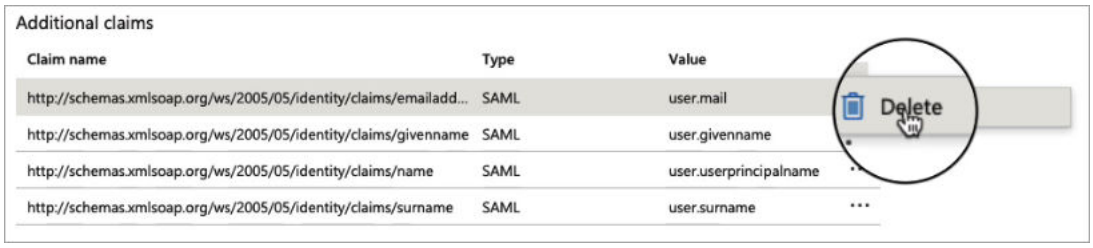

10 Click **Add new claim** and add three new claims:

.

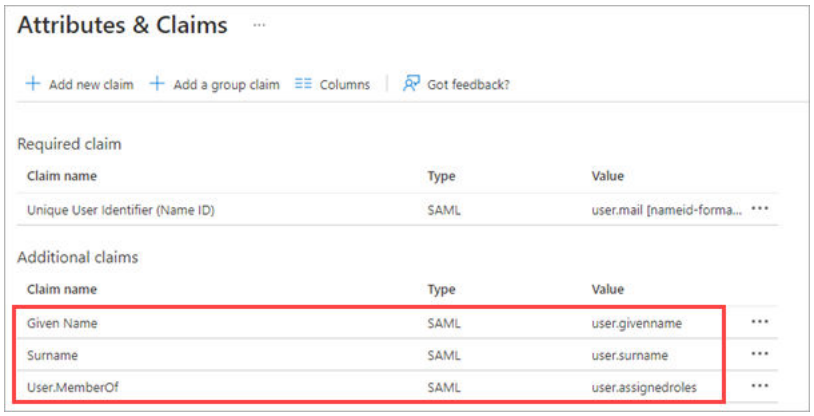

Given Name
Surname
User.MemberOf

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.

1. On the **SAML Certificates** tile, click **Download** for the **Certificate (Base64)** certificate and save the certificate to your computer.

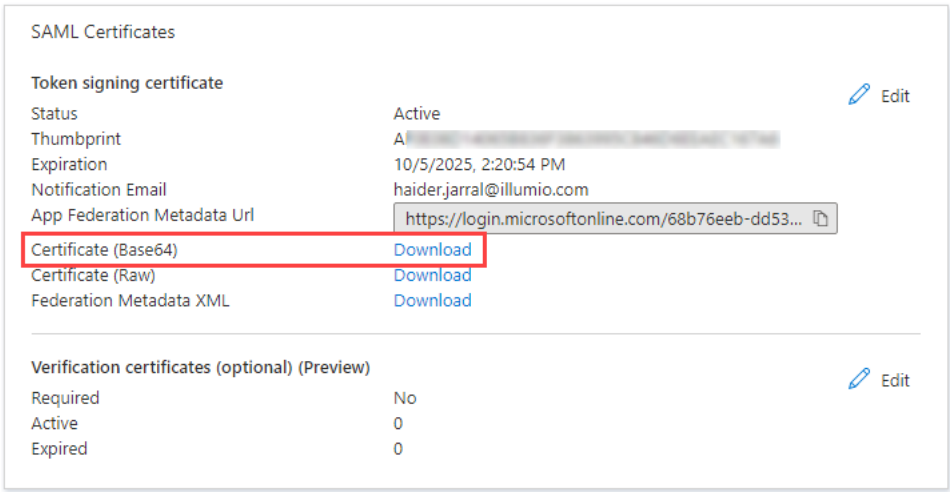

2. On the **Set up Illumio SSO** tile, copy and preserve the following URLs that you'll later paste into the Illumio PCE SAML setup.

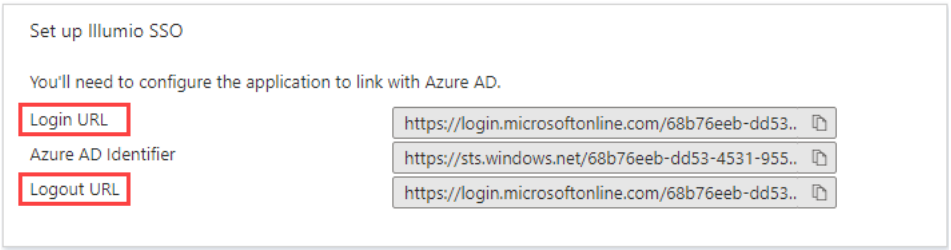

STEP 3: Create and assign a test user in Azure AD

1. In Azure, go to **Azure Active Directory**.
2. In the left pane, click **Users** and then **All users**.
3. Click **+ New user**.
4. In the **User** properties:
   a. In the **Name** field, enter a name (Example.Name).
   b. In the **User name** field, enter the user name in the form of an email address (Example.Name@exampledomain.com).
   c. Select **Show password**, and then make a note of the value that appears in the **Password** box.
5. Click **Create**.
6. Go to **Home > Azure Active Directory**.
7. Under **Overview > Manage**, click **Enterprise applications > All applications**.
8. Search for and click the **Illumio SSO** app.

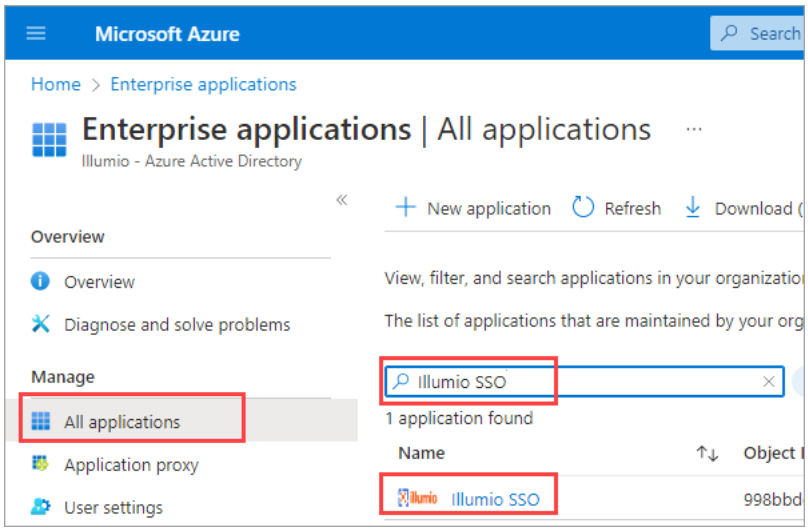

9. In the left pane under **Manage**, click **Users and Groups**.
10 Click **+ Add user/group**.
.
11. On the **Add Assignment** page, click **Users and groups**.
12. In the **Users and groups** panel that opens, click the user you created in a previous step (Example.Name).
13. Click **Select**.
14. On the **Add Assignment** page under **Select a role**, click one of the roles you created in a previous step.
15. Click **Assign**.

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:

- 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 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.

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**.

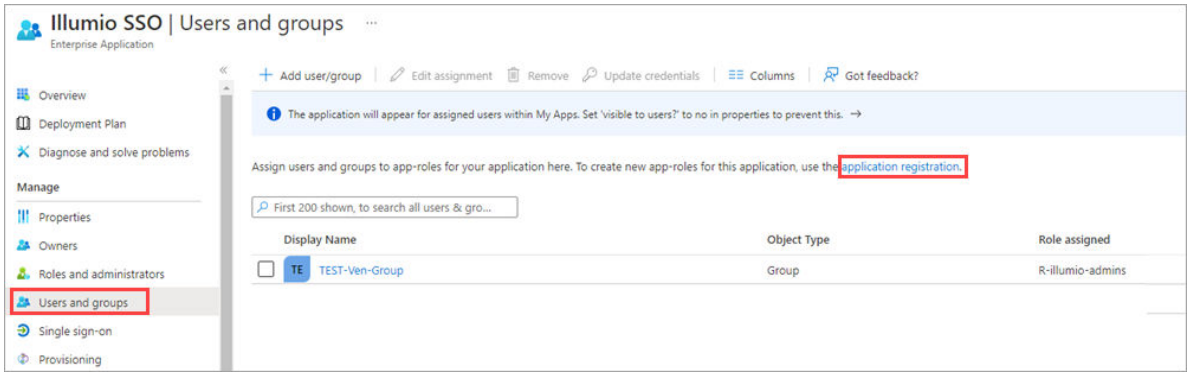

2. Create the roles you want by clicking **+ Create app role** and entering the required information for each role:
   - **Display name:** For example, enter one of the Global Roles that appear in the PCE Web Console.
   - **Value:** This must match the name you'll enter in the **Add External Groups** dialog box.
   - **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 **mslam_access**.

   Note: You first need to disable the default app role before you can delete it.
   a. Click **mslam_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 **mslam_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:

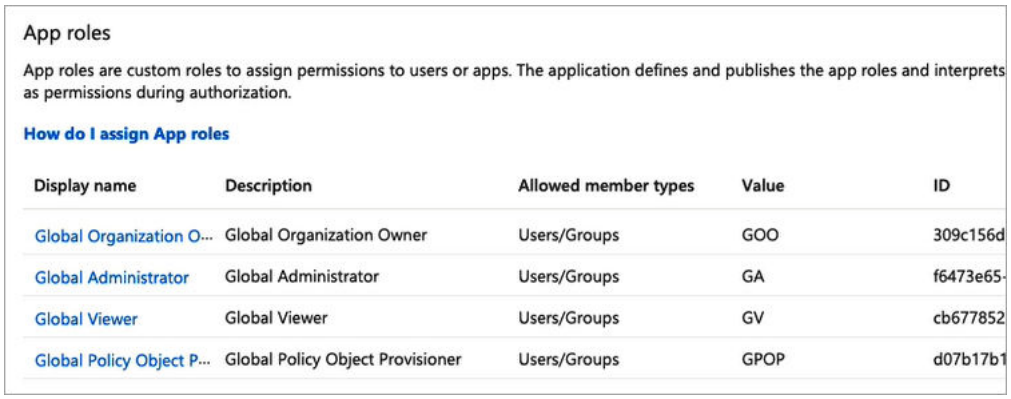

App roles

App roles are custom roles to assign permissions to users or apps. The application defines and publishes the app roles and interprets as permissions during authorization.

How do I assign App roles

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.

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.

- Click **Add**.

4. Repeat for additional groups.

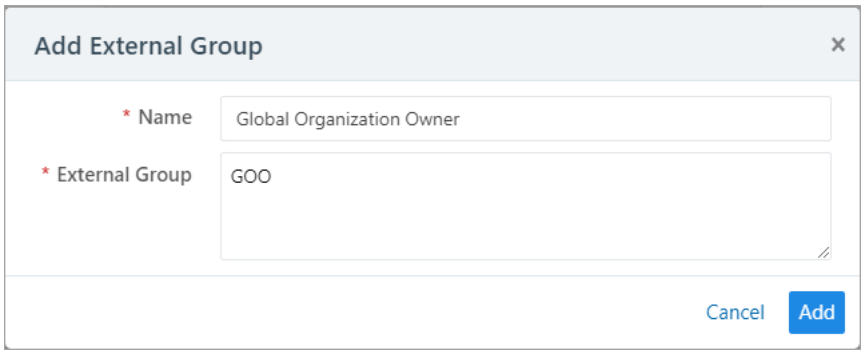

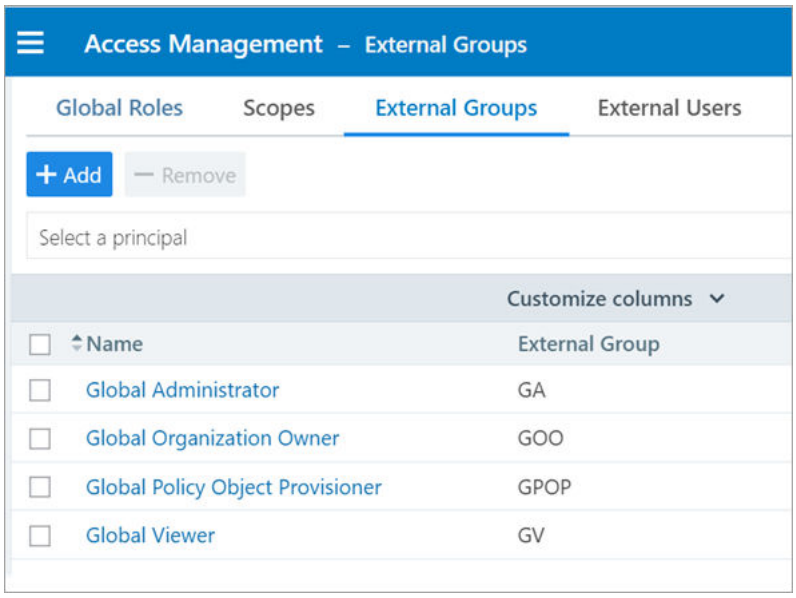

5. Click to open a group you created in the above step.
6. Click **Add Role > Add Global Role** or **Add Scoped Role**.
7. In the **Access Wizard**, select the appropriate **Role** and then click **Grant Access**.
8. Repeat for additional groups.

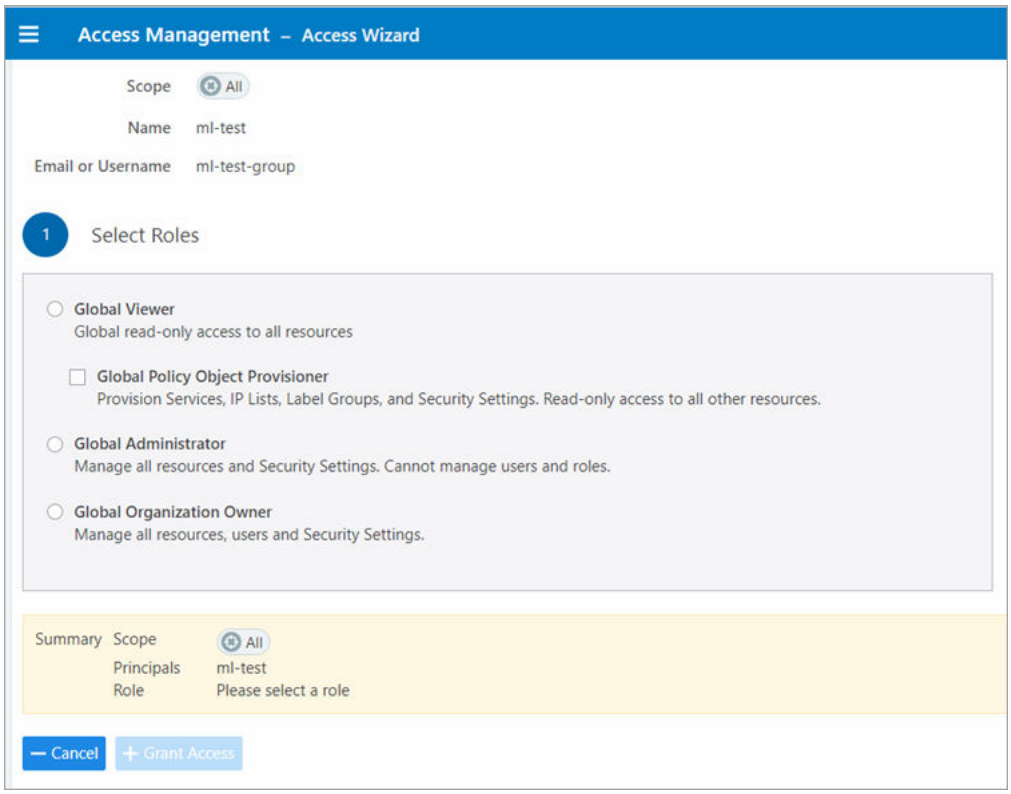

STEP 8: Turn on SAML authentication in the PCE Web Console

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**.

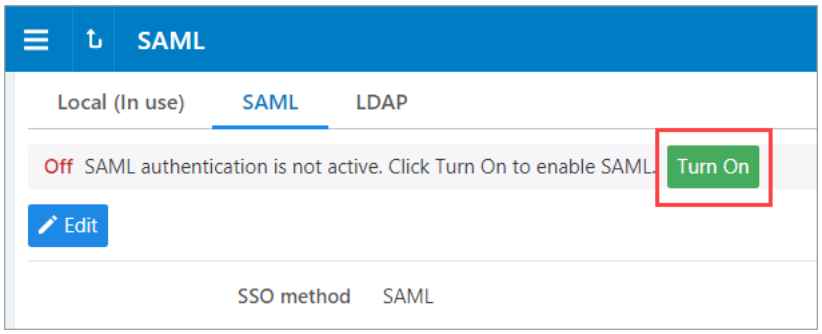

STEP 9: Test SSO

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**.

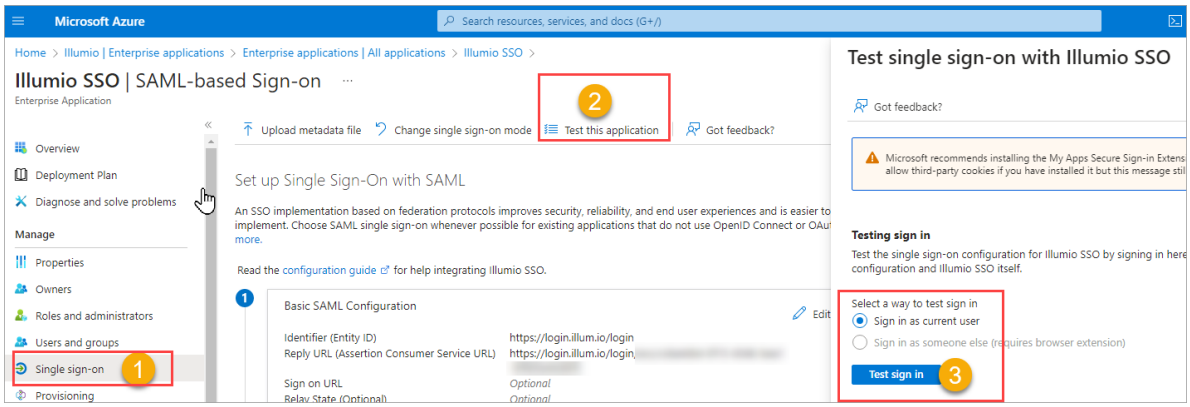

4. If the test is successful, the PCE will log you in to the **Welcome to Illumio** screen.

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 [18].)
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.

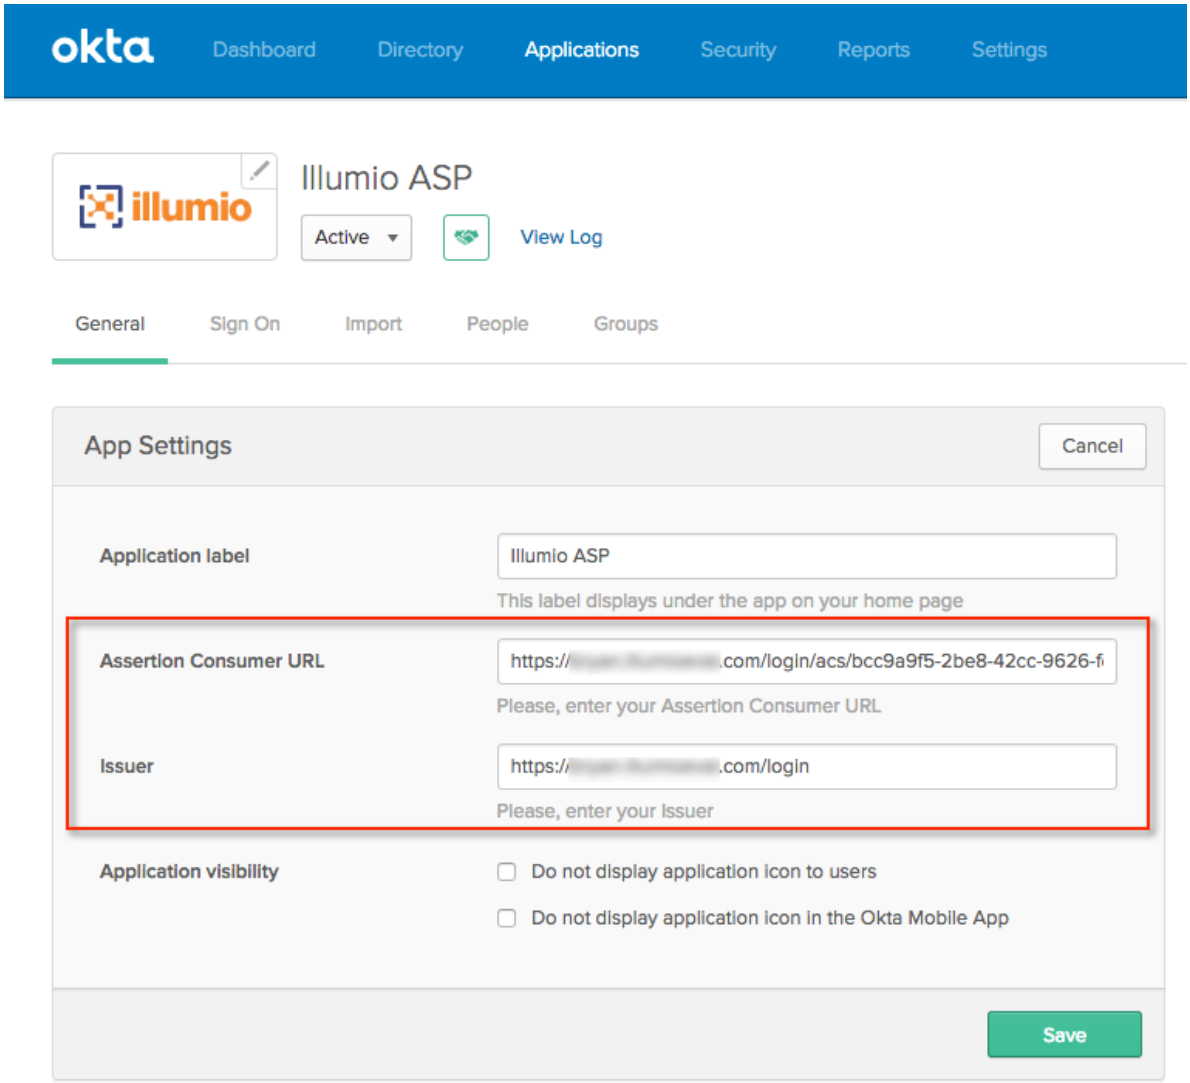

12. Click **Save**.
Your PCE is now configured to use Okta SSO for authenticating users with the PCE.

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 [18].)
6. In the Information for Identity Provider section, make note of the following fields:
   - Issuer
   - Assertion Consumer URL
   - Logout URL

   You 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.

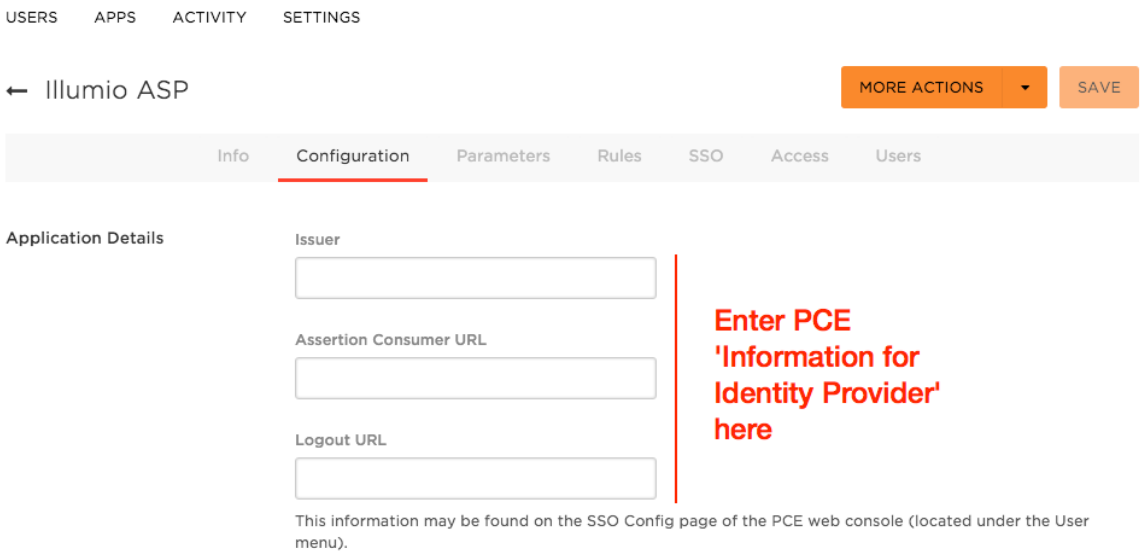

13. Click **Save**.
    Your PCE is now configured to use OneLogin SSO for authenticating users with the PCE.

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.

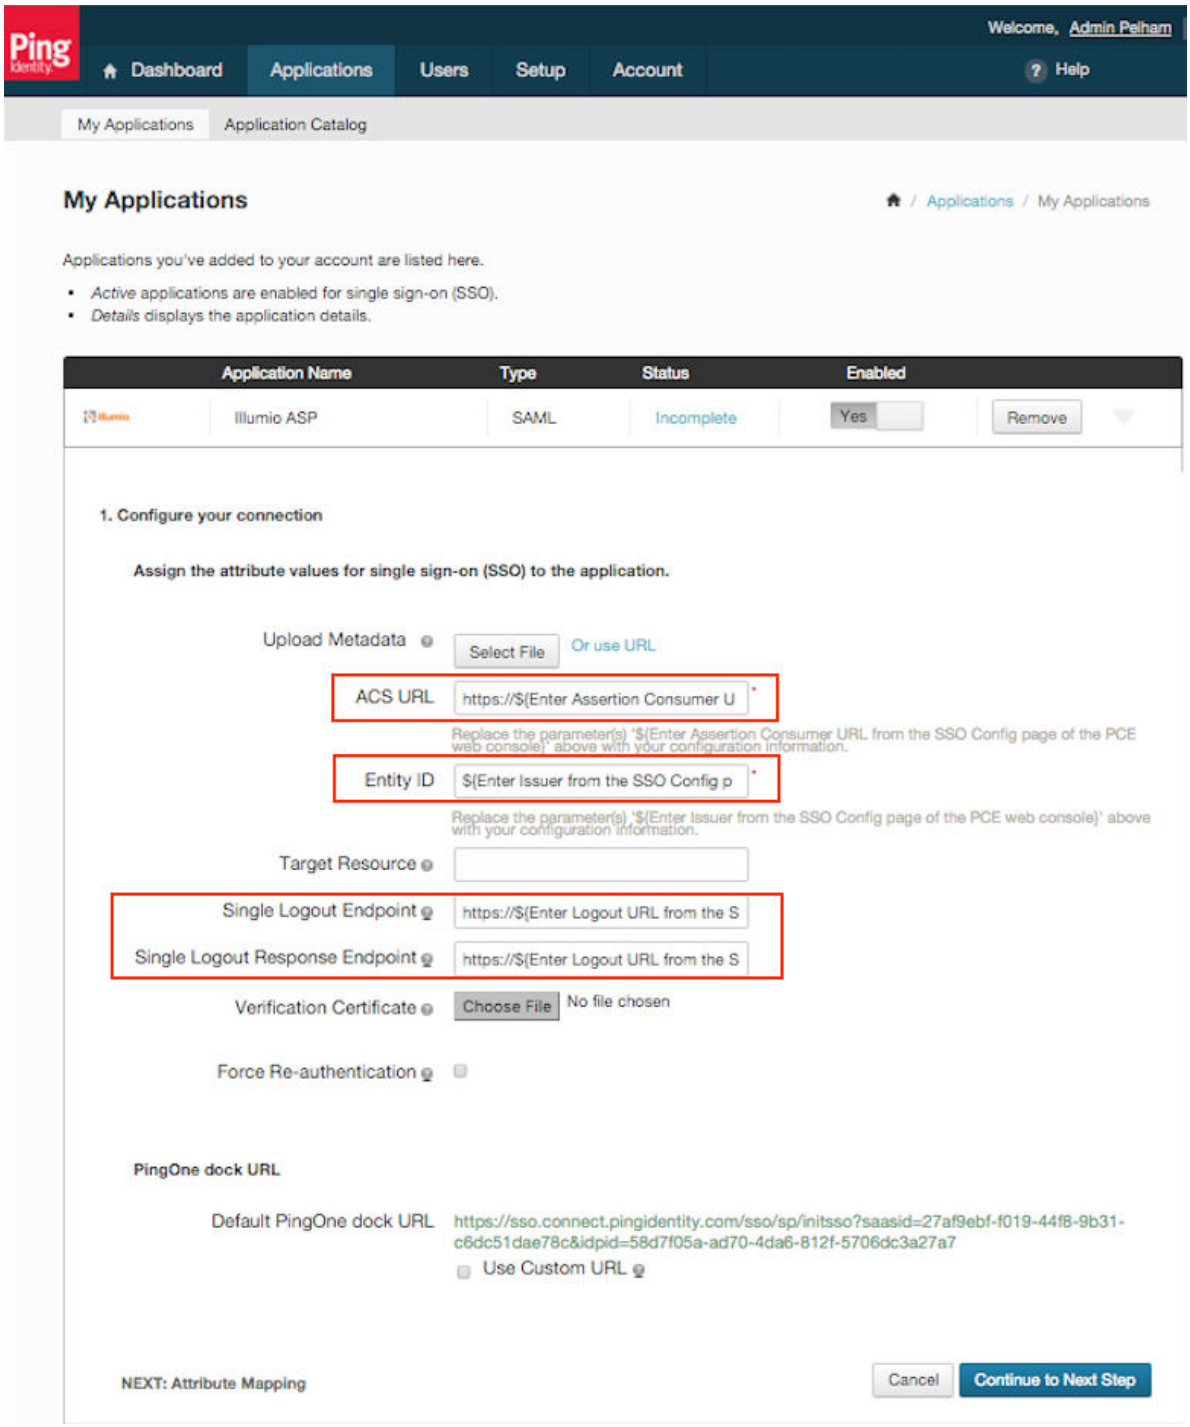

12. Click **Continue to Next Step**.
13. You will now configure the SAML_SUBJECT attribute mapping. Under Advanced Attribute Mapping, next to the Name ID Format to send to SP, select `urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress`.

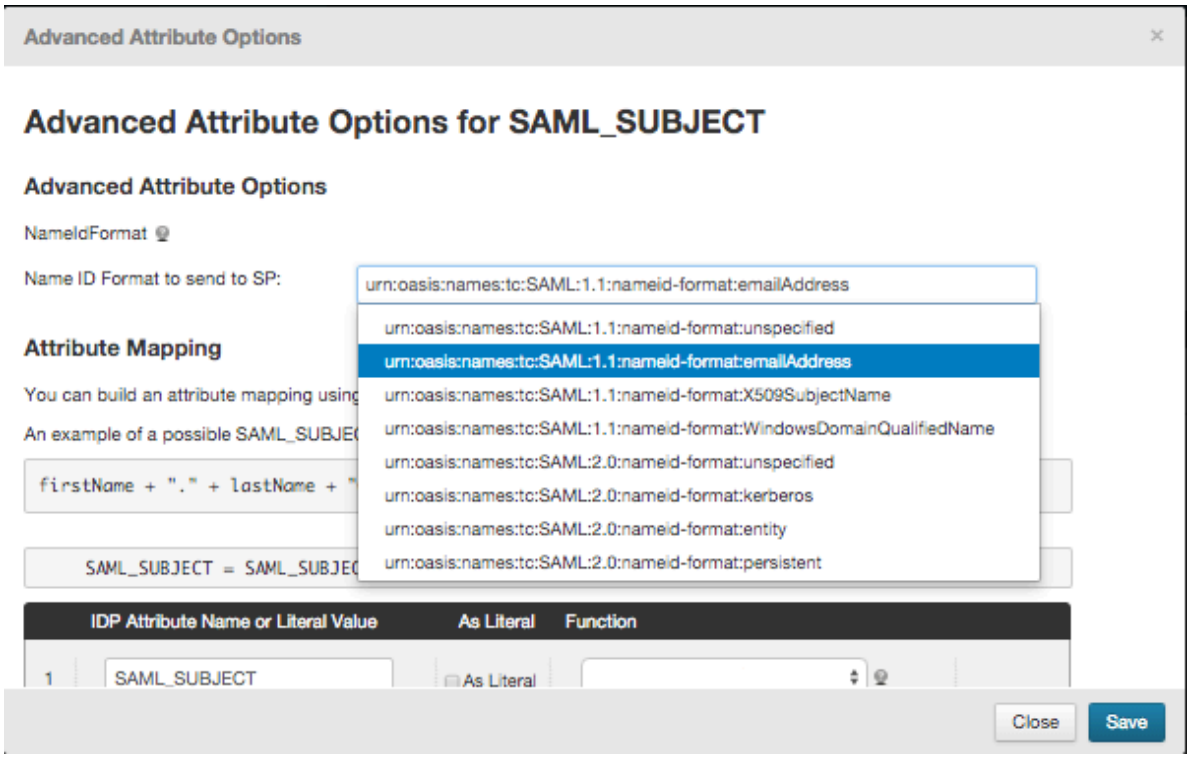

14. Click **Save**.

Your PCE is now configured to use Ping Identity SSO for authenticating users with the PCE.

Manage PCE Nodes and Clusters

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 PCE Health Monitoring [149].
  - 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 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 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 `postgresql.bak` is sometimes left behind in the `<postgresql directory>`, especially on the database master node. You can delete the directory `postgresql.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 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 --runlevel 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
`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
`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 iimnttttttt 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
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.

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.
2. Stop *all nodes* in your deployment:

   ```
   $ sudo -u ilo-pce illumio-pce-ctl stop
   ```
3. On *all nodes*, 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 *all nodes*, 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

To change the PCE FQDN:

- Backup the database and restore the database with the `change-fqdn` option.
- Configure `runtime_env` prior to the restore and make sure the web certificate has the new FQDN.

NOTE

Any VENs paired *after* PCE FQDN was changed will receive the masterconfig with the new PCE FQDN.

VENs paired to the PCE *before* the FQDN change will continue to point to the old PCE FQDN. The masterconfig of those VENs will be updated by a periodic job on the VEN or by manually restarting the VEN.

It is necessary that the old PCE FQDN still resolves to the PCE until those VENs can communicate to the PCE with the new FQDN.

Ideally, the old FQDN is used as a Subject Alternative name on the new certificate. This way, the VENs can still connect to the PCE and update the FQDN on its own configuration, which depends on the reason the FQDN is being changed.

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
```

NOTE

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.

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

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

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.
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.

   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
   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      192.0.2.242
collector_service                    192.0.2.241      192.0.2.241      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
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
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

- Install the PCE software and configure its runtime parameters.

  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 in PCE Installation and Upgrade Guide.

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.

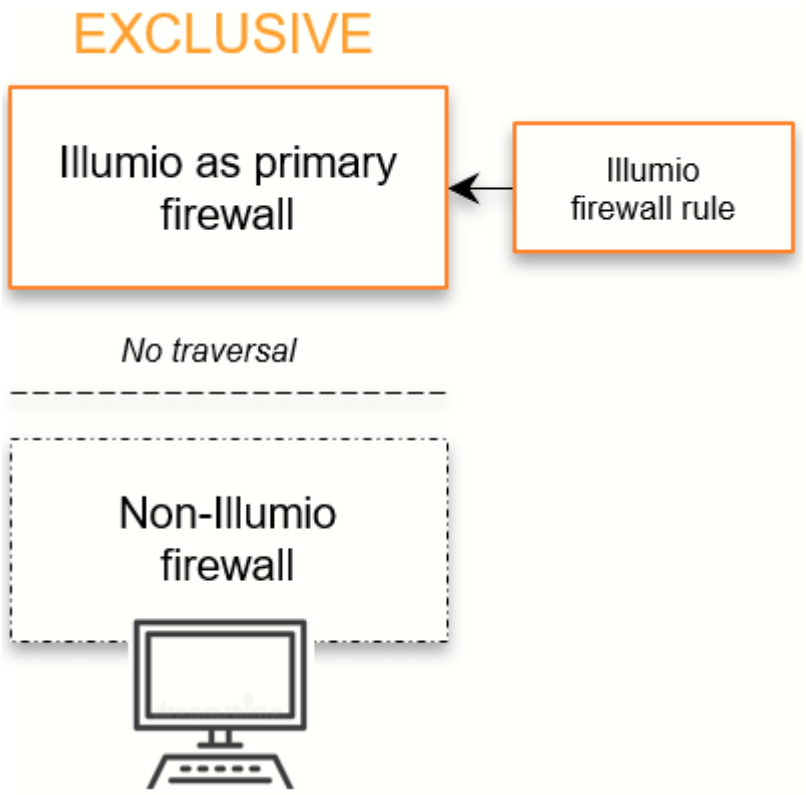

Coexistence Mode

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.

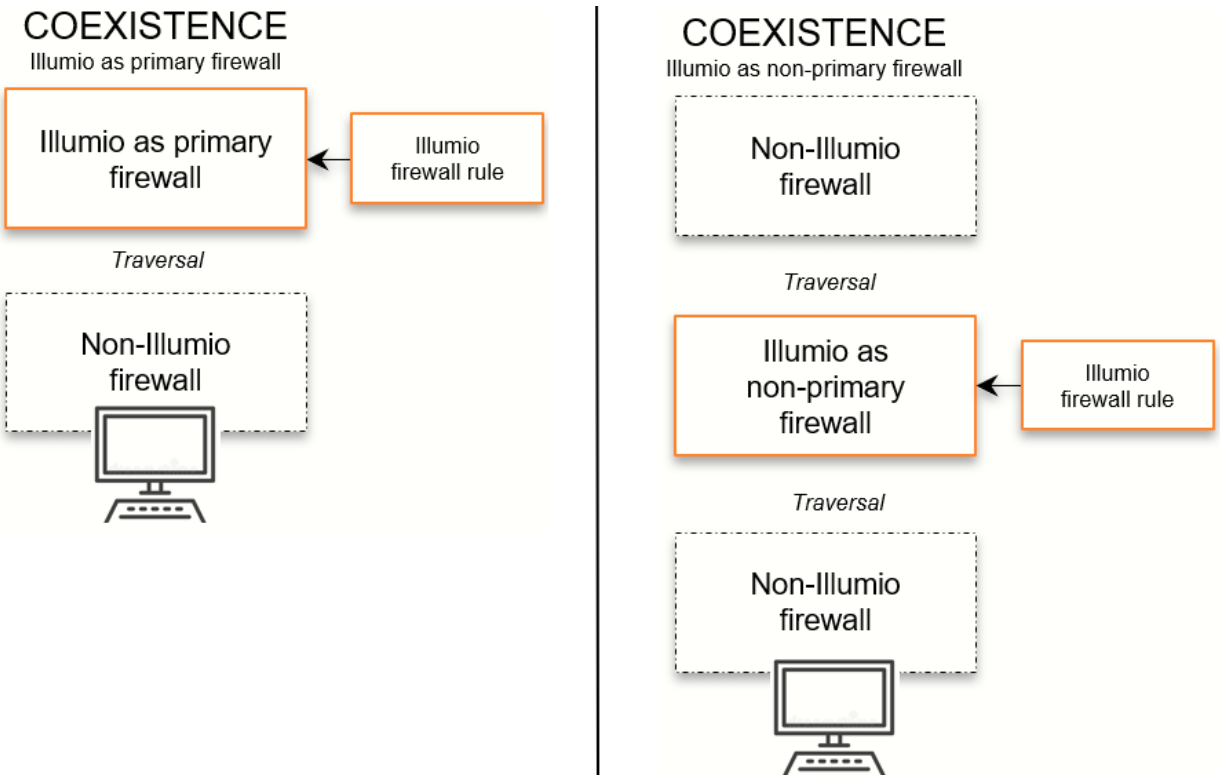

Prerequisites and Recommendations

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

WARNING

Firewall Coexistence is not supported on Solaris and AIX platforms.

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**.

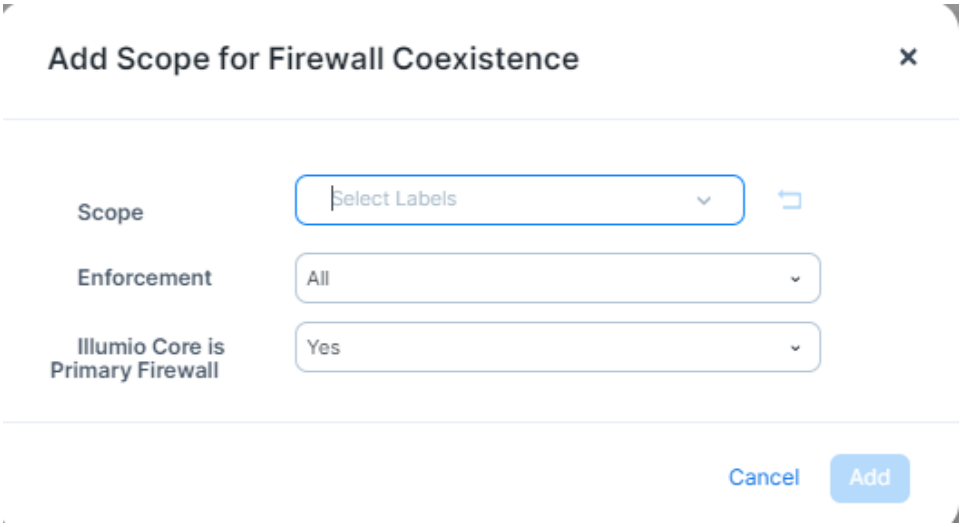

   Start adding and configuring the Scope, Enforcement, and whether it is the Primary Firewall.
2. From the Scope drop-down list, select the labels.
3. From the Enforcement drop-down list, select All, Enforced, or Illuminated.
4. In the Illumio Core is Primary Firewall, select either Yes or No.
5. Once 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 [129] 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
`last_heartbeat_on[lte]`	Allows you to search for workloads whose last heartbeat occurred before a specific time. `lte`: Less than or equal to.	String (timestamp_in_rfc3339)	No
`last_heartbeat_on[gte]`	Allows you to search for workloads whose last heartbeart 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.

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.

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 runlevel 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 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	`illumio-pce-db-management dump --file backup_filename`	Backs up the policy database.
Traffic	`illumio-pce-db-management traffic dump --file traffic_backup_filename`	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.
- 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_filenam
   ```

   Traffic database

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

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

   ```
   sudo -u ilo-pce /opt/illumio_pce/illumio-pce-db-management mask-db-dump --traffic --i
   ```

   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 umasked 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.

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 [126] 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 [132].

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
   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-du
   ```

   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 [131].

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_backu
   ```

   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.yml` 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 back
```

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 a current one. 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 restoring a backup file from a previous version of the PCE software.
- After you have completed a new PCE build and installation and initialized the database via the `Illumio-pce-db-management` setup command.

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 the PCE in runlevel 1.

NOTE

Illumio recommends waiting at least 15 minutes to restore a policy database backup after taking it. When you restore a policy database backup sooner than 15 minutes, the PCE might only apply policy correctly to some 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 the "Back Up the Policy Database" topic.)

```
$ sudo -u ilo-pce illumio-pce-db-management restore --file /path/to/policy_db_dump_fi
$ 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_d
```

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:
   - 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.

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
   ```
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.g
   ```
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.
   - 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 desribed, set `node_type` to `citus_coordinator`.
   - On each worker node, in addition to the settings already desribed, 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.
   ```
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.
    ```

    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
```

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.
```

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.

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: 250 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_actors_per_rule`)	Total number of items allowed per rule in the Providers and Consumers fields. 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.	50	200
Pairing Keys (active) (`total_active_pairing_keys`)	Total number of active pairing keys A pairing key is active when you create a pairing profile, click **Start Pairing**, and generate the key.	1200	5K

Object	Description	Soft Limit	Hard Limit
	When you click **Stop Pairing**, the pairing key becomes inactive and is no longer counted in the object limit.		
Pairing Profiles (`total_pairing_profiles`)	Total number of pairing profiles	1200	5K
RBAC Permissions (`total_org_permissions`)	Total number of RBAC permissions Each RBAC permission is a three tuple of an RBAC user or user group, role, and scope.	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_service_ports_per_rule`)	Total number of ports that can be 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.	400	500
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 x 3) + (2 x 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 number of all RBAC users and groups	1600	2000

Object	Description	Soft Limit	Hard Limit
(`total_org_auth_security_principals`)			
Adaptive User Segmentation (AUS) users (`total_security_principals`)	Total number of Adaptive User Segmentation (AUS) users used in rules	45K	50K
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: 200 2x2 (small): 10,000 2x2: 40,000 4x2: 100,000	SNC: 250 2x2(small): 12,500 2x2: 50,000 4x2: 125,000
Container workloads (`total_container_workloads`)	Total number of container workloads. The term *container workloads* 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 *Kubernetes workloads* 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 profiles (`container_workload_profiles_per_container_cluster`)	Total number of Container Workload Profiles in each container cluster.	800	1K
Container clusters (`total_container_clusters`)	Total number of container clusters.	80	100
User 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	100	125

Object	Description	Soft Limit	Hard Limit
(`total_active_sessions`)	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.		

Monitor and Diagnose PCE Health

This section describes monitoring the PCE to ensure it is operating correctly. You can view events generated by the PCE, read PCE logs, and generate reports about PCE activity.

PCE Logs

Most PCE logs are written to syslog, but some 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 primary 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_monitor	program("illumio_pce/database_monitor");
database_servicedata-base_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("haxproxy");
	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.
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 helpful system and application statistics that can help Illumio Customer Support troubleshoot 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 [149].

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 [150].
- If you are using the API, similarly, monitor the status field. For details, see Health Monitoring Using Health REST API [151].
- 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.

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=INFO`	Green	• 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=WARN`	Yellow	• One or more nodes are unreachable. • One or more optional services are missing, or one or more required services have been degraded. • 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. • Database replication lag is greater than 30 seconds. • (In a PCE Supercluster only) Supercluster replication lag is greater than 120 seconds.
Critical or `sev=ERR`	Red	• One or more required services are missing. • A health metric has exceeded its critical/error threshold.

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.

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 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/
```

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
```

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_p
```

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      10.0.6.1
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      10.0.6.1
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
fileserver_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      10.0.6.1
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.

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	• 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): Warning, Critical: If the situation persists, contact Illumio Support so that the reason for the underperformance of the auto-vacuum can be investigated.
VEN heartbeat performance	• 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. • 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	• 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. • 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 . 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	• 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. • Success rate, node: This metric is informational. However, if counts differ across core machines, ensure intra-PCE latency is within the 10ms limit. • Failure percentage ratio, node: 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.
Traffic summary	• Ingest rate, node: 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

Name	Troubleshoot
	and lose data. Adding additional `flow_analytics` daemons will distribute the work, but eventually PostgreSQL itself could become the bottleneck, requiring the use of DX. • Backlog size, node: 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. • Backlog Disk Utilization: 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.
Database Replication Lag	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.
Supercluster Replication Lag	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.

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.

NOTE

Key and value pair for lag for policy and traffic databases is:

NOTE

`policy_database_replication_lag=x seconds`

NOTE

`traffic_database_replication_lag=x seconds`

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`
		The PCE node reports disk space for the PCE application directories (configured in the `runtime_env.yml` file):
		
		• `ephemeral_data_root`
		• `runtime_data_root`
		• `log_dir`
		• `persistent_data_root` directories
		
		When all these directories are on a single mount point, the node reports: `disk=n%`
		
		When multiple mount points exist, the node reports the first discovered path by name, such as:
		
		`ephemeral_data_root=n%`
		
		`log_dir=n%`
		
		When the PCE encounters an error determining this information, the node reports: `disk=?`.
		
		`disk_space_percent_thresholds` consists of two values that determine 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.
			Default: The following thresholds trigger the following severity levels:
			
			• NOTICE: `disk_space >= 90%` or `disk_inodes >= 90%`
			• WARNING: `disk_space >= 95%` or `disk_inodes >= 95%`
			
			These default thresholds can be modified using `disk_space_percent_thresholds` or `disk_inode_percent_thresholds`.
			
			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 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: • 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) is issued when the CPU load exceeds a given percentage for a given amount of time. `cpu_max_percent` is the CPU usage percentage above which the notification timer begins.	Default: the following values trigger the following severity levels: • NOTICE: `cpu >= 95% for more than 1 minute` • WARNING: `cpu >= 95% for more than 5 minutes` These default thresholds can be modified using `cpu_max_per-`

Category	Key	Description	Severity Levels
		`cpu_tolerance_seconds` 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.	`cent` and `cpu_tolerance_seconds`.
Cluster Leader	`leader`	The IP address of the current leader, or unavailable when no leader exists or it is unreachable.	
Cluster Status	`cluster`	The overall health of the cluster, reported by the leader only: • `cluster=healthy`: Everything is operating properly and all PCE services are running. • `cluster=degraded`: The cluster is running but has unhealthy nodes. • `cluster=down`: The cluster is missing a required service < 5 minutes. • `cluster=failed`: The cluster is missing a required service for >= 5 minutes.	These status values trigger the following severity levels: • NOTICE: `cluster=degraded` (<2 minutes) • WARN: `cluster=degraded` (>=2 minutes) • WARN: `cluster=down` (<2 minutes) • ERROR: `cluster=down` (>= 2 minutes) • FATAL: `cluster=failed`
Missing Nodes	`missing`	The number of nodes that are missing from the cluster. If no nodes are missing, this metric is not reported.	
Replication Lag	`database_ replication_ lag`	The number of seconds the database replica is lagging behind the primary database. Output by database replica nodes only.	These thresholds trigger the following severity level: • WARNING: `>=30 seconds`
Disk Latency	`policy_disk_latency_milliseconds`, `traffic_disk_latency_milliseconds`	(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 iostat calculates await. Values: delay (milliseconds), disk Usefulness: Indicates Disk I/O, which is especially useful when the DB services are under heavy load.	• Normal: <= 300 • Warning: >300 <800 • Critical: >= 800
Policy Database: Size	`policy_database_size_gb`	(19.3.2 and later) Informational. 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:	`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.	• Normal: < 90 • Warning: [90 - 95] • Critical: >= 95

Category	Key	Description	Severity Levels
Disk Utilization			
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).	• 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 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.	• Normal: < 40 • Warning: 40 - 80 and current number of vacuum-ready rows is above Postgres default minimum to trigger vacuum (20% +50) • 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 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).	• Warning: average > 500ms • Critical: average > 5 sec
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	• 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

Category	Key	Description	Severity Levels
		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 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.	requests for the time window
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 spent in the load balancer queues, as queue time may be influenced by a number of other external factors. 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.	• Warning: average > 10 sec • Critical: average > 30 sec
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 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.	• Warning: > 1M req/hour

Category	Key	Description	Severity Levels
		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 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.	• 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, 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.	• 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.	• 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	`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.	• Warning: > 10% • Critical: > 50%

Category	Key	Description	Severity Levels
Size: % of Allocated			
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.	
Traffic Summary: Backlog Disk Utilization	`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 to be roughly equal or this buffered backlog will grow.	• Warning: > 10% • Critical: > 50%
Supercluster Replication Lag		(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. 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.	• 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.

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.

Generate a support bundle

1. Choose **Troubleshooting** > **VEN Support Bundles** .
2. Click the PCE tab and then **Generate**.
   Select optional content as desired.
3. Click **Generate**
   When the support report is complete, it will be available to download on that page

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

Generate a Report for the last 24 hours

Use the option "all" to get the support report for the last 24 hours from the time the command is run:

```
sudo -u ilo-pce /opt/illumio-pce/illumio/bin/support_report all wait duration=24
```

This command gets the support report 3 hours prior to the ending time. In this example, Sept 14, 2018 at 15:40:

```
sudo -u ilo-pce /opt/illumio-pce/illumio/bin/support_report all wait duration=3 to=09/14
```

Get all info in the logs

If you are not sure of the time the issue started, run this command instead:

```
sudo -u ilo-pce /opt/illumio-pce/illumio/bin/support_report all wait
```

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 `support_report` options.
`logs` (+ optional search arguments)	Includes logs and the optional search criteria described in Search the PCE Log Files [147].
`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
   ```
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 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 of 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 [149] 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	• Front end • Processing • Service and caching
Data	• 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`

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.

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

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 consist of 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 existing 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 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.
- FQDN of the front-end load balancer.
- 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`. 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 [170].

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 [170].
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 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=6efefeafe34ewroopll494934kdf
$ 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 [170].
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 [170] and Set Up a Standby PCE [173].

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.)

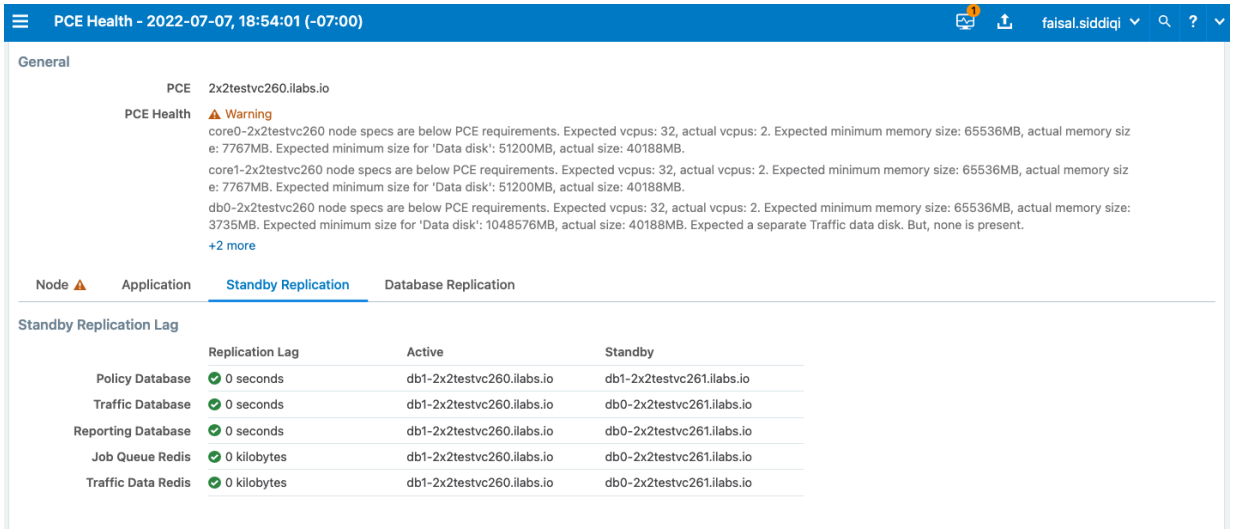

Another way that the PCE administrator can monitor replication is by watching the service discovery log for WAL segment missing errors. This error may appear when the standby traffic database service could not keep up with synchronization from the active traffic database service. When this error occurs, 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.

PCE Core Deployments and High Availability (HA)

The most common PCE Core deployments are either 2x2 or 4x2 setup.

For High Availability (HA) purpose, the PCE nodes can be deployed as 2 separate pairs (either 1core+1data or 2core+1data respectively) in separate data centers.

For high availability, the database services run in a primary replica mode with the primary service running on either of the data nodes.

NOTE

Both data nodes (data0 & data1) are always working as "active". Therefore, one of the data nodes (data1) is not on a "warm" standby that would become "active" when the primary data node has failed.

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 recover 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 behavior	The behavior of the PCE and VENs from the time the failure occurs to recovery. It can be caused by the failure itself or by the execution of recovery procedures.
Recovery	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. Includes the following items: • **Recovery type:** Can the PCE and VENs automatically recover from the failure, or is human intervention 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 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.

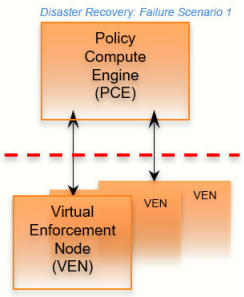

Stage	Details
Precondi-tions	None
Failure Behavior	PCE • Users cannot provision any changes to the VENs until the connection is re-established. • The information displayed in the Illumination map in the PCE web console is only as current as the last time the VENs reported to the PCE. • The PCE ignores 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. VENs • 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 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	• **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:
	  • 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.

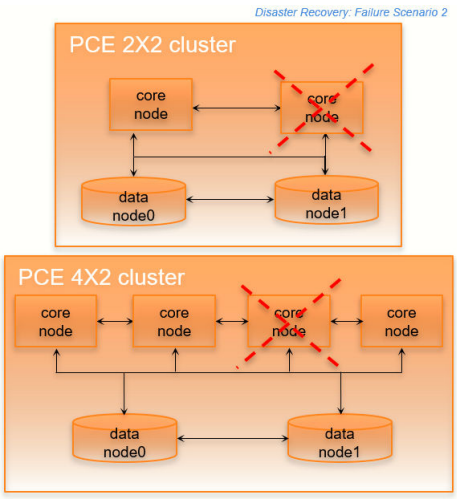

Stage	Details
Precondi-tions	None.
Failure Be-havior	**PCE** • The PCE might be temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return an HTTP 502 response and the `/node_available` API request might return an HTTP 404 error. • Other services that are dependent on the failed services might be restarted within the cluster. **VENs** • 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	• **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 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.

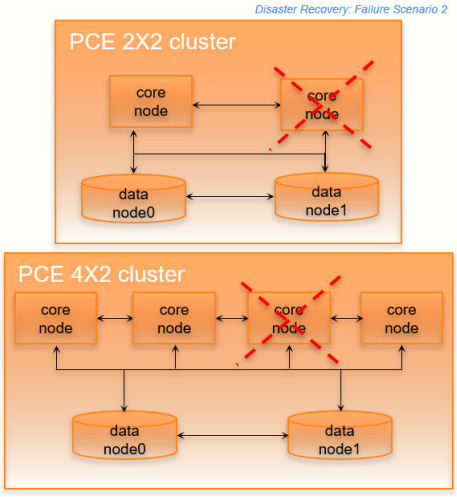

Stage	Details
Precondi-tions	The load balancer must be able to run application level health checks on each of the core nodes in the PCE cluster, so that it can be aware at all times whether a node is available. **IMPORTANT**: When you use a DNS load balancer and need to provision a new core node to recover from this failure, the `runtime_env.yml` file parameter named `cluster_public_ips` must include the IP address of your existing core nodes and the IP addresses of the replacement nodes. 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.
Failure Behavior	**PCE** • 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 `/node_available` API call might return an HTTP 404 error. • Other services that are dependent on the failed services might be restarted within the cluster. **VENs** • 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	• **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.
Full Re-covery	Either recover the failed node or provision a new node and join it to the cluster. For information, see Replace a Failed Node [124].

Data Node Failure

In this failure case, one of the data nodes completely fails.

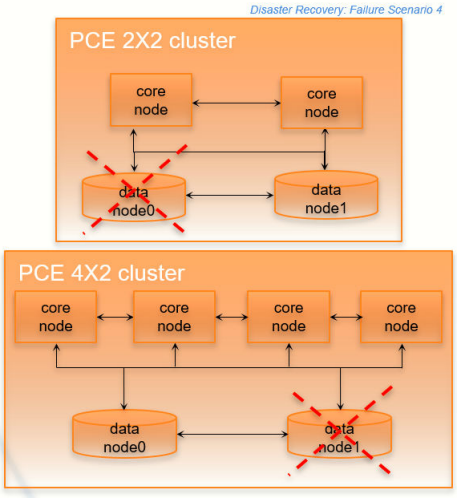
Disaster Recovery: Failure Scenario 4
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Preconditions	You should continually monitor the replication lag of the replica database to make sure it is in sync with the primary database. You can accomplish this precondition by monitoring the `illumio_pce/system_health` syslog messages or by running the following command on one of the *data nodes*: `$ sudo -u ilo-pce illumio-pce-db-management show-replication-info`
Failure Behavior	**PCE** • 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 `/node_available` 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 `set_server_redis_server` service is running on the failed data node, the VENs go into the syncing state and policy is re-computed for each VEN, even when no new policy has been provisioned. The CPU usage on the PCE core nodes might spike and stay at very high levels until policy computation is completed. **VENs** • 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	• **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:
	  • `set_server_redis_server`: 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.
	  • `database_service`: 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.
	  • `database_slave_service`: Implies the failed data node is the replica database. No 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	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. For recovery information, see Replace a Failed Node [124].

Primary Database Doesn't Start

In this failure case, the database node fails to start.

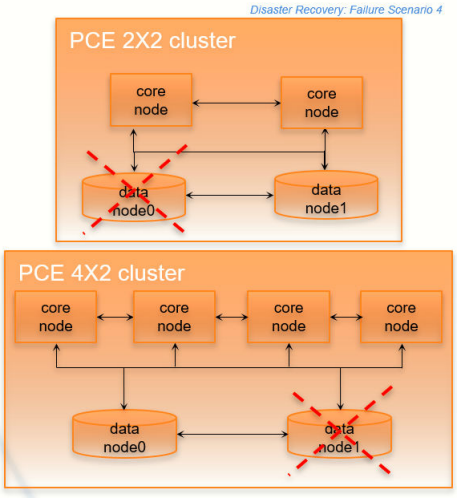

Stage	Details
Preconditions	The primary database node does not start.
Failure Behavior	The database cannot be started. Therefore, the entire PCE cluster cannot be started.
Full Recovery	**Recovery type:** Manual. You have two recovery options: • Find the root cause of the primary database failure and correct it. Contact Illumio Customer Support for assistance if needed. • Promote the replica data node to be the primary data node. **WARNING** Promoting a replica to primary risks data loss Illumio strongly recommends that this option be a last resort because of the potential for data loss. When you decide on the second option, see Configure Data1 and Core Nodes as Standalone Cluster [187]. When the PCE Supercluster is affected by this problem, you must also restore data on the promoted primary database.

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	**Recovery type:** Manual. You have two recovery options: • 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. **WARNING** **Promoting replica to primary risks data loss** Consider this option as a last resort because of the potential for data loss, depending on the replication lag. When you decide on the second option, on the *replica database node*, run the following command: `$ sudo ilo-pce illumio-pce-ctl promote-data-node <core-node-ip-address>` This command promotes the node to be the primary database for the cluster whose leader is at the specified IP address.

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 fails simultaneously for any reason.

Split Cluster Failure Involving Data1

In this failure case, data1 and half the core nodes completely fail.

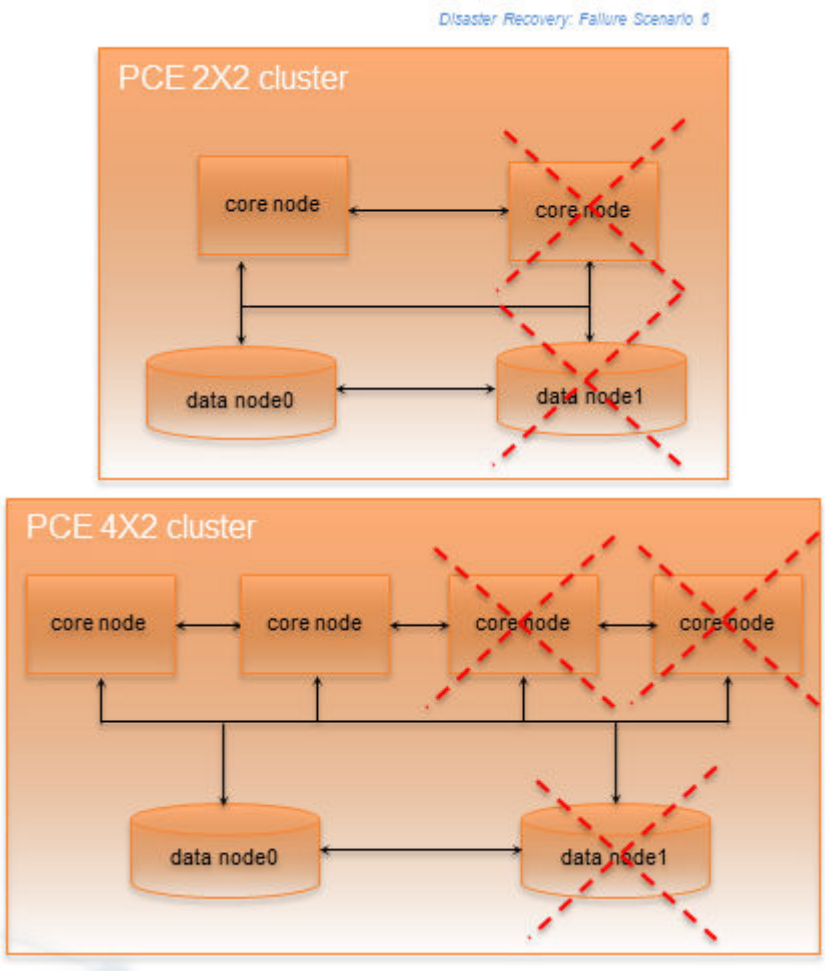

Stage	Details
Preconditions	None.
Failure Behavior	**PCE** • The PCE is temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return a HTTP 502 response and the `/node_available` API request might return am HTTP 404 error. • Other services that are dependent on the failed services might be restarted within the cluster. **VENs** • 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	• **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:
	  • `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	Either recover the failed nodes or provision new nodes and join them to the cluster. For recovery information, see Replace a Failed Node [124].

Split Cluster Failure Involving Data0

In this failure case, data0 and half of the total number of core nodes completely fail.

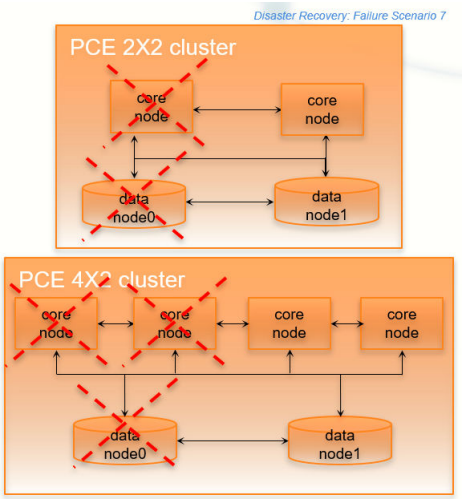

Stage	Details
Preconditions	**CAUTION** When reverting the standalone cluster back to a full cluster, you must be able to control the recovery process so that each recovered node is powered on and re-joined to the cluster one node at a time (while the other recovered nodes are powered off). Otherwise, the cluster could become corrupted and need to be fully rebuilt.
Failure Behavior	**PCE** • The PCE is unavailable because it does not have the minimum number of nodes to maintain quorum. **VENs** • 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.

Recovery	• **Recovery type:** Manual intervention is required to recover from this failure case.
	• **Recovery procedure:** See Configure Data1 and Core Node(s) to Operate as a Standalone Cluster [187] 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:
	  • `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 [188] 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> --spli
   ```

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 [130].

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.

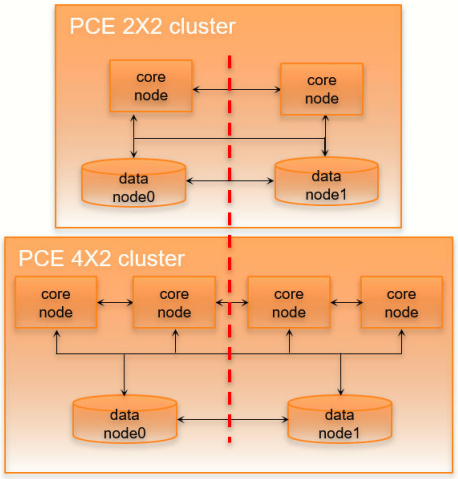

Stage	Details
Preconditions	None.
Failure behavior	**PCE** • The PCE is temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return an HTTP 502 response and the `/node_available` API request might return an HTTP 404 error. • Other services that are dependent on the failed services might be restarted within the cluster. **VENs** • 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	• **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:
	  • `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	No additional steps are required to revert the PCE 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.

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. 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 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 a 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 nailure

If the worker's primary node fails, all data-related applications might be unavailable briefly. 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 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 the 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.

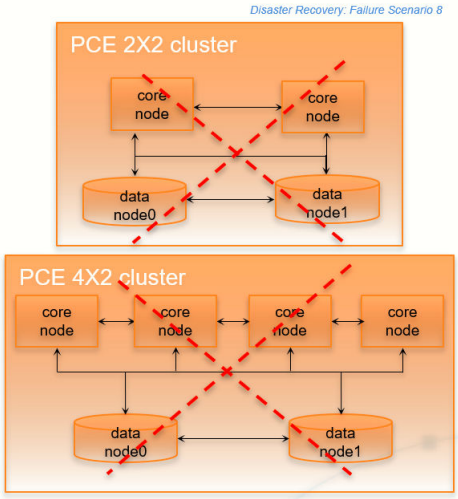

Stage	Details
Precondi-tions	Illumio assumes that you have met the following conditions before the failure occurs for this failure case. **IMPORTANT**: You must consistently and frequently back up the PCE primary database to an external storage system that can be used for restoring the primary database after this type of failure. You need access to this backup database file to recover from this failure case. The `runtime_env.yml` file parameter named `cluster_public_ips` must include the front-end IP addresses of the primary and secondary clusters. When this is not configured correctly, VENs will not have outbound rules programmed to allow them to connect to the secondary cluster in a failure case. Illumio recommends that you pre-allocate 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. • Regularly back up the PCE `runtime_env.yml` file for each node in the functioning cluster before failure. • Have a secondary PCE cluster deployed in a data center different from the primary cluster. The secondary PCE cluster can have IP addresses and hostnames that are different from the primary clusters.
Failure behavior	**PCE** • The PCE is unavailable. **VENs** • 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 and then purges the oldest data first during an extended event. • The VEN is degraded after missing 3 heartbeats (approximately 15 minutes). 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 upgrades and Support Reports.
Recovery	• **Recovery type:** Manual intervention is required to fully recover from this failure case. • **Recovery procedure:** See Complete Cluster Recovery [193] 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.
Full Recovery	See Complete Cluster Recovery [193] for full recovery information; perform all the listed steps on the restored primary cluster.

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 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
```

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 the PCE VEN Library.

- When you have at least one PCE core node unaffected 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 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 the 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 affected workloads.

Connectivity Configuration for PCE

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.

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.

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.

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.

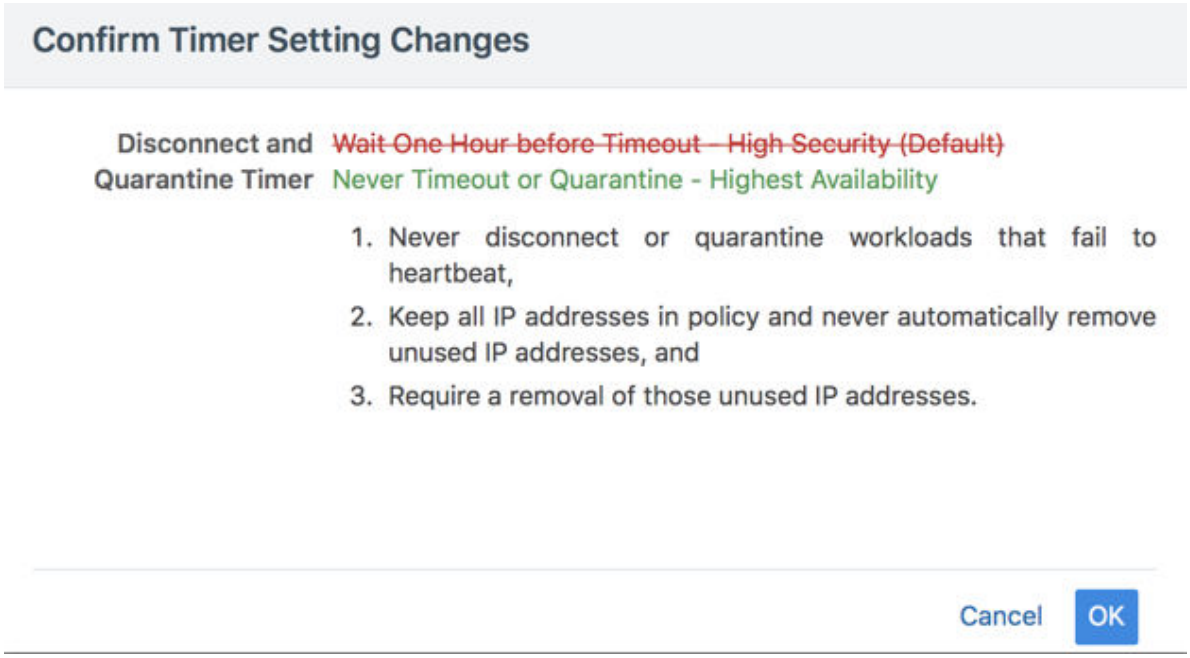

6. Click **OK** to save the new settings.

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 /
```

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
		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 default outbound access in Azure	The PCE will ignore the public IP address of the workload in policy computation. This setting is used in environments where 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 Container 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.

NOTE

Use the API call to the PCE to enable this feature so it appears in the Security menu as an option.

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, datacenter, 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.

SecureConnect Features

SecureConnect has the following key features.

Supported Platforms

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 encrypts the data between endpoint workloads.

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`, `platform.log`) from the Windows system related 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 assess the power required to process the encrypted traffic using this feature.

In practice, enabling SecureConnect for a workload will unlikely cause a significant spike in CPU processing or a decrease in network throughput. However, Illumio recommends benchmarking performance before enabling SecureConnect and comparing results after enabling it.

Prerequisites, Limitations, and Caveats

Before configuring your workloads to use SecureConnect, review the following prerequisites and limitations, and consider the following caveats.

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 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 and Build and Test Policy States

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. If misconfigured, you could inadvertently block traffic for workloads running in the Build and Test policy states.

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.

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.

Use Pre-Shared Keys with SecureConnect

SecureConnect supports 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 regener-

ate 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 Core PCE does not support multiple Organizations when you have installed the PCE in your data center.

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**.

Use PKI Certificates with SecureConnect

SecureConnect lets you use client-side PKI certificates for IKE authentication and IPsec communication between managed workloads. If you have a certificate management infrastructure, you can leverage it for IKE authentication between workloads because it provides higher security than 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.

Requirements and Caveats

- 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.
- Only use certificates obtained from trusted sources.
- 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 cannot communicate with each other.
- The certificates you deploy for PKI or IPsec must have the following properties:
  Leaf certificate X.509 field requirement
  - Version 3

- Subject Name DN must contain the Common Name
- SubjectAltName (must be the same as the Common Name)
- CN and SubjectAltName must be in one of the following formats:
  - Email Address
  - DNS
- Must contain key usage with:
  - Digital Signature
  - Key Encipherment
  - Data Encipherment
  - Key Agreement
- Must contain Extended key Usage with:
  - IPSec End System
  - IPSec User
  - TLS Web Server Authentication (optional for mac OS x compatibility)
- Must contain Authority Key Identifier

Set up Certificates on Workloads

To use PKI certificates with SecureConnect, you must set up certificates on your Windows and Linux workloads independently.

File Requirements

File	Requirements
Issuer's certificate	The global CA certificate, either root or intermediate, in PEM or DER format **NOTE** On Linux, the issuer's certificate must be readable by the Illumio user.
pkcs12 container	Archive containing the public key, private key, and identity certificate generated for the workload host. Sign the identity certificate using the global root certificate. You can password protect the container and private key but do not password protect the public key.

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

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`

Configure PKI Certificates

You can 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 or deliver certificates for your organization to your workloads.

NOTE

You must set up certificates on your Windows and Linux workloads independently. For information, see Requirements for Certificate Setup on Workloads [204].

1. Go to **Settings** > **Security Settings**.
2. Choose **Edit** > **Configure SecureConnect**.
3. In the Default IPsec Authority field, select **Certificate Authority**.
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**.

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 [205]
  - Requirements for Certificate Setup on Workloads [204]
  - Certificates for AdminConnect [207]
- • 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 [205] and Requirements for Certificate Setup on Workloads [204] 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. [207]
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 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 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=
Auth1CertMapping:                       No
Auth1ExcludeCAName:                     No
Auth1CertType:                          Intermediate
Auth1HealthCert:                        No
MainModeSecMethods:                     ECDHP384-AES256-SHA384
QuickModeSecMethods:                    ESP:SHA1-AES256+60min+100256kb
ApplyAuthorization:                     No
Ok.
```

PCE Troubleshooting

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 `postgresql.log`, look for messages like the following:

```
ERROR: database is not accepting commands to avoid wraparound data loss in database "<da
```

`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 [152].

WARNING

If you find messages that indicate a risk of transaction wraparound, immediately contact Illumio Support for assistance.

VEN Administration Guide

Overview of VEN Administration

This section describes the VEN characteristics and the VEN commands that you use to administer the VEN on the workloads in your environment after you have installed the VEN and the workloads are managed by Illumio Core.

About This Administration Guide

This guide shows you how use `illumio-ven-ctl` (for Linux, AIX, and Solaris) and `illumio-ven-ctl.ps1` (for Windows) and other commands to administer the Virtual Enforcement Node (VEN) on a managed workload for operational tasks such as start/stop, suspend, and other functions on the VEN and with the Policy Compute Engine (PCE) in an on-premise deployment.

How To Use This Guide

The VEN Administration Guide has several main divisions:

- Overview of VEN Software Architecture and Description of Components.
- VEN deployment models
- Command-line-oriented sections with syntax examples for `illumio-ven-ctl` for on-workload managing the VEN.
- Basic Theory of VEN Operations.

Before Reading This Guide

Illumio recommends that you be familiar with the following topics before you follow the procedures in this guide:

- Your organization's security goals
- The Illumio Core platform
- General computer system administration of Linux and Windows operating systems, including startup/shutdown, and common processes or services
- Linux/UNIX shell (bash) and Windows command line
- TCP/IP networks, including protocols and well-known ports

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
...
```

VEN Architecture and Components

This topic describes the basic concepts relevant to the VEN and for Illumio Core software. Additionally, it explains the VEN architecture and components.

Basic Concepts for Illumio Core Software

- A *workload* is a bare metal server, virtual machine (VM), or container.
- The *VEN* is a lightweight, multiple-process application with a minimal footprint that runs on a workload.
- *Native network interfaces* are also know as the OS's firewall platform.

The VEN manages firewalls at an OS level, so you must install a VEN on every bare-metal server or virtual machine you want to secure. However, you only need to install a single VEN to secure all the containers on a machine. A secured workload is known as a *managed workload*.

Once installed, the VEN performs the following tasks:

- Interacts with the native networking interfaces to collect traffic flow data.
- Enforces policy received from the PCE.
- Only consumes CPU as needed to calculate or optimize and apply the firewall, and so on, while remaining idle in the background as much as possible.
- Uses configurable operational modes to minimize the impact to workloads.
- Summarizes the collected traffic-flow data, then reports it to the PCE.

You control the VEN's operations through the PCE web console or from the command line on the machine with the installed VEN itself.

Activation or Pairing

The terms “activation” and “pairing” indicate the same function from different perspectives, namely putting the workload under managed control by the PCE:

- The VEN sees itself as *activated* or *deactivated*.
- The PCE sees a VEN as *paired* or *unpaired*.

Pairing and Activating the VEN

1	The VEN is installed.	The PCE remains unaware the VEN is present.
2	The VEN and the PCE are paired.	The PCE uses a pairing key (activation code) to pair with the VEN. After pairing, the PCE becomes aware of the VEN.
3	The VEN is activated.	The VEN uses an activation code generated by the PCE. After activation, the VEN is ready to function.

Unpairing or Deactivating the VEN

- When the PCE is unpaired with the VEN, the VEN is deactivated and uninstalled.
- When the VEN is deactivated, it remains installed and can be reactivated.
- Use the `illumio-ven-ctl` command to deactivate the VEN. You can't deactivate a VEN by using the PCE UI; you may only unpair it.

VEN Architectural Diagram

At startup, the VEN instantiates the following processes or services.

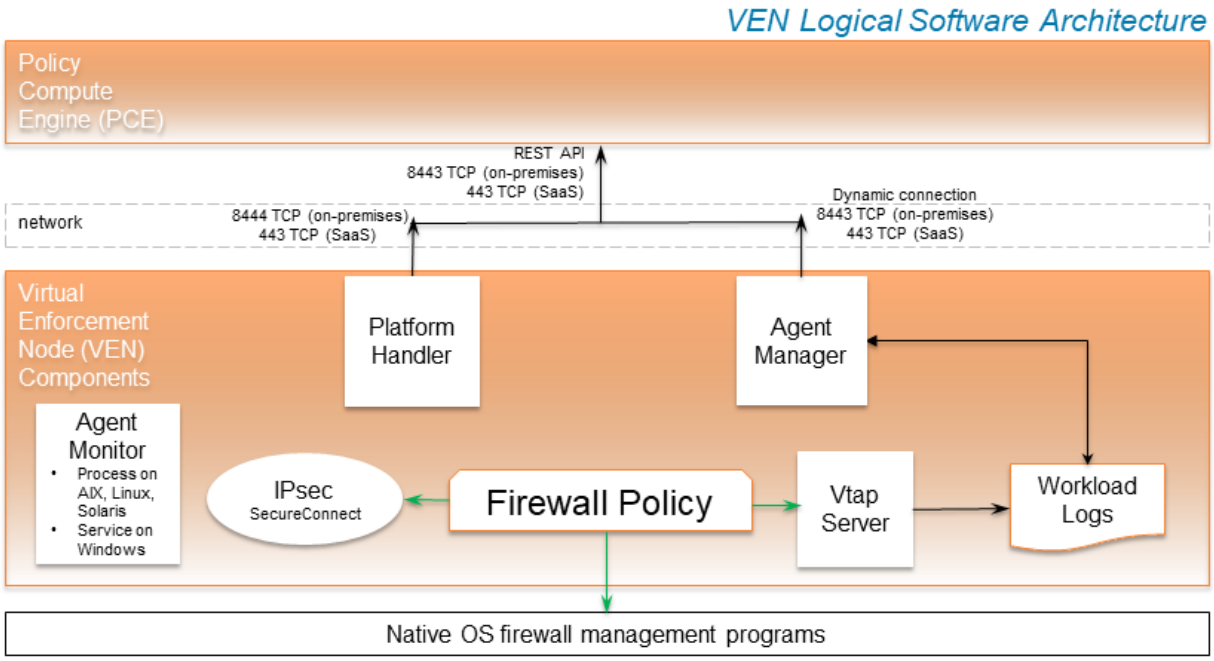

1. The VEN reports to the PCE the status of the workloads.
2. The PCE computes a unique security policy for each managed workload and transmits it to the VEN.
3. The VEN receives the policy and it programs a firewall by using the firewall platform of the OS. The VEN supports the following firewall platforms:
   a. iptables (older Linux)
   b. nftables (newer Linux)
   c. Packet Filter (newer Solaris)
   d. Ipfilter (older Solaris)
   e. Windows Filtering Platform (Windows)
4. When the VEN is finished programming a firewall for each workload, it reports back to the PCE. The PCE then considers these workloads as having a *synced* policy.

Main Components of the VEN

VEN Process	Description	Linux/ AIX/ Solaris User	Windows User
AgentManager	• Manages PCE-driven uninstallation and upgrades. • All actions relating to active service reporting. • Mines the workload's system information, such as network interfaces, and listening processes, and sends them to the PCE. • Sends heartbeats to the PCE. • Calls netstat periodically for connection status through a shell script or with a direct program call.	`root`	LOCAL SYSTEM
PlatformHandler	• Firewall configuration via native OS mechanisms. • Tamper detection and protection. • Upgrades and uninstallation.	`root`	LOCAL SYSTEM
VtapServer	• Windows: VTAP runs under the "Local System" account. • Retrieves traffic flow data from the ilowfp kernel mode driver (Windows) or firewall (other platforms) and generates flow logs in a database. • Receives events from the firewall on blocked packets and allowed connections.	`root`	LOCAL SYSTEM
AgentMonitor	• Service account: NT Authority/Local System • Monitors VEN processes or services and restarts them when necessary.	`root`	LOCAL SYSTEM

SecureConnect Architecture

Illumio's optional SecureConnect feature configures Internet Protocol Security (IPsec), a set of protocols to enforce security for IP networks. IPsec can be configured to use cryptography.

IPsec runs as root in LOCAL SYSTEM.

VEN Interactions with Files and Components

The VEN interacts with files and components for installation, root tasks, and initialization tasks. Minor tasks include working with install logs, the registry key, and read-only access to machine resources.

The VEN interacts with the following files and components:

Linux/AIX/Solaris

Function	Description	File/Location
Root file	DATA_ROOT is a variable that points to a filepath.	`/opt/illumio_ven_data` (by default)
Package repository	INSTALL_ROOT is a variable that points to a filepath.	`/opt/illumio_ven` (by default)
System initialization	Initializes system	`/etc/illumio_ven` (typically)
Persistent install log	Persistent install log	`/var/log/illumio.log`
Firewall	Dynamically adds IPs to ipsets:	Snoop on special packets.
	Strongswan IPSec system.	Snoop on Security Associations.
	Read system files (e.g., netstat).	`/proc`

Windows

Function	Description	File/Location
Runtime data files	DATAFOLDER is an installer parameter that points to a filepath.	`c:\ProgramData\Illumio` (by default)
Executable program files	INSTALLFOLDER is an installer parameter that points to a filepath.	`c:\Program Files\Illumio` (by default)
Install log	Persistent install log.	`c:Windows\Temp\illumio.log` (by default)
		`c:Windows\Temp\Illumio_VEN_Install.log` (by default)
		`c:Windows\Temp\Illumio_VEN_Uninstall.log` (by default)
System initialization	N/A	N/A
Firewall	For network filtering.	Windows Filtering Platform

Management Interfaces for the VEN and PCE

The diagram below is a logical view of the management interfaces to the PCE and VEN.

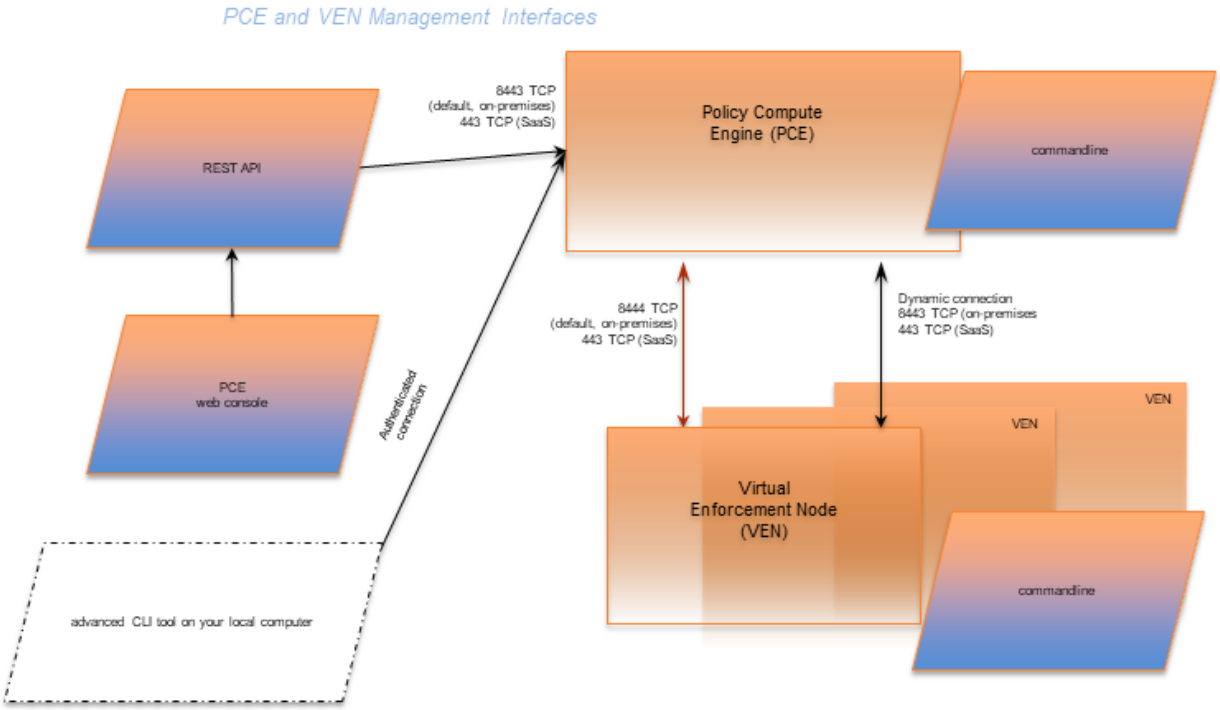

Interface	Notes	See...
PCE web console	With the PCE web console, you can perform many common tasks for managing Illumio Core.	Security Policy Guide
PCE command line	Use of the command line directly on the PCE. A primary management tool on the PCE is the command line `illumio-pce-ctl` control script. You can perform many common tasks for managing the Illumio Core on the PCE command line, including installing and updating the VEN.	PCE Administration Guide
PCE advanced command-line tool	From your own local computer, you can run the PCE advanced CLI tool for many management tasks on the PCE's resource objects: • Importing vulnerability data for analysis with Illumination®. • Importing/exporting security policy rules. • Managing security policy rules and rulesets, labels, and other resources.	PCE CLI Tool Guide
REST API	With the Illumio Core REST API, you can perform many common management tasks. One use is to 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.	REST API Developer Guide
VEN command line	A primary management tool on the VEN command line is the `illumio-ven-ctl` control script.	VEN Administration Guide

VEN Supported Interfaces

Windows

- Ethernet
- Tunnel
- WLAN (Endpoint only)
- PPP (Endpoint only)

Linux/Unix

- Ethernet
- Tunnel
- Infiniband
- GRE
- Loopback

About VEN Administration on Workloads

The following topic explains the VEN states and characteristics necessary to understand when administering the VEN on workloads.

Workload Policy States

After activation, the VEN can be in one of the following policy states. The VEN policy state determines how the rules received from the PCE affect the network communication of a workload.

Change the policy state of the VEN by modifying settings in the PCE or by making calls to the REST API.

VEN Enforcement Characteristics

Policy enforcement is managed through both enforcement states and visibility states to specify how much data the VEN collects from a workload.

The following table summarizes the key enforcement characteristics of the VEN:

Workload Enforcement State	VEN Mode	VEN Visibility Level	Log Traffic
Idle	Idle	Limited	Limited
Visibility Only	Illuminated	Off	VEN does not log traffic connection information
		Blocked	VEN logs connection information for blocked and potentially blocked traffic only
		Blocked+Allowed	VEN logs connection information for allowed, blocked, and potentially blocked traffic
		Enhanced Data Collection	VEN logs byte counts in addition to connection details for allowed, blocked, and potentially blocked traffic
Selective	Selective	Off	VEN does not log traffic connection information
		Blocked	VEN logs connection information for blocked and potentially blocked traffic only
		Blocked+Allowed	VEN logs connection information for allowed, blocked, and potentially blocked traffic
		Enhanced Data Collection	VEN logs byte counts in addition to connection details for allowed, blocked, and potentially blocked traffic
Full	Enforced	Off	VEN does not log traffic connection information
		Blocked	VEN logs connection information for blocked and potentially blocked traffic only
		Blocked+Allowed	VEN logs connection information for allowed, blocked, and potentially blocked traffic
		Enhanced Data Collection	VEN logs byte counts in addition to connection details for allowed, blocked, and potentially blocked traffic

For more information, see "Ways to Enforce Policy" in the Security Policy Guide.

VEN Features by Initial Release

The following tables list key Illumio Core features by their introductory release.

VEN Features in Release Pre-19.3.0

Feature	Initial Release
Firewall coexistence	Pre-19.3.0
illumio-ven-ctl start/stop/activate/unpair	Pre-19.3.0
illumio-ven-ctl unpair open|saved|recommended	Pre-19.3.0
illumio-ven-ctl suspend	Pre-19.3.0
IPSec (SecureConnect)	Pre-19.3.0
Kerberos PKI-based Pairing on Solaris/AIX	Pre-19.3.0
PCE Repo Upgrade	Pre-19.3.0
Process-based Policies	Pre-19.3.0
Solaris Zone Support	Pre-19.3.0
Support report	Pre-19.3.0

VEN Features in Release 19.3.x

Feature	Initial Release
Compatibility Report for IPv6 Support	19.3
Custom iptable Rules	19.3
Easy installation of VEN on container hosts	19.3
Ignored Interfaces on Windows VENs	19.3
Management of Conntrack Table Size	19.3
Modes: idle, illuminated, enforced	19.3
nftables for RHEL 8	19.3
Solaris 11.4 Support	19.3
Support Reports New Options	19.3
Faster Supercluster Full Restore	19.3.0
FQDN policy on Domain controller/DNS server	19.3.0
State Table Sizes on AIX and Solaris	19.3.0
illumio-ven-ctl deactivate	19.3.0
CRI-O Support	19.3.1
Loadbalancer TCP port 8302 and TCP+UDP port 8302 Enhancements	19.3.1
Docker/ContainerD/CRIO	19.3.1
SLES on Power Series hardware	19.3.2
Oracle Exadata Support	19.3.4
Oracle ZDLRA Support	19.3.4
FQDN-Based Rules Enhancements	19.3.5
LDAP Authentication	19.3.5
Aggressive Tampering Protection for nftables	19.3.6
Illumio Core REST API	19.3.6
Debian 11 Support	19.3.7
IBM Z Support	19.3.7

VEN Features in Release 20.x

Feature	Initial Release
Agent Monitor	20.1.0
REJECT Rules	20.1.0
Workloads and VENs Separation	20.1.0
Flow Duration Attributes	20.2.0
IPv6 for Linux and Windows VENs	20.2.0
IPv6 for VEN	20.2.0
IPv6 is Enabled by Default on Datacenter VENs	20.2.0
Software Management from PCE	20.2.0
Stopped Status	20.2.0
Tamper Detection	20.2.0
Clone Detection	20.2.0 (Edge 20.1, Core 20.2)
Selective Enforcement	20.2.0-PCE

VEN Features in Release 21.x

Feature	Initial Release
Core 21.2.0, Illumio previewed the Reports feature	21.2.0
Enforcement Boundaries	21.2.0
Linux Pairing Script Activation for Proxy Servers	21.2.0
Network-Specific Policy	21.2.0
Uninterrupted Traffic between the VEN and the PCE	21.2.0
Network_deny List	21.2.0-PCE
Adaptive User Segmentation	21.2.0-VEN
Explorer Allows Label Search of All Types	21.2.1
Open Source Package Updates for 21.2.1	21.2.1
RHEL 8 support for PCE	21.2.1
Supercluster 8-Region Support in 21.2.1	21.2.1
Syslog Forwarding Change	21.2.1
Threshold Configuration Settings	21.2.1
File Settings Option	21.2.1
VEN Package Format Changes	21.2.1
Proxy Fallback Enhancement on Windows	21.2.4
Robustness and Reliability	21.5.0
Run as a Different User with AUS on Windows	21.5.0
IBM Z with RHEL 7 and RHEL 8	21.5.11
Label-based Security Setting for IP Forwarding	21.5.11

VEN Features in Release 21.x-C (Container)

Feature	Initial Release
Containerized VEN	21.2.0-C VEN
Containerized VEN Base Image	21.2.1-C-VEN

VEN Features in Release 22.x

Feature	Initial Release
Advanced Diags (strace/tcpdump)	22.5.0
Configurable Time for Heartbeat Warning Events	22.2.0
Disable and Enable Enforcement Boundaries	22.2.0
Essential Rule Coverage in Illumination and Explorer	22.2.0
Firewall Script Logging	22.2.0
Traffic Flow Query Report	22.2.0
Wireless Connections and VPNs	22.2.0

VEN Features in Release 23.x

Feature	Initial Release
Extended RHEL 5 Support	23.2.0
Configurable enforcement node type (server or endpoint) in pairing profile	23.2.0

Major VEN Features by Supported OS

The following table lists key VEN features by supported platform.

Feature	Windows	Windows Edge	Linux	RHEL 5	C-VEN	CentOS8	AIX	Solaris	MacOS (Endpoint)
Firewall	WFP	WFP	IPtables	IPtables	IPtables	NFTables	IPFilter	IPFilter/PF	PF
Firewall coexistence	✓	✓	✓	✓	✓	✓	-	-	✓
Container support	-	-	✓	✓	✓	✓	-	-	-
IPv6	✓	✓	✓	-	✓	✓	-	✓	✓
PCE repo upgrade	✓	✓	✓	✓	-	✓	-	-	✓
Aggressive Tampering Detection	✓	✓	✓	✓	-	-	-	-	-
Process-based policies	✓	✓	-	-	-	-	-	-	-
Extended process path/args (vtap)	✓	✓	✓	✓	✓	✓	✓	✓	✓
Flow-byte counting	✓	✓	✓	-	-	-	-	-	-
Kerberos	✓	✓	✓	✓	✓	✓	✓	✓	✓
FIPS	✓	✓	✓	✓	✓	✓	✓	✓	-
FQDN Policies	✓	✓	✓	-	✓	✓	-	-	✓
FQDN Traffic reporting	✓	✓	✓	✓	✓	✓	-	-	-
IPSec (SecureConnect)	✓	✓	✓	✓	✓	✓	-	-	-
Installer	MSI; EXE	MSI; EXE	pkg	pkg	apk; rpm	pkg	bff	pkg	dmg

Feature	Windows	Windows Edge	Linux	RHEL 5	C-VEN	CentOS8	AIX	Solaris	MacOS (Endpoint)
	(from 21.2.1)	(from 21.2.1)			(from 19.3.2)				
Pairing script (oneliner from PCE UI)	✓	✓	✓	✓	✓	✓	-	-	✓
Process-based policies	✓	✓	o e-bpf	o e-bpf	-	-	-	-	o (P1) networkextension

NOTE
On RHEL 5, machine authentication is not supported.

VEN Policy Sync States

To help you administer and troubleshoot the VEN, it reports many Policy Sync states. Here are the Policy Sync states and their definitions:

- **Active (Syncing)**: Policy is currently being applied to the workload. Appears if the VEN is not currently heartbeating but the PCE has not received a goodbye event from the VN, and the disconnect & quarantine threshold timer has not yet been reached. This is appropriate because, from the PCE's point of view, the VEN status is not stopped and the policy sync status is Syncing. Compare with Syncing [225].

  NOTE
  A workload may also have a status of Active (Syncing) if there is a high rate of policy changes taking place, either from user provisioning actions or from VEN environmental policy changes (for example, new VENs being activated or old VENs being deactivated/unpaired).

- **Syncing:** Appears if the PCE has received a goodbye event from a VEN but the decommission offline timer threshold has not yet been reached. This is appropriate because the VEN, although stopped, is not yet removed from policy and therefore has not yet been marked as **Offline**. When the offline timer expires, the VEN's status transitions to **Stopped** and its IP is removed from policy. Compare with Active (Syncing) [225].
- **Active**: The most recent policy provisioning was successful, no unwanted changes to the workload's firewall have been reported, none of the configured SecureConnect connections are in an erroneous state, and all VEN processes are running correctly.
  - For more information on SecureConnect, see Security Policy Guide.
- **Staged**:The PCE has successfully sent policy to the VEN, and it is staged and scheduled to be applied at a later time. This state only appears when you have configured the Policy

Update Mode for the workload to use Static Policy. See Static Policy and Staged Policy for information. For information, see "Types of Illumio Policy" in the Security Policy Guide.

- **Error**: One of the following errors has been reported by the VEN:
  - The most recent policy provisioning has failed.
  - Unwanted changes to the workload's firewall have been reported.
  - At least one VEN process is not running correctly.
  - There is a SecureConnect or Machine Authentication policy, but leaf certificates are not set up properly.
- **Warning**: At least one SecureConnect connection is in an erroneous state, and either the most recent policy provisioning was successful or no unwanted changes to the workload's firewall have been reported.
- **Suspended**: Used by admins to debug. Rules programmed into the platform firewall (including custom iptables rules) are removed completely. No Illumio-related processes are running on the workload.

VEN Health Status on Workloads

The VEN health status on the workload's details page displays information related to the current state of VEN connectivity, the most recently provisioned policy changes to that workload, and any errors reported by the VEN.

These errors include any unwanted changes to the workload's firewall settings, any SecureConnect functionality issues, or any VEN process health errors.

To view a workload's VEN health status, view the VEN section on the **Summary** tab for the workload's details page.

VEN Process Health

The health status of the VEN can be monitored from the PCE web console. If for any reason one or more Illumio processes on the workload are not running, the VEN reports the error to the PCE. The PCE marks the workload as in an error state and adds a notification on the Workloads page. It also logs an audit event that includes the Illumio processes which were not running on the workload.

Workload Clone Alerts

Workloads can be filtered according to whether a cloned node has been detected. On Windows and Linux, when the PCE detects a cloned node, it notifies the VEN through a heartbeat. The VEN verifies that a clone exists, prevents it from being activated, and deletes it.

In the Illumio REST API, detection is done by using the `clone_detected` state. In the PCE web console UI, search the workloads list by filtering on, "clone detected." If there are workloads in the `clone_detected` state, a red banner (similar to *workloads in suspension*) is displayed at the top of the workload list page.

NOTE

Automatic Cloned VEN Remediation

For on-prem domain joined Windows workloads, cloned VENs support automatic clone remediation by detecting changes to the workload's domain Security identifier (SID). After the VEN reports such changes to the PCE, the PCE tells the clone to re-activate itself, after which the cloned VEN is remediated and becomes a distinct agent from the original VEN.

VEN Software Management from PCE

The ability to manage VEN software and install the VEN by using the PCE has been enhanced in this release in the following ways:

- You can upgrade all VENs or just a subset of VENs from the PCE.
- You can upgrade VENs by using filters, such as for labels, OSs, VEN health, IP address, current VEN version.
- When upgrading, the PCE informs you of the version the VENs will be upgraded to.
- You can monitor and troubleshoot VEN upgrade issues.
- You can perform VEN version reporting and compatibility.

Stopped VEN Status

The `stopped` status has the following affect on the PCE web console UI:

- On the Workload list page, the "Connectivity" column is replaced with "Status."
- On the Workload details pages, "VEN Connectivity" is changed to "VEN status."
- You can filter the Workload list page by the new VEN stopped status.

Aggressive Tampering Protection for nftables

Firewall changes that are not explicitly configured by the VEN are logged as tampering attempts. This feature extends Release 19.3 nftables support with the inclusion of aggressive tampering protection.

VEN Proxy Support on Linux, AIX, and Solaris

VEN proxy support includes Linux, AIX, Solaris, and Windows devices.

For information, see "VEN Proxy Support" in VEN Installation and Upgrade Guide.

Support on IBM Z With RHEL 7 and RHEL 8

In the Illumio Core 19.3 release, Illumio supports installing and operating the VEN on IBM Z systems running Red Hat Enterprise Linux 7 (RHEL 7) and RHEL 8.

Support on SLES 11 SP2

The VEN can be installed on systems running SLES 11 SP2 when the following packages are installed:

From the SLES 11 SP2 Latest Updates:

- libipset2-6.12-0.7.7.1
- ipset-6.12-0.7.7.1
- libmnl0-1.0.3-0.5.4
- kernel-default-3.0.101-0.7.17.1
- kernel-default-base-3.0.101-0.7.17.1

From the SLES 11 SP4 DVD:

- libxtables9-1.4.16.3-1.37
- libiptc0-1.4.16.3-1.37
- iptables-1.4.16.3-1.37
- libnfnetlink0-1.0.0+git1-9.5.56

VEN File Settings Option

In 21.2.1, the VEN IPFilter state table supports a new option for AIX workloads to support traffic from NFS servers:

VEN File Setting: `IPFILTER_TCPCLOSED=<value>`

Ipfilter Setting: `fr_tcpclosed=<value>`

For more information about this option, see "VEN Activate Command Reference" in the VEN Installation and Upgrade Guide.

Debian 11 Support

Starting from Release 21.2.3, Illumio supports installing and operating the VEN on the Debian 11 operating system.

Windows VEN Proxy Fallback Enhancement

Starting from Illumio Core 21.2.1 and 21.2.2, the VEN automatically detects a web proxy. However, it always attempts to connect directly to the PCE first. In this release, Illumio enhanced the heuristic in the VEN for falling back to the configured web proxy. After an attempt fails to connect to the PCE directly due to an HTTPS intercepting proxy, the VEN falls back to use the configured web proxy.

VEN Enhancements in 21.5.11

The following enhancements were added in Illumio Core 21.5.11.

Support on IBM Z With RHEL 7 and RHEL 8

In this release, the system supports installing and operating the VEN on IBM Z systems running Red Hat Enterprise Linux 7 (RHEL 7) and RHEL 8.

Label-based Security Setting for IP Forwarding

Illumio has enabled IP forwarding to hosts running Linux. A container networking solution routes the traffic to the VMs. To configure IP forwarding, use the new IP Forwarding tab in

the PCE web console. In this tab, you can use labels and label groups to enable IP forwarding for the workloads that match the label combination.

To enable this feature, contact Illumio Support. For details about how to set up IP forwarding for workloads, see "Connectivity Settings" in the PCE Administration Guide.

Uninterrupted Traffic Between the VEN and the PCE

The VEN implementation provides an extra layer of self-protection that prevents any erroneous policy from being applied to the VEN. The VEN employs a defensive approach that reviews policies before applying them. In case the VEN detects that the new policy may disrupt communications between the VEN and the PCE, the VEN automatically isolates that policy and logs an error in the event log. The VEN then continues to communicate with the PCE using the existing functional policy.

IPv6 Support and Features for the VEN

In Illumio Core 20.2.0 and later releases, the VEN supports both IPv4 and Ipv6 address versions and the IP address version appears correctly in the PCE; for example, in the Workload section of the VEN summary page in the PCE web console.

You can configure how the PCE treats IPv6 traffic from workloads. For more information, see "Allow or Block IPv6 Traffic" in the PCE Administration Guide.

The VEN supports IPv6 in the following ways.

IPv6 is Enabled by Default on Datacenter VENs

Release 20.2.0 and later support configuring inbound or outbound IPv6 traffic by organization (ORG). In previous releases, you are only able to block all, or allow all IPv6 traffic by organization.

The default settings are as follows:

- If the previous ORG-wide IPv6 policy is to *block all* IPv6 traffic, then this setting is *preserved*.
- If the previous ORG-wide IPv6 policy is to *allow all* IPv6 traffic, then this setting is *not preserved*.

IPv6 Support for Linux and Windows VENs

Beginning with Release 20.1, the Linux and Windows VENs support IPv6 rules.

VEN Compatibility Report for IPv6 Support

Illumio supports IPv6 for workloads. This includes providing a warning in the Compatibility Report. The Compatibility Report is used to detect the possible issues before moving VEN out of idle state. See "VEN Compatibility Check" in the VEN Installation and Upgrade Guide. In this release, Illumio updated the options in the Compatibility Report to increase it's usability.

The following command and command options are supported:

- On Linux and SunOS, this command option is available regardless of whether IPv6 is enabled:
  - **Ipv6_forwarding_enabled**
    - At least 1 iptables forwarding rule is detected in the IPv6 forwarding chain. VEN removes existing iptables rules in the non-Idle policy state.

- On Windows, we do not support all IPv6 transition tunnels that is a part of the IPv6 transition technology (RFC 4213). The following options are available:
  - **teredo_tunneling_enabled**
    - Teredo tunneling allows for IPv6 connectivity.
    - Teredo is an IPv6 transition tunnel.
    - We do not report on Teredo adapters.
  - **IPv6 enabled**
    - Continues to be supported.
    - Detects potential transition technology usage on Windows.

`illumio-ven-ctl` General Syntax

The `illumio-ven-ctl` is a primary tool for managing VENs on individual workloads. The script varies slightly by platform.

Set PATH Environment Variable

For easier invocation of `illumio-ven-ctl` and other control scripts, set your `PATH` environment variable to the directories where they are located:

- Linux: default location is `/opt/illumio_ven`
- Windows: default location is `C:\Program Files\Illumio`

Command Line Syntax by Platform

Platform	Command	Notes
Linux/AIX/ Solaris	`illumio-ven-ctl`	**IMPORTANT** Parameters for the subcommands are preceded by two hyphens: `--option1 var --option2 var ...`
Windows	`illumio-ven-ctl.exe`	**IMPORTANT** Parameters for the script are preceded by a single hyphen: `-option1 var -option2 var ...`

Linux/AIX/Solaris Command Line Help

```
$ illumio-ven-ctl --help
Usage:  {activate|backup|check-env|conncheck|connectivity-test|deactivate|gen-supportrep
```

Windows Command Line Help

```
illumio-ven-ctl.exe <action> <options>
```

Useful VEN and OS Commands

This topic provides is a short description of the VEN command-line tools that you commonly use for various operations, and some useful native OS commands. Syntax for the VEN-provided commands is detailed throughout this guide, and in the help of the commands themselves.

Additionally, this topic lists the availability of the VEN commands across operating systems.

Verify VEN Version Number

You can verify the version of the VEN software in several different ways:

- View the VEN version in the PCE web console.
- Run the following command on the workload:

  `# /opt/illumio_ven/illumio-ven-ctlversion 21.5.0-xxxx`
- Run the following command on a Windows workload:

  `<VEN Installation Directory>\illumio-ven-ctl.exe version`
- Examine the columns in **Add or remove programs** or Task Manager.
- Examine the **Properties** > **Details** tab of `venAgentMgr.exe` or `venPlatformHandler.exe`.
- Use the Illumio Core REST API. With the REST API, the `agent-version` key and value are returned in the payload of every response.

Commonly Used VEN Commands

Platform	Command	Description	
Linux & macOS	`/opt/illumio_ven/illumio-ven-ctl`	VEN Linux shell control script to control VEN control VEN settings and functions	
	`/opt/illumio_ven/bin/agent_status.sh`	Alternative to `illumio-ven-ctl status`	
	`ps`	Native OS command to list all system processes	
	`chkconfig`	Native OS command to update and query runlevel information for system services	
Windows	`tasklist /svc`	Native OS command to display system services	
	`wf.msc`	Native OS command to manage the Windows firewall	
	AIX/Solaris	`/opt/illumio_ven/illumio-ven-ctl`	VEN AIX/Solaris shell control script to control VEN control VEN settings and functions
`/opt/illumio_ven/bin/agent_status.sh`		Alternative to `illumio-ven-ctl status`	
`/opt/illumio_ven/bin/agent_status.sh`		Alternative to `illumio-ven-ctl status`	
`ps`		Native OS command to list all system processes	
AIX	`lssrc`	Native OS command to list OS subsystem status	
Solaris	`svcs`	Native OS command to list OS service status	

illumio-ven-ctl Command Options by OS

NOTE

Options and subcommands are not yet provided for every command listed below. However, this table may be updated periodically.

The following tables detail the **illumio-ven-ctl** usage constraints and command support by operating system.

Table 3. Usage

/opt/illumio_ven/illumio-ven-ctl <command> [command-options] <command-args>

/opt/illumio_ven/illumio-ven-ctl <command> [command-options] <subcommand> [subcommand-options]

WARNING
illumio-ven-ctl is the only supported way to manage the VEN.

Do not attempt to use any of the following directly:

- Linux **systemd systemctl** commands
- Solaris SMF **svcs** and **svcadm** commands
- Legacy **init.d** start/stop scripts
- Windows Service Control Manager

While the above usage will not break the VEN, it is only designed to work when invoked automatically by the OS at boot or shutdown time.

Table 4. Commands by Operating System

Command	Description	Windows	AIX	CentOS	Debian	RHEL & macOS	Solaris	SUSE	Ubuntu
activate	Activate VEN	Y	Y	Y	Y	Y	Y	Y	Y
check-env	Check VEN **runtime_env.yml** settings	Y	Y	Y	Y	Y	Y	Y	Y
conncheck	Query VEN policy	Y	Y	Y	Y	Y	Y	Y	Y
connectivity-test [-v] [-j] [--test-all-ips]	Test connectivity with PCE	Y	Y	Y	Y	Y	Y	Y	Y
deactivate [--maintenance-token <token>] [--notify-pce <true | false>]	Deactivate VEN without uninstalling	Y	Y	Y	Y	Y	Y	Y	Y
gen-supportreport [-y] [-f <file>] [-b] **Note:** This command does not upload VEN Support Reports to the PCE. Be sure to move VEN Support Reports off the workload as needed.	Generate VEN support reports	Y	Y	Y	Y	Y	Y	Y	Y
prepare	Prepare VEN image	Y	Y	Y	Y	Y	Y	Y	Y
restart [--maintenance-token <token>]	Restart VEN services	Y	Y	Y	Y	Y	Y	Y	Y
set-proxy <server:port> **reset-proxy** **show-proxy** **Note:** For the **set-proxy** command, **server:port** must be specified using one of the following: • IP address of the proxy (for example, `10.10.10.10:8080`) • FQDN of the proxy (for example, `proxy.example.com:8080`)	Manage VEN proxy settings	Y	Y	Y	Y	Y for RHEL No for macoS	Y	Y	Y

Command	Description	Windows	AIX	CentOS	Debian	RHEL & macOS	Solaris	SUSE	Ubuntu
• HTTP or HTTPS schema (for example, `https://proxy.example.com:8080`									
start	Start VEN services	Y	Y	Y	Y	Y	Y	Y	Y
status [-v] [-x | --stdexit]									
status connectivity	Report VEN status	Y	Y	Y	Y	Y	Y	Y	Y
status health									
status policy									
stop [--maintenance-token <token>]	Stop VEN services	Y	Y	Y	Y	Y	Y	Y	Y
suspend [--maintenance-token <token>] [-y]									
Important: The **suspend** command stops the VEN and removes all Illumio rules from the OS firewall, thereby exposing the workload. This is a step further than merely marking the VEN as suspended on the PCE console.	Suspend VEN (enter emergency state)	Y	Y	Y	Y	Y	Y	Y	Y
unpair [--maintenance-token <token>] <saved | open | recommended> [noreport]									
Subcommands:									
<saved | open | recommended>	Unpair VEN	Y	Y	Y	Y	Y	Y	Y	Y
Subcommand arguments:									
[noreport]									
unsuspend [--maintenance-token <token>] [-y]	Unsuspend VEN (exit emergency state)	Y	Y	Y	Y	Y	Y	Y	Y

Command	Description	Windows	AIX	CentOS	Debian	RHEL & macOS	Solaris	SUSE	Ubuntu
version	Display VEN version	Y	Y	Y	Y	Y	Y	Y	Y

Notes:

--maintenance-token <token>	Specify the maintenance **<token>** that will authorize the subcommand. This option is not needed if a maintenance token was not generated by the PCE.
--notify-pce	Specify whether (true) or not (false) to notify the PCE that the VEN has been deactivated. By default the PCE is always notified.
-b	Block and do not exit until this command completes. By default this command exits after work is queued in background.
-f <file>	The original support report is always saved as **/opt/illumio_ven_data/reports/illumio-agent-report.tgz**. Save another copy as the specified **<file>** (can include an absolute path).
-j	Enable JSON output.
--stdexit	Use the following exit codes: 0 = all VEN process running; 1 = error or partialy running; 3 = no VEN process running.
--test-all-ips	Instead of using default OS name resolution to test a single PCE IP address, explicitly resolve and test all IP addresses returned for the PCE FQDN.
-v	Enable verbose output.
-x	Synonym for **--stdexit**
-y	Assume **yes** for all yes/no prompts, don't prompt for confirmation. By default, this command prompts for confirmation.
saved	Subcommand used with **unpair**. Corresponds to PCE UI "Remove Illumio policy." Restore firewall as it was when VEN was installed. Dangerous if the VEN was installed long ago since old firewall is probably stale and incorrect.
open	Subcommand used with **unpair**. Corresponds to PCE UI "Open all ports." Do not block any traffic after uninstalling. User is expected to create a new firewall (current firewall won't survive reboot).
recommended	Subcommand used with **unpair**. Corresponds to PCE UI "Close all ports except remote management." User is expected to create a new firewall (current firewall won't survive reboot). Remote management includes SSH, RDP, and WinRM.
noreport	Subcommand argument used with **unpair**. Do not generate a support report before uninstalling.

VEN State

This section describes all the VEN's states and how you can manage them. VEN state refers to the active state of the VEN on a workload; basically, is it running, stopped, enabled, disabled, or suspended.

VEN Startup and Shutdown

This topic provides information on starting and stopping VENs.

VEN Startup and Shutdown (illumio.com)

- AIX and Solaris: Start up the VEN.
- AIX and Solaris: Shut down the VEN and send a Goodbye message.

Start Up VENs

The VEN starts when the workload is booted from the system boot files. The VEN can also be started manually.

Automatic Startup

The VEN starts when the workload is booted from system boot files:

Platform	Command	Notes

Linux/AIX/Solaris	`/etc/rc.d/init.d/illumio-ven`
	Or
	`/etc/init.d/illumio-ven`
		Installs firewall kernel modules if necessary, sets firewall to the desired state.

	CentOS/RHEL 7+, starting from 19.3.2 `/usr/lib/systemd/system/illumioven.service`	Initializes and starts the daemon processes needed for VEN operation. **IMPORTANT** This command is only supported in Illumio Core 19.3.2-VEN and later.
Windows	None needed.	The Service Control Manager (SCM) starts all VEN services at boot.

Manual Startup

The VEN can also be started manually with `illumio-ven-ctl start`.

Platform	Command
Linux/AIX/Solaris/RHEL/CentOS	`/opt/illumio_ven/illumio-ven-ctl start`
Windows	`C:\Program Files\Illumio\illumio-ven-ctl.ps1 start`

Shut Down VENs

At shutdown, the VEN sends a “goodbye” message to the PCE. The PCE marks the workload as offline and initiates a policy recomputation. After the new policy is distributed throughout the network, the workload without the VEN is effectively isolated from the network.

Linux/AIX/Solaris Workload Shutdown

Platform	Command	Notes
Linux/AIX/Solaris/RHEL/ CentOS	`illumio-ven-ctl stop`	• Stops all VEN processes. • The VEN sends a “goodbye” message to the PCE.
Windows	None needed.	• Service Control Manager (SCM) stops all VEN services. • The VEN sends a “goodbye” message to the PCE.

Disable and Enable VENs (Windows only)

If you want to install the VEN but activate it later, you can disable the VEN after you first install it. This is only available on the Windows platform.

For example, you can load the VEN on machine image and disable the VEN. See considerations regarding preparing a "Golden Master" in the VEN Installation and Upgrade Guide.

Platform	Action	Command

Windows	• Enable
	• Disable:
		`PS C:\Program Files\Illumio> .\illumio-ven-ctl.ps1 enable`
		
		`PS C:\Program Files\Illumio> .\illumio-ven-ctl.ps1 disable`

VEN Suspension

If users are not able to reach an app on a workload, you can suspend the VEN to see if the VEN was causing the issue. The VEN suspension feature allows you to isolate a VEN on a workload to troubleshoot any communication issues with that workload, and to determine if the VEN is the cause of the anomalous behavior.

IMPORTANT

Security Implications: When the VEN is suspended, the workload firewall rules are removed leaving the VEN open and all traffic is allowed.

About VEN Suspension

When a VEN is suspended, the following is true:

- Any rules programmed into the workload's iptables (including Custom iptables rules), Windows Filtering Platform (WFP), or ipfilter, or pf firewalls are removed completely, and all VEN software processes are shut down.
- The VEN connectivity and policy sync status are changed to **Suspended**.

- The VEN informs the PCE that it is in the suspended state. If the PCE does not receive this notification, you must mark the workload as **Suspended** in the PCE web console.
- If the PCE does not receive the VEN suspension notification and you do not mark the VEN as suspended in the PCE, after one hour, the PCE assumes the workload is offline and removes it from the policy, which effectively isolates the workload from the network. For example, users will not be able to reach apps on the workload.
- Workloads communicating with the suspended VEN continue to have their rules programmed into iptables or WFP.
- The SecureConnect policy continues to be in effect while the VEN is suspended.
- An organization event (`server_suspended`) is logged. This event is exportable to CEF/LEEF and has a severity of WARNING.

Properties of a suspended VEN:

- The workload continues to appear in the PCE in the workloads list page and Illumination map.
- You can unpair a workload while its VEN is suspended.
- You can change the policy state of the workload in the PCE Web Console while the VEN is suspended.
- When the VEN is unsuspended, the new policy state is applied.
- Heartbeats or other communication is not expected, but if one is received, any communication is logged by the PCE.
- If the PCE is rebooted, the VEN remains suspended.

When a VEN is unsuspended:

- The PCE is informed that the VEN is no longer suspended and can now receive policy from the PCE.
- If existing Rules affect the unsuspended workload, the PCE will reprogram those Rules.
- An organization event (`server_unsuspended`) is logged. This event is exportable to CEF/LEEF and has a severity of WARNING.
- The workload will revert to its policy state prior to Suspended.
- Custom iptables Rules are configured back into the iptables.

You can manage VEN suspension by using these features of the Illumio Core:

- The REST API

  For more information on this method, see "VEN Operations" in the REST API Developer Guide.
- The command line
- The PCE web console

  For more information, see Mark VEN as Suspended Using the PCE Web Console [240] in this topic.

Linux VEN: Back Up Custom iptables/NAT Rules

NOTE

Before suspending a Linux VEN, back up the workload PCE custom iptables filter or NAT rules.

After a workload is suspended, restore the rules on the workload because all custom iptables filter or NAT rules will have been removed from the workload.

Suspend and Unsuspend Commands

Platform	Action	Command	Notes

Linux/ Unix	• Suspend
	• Unsuspend
		`$ illumio-ven-ctl suspend`
		`Suspending the VEN... The VEN has been suspended. PCE was notified.`
		
		`$ illumio-ven-ctl unsuspend`
		`Unsuspending the VEN... The VEN has been unsuspended. PCE was notified.`
			On Linux, be sure to backup your custom configuration.
			
			See Linux VEN: Back Up iptables/ NAT Rules [239].

Windows	• Suspend
	• Unsuspend
		`<VEN Installation Directory>\illumio-ven-ctl.exe suspend`
		`Suspending the VEN... The VEN has been suspended. PCE was notified.`
		
		`<VEN Installation Directory>\illumio-ven-ctl.exe unsuspend`
		`Unsuspending the VEN... The VEN has been unsuspended. PCE was notified.`	

Mark VEN as Suspended Using the PCE Web Console

In addition to using the command explained in the previous section, you can mark a workload as **Suspended** using the PCE web console.

NOTE

Marking a workload as **Suspended** in the PCE web console does **not** actually suspend the VEN. It should only be used if the VEN went offline before it could be suspended. Marking the workload as **Suspended** is a way to keep the PCE from removing the VEN from the policy and isolating it from the rest of the network.

To mark a VEN Suspended:

1. Go to **Servers & Endpoints > Workloads**.
2. Click the **VENs** tab.
3. Click the name of the VEN you want to mark as suspended.
4. On the VEN's detail page, click **Mark as Suspended**.
5. Click **Suspend** to confirm the VEN suspension.

The number of suspended workloads is displayed at the top of the page and the suspended workload is displayed on the Workloads page with a red "Suspended" icon.

To clear a VEN's Suspension status:

1. Go to **Servers & Endpoints > Workloads**.
2. Click the **VENs** tab.
3. Click the name of a VEN marked as suspended that you want to mark as unsuspended.
4. On the VEN's detail page, click **Clear Suspension**.
5. Click **Clear** to confirm.

Disable VEN Suspension on Workloads

You can disable the ability to suspend a VEN on a workload. To disable the VEN suspension feature, define the following environment variable for the VEN. How you set the variable varies by VEN platform. See the procedures to set the environment variable for each platform.

Environment Variable	Values
VEN_NO_SUSPEND	1 – Disable VEN suspension 0 – VEN suspension is enabled

NOTE

Disabling VEN suspension is not supported for Illumio Secure Cloud customers.

Linux VENs

Before installing or upgrading the Linux VEN, enter the following command line syntax to set the environment variable:

```
# VEN_NO_SUSPEND=1 <ven_install_or_upgrade_command>
```

Examples:

```
# VEN_NO_SUSPEND=1 rpm -i <illumio-ven-pkg>.rpm

# VEN_NO_SUSPEND=1 dpkg -i <illumio-ven-pkg>.deb

# VEN_NO_SUSPEND=1 rpm -U <illumio-ven-pkg>.rpm
```

Windows VENs

Disabling the suspend command:

```
<ven_installation_filename>.exe <options> VEN_NO_SUSPEND=1
```

Available options include:

- `/install`
- `/log logfile.log`
- `/quiet`

Example:

```
ven_install_filename.exe /install EN_NO_SUSPEND=1
```

AIX VENs

Before installing or upgrading the AIX VEN, enter the following command line syntax to set the environment variable:

```
# VEN_NO_SUSPEND=1 <ven_install_or_upgrade_command>
```

Example:

```
# VEN_NO_SUSPEND=1 installp -acXgd <path_to_bff_package> illumio-ven
```

Solaris VENs

When you install the Solaris VEN by interactively responding to installer prompts, enter `n` at the following prompt:

```
"Do you want to disable VEN suspend? [y,n] ", enter as required : y - disable, n - defau
```

When you use the template file in the VEN package to pre-load responses to installer prompts, copy the following file:

```
illumio-ven/root/opt/illumio_ven/etc/templates/response
```

Change the copied file in the following way:

```
/usr/xpg4/bin/sed 's/^VEN_NO_SUSPEND=0/VEN_NO_SUSPEND=1/g’ \
< illumio-ven/root/opt/illumio_ven/etc/templates/response \
> illumio-ven/root/opt/illumio_ven/etc/templates/response.custom
```

Deactivate and Unpair VENs

VEN Deactivation and Unpairing

This section describes all the ways that you can change the VEN software running on a workload, from reverting it to an earlier release, deactivating the software, or uninstalling it completely.

This section describes how to deactivate and unpair the VEN software.

Deactivate and Unpair VENs

This topic describes how to deactivate and unpair VENs by operating system. Additionally, it explains the security implications for performing these tasks and makes recommendations on how to properly deactivate and unpair VENs.

See VEN Unpairing Details [245].

Deactivate Using VEN Command Line

To deactivate the VEN, you must use the `illumio-ven-ctl` command.

`deactivate` breaks the PCE-to-workload connection but doesn't uninstall the VEN software (as `unpair` would).

After deactivation, the workload reverts to its pre-Illumio native firewall settings.

Linux/AIX/Solaris

```
# /opt/illumio_ven/illumio-ven-ctl deactivate
```

Windows

```
<VEN Installation Directory>\illumio-ven-ctl.exe deactivate
```

Unpair Using VEN Command Line

The `unpair` command breaks the PCE-to-workload connection, and uninstalls the VEN software. The `unpair` command gives you control over the post-unpair state, as described below.

Linux/AIX/Solaris

With `illumio-ven-ctl unpair`, specify the post-unpair state for the VEN:

```
# /opt/illumio_ven/illumio-ven-ctl unpair [recommended | saved | open]
```

NOTE

On Linux, the `unmanaged` option is not available.

Unpair Options on Linux/AIX/Solaris

- `recommended`: Uninstalls the VEN and temporarily allows only SSH/22 until reboot.

IMPORTANT

Security Implications: When the workload is running a production application, it could break because this workload will no longer allow any connections to it other than SSH on port 22.

- `saved`: Uninstalls the VEN and reverts to pre-Illumio policy to the state before the VEN was first installed. Revert the state of the workload's iptables to the state before the VEN was installed. The dialog displays the amount of time that has passed since the VEN was installed.

  IMPORTANT

  Security Implications: Depending on how old the iptables configuration is on the workload, VEN removal could impact the application.

- `open`: Uninstalls the VEN and leaves all ports on the workload open.

  IMPORTANT

  Security Implications: When iptables or Illumio are the only security being used for the workload, the workload is open to anyone and becomes vulnerable to attack.

Windows

With `illumio-ven-ctl.ps1` unpair, specify the post-deactivation state for the VEN:

```
<VEN Installation Directory>\illumio-ven-ctl.exe unpair [recommended | saved | open | un
```

Unpair Options on Windows

- `recommended`: Temporarily allow only RDP/3389 and WinRM/5985,5986 until reboot.

  IMPORTANT

  Security Implications: If the workload is running a production application, the application could break because the workload no longer allows any connections to it.

- `saved`: Restores firewall rules and configuration to the state it was in at the time the workload was paired. Reverts the state of the firewall to before Illumio was installed.

  IMPORTANT

  Security Implications: Depending on how old the WFP configuration was on the workload, VEN removal could impact the application.

- `open`: Uninstalls the VEN and leaves all ports on the workload open.

  IMPORTANT

  Security Implications: When WFP or the PCE are the only security being used for the workload, the workload is open to anyone and becomes vulnerable to attack.

- `unmanaged`: Uninstalls the VEN and reverts to the workload's currently configured Windows Firewall policy.

Unpair Using System Commands

You can use the `illumio-ven-ctl` (Linux/AIX/Solaris) or `illumio-ven-ctl.ps1` (Windows) to unpair the VEN.

IMPORTANT

As an alternative, you can use the `system uninstall` command to unpair the VEN, however it is not recommended. This command should only used as a fallback if there are issues with unpairing with `illumio-ven-ctl` or `illumio-ven-ctl.ps1`.

Linux

- RPM: `rpm -e illumio-ven`
- DPKG: `dpkg -P illumio-ven`

Windows

- Use the Control Panel to uninstall the VEN.

AIX

- `installp -u illumio-ven`

Solaris

- `pkgrm illumio-ven`

VEN Unpairing Details

During unpairing, the VEN performs the following actions. These actions are specific to the workload operating system.

Linux/AIX/Solaris

- Unpairs the VEN from the PCE.
  - Sends a "deactivate" message to the PCE.
- Restores the host firewall state to the requested or open state if no state is specified. Possible values of the state are:
  - Open: All ports are open after VEN uninstalls.
  - Saved: The firewall is restored to its state just before the VEN was installed.
- Uninstalls the `illumio-ven` package.
  - Removes program and data files.
  - Removes repo and GPG files and package.

Windows

- Unpairs the VEN from the PCE.
  - Sends a “deactivate” message to PCE.
- Stops all VEN services.
- Unregisters services from Service Control Manager.
- Restores Windows Firewall to requested state.
  - Open: All ports are open after VEN uninstalls.
  - Saved: Restore the firewall to its state just before the VEN was installed.
- Removes Program Files and `ProgramData` directories.
- Removes VEN registry keys.
- Removes Certificate.
- Unregisters VEN Event provider.

Support Report During Unpairing

When you unpair a workload, the VEN creates a local Support Report for diagnostic purposes in case you need a record of the VEN after it is uninstalled.

On Linux/Unix, the generated Support Report is saved to the `/tmp` directory. On Windows, the generated Support Report is saved to the `C:\Windows\Temp` directory. If a there is an existing Support Report in this directory, it will be overwritten with the new one.

Monitor and Diagnose VEN Status

This section provides you with the necessary information to monitor VEN status on your workloads and to troubleshoot any problems that might occur.

VEN-to-PCE Communication

This topic discusses how the VEN communicates with the PCE for both Illumio Core Cloud customers and Illumio Core On-Premises customers.

Details about VEN-to-PCE Communication

On Prem

The VEN, by default, communicates with the PCE when installed in customers data centers (On-Premises) over the following ports:

- Port 8443 - HTTPS requests
- Port 8444 - long-lived TLS-over-TCP connection

SaaS

The VEN communicates with the Illumio Core Cloud PCE over Port 443 for both HTTPS requests and the long-lived TLS-over-TCP connection.

The VEN uses Transport Level Security (TLS) to connect to the PCE. The PCE certificate must be trusted by the VEN before communication can occur.

The VEN sends the following details to the PCE:

- Regular heartbeat with the latest hostname and other properties of the workload
- Traffic log
- Network interfaces
- Processes
- Open ports
- Interactive users (Windows only)
- Container workload information (C-VEN only)

The VEN receives the following details from the PCE:

- Firewall policy
- Lightning bolts/heartbeat responses with action to perform, such as sending a support report

PCE Certificate Verification

Keep in mind the following:

- The VEN requires that the full certificate chain, *up to but not including a self-signed root certificate trusted by the OS*, be sent as part of the TLS handshake with the PCE.
- The PCE will always send the full certificate chain, *minus the root certificate*.
- If a “Man In The Middle" (MITM) device with TLS inspection capability is deployed on a path between the VEN and the PCE, Illumio recommends bypassing such capabilities for VEN-to-PCE communication:
  - Some MITM devices that forge the PCE certificate will not send the full certificate chain, resulting in a TLS failure with some VEN and OS combinations.
  - Illumio does not test coexistence with any MITM devices. With respect to compatibility with partial certificate chains in the TLS handshake, the behavior of the VEN and the behavior of the MITM device may change at any time without notice on either side.

Configurable Time for Heartbeat Warning

You can change the threshold for the time the VEN goes without a heartbeat and goes into the Warning state. To change the 15-minute threshold in the PCE interface:

1. Go to **Settings** > **Offline Timers**.
2. Click **Edit**.
3. In the **Disconnect and Quarantine** section, select **Custom Timeout**.
4. Specify a wait time.
5. Click **Save**.

VEN Connectivity

- **Online:** The workload is connected to the network and can communicate with the PCE.
- **Offline:** The workload is *not* connected to the network and cannot communicate with the PCE.
- **Suspended:** The VEN is in the suspended state and any rules programmed into the workload's IP tables (including custom iptables rules) or Windows filtering platform firewalls are removed completely. No Illumio-related processes are running on the workload.

VEN Support for IPv6 Traffic

You can configure how VENs support IPv6 traffic. Go to **Settings > Security** and click the General tab:

For VEN releases 20.2.0 and later, choose one of these options:

- Allow IPv6 traffic according to your policy
- Block IPv6 traffic only when in Full Enforcement. (Traffic will always be allowed on AIX and Solaris workstations.)

For VEN releases pre-20.2.0, choose one of these options:

- Allow all IPv6 traffic
- Block IPv6 traffic only when in Full Enforcement. (Traffic will always be allowed on AIX and Solaris workstations.)

Communication Frequency

The following table shows the frequency of communications to the PCE for common VEN operations. See PCE Administration Guide for more details about these intervals and their effects.

Function	Frequency	Notes
Firewall policy updates	Real-time if lightning bolts are enabled.	If lightning bolts are displayed or the channel is not functional, policy updates are communicated to the VEN by a heartbeat action.
Active service reporting	See note.	• AgentManager performs all active service reporting tasks. • At start-up, a snapshot of processes and ports is sent to the PCE. • Every 24 hours, a snapshot of *all* listening processes is taken and sent to the PCE.
Interface reports and changes	Event driven.	Only if there are changes to the interfaces; otherwise, no data are sent.
Traffic flow log	Every 10 minutes.	• The VEN checks if there are logs, and if so, sends them to the PCE. • If the PCE is inaccessible, the VEN retains flow summaries for the previous 24 hours but purges logs that are older than 24 hours, with the oldest log at every 24-hour mark. • When logs are purged, the VEN locally logs an alert, which is posted to the PCE as an event when connectivity is restored.
Heartbeat	Every 5 minutes.	If the PCE does not receive three consecutive heartbeats, an event is written to the PCE's event log. See also VEN Heartbeats and Lost Agents [248].
Dead-peer interval	Configurable	Default is 60 minutes (or 12 heartbeats). See also VEN Offline Timers and Isolation [249].
VEN tampering detection	Within a few seconds on Windows and Linux.	For more information, see Host Firewall Tampering Protection [262].

VEN Heartbeats and Lost Agents

The VEN sends a heartbeat message every five minutes to the PCE to inform the PCE that it is up and running. If the VEN fails to send a heartbeat, check the workload where the VEN

is installed and investigate any connectivity issues. If the VEN continues to fail to send a heartbeat, it eventually is marked Offline, which means it can no longer communicate with the PCE or other managed workloads.

PCE down or network issue and the VEN degraded state

- If the VEN cannot connect to the PCE either because the PCE is down or because of a network issue, the VEN continues to enforce the last-known-good policy while it tries to reconnect with the PCE.
- After missing three heartbeats, the VEN enters the *degraded state*. In the degraded state, the VEN ignores all the asynchronous commands received as lightning bolts from the PCE, except the commands for software upgrades and support reports.
- After connectivity to the PCE is restored, the VEN comes out of the degraded state after three successful heartbeats.

Failed authentication and the VEN minimal state

- If the VEN enters the degraded state because of failed authentications, the VEN enters a state called *minimal*. In the minimal state, the VEN only attempts to connect with the PCE every four hours through a heartbeat.
- If the authentication failure was temporary, the VEN exits the minimal state after its first successful connection to the PCE. Whenever the VEN enters the minimal state, it stops the VTAP service. VTAP is then restarted when the VEN exits the minimal state.
- If Kerberos authentication is used, the VEN attempts to refresh the agent token with a new Kerberos ticket before sending a heartbeat. If the authentication error is not recovered after four hours, the VEN sends a lost-agent message to the PCE which then logs a message in the Organization Events. The message informs the user that the VEN needs to be uninstalled or reinstalled manually on this workload.

VEN Offline Timers and Isolation

When the VEN on a workload is stopped, the VEN makes a "best effort" REST API goodbye call to the PCE. After a delay specified by the "workload goodbye timer" (a default of 15 minutes), the PCE marks the workload offline and removes it from the policy.

If the REST API call (goodbye) fails, or if the workload goes offline abruptly (for example, due to a power outage), the PCE stops receiving heartbeats from the workload. After the period of time configured in the PCE web console **Settings > Offline Timers** elapses, the PCE marks the workload offline and recomputes policies for the peer workloads to isolate the offline workload. If no time period has been configured, the default is 60 minutes, or 12 heartbeats.

The `system_task.agent_missed_heartbeats_check` alert triggers an alert to be sent at 25% of the time configured in the offline timer. For example, if the offline timer is configured to 1 hour, an alert is sent after the VEN has not sent a heartbeat for 15 minutes; if the offline timer is configured to 4 hours, an alert is sent after the VEN hasn't sent a heartbeat for 1 hour. If a user has customized the timer, the event will show up when 25% of the timer has elapsed.

Sampling Mode for VENs

If the VEN receives a sustained amount of high traffic per second from many individual connections, the VEN enters Sampling Mode to reduce the load. Sampling Mode is a protection mechanism to ensure that the VEN does not contribute to the consumption of CPU. In Sampling Mode, not every flow is reported. Instead, flows are periodically sampled and logged.

After CPU usage on the VEN decreases, Sampling Mode is disabled and each connection is reported to the VEN. The entry and exit from sampling-mode is automatically performed by the VEN depending on the load on the VEN.

Details about entering and exiting Sampling Mode are captured in `/opt/illumio_ven_data/log/vtap.log`. Look for `Entering` and `Exiting throttle state`.

Linux TCP Timeout Variable

For VENs installed on Linux workloads, the VEN relies on conntrack to manage the `nf_conntrack_tcp_timeout_established` variable.

By default, as soon as the VEN is installed, it sets the `nf_conntrack_tcp_timeout_established` frequency to eight hours (28,800 seconds). Setting this frequency manages workload memory by removing unused connections from the table and thereby increasing performance.

If you change the frequency via `sysctl`, it is reverted the next time the workload is rebooted or the next time the VEN's configuration file is read.

Wireless Connections and VPNs

The Illumio Core VEN supports wireless connections for VENs installed on endpoints in the Illumio Core.

For more information about installing the VEN on an endpoint, and supporting a wireless network connection, see the *Endpoint Installation and Usage Guide*.

NOTE

Wireless network support is only available for endpoints in Illumio Core. It is not available for other support server types, such as bare-metal servers, virtual machines (VMs), or container hosts.

Show Amount of Data Transfer

The operation of 'show amount of data transfer' capability on the PCE is a preview feature available with the 20.2.0 release. The PCE now reports amount of data transferred in to and out of workloads and applications in a datacenter. The number of bytes sent by and received by the provider of an application are provided separately. These values can be seen in traffic flow summaries streamed out of the PCE. This capability can be enabled on a per-workload basis in the Workload page. It can also be enabled in the pairing profile so that workloads are directly paired into this mode.

After the feature is enabled, the VEN starts reporting the number of bytes transferred over the connections. The PCE collects this data, adds relevant information, such as, labels and sends the traffic flow summaries out of the PCE.

The direction reported in flow summary is from the viewpoint of the provider of the flow.

- Destination Total Bytes Out (`dst_tbo`): Number of bytes transferred out of provider (Connection Responder)
- Destination Total Bytes In (`dst_tbi`): Number of bytes transferred in to provider (Connection Responder)

The number of bytes includes:

1. L3 and L4 header sizes of each packet (IP Header and TCP Header)
2. Sizes of multiple headers that may be included in communication (when SecureConnect is enabled)
3. Retransmitted packets.

   The bytes transferred in the packets of a connection are included in measurement. This is similar to various networking products such as firewalls, span-port measurement tools, and other network traffic measurement tools that measure network traffic.

Term	Description
dst_tbi	**Destination Total Bytes** In Total bytes received till now by the destination over the flows included in this flow-summary in the latest sampled interval. This is the same as bytes sent by the source. Present in 'A', 'C', and 'T' flow-summaries. source = client = connection initiator, destination = server = connection responder.
dst_tbo	**Destination Total Bytes** Out Total bytes sent till now by the destination over the flows included in this flow-summary in the latest sampled interval. This is the same as bytes received by the source. Present in 'A', 'C', and 'T' flow-summaries. source = client = connection initiator, destination = server = connection responder.
dst_tbi	**Destination Delta Bytes** In Number of bytes received by the destination in the latest sampled interval, over the flows included in this flow-summary. This is the same as bytes sent by the source. Present in 'A', 'C', and 'T' flow-summaries. source = client = connection initiator, destination = server = connection responder.
dst_dbo	**Destination Delta Bytes** Out Number of bytes sent by the destination in the latest sampled interval, over the flows included in this flow-summary. This is the same as bytes received by the source. Present in 'A', 'C', and 'T' flow-summaries. source = client = connection initiator, destination = server = connection responder.
interval_sec T	**Time Interval in Seconds** Duration of latest sampled interval over which the above metrics are valid.

Connection State	Description
A	**Active:** The connection is still active at the time the record was posted. Typically observed with long-lived flows on source and destination side of communication.
T	**Timed Out:** Flow does not exist any more. It has timed out. Typically observed on destination side of communication.
C	**Closed:** Flow does not exist any more. It has been closed. Typically observed on source side of communication.
S	**Snapshot:** Connection was active at the time VEN sampled the flow. Typically observed when the VEN is in Idle state.

VEN Status Command and Options

This topic describes various commands for determining the status of a VEN. Log in as root to run these commands.

Command

The `status` command returns the status of the VEN on the workload.

```
illumio-ven-ctl status
```

Linux/AIX/Solaris VENs

```
# /opt/illumio_ven/illumio-ven-ctl status
```

Windows VENs

```
C:\Program Files\Illumio\illumio-ven-ctl status
```

Return parameters

Linux

```
Status for illumio-control:
 - Environment Illumio VEN Environment is setup
 - venAgentMgr venAgentMgr (pid 23598) is running...
 - IPSec IPSec feature not enabled
 - venPlatformHandler venPlatformHandler (pid 23676) is running...
 - venVtapServer venVtapServer (pid 23737) is running...
 - venAgentMonitor active(running)

Agent state: enforced
```

Windows

```
Service venAgentMgrSvc:         Running
Service venPlatformHandlerSvc:  Running
Service venVtapServerSvc:       Running
```

```
Service venAgentMonitorSvc:     Running
Service venAgentMgrSvc:         Enabled
Service venPlatformHandlerSvc:  Enabled
Service venVtapServerSvc:       Enabled
Service venAgentMonitorSvc:     Enabled
```

Field definitions

Linux/AIX/Solaris

Name	Definition
Environment	Whether or not the Illumio VEN environment is setup
venAgentMgr	**venAgentMgr** status, and if running its pid
IPSec	Whether or not the **IPSec** feature is enabled
venPlatformHandler	**venPlatformHandler** status, and if running its **pid**
venVtapServer	**venVtapServer status**, and if running its **pid**
venAgentMonitor	**venAgentMonitor** status
Agent state	For example, **enforced**QQ

Options

This section describes these options:

- Policy
- Health
- Connectivity

Policy option

```
illumio-ven-ctl status policy
```

Th `policy` option returns the timestamp, ID, and state of the current security policy the VEN received from the PCE.

Linux/AIX/Solaris

```
# /opt/illumio_ven/illumio-ven-ctl status policy
```

Windows

```
C:\Program Files\Illumio> .\illumio-ven-ctl.ps1 status policy
```

Return parameters

Windows

```
{
   "timestamp" : "2019-06-14T00:41:41Z",
```

```
   "id" : "xxxxxxxx940d0f4c2531b0d44400523dae055674-xxxxxxxx7a6796c210fb846b0321847bc22d
   "state" : "enforced"
}
```

Field definitions

Linux/AIX/Solaris

Policy Field Name	Definition
`timestamp`	Time the policy was received from the PCE (Local time + UTC offset)
`id`	ID of the security policy (computed locally)
`state`	Policy state (for example, `enforced`)

Health option

```
illumio-ven-ctl status health
```

The `health` option shows whether or not the VEN can write logs locally.

NOTE

This is not the same as PCE health.

Linux/AIX/Solaris VENs

```
# /opt/illumio_ven/illumio-ven-ctl status health
```

Windows

```
<VEN Installation Directory>\illumio_ven\illumio-ven-ctl status health
```

Return parameters

Windows

```
{
  "results": [
    {
      "test": "VEN has write access to the log directory",
      "result": "pass"
    }
  ],
  "state": "healthy"
}
```

Field definitions

Linux/AIX/Solaris

Field Name	Definition
`results`	Array of test results
`test`	VEN has write access to the log directory
`result`	"`pass`" or an error
`state`	VEN health status ("healthy" or "unhealthy"); "`healthy`" means the VEN can write logs locally

Connectivity option

The `connectivity` option returns the status of the VEN connectivity with the PCE.

```
illumio-ven-ctl status connectivity
```

Linux/AIX/Solaris

```
/opt/illumio_ven/illumio-ven-ctl status connectivity
```

Windows

```
C:\Program Files\Illumio\illumio-ven-ctl status connectivity
```

Return parameters

```
{
   "connectivity" : {
      "ips_returned" : 1,
      "pce" : "someName.someDomain",
      "port" : 8443,
      "results" : [
         {
            "ip" : "xx.xx.xxx.xxx",
            "result" : "pass",
            "http_code" : 204
         }
      ]
   },
   "last_successful_hb" : "2019-06-14T04:10:28Z",
   "time_now" : "2019-06-14T04:14:06Z"
}
```

VEN Logging

The VEN captures logs of its operation and traffic flow summaries locally on the workload. There are several different application log files, each with one backup. Application logs are rotated from primary to backup when their size reaches 15 MB. Application log files are preserved at reboot, because application logs are stored in files on a workload.

VEN Traffic Logging

The VEN stores traffic flow summaries, rather than each individual traffic flow. For each connection, the traffic flow summary includes:

- Source IP
- Destination IP
- Destination Port
- Protocol
- Number of connections

Querying Flow Log Databases

The `sqlite` command-line tool, which comes with the VEN, is used to query the flow log databases.

Linux/AIX/Solaris Database Query Examples

Query Type	Example
Non-aggregated accepted flows	`/opt/illumio_ven/bin/sqlite3 /opt/illumio_ven_data/log/flow.db "select * from flow_view"`
Non-aggregated dropped flows	`/opt/illumio_ven/bin/sqlite3 /opt/illumio_ven_data/log/flow.db "select * from drop_flow_view"`
Aggregated accepted flows	`/opt/illumio_ven/bin/sqlite3 /opt/illumio_ven_data/log/flowsum.db "select * from flow_view"`
Aggregated dropped flows	`/opt/illumio_ven/bin/sqlite3 /opt/illumio_ven_data/log/flowsum.db "select * from drop_flow_view"`

Window Database Query Examples

Query Type	Example
Non-aggregated accepted flows	`"c:\Program Files\Illumio\bin\sqlite.exe" c:\ProgramData\Illumio\log\flow.db "select * from flow_view"`
Non-aggregated dropped flows	`"c:\Program Files\Illumio\bin\sqlite.exe" c:\ProgramData\Illumio\log\flow.db "select * from drop_flow_vi`
Aggregated accepted flows	`"c:\Program Files\Illumio\bin\sqlite.exe" c:\ProgramData\Illumio\log\flowsum.db "select * from flow_view`
Aggregated dropped flows	`"c:\Program Files\Illumio\bin\sqlite.exe" c:\Program Data\Illumio\log\flowsum.db "select * from drop_flo`

List of Local Processes

The names of local process are captured in traffic flow data and stored in the PCE.

OS	Description
Windows	Indicates whether auto resize of the Conntrack table is required.
Linux, AIX, and Solaris	The VEN monitors the list of all processes with listening ports on TCP and UDP inbound connections, then matches process names to the list. Refreshes occur every 30 seconds. This process allows for a lower impact on the CPU.

The data can be exported in near-real-time to a Security Information and Event Management (SIEM) or another collector.

VEN Firewall Script Logging

The Illumio firewall scripts log all errors and other key information into the `platform.log` file. This log file can help Illumio debug issues.

Traffic Flow Query Report

You can generate, schedule, and email reports which are based off saved and recent filters from Explorer for reporting. The CSV report is downloadable and can be emailed to the user.

Tuning the IPFilter State Table (AIX/Solaris)

In versions 11.3 and earlier, you can tune the IPFilter state table for AIX and Solaris workloads. Solaris versions before 11.4, you must tune the IPFilter state table. In version 11.4 and after, you must tune the packet filter.

About State Table Tuning

In most environments, the state table default values are sufficient to handle the number of network connections encountered by Solaris and AIX workloads. However, if your system has a very large number of network connections, you might need to tune the state table. You can do so either before or after VEN activation. Tuning the state table values persists through rebooting, restarting, and suspending the VEN.

By default, Solaris and AIX VENs are installed with the following state table values:

- `fr_statemax`: 1,000,000
- `fr_statesize`: 250,007
- `fr_state_maxbucket`: 256
- `fr_tcpclosed`: 120

Set a Custom IPFilter State Table Size

1. Create the following file on your Solaris or AIX workload as `root` or the Illumio VEN user, `ilo-ven`.

   

   NOTE

   The following file that must be created by the `root` user or the Illumio VEN user `ilo-ven`: `/etc/default/illumio-agent`.

   This file cannot be world-readable or -writeable.

2. Add the following settings and values to the file. Do not include spaces in the settings or values.

VEN File Setting	ipfilter Setting	Description
`IPFIL-TER_STATE_MAX=<value>`	`fr_statemax`	Maximum number of network connections stored in the state table. You must also set `IPFILTER_STATE_SIZE`.
`IPFIL-TER_STATE_SIZE=<value>`	`fr_statesize`	Size of the hash table. Must be a prime number. You must also set `IPFIL-TER_STATE_MAX`. **Recommended:** Set the hash table size to 1/4 of the number in `fr_statemax`. This setting allows each hash bucket to contain about 4 states.
`IPFILTER_STATE_MAXBUCK-ET=<value>`	`fr_state_max-bucket`	Number of allowed hash collisions before the VEN starts dropping network connections **Recommended:** Increase this value beyond the default value to avoid dropping network connections.
`IPFIL-TER_TCPCLOSED=<value>`	`fr_tcpclosed`	Option introduced and supported for Illumio Core 21.2.1 VEN and later. To support NFS traffic so that the workload does not drop this traffic even when a rule exists in the PCE allowing the traffic. This issue occurs due to TCP port number reuse. **Recommended:** Illumio customers have found that setting the value for the IPFILTER_TCPCLOSED option to 2 (2 equals 1 second) resolved the issue.

NOTE

If you set `IPFILTER_STATE_MAX`, you must also set `IPFILTER_STATE_SIZE`. If you add only one of these settings in the `illumio-agent` file, the VEN ignores the value and uses default values for both settings.

3. This step depends on whether the VEN has been activated.
   - If the VEN has not yet been activated, skip this step.
   - If the VEN has been activated, restart the VEN by entering the following command:

   ```
   /opt/illumio_ven/illumio-ven-ctl restart
   ```

4. Enter the following command to confirm the new values are configured for the state table:

   ```
   /usr/sbin/ipf -T fr_statemax,fr_statesize,fr_state_maxbucket
   ```

   The command output displays the values from the state table. In this example, the settings are still at the default values:

   ```
   fr_statemax min 0x1 max 0x7fffffff current 1000000
   fr_statesize min 0x1 max 0x7fffffff current 250007
   fr_state_maxbucket min 0x1 max 0x7fffffff current 256
   ```

Manage Conntrack Table Size (Linux)

This topic explains how to manage the kernel firewall state table.

About Managing the State Table

Conntrack is only supported on Linux systems, and IPFilter is supported on AIX and Solaris before version 11.4. Both are system-specific names for the *Kernel Firewall State Table*.

- Linux workloads: Manage the Conntrack table.
- AIX or Solaris workloads, versions 11.3 and earlier: Manage the IPFilter state table.
  For more information about AIX and Solaris, see Tuning the IP Filter State Table (AIX/Solaris) [258].

On Linux workloads, the VEN automatically increases and decreases the size of the Conntrack table as needed based on the number of active connections on the workload.

The VEN automatically increases the size to minimize the possibility of the workload running out of space in the Conntrack table and blocking valid connections.

The VEN uses the following behavior to manage the Conntrack table size:

- By default, the size of the Conntrack table starts at 1M. This is the baseline value. The baseline value is used as the starting point for automatically resizing the Conntrack table.
- Every 10 seconds, the VEN polls the table size to check the fill percentage.
- When the table reaches 80% of the maximum size, the VEN doubles the value set for the maximum size.
- The VEN doubles the maximum size value only 3 times (8x of the baseline value).
- For a 1M baseline value, the maximum table size after adjustment is 8M.

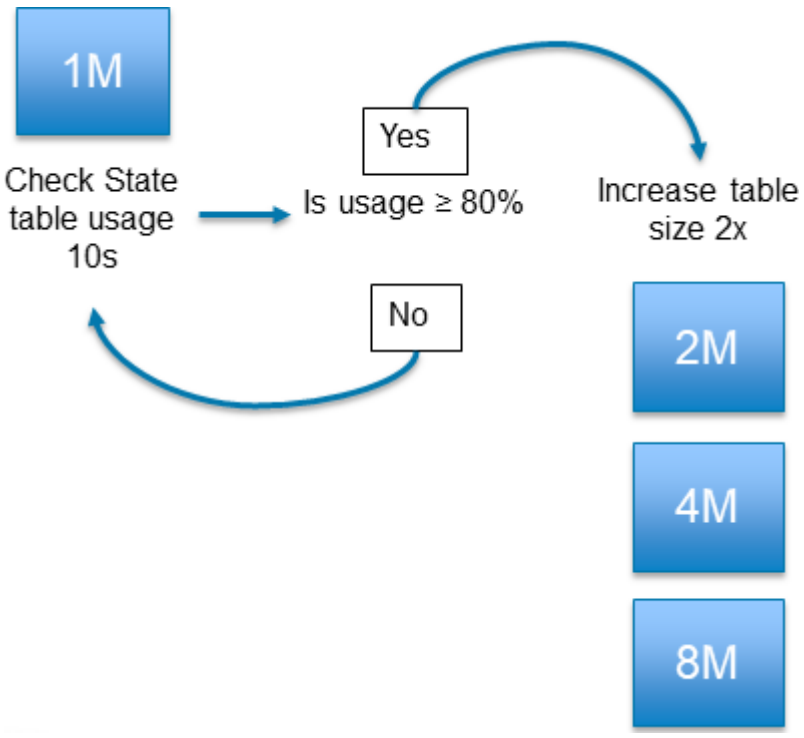

Customizing the VEN Adjustment Behavior

If the Conntrack table is experiencing issues with the size limit, you can adjust the way by which the VEN automatically manages the table size. Adjust the VEN behavior by setting the following values in the VEN configuration file `/etc/default/illumio-agent`.

Setting	Default	Description
`FW_STATE_TABLE_AUTO_RESIZE`	True	Indicates whether auto resize of the Conntrack table is required.
`CONNTRACK_MAX`	1000000	• Defines the maximum number of Conntrack table entries. • Configures the system value for `/proc/sys/net/nf_conntrack_max`
`CONNTRACK_HASH_SIZE`	256000	• Defines the starting size of the Conntrack hash table. • Configures the system value for `/sys/module/nf_conntrack/parameters/hashsize`

NOTE

When you install a VEN on a Linux workload, this feature is enabled by default using the default values. If you customize the values in the `illumio-agent` configuration file before installing the VEN, the custom values will apply on installation. If you customize the values after installing the VEN, you must restart the VEN for the values to take effect in runtime.

Restrictions for VEN Adjustment

Customizing the VEN adjustment behavior has the following restrictions:

- The value you set for `CONNTRACK_HASH_SIZE` should be 25% of the value of `CONNTRACK_MAX`.
- You must set the values to 512 or higher. If you set a value below 512, the Linux kernel will automatically adjust the value to 512.

VEN Firewall Tampering Detection

The PCE distributes the latest policy applicable to each workload to ensure that the VEN receives the latest policy updates. The VEN internally creates and maintains a set of meta information of these rules, which it uses to detect tampering.

Automatic History of Firewall Changes

Changes to the firewall on a workload are historically recorded for an audit trail. Up to 10 changes to the firewall history are saved. The history is viewable via the PCE Support Reports.

Host Firewall Tampering Protection

During periodic tampering detection (default: every 10 mins), the VEN checks whether certain static configurations have changed, including whether the runtime IPSec policy is identical to the policy that the PCE generated.

If a host firewall is tampered with, firewall tampering protection start firewall validation procedure. If the outcome detects any of the Illumio-added rules have been tampered, then the restoration procedure starts.

The procedure attempts to fetch a new security policy from the PCE, but if it fails due to a network connectivity issue, you can try to recover your last known good copy of a policy stored locally. The last step is validating the policy against the meta information of the policy. The tampering attempt is reported to the PCE as an `agent.tampering` event.

A host firewall tampering event occurs when another administrator or an attacker:

- Adds a firewall rule to the Illumio firewall compartment.
- Modifies a firewall rule added by Illumio.
- Deletes a firewall rule added by Illumio.
- Deletes all firewall rules (flush) added by Illumio.

The norm is that Illumio tries to detect tampering attempts only to Illumio firewall policy only and not to others.

Workload OS	Tampering Detection
Linux	The VEN monitors any underlying iptables, ipset, and IPsec changes. Once the VEN detects a tampering attempt, it validates the snapshot of iptables/ipset/IPsec against the firewall policy validation meta information.
Windows	The VEN monitors any changes in the Windows Filtering Platform (WFP) layer and the runtime IPsec policy. If it detects a change, it starts the validation and restore procedure.
AIX/Solaris	• On AIX (all versions) and Solaris (versions before 11.4) , the VEN monitors any underlying ipfilter changes. If the VEN detects a tampering attempt, it validates the snapshot of the ipfilter against the firewall policy validation meta information. • On Solaris versions 11.4 and later, the VEN checks packet filter. • On AIX and Solaris, the feature is enabled by default and updated every 10 minutes. • On AIX, the VEN monitors any changes in the runtime IPsec policy. If it detects a change, it starts the validation and restore procedure.

Host Firewall Tampering Alerts

Host firewall tampering alerts can be viewed:

- On the host VEN.
- In the PCE web console.
- In the return from a call to the `/events`Illumio Core REST API.
- In the return from a query in Splunk or other SIEM software.

View Tampering Alerts on VEN Host

Work-load OS	Procedure

Linux	As root, separately execute the following commands:
	
	Tail the VEN log file to see suspected tampering events and hash comparisons:
	
	`$ tail -f /opt/illumio_ven_data/log/platform.log`
	
	`INFO: Possible tamper detected...`
	`INFO: FW iptables checksums ... (compares security policy hashes to see if anything changed)`

Win-dows	Check `\programdata\illumio\log\platform.log` and search "`!!!Tampering detected`" **NOTE** This alter displays "Filtering Platform Policy Change" when a tampering event is detected. Double-click the alert for detailed information.

View Tampering Alerts Sent to PCE

PCE Web Console

To view `agent.tampering` events in the PCE web console, navigate to **Troubleshooting > Events**.

Double-click an `agent.tampering` event to see its details.

Illumio Core REST APIs

To return all tampering events for an organization, execute the following command using your organization URI. For more information, see Events in the *REST API Developer Guide*.

Example Curl Command to Get Information for All `agent.tampering` Events:

```
$ curl -i -X GET https://pce.example.com:8443/api/v2/orgs/1/events/?event_type=agent.tam
```

Example Curl Command to Get Information for a Specific `agent.tampering` Event:

```
$ curl -i -X GET https://pce.example.com:8443/api/v2/orgs/1/events/some_event_ID -H "Acc
```

Example JSON Response Body from Getting an `agent.tampering` Event:

```
{
    "href": "/orgs/1/events/some_event_ID",
    "timestamp": "2019-06-17T05:42:10.419Z",
    "pce_fqdn": "someName.someDomain",
    "created_by": {
```

```
        "agent": {
            "href": "/orgs/1/agents/xxxxx",
            "hostname": "someHostname"
        }
    },
    "event_type": "agent.tampering",
    "status": "success",
    "severity": "err",
    "action": {
        "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
        "api_endpoint": "FILTERED",
        "api_method": "PUT",
        "http_status_code": 204,
        "src_ip": "xx.xxx.xx.xx"
    },
    "resource_changes": [],
    "notifications": [
        {
            "uuid": "yyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyyyy",
"notification_type": "workload.oob_policy_changes",
            "info": {
                "tampering_revert_succeeded": true,
                "beginning_timestamp": "2019-06-17T05:42:10Z",
                "ending_timestamp": "2019-06-17T05:42:10Z",
                "num_events": 1
            }
        }
    ]
}
```


Splunk or Other SIEM Software

If you send VEN events received by the PCE to Splunk or other SIEM software, query for `agent.tampering` events in accordance with the SIEM vendor's query procedures.

VEN Tampering Protection

In Illumio Core and Illumio Endpoint 22.5.10 and later releases, you can protect the following types of VENs from unintended actions and tampering:

- Windows and Linux VENs running on servers
- Windows VENs running on endpoints

This feature protects the VEN itself from tampering versus protecting the workload host that the VEN is running on from being tampered with. For information about how the VEN detects tampering with the host firewall, see VEN Firewall Tampering Detection [261].

About Tampering Protection

NOTE

Before using this feature, complete the tasks in Requirements for Using Tampering Protection [265].

This feature protects VENs from unintended, accidental invocation of VEN CLI actions and installer commands that impact VEN functionality, and malicious attempts (including from System Administrators) to disable or uninstall the VEN, or otherwise render the VEN unusable.

Using this feature, you control the ability to run the following VEN administrative actions with the VEN CLI:

- Stopping the VEN; for information see Shut Down VENs. [237]
- Restarting the VEN; for information see Start Up VENs. [237]
- Suspending the VEN; for information, see VEN Suspension. [238]
- Deactivating the VEN; for information, see Deactivate Using VEN Command Line. [243]
- Unpairing the VEN from the PCE; for information, see Unpair Using VEN Command Line. [243]
- Upgrading the VEN on the server or endpoint; see the topics for managing the VENs using the CLI.

NOTE

Providing a maintenance token is not required when upgrading VENs by using the PCE web console.

- Uninstalling the VEN from the server or endpoint; see the topics for managing the VENs using the CLI.

NOTE

Providing a maintenance token is not required when uninstalling VENs from workloads by using the PCE web console.

This tampering protection restricts VEN CLI commands issued by all users, including the users who have administrative or root access to the VEN hosts (servers and endpoints).

Requirements for Using VEN Tampering Protection

To use this feature, you must complete the following requirements:

1. Enable the feature for your organization. See Enable VEN Tampering Protection [266].
2. Generate a maintenance token for all VENs or for specific VENs that you want protected. See Generate VEN Maintenance Token. [266]

   To generate this token, users must be part of one of the following Illumio Authorization roles:
   - Global Organization Owner
   - Global Administrators
   - Workload Managers (only for the workloads to which the users have access)

     When you are part of the Workload Manager role, you can set up tampering protection for the VENs you have access to. See "Workload Manager Role" in the PCE Administration Guide for information.
3. Include the token when running VEN CLI commands. See Manage VEN When Tampering Protection Enabled [267].

Enable VEN Tampering Protection

Before you can generate maintenance tokens for VENs or use the tampering protection feature, you must enable it in the PCE web console for your organization.

1. From the PCE web console main menu, go to **Settings** > **VEN Operations**.

   IMPORTANT

   To access the Setting page for VEN Operations, you must be a memember of the Global Organization Owner role. You cannot enable the VEN tampering protection feature without this level of Illumio authorization.

2. Click **Edit**.
3. In the Tampering Protection section, select **Yes** to require a maintenance token when running VEN commands on the VEN CTL.
4. Click **Save**.

Generate a VEN Maintenance Token

NOTE

Before you generate a VEN and Endpoint maintenance token, you must enable the feature for your organization.

You can generate maintenance tokens for all your VENs or for a specific VEN.

To generate a maintenance token:

1. Go to **Workloads** and click the **VENs** tab.
   - To generate support tokens for all of the VENs, click **Generate Maintenance Token**.
   - To generate a token for a specific VEN, click the name of a VEN to open the details page for that VEN, and then click **Generate Maintenance Token**.

     A **Generate Maintenance Token** dialog box appears where you can generate tokens for all VENs or the specific VEN you selected.

     NOTE

     If the tampering protection feature is enabled for the PCE, the page includes a **Generate Maintenance Token** button. If the page does not include this button, you must enable the feature for your PCE. See Enable VEN Tampering Protection [266].

2. Specify the time period for the token: unlimited (will never expire or need to be regenerated) or a set time period. By default, the dialog box specifies 7 days for the time period.
3. Click **Generate**.

   When ready, the dialog refreshes with the text string for the maintenance token and the timestamp for then the token was generated.
4. Copy the text string for the token and store it in a secure location. You will need to provide this string on the command line when you run VEN commands using the VEN CLI.

5. Click **Done** to close the dialog box.

Manage a VEN when Tampering Protection Enabled

When you've enabled tampering protection for a VEN, you must include the new parameter `maintenance-token <token>` on the VEN command line after the action you want to run. See the following examples. On Windows, include one dash with the parameter (`-maintenance-token <token>`); on Linux, include two dashes (`--maintenance-token <token>`) to run the parameter.

When enabled, running the VEN actions without specifying the token will fail.

NOTE

Not all VEN actions support using a maintenance token for tampering protection. See About Tampering Protection [264] for the list of supported actions.

When enabled, the VEN validates the maintenance token and the token expiration date, and runs the commands as usual.

When the token expires, you can regenerate it in the PCE web console.

Example: Windows Command Line to Run Protected VENs

```
<VEN Installation Directory>\illumio-ven-ctl.exe stop
Maintenance token is required for this operation.
<VEN Installation Directory>\illumio-ven-ctl.exe stop
-maintenance-token eyJhY3Rpb25zIjpudWxsLCJleHBpcmVzX2F0IjpudWxsLCJhZ2VudF9p
ZHMiOm51bGwsIm9yZ19pZCI6MX0=.MGUCMHSfLNS8yGHgFY0D3CuFvi+L8m6VUVI9FHRzT31sn37F+
GsKecpSnbR8abYuSoz2wgIxALhrtjAXZNN8unxLuN8WO/kcLONz7gwboRCT/Sc2FdwXAkLvioh+9
jyU8OBeAj5poA==Stopping venAgentMonitorSvc
Stopping venPlatformHandlerSvc
Stopping venVtapServerSvc
Stopping venAgentMgrSvc
Success
<VEN Installation Directory>\Illumio>
```

Example: Linux Command Line to Run Protected VENs

```
[root@localhost illumio_ven]# ./illumio-ven-ctl unpair open noreport
Maintenance token is required for this operation.
[root@localhost illumio_ven]# ./illumio-ven-ctl unpair --maintenance-token
eyJhY3Rpb25zIjpudWxsLCJleHBpcmVzX2F0IjpudWxsLCJhZ2VudF9pZHMiOm51bGwsIm9yZ19pZCI
6MX0=.MGUCMHSfLNS8yGHgFY0D3CuFvi+L8m6VUVI9FHRzT31sn37F+GsKecpSnbR8abYuSoz2wgIxAL
hrtjAXZNN8unxLuN8WO/kcLONz7gwboRCT/Sc2FdwXAkLvioh+9jyU8OBeAj5poA== open noreport
Stopping venAgentMonitor:    ...done.
Stopping venVtapServer:    ...done.
Stopping IPSec:    ...done.
Stopping venPlatformHandler:    ...done.
Stopping venAgentMgr:    ...done.
```

```
Checking agent state
   ...done.
 * Flush IPv4   ...done.
   ...done.
Unloading modules   ...done.Illumio VEN is being uninstalled...
2023-01-17T12:51:01-0800 Uninstalling Illumio ............
2023-01-17T12:51:04-08:00 Stopped all daemons
2023-01-17T12:51:04-08:00 Init scripts disabled
2023-01-17T12:51:04-08:00 VEN state on uninstall: enforced
2023-01-17T12:51:04-0800 Deactivating Illumio VEN .......
2023-01-17T12:51:05-0800 Agent 15 Org 1 successfully deactivated
2023-01-17T12:51:05-0800 Deactivation complete
2023-01-17T12:51:05-08:00 /opt/illumio_ven/system/etc/init.d/illumio-firewall
disable -w workload/c3364c6d-43f7-43fd-a4e4-9eb6258808b4/current
2023-01-17T12:51:07-08:00 Firewall Rules successfully restored
2023-01-17T12:51:07-08:00 Removed ilo-ven user entries
2023-01-17T12:51:07-08:00 Removed data distribution tree from /opt
2023-01-17T12:51:07-08:00 Removed binary distribution tree from /opt
2023-01-17T12:51:07-0800 Uninstall successful
VEN has been SUCCESSFULLY unpaired with Illumio
[root@localhost illumio_ven]#
```

Windows VEN Installer Changes

When you enable the VEN tampering protection feature, the Windows VEN installer can include the new `MAINTENANCE_TOKEN` parameter for the `upgrade`, `uninstall`, and `repair` commands, as shown in the following examples.

Upgrade a VEN

```
ven_installer.exe /install /quiet /log ven_install.log MAINTENANCE_TOKEN=xxx
```

Uninstall a VEN

```
ven_installer.exe /uninstall /quiet /log ven_uninstall.log MAINTENANCE_TOKEN=xxx
```

Repair a VEN

```
ven_installer.exe /repair /quiet /log ven_repair.log MAINTENANCE_TOKEN=xxx
```

VEN Support Reports

A workload's support report provides diagnostic information for selected workloads. To troubleshoot issues with your workloads, you can generate a support report and send it to Illumio support.

NOTE

Your PCE user account must have the Organization Owner or Admin user role to perform this task and the workload should be an active, managed workload.

Generate a VEN Support Report from the PCE UI

1. In the PCE web console, go to **Workloads**.
2. Click the **VENs** tab.
3. Click the name of a VEN to go to its details page.
4. Click **Generate Support Bundle**. Generating the bundle may take up to 10 minutes.
5. When the bundle is finished generating, click **Download**.

Generate Linux/AIX/Solaris Support Report Using CLI

If you need to troubleshoot VEN issues, you can generate a VEN support report from the command line for any workload and then send the report to Illumio support.

On Linux, AIX, and Solaris, the generated report is saved to the `/tmp` directory and overwrites any previously generated copy of the same report.

NOTE

You must have root privileges on the workload to run the support report command.

You can also run a VEN support report when you unpair a workload.

To generate a VEN support report for a Linux workload:

1. Establish a secure shell connection (SSH) to the Linux workload.
2. Execute the following command as root to generate the support report.

   ```
   /opt/illumio_ven/illumio-ven-ctl gen-supportreport
   ```
3. Type Y when asked if you want to run the report.
4. Optionally, if you want to bypass the confirmation prompt, you can execute the script with a `-y` or `-Y` option:

   ```
   /opt/illumio_ven/illumio-ven-ctl gen-supportreport -y
   ```
5. To view the report generation log, enter the following command:

   ```
   more -n 10 -f /opt/illumio_ven_data/log/report.log
   ```
6. The support report generation is complete when "Successfully created report" or "Failed to create report" is logged. After the report is successfully generated, the report is sent to the PCE.

Generate Windows Support Report Using CLI

If you need to troubleshoot VEN issues, you can generate a VEN support report from the command line for any workload and then send the report to Illumio Customer Support.

On Windows, the generated report is saved to the `C:\Windows\Temp` directory and overwrites any previously generated copy of the same report.

You can also run a support report when you unpair a workload.

To generate a VEN support report

```
illumo-ven-ctl.exe gen-supportreport
```

To bypass the confirmation prompt

```
illumo-ven-ctl.exe gen-supportreport -noprompt yes
```

VEN Troubleshooting

This topic describes some important system administration considerations on Windows, useful tools, and a generalized set of actions to troubleshoot VEN operations.

Windows: Enable Base Filtering Engine (BFE)

Windows BFE is a Windows subsystem that determines which packets should be allowed to the network stack. BFE is enabled by default. If you disable BFE on your Windows workload, all packets are sent to the TCP/IP stack bypassing BFE which can result in different behavior from one system to another. The worst case scenario is all the ingress and egress packets get dropped.

If you have disabled BFE on your Windows workload, re-enable it.

Linux: ignored_interface

The Linux `ignored_interface` inhibits PCE policy updates.

Transitioning an enforced workload's interface from or to `ignored_interface` might drop the dynamic, long-lived connections maintained by the system.

When a VEN interface is placed in the `ignore_interface` list, the any flow state over the interface won't be kept by conntrack an longer. (The conntrack table on Linux stores information on network connections.) If the connection on TCP port 8444 to the PCE is reinitialized, any arriving packets from the PCE are dropped, because the packets do not have any state in conntrack.

The VEN heartbeat eventually restores connections, but meanwhile the VEN implements any policy sent by lightning bolt from the PCE.

VEN Troubleshooting Tools

Illumio provides the following tools for VEN connectivity checking and troubleshooting VEN issues on workloads:

- A VEN connectivity checking tool called `venconch` for workloads is available on the Illumio Support site.
- A VEN compatibility checking feature is available in the PCE web console for paired workloads.

Commands to Obtain Firewall Snapshot

Run the following commands on the workload to get a copy of the logs and configured firewall settings.

Linux

- `iptables-save`
- `ipset –L`

Windows

- `netsh wfp show state`

Solaris

```
ipfstat -ionv
```

AIX

```
ipfstat -ionv
```

Troubleshooting Tips

Connectivity Issues

Perform the following actions to identify why a workload is unreachable, cannot reach other workloads, or cannot communicate with the PCE:

- Determine if all workloads are unable to communicate or just a subset of the workloads are reported as disconnected. If the PCE reports that all workloads are offline, check if PCE is reachable from workloads.
- If a subset of workloads are down, check if there are differences in network configuration between those and the workloads that are connected, and if they are contributing to PCE being unreachable.
- Check if any workloads that are unable to communicate are located behind NAT devices, firewalls, or remote data centers.
- Ensure the following port configuration:
  - On Prem
    - Port 8443 - HTTPS requests
    - Port 8444 - long-lived TLS-over-TCP connection
  - SaaS
    - Port 443 for both HTTPS requests and the long-lived TLS-over-TCP connection
- If running in a public cloud instance:
  - For AWS, ensure security groups permit TCP port 443.
  - For Azure, ensure that Endpoints are configured to allow traffic.

VEN Process Issues

Check the status of the VEN-specific processes and ensure that they are running and active:

- **Linux:** Run `/opt/illumio/illumio-ven-ctl status`
- **Windows:** Execute `tasklist`

Ensure the following processes are running and active:

- **Linux:** `venAgentManager`, `venPlatformHandler`, `venAgentLManager`, `VtapServer`, and `AgentMonitor`
- **Windows:** `venAgentLogMgrSvc`, `venPlatformHandler`, `venVtapServerSvc`, and `ilowfp`

Errors in the VEN Logs

Review the VEN log files to find any errors generated by the system (`sudo` required):

- Logs in `Data_Dir/log` directory

  To look for any errors in the log files, execute `grep –ir ERROR *`

  To check for firewall updates, view the `platform.log` file. Look for logs related to firewall updates; for example:

```
2014-07-26T22:20:41Z INFO:: Enforcement mode is: XXXX
2014-07-26T22:20:41Z INFO:: Is fw update yes
2014-07-26T22:20:41Z INFO:: Is ipset update yes
2014-07-26T22:20:41Z INFO:: saved fw-json
```

- Check heartbeat logs for records related to update messages from the PCE. See the following example heartbeats:

```
2014-07-26T22:43:12Z Received HELLO from EventService.
2014-07-26T22:43:12Z Sent ACK to EventService.
Events – f/w updates etc.
014-07-26T22:34:11Z Received EVENT from EventService.
2014-07-26T22:34:11Z Added EVENT from EventService to PLATFORM handler thread message
```

```
iptables-save | grep 443 | grep allow_out
    -A tcp_allow_out -d 54.185.43.60/32 -p tcp -m multiport --dports 443 -m conntrack
    -A tcp_allow_out -d 54.185.43.60/32 -p tcp -m multiport --dports 443 -m conntrack
    -A tcp_allow_out -d 204.51.153.0/27 -p tcp -m multiport --dports 443 -m conntrack
    -A tcp_allow_out -d 204.51.153.0/27 -p tcp -m multiport --dports 443 -m conntrack
iptables-save | grep 444 | grep allow_out
    -A tcp_allow_out -d 54.185.43.60/32 -p tcp -m multiport --dports 444 -m conntrack
    -A tcp_allow_out -d 54.185.43.60/32 -p tcp -m multiport --dports 444 -m conntrack
```

Policy Sync Might Require Reboot

Persistent errors with policy sync on a workload can be cleared by rebooting the VEN.

Event Viewer Stops Logging

After you upgrade the VEN, **Event Viewer** can stop logging so that the support report does not include `windows_evt_application`, `windows_evt_system`, and the system directory (e.g.: `msinfo32`). To correct the issue, close **Event Viewer** before upgrading the VEN. Then reopen **Event Viewer**.

Events Administration and REST APIs

Overview of Events Administration

This section describes how to do typical administration tasks related to PCE events.

Before You Begin

Illumio recommends that you be familiar with the following technology:

- Solid understanding of Illumio Core
- Familiarity with syslog
- Familiarity with your organizations' Security Information and Event Management (SIEM) systems

About This Guide

This guide provides the following information to administer your PCE deployment:

- An overview of events and SIEM integration
- Events setup considerations
- Event record formats, types, and common fields
- Event types by resource
- SIEM integration considerations and recommendations

See also the following related documentation:

- U.S. National Institute for Standards and Technology's NIST 800-92 Guide to Computer Security Log Management
- U.S. Department of Homeland Security National Cybersecurity Center

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
...
```

Events Framework

The Illumio events framework provides an information-rich, deep foundation for actionable insights into the operations of Illumio Core.

Overview of the Framework

Auditable events are records of transactions collected from the following management interfaces:

- PCE web console
- REST API
- PCE command-line tools
- VEN command-line tools

All actions that change the configuration of the PCE, security policy, and the VENs are recorded, including workload firewall tampering.

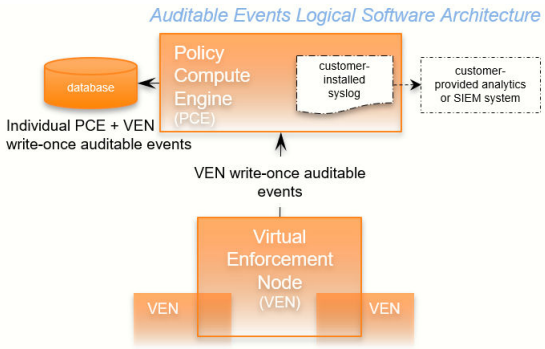

As required by auditing standards, every recorded change includes a reference to the program that made the change, the change's timestamp, and other fields. After recording, the auditable events are read-only.

Auditable events comply with the Common Criteria Class FAU Security Audit requirements standard for auditing.

Auditing Needs Satisfied by Framework

Need	Description	See topic...
Audit and Compliance	Evidence to show that resources are managed according to rules and regulatory standards.	Events Record Information [277]
Resource Lifecycle Tracking	All information necessary to track a resource through creation, modification, and deletion.	Events Lifecycle for Resources [275]
Operations	Trace of recent changes to resources.	Events Lifecycle for Resources [275]
Security	Evidence to show which changes failed, such as incorrect user permissions or failed authentication.	User Password Update Failed (JSON) [296]

Benefits of Events Framework

The events framework in the Illumio Core provides the following benefits:

- Exceeds industry standards
- Delivers complete content
  - Comprehensive set of event types
  - Includes more than 200 events
  - Additional notable system events are generated
- Easily accessible interfaces to capture events:
  - Event Viewer in the PCE web console
  - REST API with filtering
  - SIEM intregration
  - Events are the same across all interfaces
- Designed for customer ease of use
  - Flattened, common structure for all events
  - Eliminates former duplicate or multiple events for single actions
  - Streamed via syslog in JSON, CEF, or LEEF format
  - Create/Update/Delete REST APIs recorded as events

    Read APIs/GET requests are not recorded, because they do not change the Illumio Core.

Events Lifecycle for Resources

Illumio resources progress through the lifecycle stages (creation, updating, deletion) and Illumio Core records them with the appropriate event types.

About the Lifecycle

Many resources have a lifecycle from creation through update to deletion. For example, the events related to a security policy rule (identified by the resource name `sec_rule`) are recorded with the following event types.

- `sec_rule.create`
- `sec_rule.update`: Update events record with the values of the resource object both before and after the event for a lifecycle audit trail.
- `sec_rule.delete`

Auditable Events: Lifecycle of a Resource

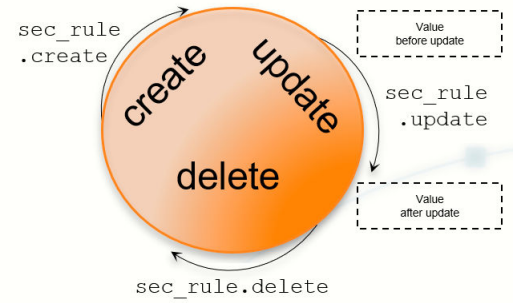

Other Resource Lifecycles

Some resources have unique characteristics and do not follow the create-update-delete pattern. For example, workloads have the following event types:

- `workload.update`
- `workload.upgrade`
- `workload.redetect_network`
- `workload.recalc_rules`
- `workload.soft_delete`
- `workload.delete`
- `workload.undelete`

Events Described

This section describes the concepts and types of PCE events.

Event Types, Syntax, and Record Format

When working with events, it is important to recognize their type, REST API schema, syntax, and record information.

Types of Events

The Illumio Core includes the following general categories of auditable events:

- Organizational events: Organizational events are further grouped by their source:
  - API-related events: Events occurring from a use of the REST API, including the PCE web console
  - System-related events: Events caused by some system-related occurrence
- Traffic events

Anonymized Database Dumps

To troubleshoot customer-reported issues, Illumio Customer Support sometimes requests that you supply an anonymized dump of the PCE database.

To safeguard your organization's privacy, the event information is not included in the anonymized database dump.

REST API Events Schema

The Events schema in JSON is downloadable from this documentation portal in the zipfile of the REST API schemas. From the documentation portal Home page, go to the **Develop** category > **REST API Public Schemas (Archive File)**.

Event Syntax

The names of recorded auditable events in have the following general syntax:

```
resource.verb[.success_or_failure]
```

Where:

- `resource` is a PCE and VEN object, such as PCE `user` or VEN `agent` component.

- `verb` describes the action of the event on that resource.
- In CEF and LEEF formats, the success or failure of the verb is included in the recorded event type. This indicator is not needed in the JSON format.

Events Record Information

The following information is included in a event record, which answers the who, what, where, how, and when:

Type of information	Description
Who	• VEN identified by hostname and agent href, and after Release 22.3, VEN href • User identified by username and href • PCE system identified by "system"

What	The action that triggered the event, including the following data:
	
	• Resource type + operation + success or failure
	• Application Request ID
	• Status of successful events and failed events:
	  • In case of failure, exception type and exception message.
	  • All failures related to security, such as authentication and authorization.
	  • Severity as INFO, WARNING, ERROR.
	• The pre-change and post-change values of the affected resources.

Where	The target resource of the action, composed of the following data:
	
	• Identifier of the target resource (primary field).
	• Friendly name for the target resource. For example:
	  • workload/VEN: `hostname`
	  • user.username
	  • ruleset, label, service, etc: name, key/value

How	API endpoint, method, HTTP status code, and source IP address of the request.
When	Timestamp of the event's occurrence. This timestamp is *not* the time the event was recorded.

Event Record Structure

Regardless of export format (JSON, CEF, or LEEF), the records and fields for all events share a common structure. This common structure of composite events makes post-processing of event data easier.

Bulk change operations on many resources simultaneously are recorded as individual operations on the resource within a single composite event. Failed attempts to change a configuration, such as incorrect authentication, are also collected.

Common Fields

Field Name	Description
`href`	Unique event identifier; contains a UUID.
`timestamp`	Exact time that the event occurred in RFC 3339 format with fractional seconds.
`pce_fqdn`	The fully qualified domain name of the PCE; especially useful for Supercluster deployments or if there are multiple PCEs sending data to the SIEM server.
`created_by`	Identifies creator of the event; could be a user, the system, or a workload.
`event_type`	Name of the event; for more information, see the List of Event Types [279] table.
`status`	"Success" or "failure;" if the status is null, the event is for information only and doesn't indicate success or failure.
`severity`	"Informational," "warning," or "error" indicating the severity of the event.
`version`	Schema version for events.

Events Displayed in PCE Web Console

The PCE web console provides an ongoing log of all Organization events that occur in the PCE. For example, Organization events capture actions such as users logging in and logging out, and failed login attempts; when a system object is created, modified, deleted, or provisioned; when a workload is paired or unpaired; and so on.

From the platform and API perspective, Organization events are referred to internally as `auditable_events` and are generated by the `auditable_events_service`.

You can use the filter at the top of the page to search for events by type of event, event severity level, and when the event occurred.

Cross-Site Request Forgery Protection

A cross-site request forgery (CSRF) is an attack that involves forcing a victim to send an HTTP request to a target destination without their knowledge or intent in order to perform an action as the victim. The underlying cause is an application functionality using predictable URL or form actions in a repeatable way. The nature of the attack is that CSRF exploits the trust that a website has for a user.

For more details on this attack, see the CSRF article on the Web Application Security Consortium website.

Illumio Core can notify you of this type of attack in the following ways:

- The PCE web console logs the attack as an Organization Event called "CSRF token validation failure."
- The event is logged in the Illumio Core REST API as `authz_csrf_validation_failure` in the `audit_log_events_get.schema`.
- The event `authz_csrf_validation_failure` appears in the PCE syslog output if you have deployed the PCE as a software.

IMPORTANT

When you see this event occur, you should immediately investigate the issue because the request might not have originated from a valid user.

List of Event Types

The following table provides the types of JSON events generated and their description. For each of these events, the CEF/LEEF success or failure events generated are the event name followed by `.success` or `.failure`.

For example, the CEF/LEEF success event for `agent.activate` is `agent.activate.success` and the failure event is `agent.activate.failure`.

Each event can generate a variety of notification messages. See Notification Messages in Events [290].

JSON Event Type	Description
access_restriction.create	Access restriction created
access_restriction.delete	Access restriction deleted
access_restriction.update	Access restriction updated
agent.activate	Agent paired
agent.activate_clone	Agent clone activated
agent.clone_detected	Agent clone detected
agent.deactivate	Agent unpaired
agent.generate_maintenance_token	Generate maintenance token for any agent
agent.goodbye	Agent disconnected
agent.machine_identifier	Agent machine identifiers updated
agent.refresh_token	Agent refreshed token
agent.reguest_policy	Policy request sent
agent.request_upgrade	VEN upgrade request sent
agent.service_not_available	Agent reported a service not running
agent.suspend	Agent suspended
agent.tampering	Agent firewall tampered
agent.unsuspend	Agent unsuspended
agent.update	Agent properties updated.
agent.update_interactive_users	Agent interactive users updated
agent.update_iptables_href	Agent updated existing iptables href
agent.update_running_containers	Agent updated existing containers
agent.upload_existing_ip_table_rules	Agent existing IP tables uploaded
agent.upload_support_report	Agent support report uploaded
agent_support_report_request.create	Agent support report request created
agent_support_report_request.delete	Agent support report request deleted
agents.clear_conditions	Condition cleared from a list of VENs
agents.unpair	Multiple agents unpaired
api_key.create	API key created

JSON Event Type	Description
api_key.delete	API key deleted
api_key.update	API key updated
auth_security_principal.create	RBAC auth security principal created
auth_security_principal.delete	RBAC auth security principal deleted
auth_security_principal.update	RBAC auth security principal updated
authentication_settings.update	Authentication settings updated
cluster.create	PCE cluster created
cluster.delete	PCE cluster deleted
cluster.update	PCE cluster updated
container_workload.update	Container workload updated
container_cluster.create	Container cluster created
container_cluster.delete	Container cluster deleted
container_cluster.update	Container cluster updated
container_cluster.update_label_map	Container cluster label mappings updated all at once
container_cluster.update_services	Container cluster services updated, created, or deleted by Kubelink
container_workload_profile.create	Container workload profile created
container_workload_profile.delete	Container workload profile deleted
container_workload_profile.update	Container workload profile updated
database.temp_table_autocleanup_started	DB temp table cleanup started
database.temp_table_autocleanup_completed	DB temp table cleanup completed
domain.create	Domain created
domain.delete	Domain deleted
domain.update	Domain updated
enforcement_boundary.create	Enforcement boundary created
enforcement_boundary.delete	Enforcement boundary deleted
enforcement_boundary.update	Enforcement boundary updated
event_settings.update	Event settings updated

JSON Event Type	Description
firewall_settings.update	Global policy settings updated
group.create	Group created
group.update	Group updated
ip_list.create	IP list created
ip_list.delete	IP list deleted
ip_list.update	IP list updated
ip_lists.delete	IP lists deleted
ip_tables_rule.create	IP tables rules created
ip_tables_rule.delete	IP tables rules deleted
ip_tables_rule.update	IP tables rules updated
job.delete	Job deleted
label.create	Label created
label.delete	Label deleted
label.update	Label updated
label_group.create	Label group created
label_group.delete	Label group deleted
label_group.update	Label group updated
labels.delete	Labels deleted
ldap_config.create	LDAP configuration created
ldap_config.delete	LDAP configuration deleted
ldap_config.update	LDAP configuration updated
ldap_config.verify_connection	LDAP server connection verified
license.delete	License deleted
license.update	License updated
login_proxy_ldap_config.create	Interservice call to login service to create LDAP config
login_proxy_ldap_config.delete	Interservice call to login service to delete LDAP config
login_proxy_ldap_config.update	Interservice call to login service to update LDAP config
login_proxy_ldap_config.verify_connection	Interservice call to login service to verify connection to the LDAP server

JSON Event Type	Description
login_proxy_msp_tenants.create	New MSP tenant created
login_proxy_msp_tenants.delete	MSP tenant deleted
login_proxy_msp_tenants.update	MSP tenant updated
login_proxy_orgs.create	New managed organization created
login_proxy_orgs.delete	Managed organization deleted
login_proxy_orgs.update	Managed organization updated
lost_agent.found	Lost agent found
network.create	Network created
network.delete	Network deleted
network.update	Network updated
network_device.ack_enforcement_instructions_applied	Enforcement instruction applied to a network device
network_device.assign_workload	Existing or new unmanaged workload assigned to a network device
network_device.create	Network device created
network_device.delete	Network device deleted
network_device.update	Network device updated
network_devices.ack_multi_enforcement_instructions_applied	Enforcement instructions applied to multiple network devices
network_endpoint.create	Network endpoint created
network_endpoint.delete	Network endpoint deleted
network_endpoint.update	Network endpoint updated
network_enforcement_node.activate	Network enforcement node activated
network_enforcement_node.clear_conditions	Network enforcement node conditions cleared
network_enforcement_node.deactivate	Network enforcement node deactivated
network_enforcement_node.degraded	Network enforcement node failed or primary lost connectivity to secondary
network_enforcement_node.missed_heartbeats	Network enforcement node did not heartbeat for more than 15 minutes
network_enforcement_node.missed_heartbeats_check	Network enforcement node missed heartbeats check

JSON Event Type	Description
network_enforcement_node.network_devices_network_endpoints_workloads	Workload added to network endpoint
network_enforcement_node.policy_ack	Network enforcement node acknowledgment of policy
network_enforcement_node.request_policy	Network enforcement node policy requested
network_enforcement_node.update_status	Network enforcement node reports when switches are not reachable
network_enforcement_nodes.clear_conditions	A condition was cleared from a list of network enforcement nodes
nfc.activate	Network function controller created
nfc.delete	Network function controller deleted
nfc.update_discovered_virtual_servers	Network function controller virtual servers discovered
nfc.update_policy_status	Network function controller policy status
nfc.update_slb_state	Network function controller SLB state updated
org.create	Organization created
org.recalc_rules	Rules for organization recalculated
org.update	Organization information updated
pairing_profile.create	Pairing profile created
pairing_profile.create_pairing_key	Pairing profile pairing key created
pairing_profile.delete	Pairing profile deleted
pairing_profile.update	Pairing profile updated
pairing_profile.delete_all_pairing_keys	Pairing keys deleted from pairing profile
pairing_profiles.delete	Pairing profiles deleted
password_policy.create	Password policy created
password_policy.delete	Password policy deleted
password_policy.update	Password policy updated
permission.create	RBAC permission created
permission.delete	RBAC permission deleted
permission.update	RBAC permission updated
radius_config.create	Create domain RADIUS configuration
radius_config.delete	Delete domain RADIUS configuration

JSON Event Type	Description
radius_config.update	Update domain RADIUS configuration
radius_config.verify_shared_secret	Verify RADIUS shared secret
request.authentication_failed	API request authentication failed
request.authorization_failed	API request authorization failed
request.internal_server_error	API request failed due to internal server error
request.service_unavailable	API request failed due to unavailable service
request.unknown_server_error	API request failed due to unknown server error
resource.create	Login resource created
resource.delete	Login resource deleted
resource.update	Login resource updated
rule_set.create	Rule set created
rule_set.delete	Rule set deleted
rule_set.update	Rule set updated
rule_sets.delete	Rule sets deleted
saml_acs.update	SAML assertion consumer services updated
saml_config.create	SAML configuration created
saml_config.delete	SAML configuration deleted
saml_config.pce_signing_cert	Generate a new cert for signing SAML AuthN requests
saml_config.update	SAML configuration updated
saml_sp_config.create	SAML Service Provider created
saml_sp_config.delete	SAML Service Provider deleted
saml_sp_config.update	SAML Service Provider updated
sec_policy.create	Security policy created
sec_policy_pending.delete	Pending security policy deleted
sec_policy.restore	Security policy restored
sec_rule.create	Security policy rules created
sec_rule.delete	Security policy rules deleted
sec_rule.update	Security policy rules updated

JSON Event Type	Description
secure_connect_gateway.create	SecureConnect gateway created
secure_connect_gateway.delete	SecureConnect gateway deleted
secure_connect_gateway.update	SecureConnect gateway updated
security_principal.create	RBAC security principal created
security_principal.delete	RBAC security principal bulk deleted
security_principal.update	RBAC security principal bulk updated
security_principals.bulk_create	RBAC security principals bulk created
service.create	Service created
service.delete	Service deleted
service.update	Service updated
service_account.create	Service account created
service_account.delete	Service account deleted
service_account.update	Service account updated
service_binding.create	Service binding created
service_binding.delete	Service binding created
service_bindings.delete	Service bindings deleted
service_bindings.delete	Service binding deleted
services.delete	Services deleted
settings.update	Explorer settings updated
slb.create	Server load balancer created
slb.delete	Server load balancer deleted
slb.update	Server load balancer updated
support_report.upload	Support report uploaded
syslog_destination.create	syslog remote destination created
syslog_destination.delete	syslog remote destination deleted
syslog_destination.update	syslog remote destination updated
system_task.agent_missed_heart-beats_check	Agent missed heartbeats

JSON Event Type	Description
system_task.agent_missing_heartbeats_after_upgrade	VEN missing heartbeat after upgrade
system_task.agent_offline_check	Agents marked offline
system_task.agent_self_signed_certs_check	VEN self signed certificate housekeeping check
system_task.agent_settings_invalidation_error_state_check	VEN settings invalidation error state check
system_task.agent_uninstall_timeout	VEN uninstall timeout
system_task.clear_auth_recover_condition	Clear VEN authentication recovery condition
system_task.compute_policy_for_unmanaged_workloads	Compute policy for unmanaged workloads
system_task.delete_expired_service_account_api_keys	An expired service account api_key was successfully deleted
system_task.delete_old_cached_perspectives	Delete old cached perspectives
system_task.endpoint_offline_check	Endpoint marked offline
system_task.provision_container_cluster_services	Container cluster services provisioned
system_task.prune_old_log_events	Event pruning completed
system_task.remove_stale_zone_subsets	Stale zone subnets removed
system_task.set_server_sync_check	Set server synced
system_task.vacuum_deactivated_agent_and_deleted_workloads	Deactivated and deleted workloads have been vacuumed
traffic_collector_setting.create	Traffic collector setting created
traffic_collector_setting.delete	Traffic collector setting deleted
traffic_collector_setting.update	Traffic collector setting updated
trusted_proxy_ips.update	Trusted proxy IPs created or updated
user.accept_invitation	User invitation accepted
user.authenticate	User authenticated
user.create	User created
user.delete	User deleted
user.invite	User invited
user.login	User logged in

JSON Event Type	Description
user.login_session_terminated	User login session terminated
user.logout	User logged
user.pce_session_terminated	User session terminated
user.reset_password	User password reset
user.sign_in	User session created
user.sign_out	User session terminated
user.update	User information updated
user.update_password	User password updated
user.use_expired_password	User entered expired password
user.verify_mfa	User verified MFA
users.auth_token	Auth token returned for user authentication on PCE
user_local_profile.create	User local profile created
user_local_profile.delete	User local profile deleted
user_local_profile.reinvite	User local profile reinvited
user_local_profile.update_password	User local password updated
ven_settings.update	VEN settings updated
ven_software.upgrade	VEN software release upgraded
ven_software_release.create	VEN software release created
ven_software_release.delete	VEN software release deleted
ven_software_release.deploy	VEN software release deployed
ven_software_release.update	VEN software release updated
ven_software_releases.set_default_version	Default VEN software version set
virtual_server.create	Virtual server created
virtual_server.delete	Virtual server created
virtual_server.update	Virtual server updated
virtual_service.create	Virtual service created
virtual_service.delete	Virtual service deleted
virtual_service.update	Virtual service updated

JSON Event Type	Description
virtual_services.bulk_create	Virtual services created in bulk
virtual_services.bulk_update	Virtual services updated in bulk
vulnerability.create	Vulnerability record created
vulnerability.delete	Vulnerability record deleted
vulnerability.update	Vulnerability record updated
vulnerability_report.delete	Vulnerability report deleted
vulnerability_report.update	Vulnerability report updated
workload.create	Workload created
workload.delete	Workload deleted
workload.online	Workload online
workload.recalc_rules	Workload policy recalculated
workload.redetect_network	Workload network redetected
workload.undelete	Workload undeleted
workload.update	Workload settings updated
workload.upgrade	Workload upgraded
workload_interface.create	Workload interface created
workload_interface.delete	Workload interface deleted
workload_interface.update	Workload interface updated
workload_interfaces.update	Workload interfaces updated For example, IP address changes, new interface added, and interface shut down.
workload_service_report.update	Workload service report updated
workload_settings.update	Workload settings updated
workloads.apply_policy	Workloads policies applied
workloads.bulk_create	Workloads created in bulk
workloads.bulk_delete	Workloads deleted in bulk
workloads.bulk_update	Workloads updated in bulk
workloads.remove_labels	Workloads labels removed
workloads.set_flow_reporting_frequency	Workload flow reporting frequency changed

JSON Event Type	Description
workloads.set_labels	Workload labels applied
workloads.unpair	Workloads unpaired
workloads.update	Workloads updated

Notification Messages in Events

Events can generate a variety of notifications that are appended after the event type:

- agent.clone_detected
- agent.fw_state_table_threshold_exceeded
- agent.missed_heartbeats
- agent.missing_heartbeats_after_upgrade
- agent.policy_deploy_failed
- agent.policy_deploy_succeeded
- agent.process_failed
- agent.service_not_available
- agent.upgrade_requested
- agent.upgrade_successful
- agent.upgrade_time_out
- container_cluster.duplicate_machine_id
- container_cluster.region_mismatch
- container_workload.invalid_pairing_config
- container_workload.not_created
- database.temp_table_autocleanup_completed
- database.temp_table_autocleanup_started
- hard_limit.exceeded
- pce.application_started
- pce.application_stopped
- remote_syslog.reachable
- remote_syslog.unreachable
- request.authentication_failed
- request.authorization_failed
- request.internal_server_error
- request.invalid
- request.service_unavailable
- request.unknown_server_error
- sec_policy.restore
- soft_limit.exceeded
- system_task.event_pruning_completed
- system_task.hard_limit_recovery_completed
- user.csrf_validation_failed
- user.login_failed
- user.login_failure_count_exceeded
- user.login_session_created
- user.login_session_terminated
- user.pce_session_created
- user.pce_session_terminated

- `user.pw_change_failure`
- `user.pw_changed`
- `user.pw_complexity_not_met`
- `user.pw_reset_completed`
- `user.pw_reset_requested`
- `virtual_service.not_created`
- `workload.duplicate_interface_reported`
- `workload.nat_rules_present`
- `workload.offline_after_ven_goodbye`
- `workload.online`
- `workload.oob_policy_changes`
- `workload.partial_policy_delivered`
- `workload.update_mismatched_interfaces`
- `workloads.flow_reporting_frequency_updated`

Common Criteria Only Events

The following table lists the types of JSON events that are generated and their descriptions.

For each of these events, the CEF/LEEF success or failure events generated are the event name followed by `.success` or `.failure`.

For example, the CEF/LEEF success event for `agent.update` is `agent.update.success` and the failure event is `agent.update.failure`.

JSON Event Type	Description
`pce.application_started`	PCE application started
`pce.application_stopped`	PCE application stopped
`remote_syslog.reachable`	Remote syslog destination reachable
`remote_syslog.unreachable`	Remote syslog destination not reachable
`tls_channel.establish`	TLS channel established
`tls_channel.terminate`	TLS channel terminated

View and Export Events

By default, you can view events in the PCE web console or by using the PCE command line. You can then export Organization events using the PCE web console.

View Events in PCE Web Console

By default, the PCE web console shows events that occur in your organization, such as when a workload is paired, if a pairing failed, when a user logs in or logs out, when a user fails to authenticate, and so on.

If you want to see only certain events you can filter by event type to see events that interest you most. You can also search for Organization events by their universally unique identifier (UUID), and filter events by their severity.

You can also export the list of organization events as a CSV file.

To view Organization events:

1. From the PCE web console menu, choose **Troubleshooting** > **Events**.
2. As the top of the page, you can use the Event Filter to filter the list by event type.

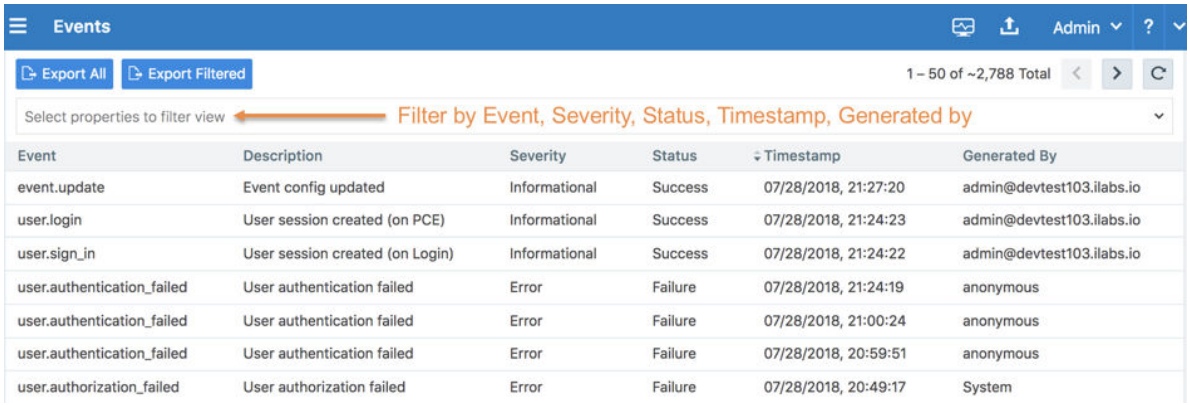

NOTE

In the Events Viewer, the suggested values for the filters are generated from all possible values. For example, the “Generated By” filter shows all users on the system. However, the actual results displayed by that filter might not contain any data.

VEN Event Not Displayed in PCE Web Console

The following events related to VENs are not currently viewable in the PCE web console. This is a two-column list of event names.

VEN Events not shown in PCE Web Console	
fw_tampering_revert_failure	lost_agent
fw_tampering_reverted	missing_os_updates
fw_tampering_subsystem_failure	pce_incompat_api_version
invoke_powershell_failure	pce_incompat_version
ipsec_conn_state_change	pce_reachable
ipsec_conn_state_failure	pce_unreachable
ipsec_monitoring_failure	proc_config_failure
ipsec_monitoring_started	proc_envsetup_failure
ipsec_monitoring_stopped	proc_init_failure
ipsec_subsystem_failure	proc_malloc_failure
ipsec_subsystem_started	proc_restart_failure
ipsec_subsystem_stopped	proc_started
refresh_token_failure	proc_stopped
refresh_token_success	

VEN href Added to Events Information

After the 22.3.0 upgrade, all events created by a VEN includes the VEN href as well as the previously included Agent href. The VEN href can be used to query the VEN API, obtain the workload record, and execute various operations on the VEN from the PCE.

View Events Using PCE Command Line

Run this 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 this command at any runlevel to display:

- The amount of disk space used by events
- The total number of events

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

Export Events Using PCE Web Console

You can export all Organization events, or export a filtered list organization events to a CSV file.

To export events:

1. From the PCE web console menu, choose **Troubleshooting** > **Events**.
   You see a list of events based on the activities performed.
2. Click **Export** > **Export All** to export all Organization events.
3. To export a filtered list of a events, filter the list and then click **Export** > **Export Filtered** to export only the filtered view.
4. To search for events based on event type, severity, status, timestamp, and who generated them, use the search filter:

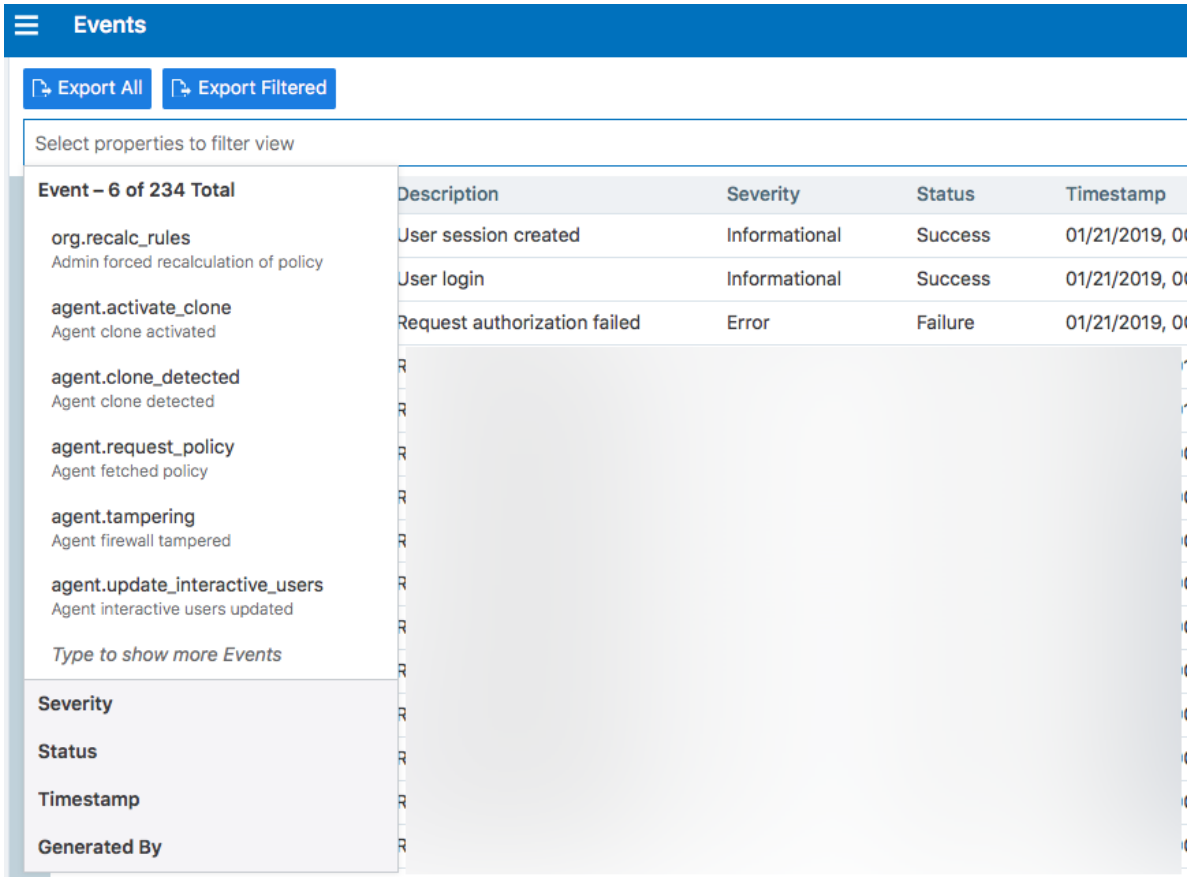

5. For a faster filtering via the browser, use the following field:

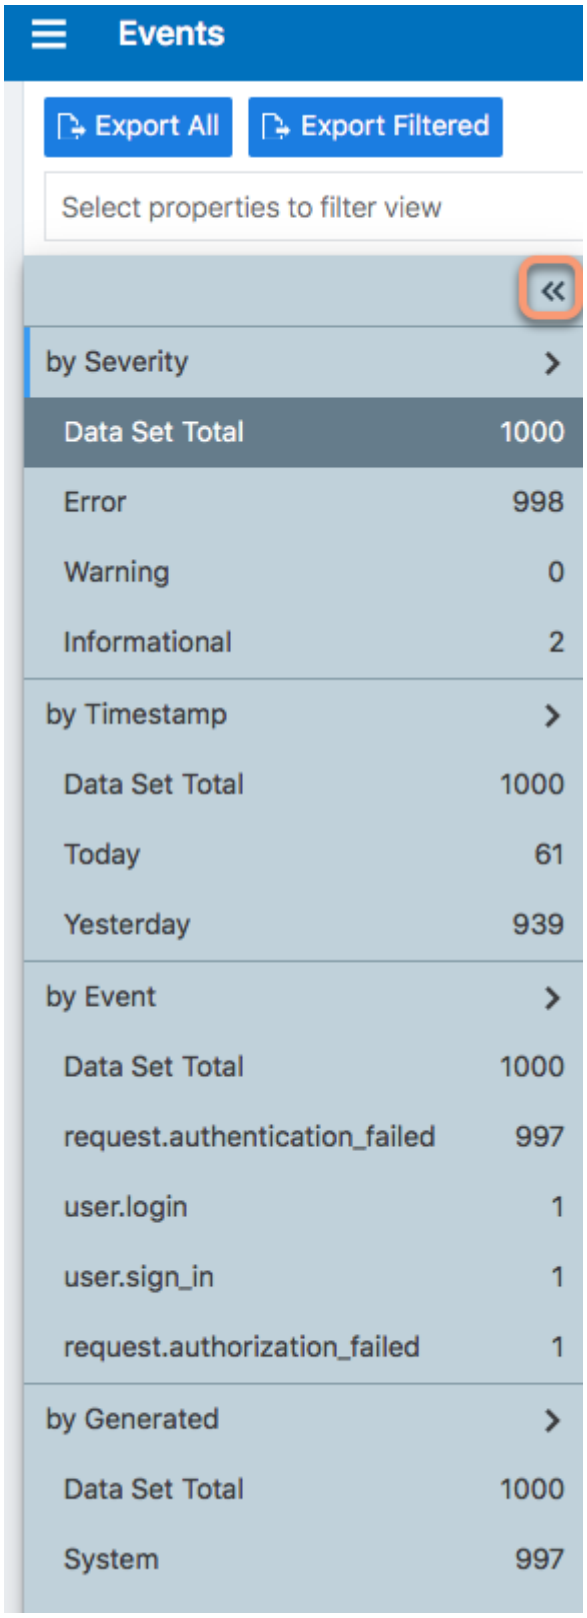

Examples of Events

This section presents examples of recorded events in JSON, CEF, and LEEF for various auditing needs.

User Password Update Failed (JSON)

This example event shows a user password change that failed validation. Event type `user.update_password` shows `"status": "failure"`, and the notification shows that the user's attempted new password did not meet complexity requirements.


```
{
        "href": "/orgs/1/events/xxxxxxxx-39bd-43f1-a680-cc17c6984925",
        "timestamp": "2018-08-29T22:07:00.978Z",
        "pce_fqdn": "pce1.bigco.com",
        "created_by": {
               "system": {}
        },
        "event_type": "user.update_password",
        "status": "failure",
        "severity": "info",
        "action": {
               "uuid": "xxxxxxxx-a5f7-4975-a2a5-b4dbd8b74493",
               "api_endpoint": "/login/users/password/update",
               "api_method": "PUT",
               "http_status_code": 302,
               "src_ip": "10.3.6.116"
        },
        "resource_changes": [],
        "notifications": [{
               "uuid": "xxxxxxxx-7b8e-4205-a62a-1f070d8a0ee2",
               "notification_type": "user.pw_complexity_not_met",
               "info": null
        }, {
               "uuid": "xxxxxxxx-9721-4971-b613-d15aa67a4ee7",
               "notification_type": "user.pw_change_failure",
               "info": {
                      "reason": "Password must have minimum of 1 new character(s)"
               }
        }],
        "version": 2
}
```


Resource Updated (JSON)

This example shows the before and after values of a successful update event `rule_set.update`. The name of the ruleset changed from `"before": "rule_set_2"` to `"after": "rule_set_3"`.


```
{ "href": "/orgs/1/events/xxxxxxxx-8033-4f1a-83e9-fde57c425807",
"timestamp": "2018-08-29T22:04:04.733Z",
"pce_fqdn": "pce1.bigco.com",
"created_by": {
"user": {
"href": "/users/1",
"username": "albert.einstein@bigco.com"
}
},
"event_type": "rule_set.update",
"status": "success",
"severity": "info",
```

```
"action": {
"uuid": "xxxxxxxx-7488-480b-9ef9-0cd2a8496004",
"api_endpoint": "/api/v2/orgs/1/sec_policy/draft/rule_sets/6",
"api_method": "PUT",
"http_status_code": 204,
"src_ip": "10.3.6.116"
},
"resource_changes": [{
"uuid": "xxxxxxxx-1d13-4e5e-8f0b-e0e8bccc44e0",
"resource": {
"rule_set": {
"href": "/orgs/1/sec_policy/draft/rule_sets/6",
"name": "rule_set_3",
"scopes": [
[{
"label": {
"href": "/orgs/1/labels/19",
"key": "app",
"value": "app2"
}
}, {
"label": {
"href": "/orgs/1/labels/20",
"key": "env",
"value": "env2"
}
}, {
"label": {
"href": "/orgs/1/labels/21",
"key": "loc",
"value": "loc2"
}
}]
]
}
},
"changes": {
"name": {
"before": "rule_set_2",
"after": "rule_set_3"
}
},
"change_type": "update"
}],
"notifications": [],
"version": 2
}
```


Security Rule Created (JSON)

In this example of a successful `sec_rule` composite event, a new security rule is created. Because this is a creation event, the `before` values are `null`.


```
{ "href": "/orgs/1/events/xxxxxxxx-6d29-4905-ad32-ee863fb63697",
"timestamp": "2018-08-29T21:48:28.954Z",
"pce_fqdn": "pce24.bigco.com",
```

```
"created_by": {
"user": {
"href": "/users/1",
"username": "albert.einstein@bigco.com"
}
},
"event_type": "sec_rule.create",
"status": "success",
"severity": "info",
"action": {
"uuid": "xxxxxxxx-165b-4e06-aaac-60e4d8b0b9a0",
"api_endpoint": "/api/v2/orgs/1/sec_policy/draft/rule_sets/1/sec_rules",
"api_method": "POST",
"http_status_code": 201,
"src_ip": "10.6.1.156"
},
"resource_changes": [{
"uuid": "9fcf6feb-bf25-4de8-a68a-a50598df4cf6",
"resource": {
"sec_rule": {
"href": "/orgs/1/sec_policy/draft/rule_sets/1/sec_rules/5"
}
},
"changes": {
"rule_list": {
"before": null,
"after": {
"href": "/orgs/1/sec_policy/draft/rule_sets/1"
}
},
"description": {
"before": null,
"after": "WinRM HTTP/HTTPS and RDP"
},
"type": {
"before": null,
"after": "SecRule"
},
"resolve_labels": {
"before": null,
"after": "1010"
},
"providers": {
"created": [{
"provider": true,
"actors": "ams"
}]
},
"consumers": {
"created": [{
"provider": false,
"actors": "ams"
}, {
"provider": false,
"ip_list": {
```

```
"href": "/orgs/1/sec_policy/draft/ip_lists/1"
}
}]
},
"ingress_services": {
"created": [{
"href": "/orgs/1/sec_policy/draft/services/7",
"name": "WinRM HTTP/HTTPS and RDP"
}]
}
},
"change_type": "create"
}],
"notifications": [],
"version": 2
}
```


User Logged In (JSON)


```
[
{
  "href": "/orgs/1/events/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "timestamp": "2019-06-25T23:34:12.948Z",
  "pce_fqdn": "someFullyQualifiedDomainName",
  "created_by": {
    "user": {
      "href": "/users/1",
      "username": "someUser@someDomain"
    }
  },
  "event_type": "user.sign_in",
  "status": "success",
  "severity": "info",
  "action": {
    "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "api_endpoint": "/login/users/sign_in",
    "api_method": "POST",
    "http_status_code": 302,
    "src_ip": "xxx.xxx.xx.x"
  },
  "resource_changes": [
    {
      "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
      "resource": {
        "user": {
          "href": "/users/1",
          "type": "local",
          "username": "someUser@someDomain"
        }
      },
      "changes": {
        "sign_in_count": {
          "before": 4,
          "after": 5
        }
      },
```

```
      "change_type": "update"
    }
  ],
  "notifications": [
    {
      "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
      "notification_type": "user.login_session_created",
      "info": {
        "user": {
          "href": "/users/1",
          "type": "local",
          "username": "someUser@someDomain"
        }
      }
    }
  ]
},
{
  "href": "/orgs/1/events/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "timestamp": "2019-06-25T23:34:15.147Z",
  "pce_fqdn": "someFullyQualifiedDomainName",
  "created_by": {
    "user": {
      "href": "/users/1",
      "username": "someUser@someDomain"
    }
  },
  "event_type": "user.login",
  "status": "success",
  "severity": "info",
  "action": {
    "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "api_endpoint": "/api/v2/users/login",
    "api_method": "GET",
    "http_status_code": 200,
    "src_ip": "xxx.xxx.xx.x"
  },
  "resource_changes": [

  ],
  "notifications": [
    {
      "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
      "notification_type": "user.pce_session_created",
      "info": {
        "user": {
          "href": "/users/1",
          "username": "someUser@someDomain"
        }
      }
    }
  ]
}
]
```

User Logged Out (JSON)


```
[
{
  "href": "/orgs/1/events/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "timestamp": "2019-06-25T23:35:16.636Z",
  "pce_fqdn": "someFullyQualifiedDomainName",
  "created_by": {
    "user": {
      "href": "/users/1",
      "username": "someUser@someDomain"
    }
  },
  "event_type": "user.sign_out",
  "status": "success",
  "severity": "info",
  "action": {
    "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "api_endpoint": "/login/logout",
    "api_method": "GET",
    "http_status_code": 302,
    "src_ip": "xxx.xxx.xx.x"
  },
  "resource_changes": [

  ],
  "notifications": [
    {
      "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
      "notification_type": "user.login_session_terminated",
      "info": {
        "reason": "user_logout",
        "user": {
          "href": "/users/1",
          "username": "someUser@someDomain"
        }
      }
    }
  ]
},
{
  "href": "/orgs/1/events/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "timestamp": "2019-06-25T23:35:16.636Z",
  "pce_fqdn": "someFullyQualifiedDomainName",
  "created_by": {
    "user": {
      "href": "/users/1",
      "username": "someUser@someDomain"
    }
  },
  "event_type": "user.sign_out",
  "status": "success",
  "severity": "info",
  "action": {
    "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
```

```
    "api_endpoint": "/login/logout",
    "api_method": "GET",
    "http_status_code": 302,
    "src_ip": "xxx.xxx.xx.x"
  },
  "resource_changes": [

  ],
  "notifications": [
    {
      "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
      "notification_type": "user.login_session_terminated",
      "info": {
        "reason": "user_logout",
        "user": {
          "href": "/users/1",
          "username": "someUser@someDomain"
        }
      }
    }
  ]
}
]
```

Login Failed — Incorrect Username (JSON)

```
{
  "href": "/orgs/1/events/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "timestamp": "2019-06-25T23:35:41.560Z",
  "pce_fqdn": "someFullyQualifiedDomainName",
  "created_by": {
    "system": {
    }
  },
  "event_type": "user.sign_in",
  "status": "failure",
  "severity": "info",
  "action": {
    "uuid": "someFullyQualifiedDomainName",
    "api_endpoint": "/login/users/sign_in",
    "api_method": "POST",
    "http_status_code": 200,
    "src_ip": "xxx.xxx.xx.x"
  },
  "resource_changes": [

  ],
  "notifications": [
    {
      "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
      "notification_type": "user.login_failed",
      "info": {
        "associated_user": {
          "supplied_username": "invalid_username@someDomain"
        }
      }
```

```
    }
  ]
}
```

Login Failed — Incorrect Password (JSON)

```
{
  "href": "/orgs/1/events/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "timestamp": "2019-06-25T23:35:27.649Z",
  "pce_fqdn": "someFullyQualifiedDomainName",
  "created_by": {
    "system": {
    }
  },
  "event_type": "user.sign_in",
  "status": "failure",
  "severity": "info",
  "action": {
    "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "api_endpoint": "/login/users/sign_in",
    "api_method": "POST",
    "http_status_code": 200,
    "src_ip": "xxx.xxx.xx.x"
  },
  "resource_changes": [

  ],
  "notifications": [
    {
      "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
      "notification_type": "user.login_failed",
      "info": {
        "associated_user": {
          "supplied_username": "someUser@someDomain"
        }
      }
    }
  ]
}
```

User Log Out (CEF)

This example of an event record in CEF shows a successful user log out.

```
CEF:0|Illumio|PCE|19.3.0|user.logout.success|User Logout Success|1|rt=Mar 06 2020
18:38:59.900 +0000 dvchost=mypce.com duser=system dst=10.6.5.4 outcome=success
cat=audit_events request=/api/v2/users/logout_from_jwt requestMethod=POST reason=204
 cs2= cs2Label=resource_changes cs4=[{"uuid":"b5ba8bf0-7ca8-47fc-870f-6c61ddc1648d",
"notification_type":"user.pce_session_terminated","info":{"reason":"user_logout",
"user":{"href":"/users/1","username":"testuser@mypce.com"}}}] cs4Label=notifications
cn2=2 cn2Label=schema-version cs1Label=event_href cs1=/system_events/
e97bd255-4316-4b5e-a885-5b937f756f17
```

Workload Security Policy Updated (LEEF)

This example of an event record in LEEF shows a successful update of security policy for a workload's Ethernet interfaces.

```
LEEF:2.0|Illumio|PCE|18.2.0|interface_status.update.success|src=xx.xxx.xxx.xxx
cat=organizational devTime=someUTCdatetime devTimeFormat=yyyy-mm-dd'T'HH:mm:ss.ttttttZ
sev=1
usrName=albert.einstein url=/orgs/7/agents/someUUID version=2 pce_fqdn=someFQDN
created_by={"agent":{"href":"/orgs/7/agents/someUUID","hostname":"someHostname"}}
action={"uuid":"someUUID",
"api_endpoint":"/api/v6/orgs/7/agents/xxxxxx/interface_statuses/update",
"api_method":"PUT","http_status_code":200,"src_ip":"someIP"}
resource_changes=[{"uuid":"someUUID",
"resource":{"workload":{"href":"/orgs/7/workloads/someUUID","name":null,
"hostname":"someHostname",
"labels":[{"href":"/orgs/7/labels/xxxxxx","key":"loc","value":"test_place_1"},
{"href":"/orgs/7/labels/xxxxxx","key":"env","value":"test_env_1"},
{"href":"/orgs/7/labels/xxxxxx","key":"app","value":"test_app_1"},
{"href":"/orgs/7/labels/xxxxxx","key":"role","value":"test_access_1"}]}},
"changes":{"workload_interfaces":
{"updated":[{"resource":
{"href":"/orgs/7/workloads/someUUID/interfaces/eth1","name":"eth0","
address":{"family":2,"addr":xxxxxxxxx,"mask_addr":someMask}},
"changes":{"address":{"before":null,"after":
{"family":2,"addr":xxxxxxxxx,"mask_addr":someMask}},
"cidr_block":{"before":null,"after":16},"default_gateway_address":
{"before":null,"after":{"family":2,"addr":someGateway,"mask_addr":someMask}},
"link_state":{"before":"unknown","after":"up"},
"network":{"before":null,"after":{"href":"/orgs/7/networks/xx"}},
"network_detection_mode":{"before":null,"after":"single_private_brn"}}},
{"resource":{"href":"/orgs/7/workloads/someUUID/interfaces/eth1",
"name":"eth1","address":{"family":2,"addr":someAddress,"mask_addr":someMask}},
"changes":{"address":{"before":null,"after":{"family":2,"addr":someAddress,
"mask_addr":someMask}},
"cidr_block":{"before":null,"after":16},"link_state":{"before":"unknown","after":"up"},
"network":{"before":null,"after":{"href":"/orgs/7/networks/xx"}},
"network_detection_mode":{"before":null,"after":"single_private_brn"}}}]}},
"change_type":"update"}] notifications=[] event_href=/orgs/7/events/someUUID
```


Differences from Previous Releases

The following table indicates which event names changed in the Illumio Core 18.2 release. If you are upgrading from a release prior to 18.2, be sure to use the current event name in your alert monitoring system.

Changed VEN Event Names

This table lists the names of VEN-related events prior to the Illumio Core 18.2 release and the names they were changed to in the 18.2 release.

Old Name Prior to 18.2	New Name as of 18.2
fw_config_change	agent.firewall_config
activation_success	agent.activate
activation_failure	
deactivation_success	agent.deactivate
deactivation_failure	

Events Monitoring Best Practices

The Illumio Core generates a rich stream of structured messages that provide the following information:

- Illumio PCE system health
- Illumio PCE notable activity
- Illumio VEN notable activity

Illumio Core events are structured and actionable. Using the event data, you can identify the severity, affected systems, and what triggered the event. Illumio Core sends the structured messages using the syslog protocol to remote systems, such as Splunk and QRadar. You can set up your remote systems to automatically process the messages and alert you.

Monitoring Operational Practices

In addition to setting up an automated system, Illumio recommends implementing the following operational practices:

1. Determine the normal quantity of events from the Illumio Core and monitor the trend for changes; investigate spikes or reductions in the event generation rate.
2. Implement good operational practices to troubleshoot and investigate alerts, and to recover from events.
3. Do not monitor Illumio Core events in isolation. Monitor them as part of your overall system. Understanding the events in the context of your overall system activity can provide as much information as the events themselves.

Recommended Events to Monitor

As a best practice, Illumio recommendations you monitor the following events at a minimum.

Events	Description
Program name = `Illumio_pce/system_health` Severity = Warning, Error, or Fatal	Provides multiple systems metrics, such as CPU and memory data, for each node in a PCE cluster. The PCE generates these events every minute. The Severity field is particularly important. When system metrics exceed thresholds, the severity changes to warning, error, or fatal. For more information about the metrics and thresholds, see the PCE Administration Guide. **Recommendation:** Monitor `system_health` messages with a severity of warning or higher and correlate the event with other operational monitoring tools to determine if administrative intervention is required.
`event_type="lost_agent.found"`	Contains the information necessary to identify workloads with lost agents. A lost agent occurs when the PCE deletes a workload from its database but that workload still has a VEN running on it. **Recommendation:** Monitor `lost_agent.found` events and send alerts in case you need to pair the workloads' VENs with the PCE again.
`event_type="system_task.agent_missed_heartbeats_check"`	Lists the VENs that missed three heartbeats (usually 15 minutes). Typically, this event precedes the PCE taking the VENs offline to perform internal maintenance. This event triggers an alert to be sent at 25% of the time configured in the offline timer. For example, if the offline timer is configured to 1 hour, an alert is sent after the VEN has not sent a heartbeat for 15 minutes; if the offline timer is configured to 4 hours, an alert is sent after the VEN hasn't sent a heartbeat for 1 hour. **Recommendation:** Monitor these events for high-value workloads because the PCE can take these workloads offline when the VENs miss 12 heartbeats (usually 60 minutes).
`event_type="system_task.agent_offline_check"`	Lists VENs that the PCE has marked offline, usually because they missed 12 heartbeats. The VENs on these workloads haven't communicated with the PCE for an hour and it removed the workloads from policy. **Recommendation:** Monitor these events for high-value workloads because they indicate change in the affected workloads' security posture.
`event_type="agent.suspend"`	Indicates that the VEN is suspended and no longer protecting the workload. If you did not intentionally run the VEN suspend command on the workload, this event can indicate the workload is under attack. **Recommendation:** Monitor these events for high-value workloads.
`event_type="agent.tampering"`	Indicates tampering of the workload's Illumio managed firewall and that the VEN recovered the firewall. Firewall tampering is one of the first signs that a workload is compromised. During a tampering attempt, the VEN and PCE continue to protect the workload; however, you should investigate the cause of the event. **Recommendation:** Monitor these events for high-value workloads.
`event_type="agent.update"`	Contains the state data that the VEN regularly sends to the PCE. Typically, these events contain routine information; however, the VEN can attach a notice indicating the following issues: • Processes not running • Policy deployment failure

Events	Description
	Recommendation: Monitor `agent.update` events that include notifications because they indicate workloads that might require administrative intervention.

`event_type="rule_set.create"`
`event_type="rule_set.update"`
`event_type="rule_sets.delete"`
	Contains the labels indicating the scope of a draft ruleset. Illumio Core generates these events when you create, update, or delete a draft ruleset. When you include "All Applications," "All Environments," or "All Locations" in a ruleset scope, the PCE represents that label type as a null HREF. Ruleset scopes that are overly broad affect a large number of workloads. Draft rulesets do not take effect until they are provisioned.
	
	Recommendation: Monitor these events to pinpoint ruleset scopes that are unintentionally overly broad.

`event_type="sec_rule.create"`
`event_type="sec_rule.update"`
`event_type="sec_rule.delete"`
	Contains labels indicating when all workloads affected, all services, or a label/label-group are used as a rule provider or consumer. Illumio Core generates these events when you create, update, or delete a draft ruleset. The removed or added labels could represent high-value applications or environments.
	
	Recommendation: Monitor these events for high-value labels.

`event_type="sec_policy.create"`	**[NEW in Illumio** Core **19.3.0]** Contains the `workloads_affected` field, which includes the number of workloads affected by a policy. Illumio Core generates this event when you provision draft policy that updates the policy on affected workloads. The number of affected workloads could be high or a significant percentage of your managed workloads. **Recommendation:** Monitor the `workloads_affected` field for a high number of affect workloads. If the number exceeds an acceptable threshold, investigate the associated the policy.
`event_type="agent.clone_detected"`	The PCE detects cloned VENs based on clone token mismatch. This is a special alert from the Illumio Core release 19.3.2 onwards, as clones have become a higher priority. Volume of these events make the severity level important and not the fact that these events occurred. **Recommendation:** If severity is 1 or 'error', some intervention may be needed.
	 **NOTE** **Automatic Cloned VEN Remediation** For on-prem domain joined Windows workloads, cloned VENs support automatic clone remediation by detecting changes to the workload's domain Security identifier (SID). After the VEN reports such changes to the PCE, the PCE tells the clone to re-activate itself, after which the cloned VEN is remediated and becomes a distinct agent from the original VEN.

Events Setup

This section describes PCE settings related to events and how to use them to configure PCE behavior.

Requirements for Events Framework

To use the events framework, ensure that you allocate enough disk space for event data, and be familiar with the disk capacity requirements.

Database Sizing for Events

Disk space for a single event is estimated at an average 1,500 bytes.

CAUTION

As the number of events increases, the increase in disk space is not a straight line. The projections below are rough estimates. Disk usage can vary in production and depends on the type of messages stored.

Number of Events	Disk Space
25 million	38GB
50 million	58GB

Data and Disk Capacity for Events

For Illumio Core Cloud customers, Illumio Operations manages all data and disk capacity requirements and configuration for events; including the default events data retention period, database dumps with and without events data, and disk compacting.

For more information, contact your Illumio Support representative.

Events Preview Runtime Setting

If you participated in the preview of Events in 18.1.0, the preview was enabled by configuring a setting in your PCE `runtime_env.yml` file.

WARNING

Remove preview parameter from runtime_env.yml

Before you upgrade to the latest release, you must remove `v2_auditable_events_recording_enabled: true` from `runtime_env.yml`. Otherwise, the upgrade does not succeed.

Removing this preview parameter does not affect the collection of “organization events” records, which continue to be recorded.

To remove the Events preview setting:

1. Edit the `runtime_env.yml` file and remove the line `v2_auditable_events_recording_enabled` :

   ```
   v2_auditable_events_recording_enabled: true
   ```

   If you are not participating in any other previews, you can also remove the line `enable_preview_features`.
2. Save your changes.

Events Settings

The following section describes how to configure the Events Settings in the PCE web console.

Events Are Always Enabled

Events are enabled by default in the PCE and cannot be disabled, in accordance with Common Criteria compliance.

Use the PCE web console to change event-related settings and the PCE `runtime_env.yml` for traffic flow summaries.

Event Settings in PCE Web Console

From the PCE web console, you can change the following event-related settings:

- **Event Severity:** Sets the severity level of events to record. Only messages at the set severity level and higher are recorded. The default severity is "Informational."
- **Retention Period:** The system retains event records for a specified number of days; from 1 day to 200 days with the default period being 30 days.
- **Event Pruning:** The system automatically prunes events based on disk usage and the age of events; events older than the retention period are pruned. When pruning is complete, the `system_task.prune_old_log_events` event is recorded.
- **Event Format:** Sets the message output to one of the three formats. The selected message output format only applies to messages that are sent over syslog to a SIEM. The REST API always returns events in JSON.
  - JavaScript Object Notation (JSON): The default; accepted by Splunk and QRadar SIEMs
  - Common Event Format (CEF): Accepted by ArcSight
  - Log Event Extended Format (LEEF): Accepted by QRadar

Event Severity Levels

Severity	Description
Emergency	System is unusable
Alert	Should be corrected immediately
Critical	Critical conditions
Error	Error conditions
Warning	Might indicate that an error will occur if action is not taken
Notice	Events that are unusual, but not error conditions
Informational	Normal operational messages that require no action
Debug	Information useful to developers for debugging the application

Output Format Change

The output format can be changed in the PCE web console:

- JSON (default)
- CEF
- LEEF

Records are in JSON format until you change to one of the other formats. Then, the new events are recorded in the new format; however, the earlier events are not changed to the selected format and they remain recorded in JSON.

Set Event Retention Values

You can set the event retention values depending on the specific conditions described below.

If you are using a SIEM, such as Splunk as the primary long-term storage for events and traffic in a dynamic environment, consider setting the event retention period to 7 days. On setting it to 7 days, you can use the PCE Troubleshooting or Events Viewer to quickly troubleshoot and diagnose events. The benefit of setting 7 days is that if an issue occurs on a Friday, it can still be diagnosed on the following Monday. A large number of events are generated in a dynamic environment, which increases the data stored (disk space used), backup size, and so on. The period of 7 days provides a good balance between disk usage and the ability to troubleshoot.

NOTE

A dynamic environment is when applications and infrastructure are subject to frequent changes; for example, usage of APIs, ETL, Containers, and so on.

If you are using a SIEM in a non-dynamic environment, consider setting the event retention period to 30 days. A smaller number of events are generated, and less disk space is used in a non-dynamic environment.

If you are not using a SIEM such as Splunk and the PCE is the primary storage for the events data used for reporting, diagnosis, and troubleshooting, set the event retention period as per the organization's record retention policy, such as 30 days. If you generate quarterly reporting using events, set the event retention period to 90 days.

SIEM	Consideration	Value
Yes: Primary storage for events	If primary storage of events is not on the PCE	7 days (PCE troubleshooting) 1 day (minimum)
No: Not primary storage for events	If primary storage of events is on the PCE, consider the organization's record retention policy as well as the available disk and event growth pattern	30 days (default)
No	• If the organization's record retention is more than 30 days • If disk monitoring is not set up, it is required to set up disk monitoring	As per your record retention policy 200 days (maximum)
Not applicable	If events data is not needed for reporting or troubleshooting	1 day (minimum)

If disk space availability and event growth projections indicate that the desired retention period cannot be safely supported, consider using a SIEM because the PCE might not store events for the desired period.

NOTE

Running the `illumio-pce-db-management events-db` command provides an output of the average number of events and the storage used.

Configure Events Settings in PCE Web Console

1. From the PCE web console menu, choose **Settings** > **Event Settings** to view your current settings.
2. Click **Edit** to change the settings.
   - For Event Severity, select from the following options:
     - Error
     - Warning
     - Informational
   - For Retention Period, enter the number of days you want to retain data.
   - For Event Format, select from the following options:
     - JSON
     - CEF
     - LEEF
3. Click **Save** once you're done.

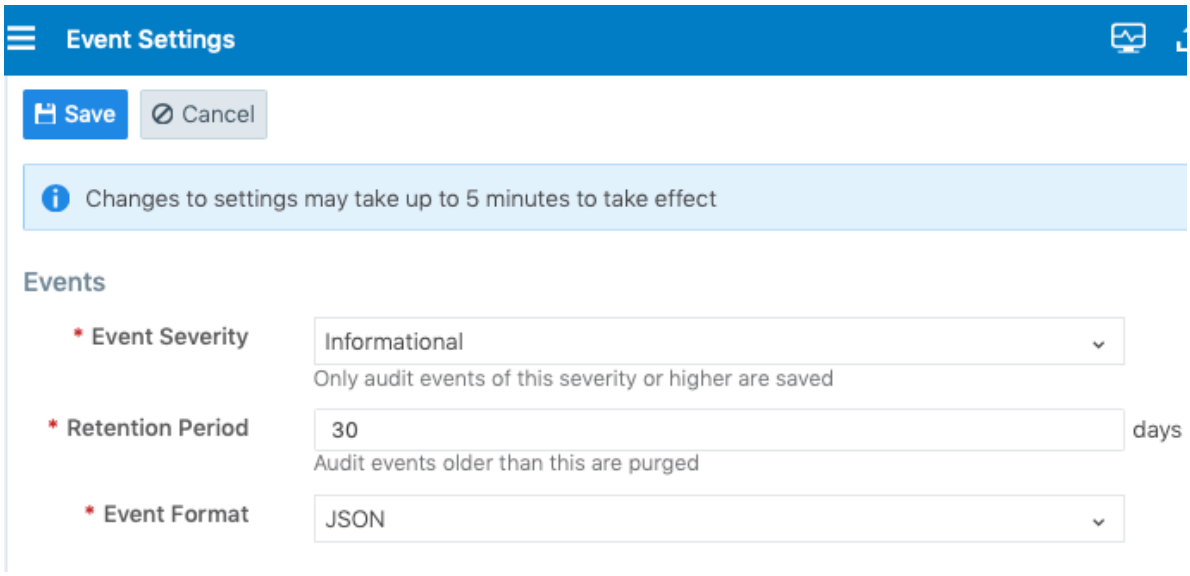

Limits on Storage

From the Illumio Core 19.3.1 release onwards, the PCE will automatically limit the maximum number of events stored. The limits are set on the volume of events stored locally in the PCE database, so that the events recorded in the database do not fill up the disk. The limit is a percentage of the disk capacity, cumulative for all services that store events on the disk.

IMPORTANT

To change the default limits, contact Illumio Support.

The configuration limit includes both hard and soft limits.

- Soft limit: 20% of disk used by event storage

  Aggressive pruning is triggered when the soft limit is reached. However, new events are still recorded while pruning. On the Events list page of the PCE Web Console, the `system_task.prune_old_log_events` event is displayed with the "Object creation soft limit exceeded" message and 'Severity: Informational'.
- Hard limit: 25% of disk used by event storage.

  More aggressive pruning is triggered when the hard limit is reached. New events are not recorded while pruning. On the Events list page of the PCE Web Console, the `system_task.prune_old_log_events` event is displayed with the message "Object creation hard limit exceeded" message and 'Severity: Error'. The pruning continues until the soft limit level of 20% is reached. When this occurs, a `system_task.hard_limit_recovery_completed` event occurs, and the PCE starts to behave as it did for the soft limit conditions.

SIEM Integration for Events

For analysis or other needs, event data can be sent using syslog to your own analytics or SIEM systems.

About SIEM Integration

This guide also explains how to configure the PCE to securely transfer PCE event data in the following message formats to some associated SIEM systems:

- JavaScript Object Notation (JSON), needed for SIEM applications, such as Splunk®.
- Common Event Format (CEF), needed for SIEM applications, such as Micro Focus ArcSight®.
- Log Event Extended Format (LEEF), needed for SIEM applications, such as IBM QRadar®.

Syslog Forwarding

The PCE can export logs to syslog. You can also use the PCE's own internal syslog configuration.

Identify Events in Syslog Stream

Event records from the syslog stream are identified by the following string:

```
"version":2
```

```
AND
```

```
'"href":\s*"/orgs/[0-9]*/events'  OR  '"href":\s*"/system_events/'
```

Forward Events to External Syslog Server

The PCE has an internal syslog repository, "Local" where all the events get stored. You can control and configure the relaying of syslog messages from the PCE to multiple external syslog servers.

To configure forwarding to an external syslog server:

1. From the PCE web console menu, choose **Settings** > **Event Settings**.
2. Click **Add**.
   The Event Settings - Add Event Forwarding page opens.
3. Click **Add Repository**.

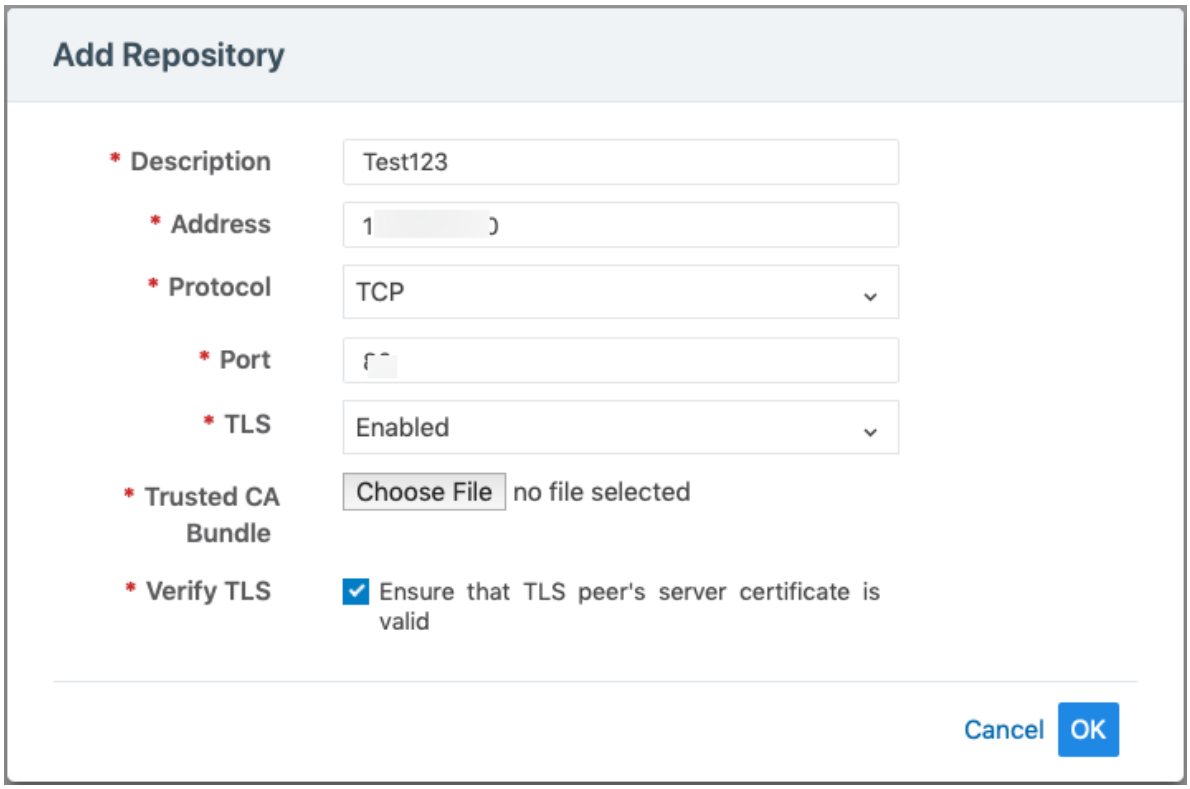

4. In the Add Repository dialog:
   - Description: Enter name of the syslog server.
   - Address: Enter the IP address for the syslog server.
   - Protocol: Select TCP or UDP. If you select UDP, you only need to enter the port number and click **OK** to save the configuration.
   - Port: Enter port number for the syslog server.
   - TLS: Select Disabled or Enabled. If you select Enabled, click “Choose File” and upload your organization's “Trusted CA Bundle” file from the location it is stored on.

     The Trusted CA Bundle contains all the certificates that the PCE (internal syslog service) needs to trust the external syslog server. If you are using a self-signed certificate, that certificate is uploaded. If you are using an internal CA, the certificate of the internal CA must be uploaded as the “Trusted CA Bundle”.
   - Verify TLS: Select the check-box to ensure that the TLS peer’s server certificate is valid.
5. Click **OK** to save the event forwarding configuration.

After ensuring that the events are being forwarded as configured to the correct external syslog servers, you can choose to stop using the “Local” server by editing the local server setting and deselect all message types.

NOTE
You cannot delete the “Local” server.

Disable Health Check Forwarding

PCE system health messages are useful for PCE operations and monitoring. You can choose to forward them if they are needed on the remote destination.

For example, IBM QRadar is usually used by security personnel, who might not need to monitor the PCE system health. The Illumio App for QRadar does not process the PCE system health messages.

The PCE system health messages are only provided in key/value syslog format. They are not translatable into CEF, LEEF, or JSON formats. If your SIEM does not support processing key/value messages in syslog format, do not forward system health messages to those SIEMs. For example, IBM QRadar and Micro Focus ArcSight do not automatically parse these system health messages.

To disable syslog forwarding of health check messages:

1. From the PCE web console menu, choose **Settings** > **Event Settings**.
2. Click the Event listed under the **Events** column.

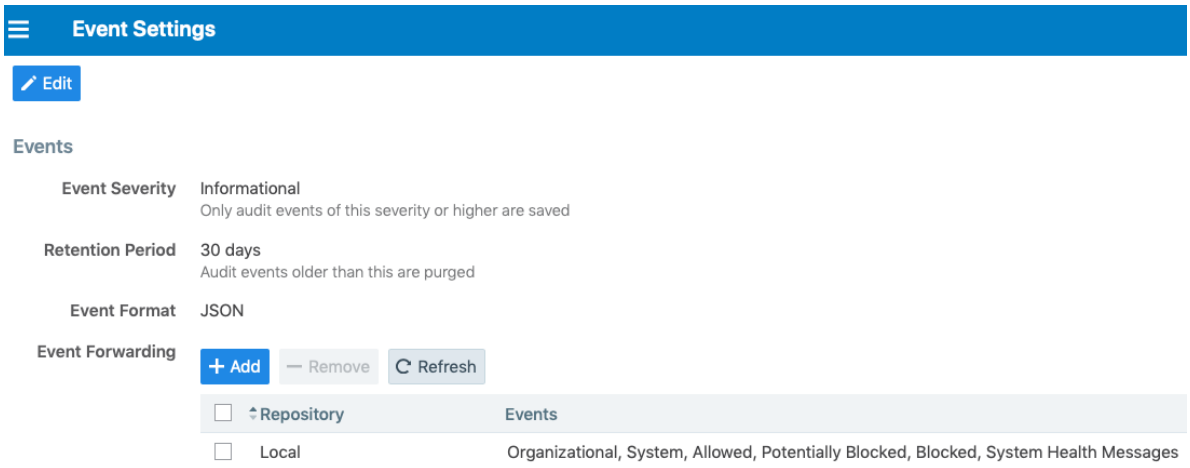

3. Under the Events block, for the Status Logs entry, deselect **System Health Messages**. System health check is only available in key-value format. Selecting a new event format does not change the system health check format to CEF or LEEF.

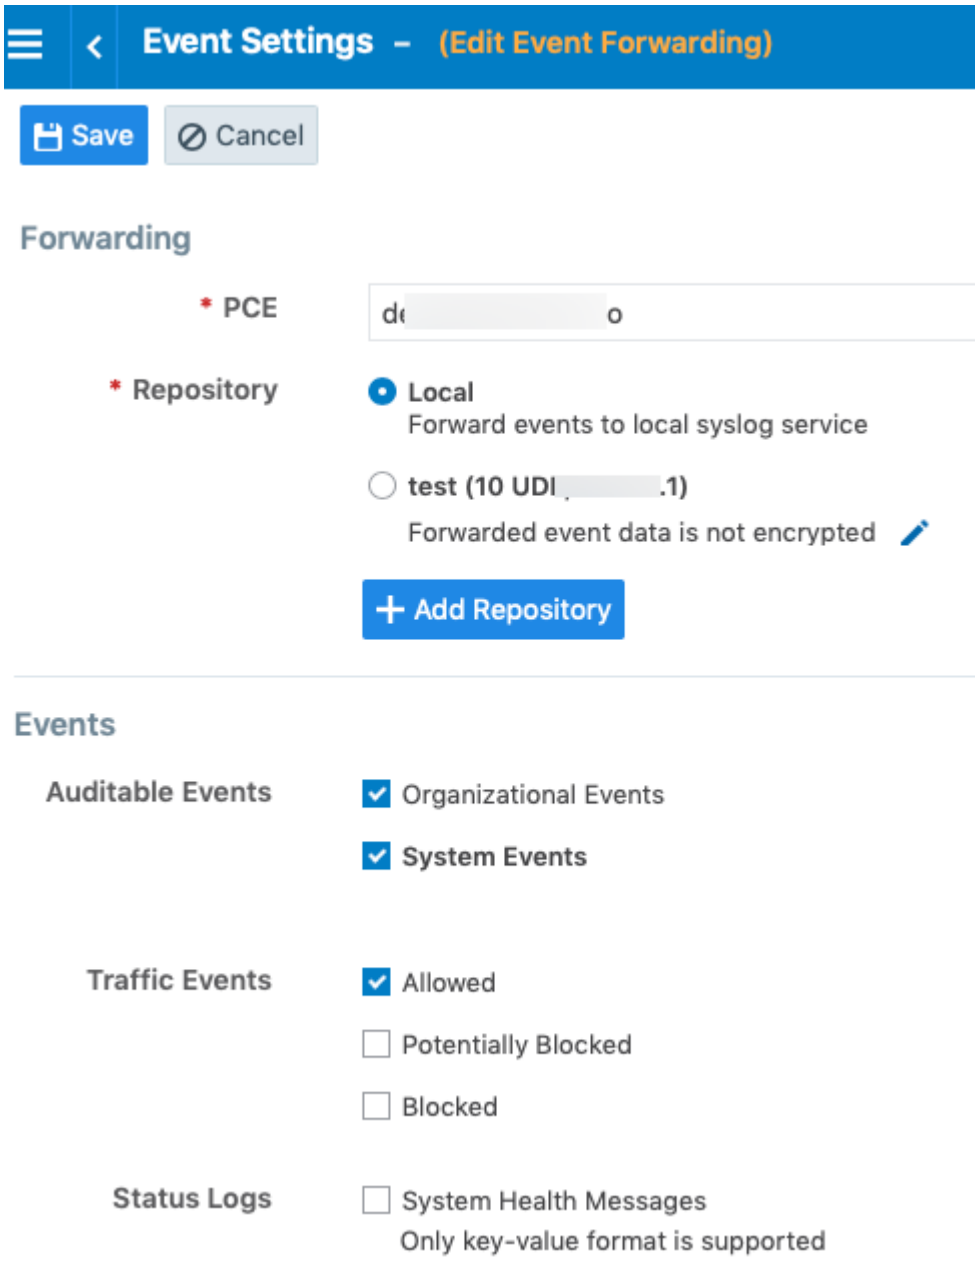

4. Click **Save**.

> **NOTE**
> IBM QRadar and HP ArcSight do not support system health messages. If you are using either of these for SIEM, make sure that you do not select the System Health Messages checkbox.

Traffic Flow Summaries

This section describes traffic flow summaries.

After you install a VEN on a workload and pair the VEN with the PCE, the VEN monitors each workload's traffic flows and sends the traffic flow summaries to the PCE.

Traffic summaries can be exported to syslog or Fluentd. If traffic data is configured for export, the PCE processes the received traffic flow summaries from each VEN and immediately sends them to syslog or Fluentd.

Traffic Flow Types and Properties

The Illumio Core logs traffic flows based on the Visibility setting. Events have attributes that can be Allowed, Blocked, or Potentially Blocked and might not appear in the traffic flow summary.

Visibility Settings

The table below indicates whether or not a traffic summary is logged as Allowed, Potentially Blocked, or Blocked depending on a workload's policy state.

NOTE

Traffic from workloads in the "Idle" policy state is not exported to syslog from the PCE.

Visibility	Logged in Traffic Flow Summary
Off	VEN does not log traffic connection information
Blocked - Low Detail	VEN logs connection information for blocked and potentially blocked traffic only
Blocked + Allowed - High Detail	VEN logs connection information for allowed, blocked, and potentially blocked traffic
Enhanced Data Collection	VEN logs byte counts in addition to connection details for allowed, blocked, and potentially blocked traffic

Event Types

In a traffic flow summary, the event type is designated by Policy Decision (`pd`).

NOTE

An asterisk (*) indicates the attribute might not appear in the summary.

Event Attributes	Allowed (pd=0)	Potentially Blocked (pd=1)	Blocked (pd=2)	Unknown (pd=3)
version	✓	✓	✓	✓
count	✓	✓	✓	✓
interval_sec	✓	✓	✓	✓
timestamp	✓	✓	✓	✓
dir	✓	✓	✓	✓
src_ip	✓	✓	✓	✓
dst_ip	✓	✓	✓	✓
proto	✓	✓	✓	✓
dst_prt	✓	✓	✓	✓
state	✓	✓	✓	✓
pd	✓	✓	✓	✓
code*	✓	✓	✓	✓
type*	✓	✓	✓	✓
dst_vulns*	✓	✓	✓	✓
fqdn*	✓	✓	✓	✓
un*	✓	✓	x	✓
pn*	✓	✓	x	✓
sn*	✓	✓	x	✓
src_labels*	✓	✓	✓	✓
dst_labels*	✓	✓	✓	✓
src_hostname*	✓	✓	✓	✓
dst_hostname*	✓	✓	✓	✓
src_href*	✓	✓	✓	✓
dst_href*	✓	✓	✓	✓

Show Amount of Data Transfer

The JSON, CEF, and LEEF for the accurate byte count work events are related to the 'Show Amount of Data Transfer' preview feature available with the Illumio Core 20.2.0 release.

The PCE now reports amount of data transferred in to and out of workloads and applications in a datacenter. The number of bytes sent by and received by the provider of an application

are provided separately. These values can be seen in traffic flow summaries streamed out of the PCE. This capability can be enabled on a per-workload basis in the Workload page. It can also be enabled in the pairing profile so that workloads are directly paired into this mode.

The direction reported in flow summary is from the viewpoint of the provider of the flow:

Destination Total Bytes Out: Number of bytes transferred out of provider:

```
dst_tbo
```

Destination Total Bytes In: Number of bytes transferred in to provider.

```
dst_tbi
```

To activate the 'Show Amount of Data Transfer' capability on the PCE, contact your Illumio representative.

LEEF Mapping

- LEEF field `X` contains JSON field `Y`
- `srcBytes` contains `dst_tbo`
- `dstBytes` contains `dst_tbi`
- `dbi` contains `dst_dbi`
- `dbo` contains `dst_dbo`

CEF Mapping

- CEF field cn2 is `dst_dbi` with cn2Label is “dbi”
- CEF field cn3 is `dst_dbo` with cn3Label is “dbo”
- CEF field “in” is `dst_tbi`
- CEF field “out” is `dst_tbo`

Manage Traffic Flows Using REST API

You can use the following properties to manage traffic flows using the REST API.

NOTE

You should ignore and *not* use any extra properties that are not described in this document, such as `tbi`, `tbo`, `dbi`, and `dbo`.

Property	Description	Type	Required	Possible Values
`version`	The version of the flow summary schema.	Integer	Yes	4
`timestamp`	Indicates the time (RFC3339) when the first flow in the summary was created, represented in UTC. Format: `yyyy-MM-dd'T'HH:mm:ss.SSSSSSZ`	String	Yes	
`interval_sec`	Sample duration for the flows in the summary. Default is approximately 600 seconds (10 minutes), depending on the VEN's ability to report traffic and PCE's current load.	Integer	Yes	
`dir`	Direction of the first packet: in or out (I, O).	String	Yes	I, O
`src_ip`	Source IP of the flows.	String	Yes	
`dst_ip`	Destination IP of the flows.	String	Yes	
`proto`	Protocol number (0-255).	Integer	Yes	Minimum=0 Maximum=255
`type`	The ICMP message type associated with the first flow in the summary. This value exists only if protocol is ICMP (1). **NOTE** This information is included in blocked flows for VEN versions lower than 19.1.0. It is included in all flows for VEN version 19.1.0 and later. Example: 3 for "Destination Unreachable."	Integer	No	Minimum=0 Maximum=255
`code`	The ICMP message code (subtype) associated with the first flow in the summary. This value exists only if protocol is ICMP (1). **NOTE** This information is included in blocked flows for VEN versions lower than 19.1.0. It is included in all flows for VEN version 19.1.0 and later. Example: 1 for "Destination host unreachable."	Integer	No	Minimum=0 Maximum=255
`dst_port`	Destination port.	Integer	Yes	Minimum=0

Property	Description	Type	Required	Possible Values
	This value exists only if protocol is not TCP (6) or UDP (17).			Maximum=65535
`pd`	Policy decision value, which indicates if the flow was allowed, potentially blocked (but allowed), blocked, or unknown. Possible values: • **0** – Allowed traffic • **1** – Allowed traffic but will be blocked after policy enforcement • **2** – Blocked traffic • **3** - Unknown **NOTE** Policy decision is "unknown" in the following cases: • Flows uploaded using existing bulk API (`/orgs/<org_id>/agents/bulk_traffic_flows`). • Flows uploaded using Network Flow Ingest Application (`/orgs/<org_id>/traffic_data`). • Traffic reported by idle VENs and specifically those that have been reported with "s" state (snapshot).	Integer	Yes	Minimum=0 Maximum=3
`count`	Count of the number of flows in the flow summary.	Integer	Yes	
`state`	Session state for the traffic flows in the flow summaries. Possible values: • **Active (A):** Connection was still open at the time the flow summary was logged. Applies to allowed and potentially blocked flows. • **Closed (C):** (Linux only) Connection closed at the time the flow summary was logged. Applies to allowed and potentially blocked flows. • **Timed out (T):** Connection timed out at the time the flow summary was logged. Applies to allowed and potentially blocked flows. Due to a limitation of WFP, a Windows VEN will report "T" even when the connection is closed at the time the flow summary was logged. • **Snapshot (S):** Snapshot of current connections to and from the VEN, which applies only to workloads whose policy state is set to Idle. Applies to allowed and potentially blocked flows. • **New connection (N):** Dropped TCP packet contains a SYN and is associated with a new connection. Applies to blocked TCP flows. The value is empty for blocked UDP flows.	String	No	A, C, T, S, N
`pn`	The program name is associated with the first flow of the summary. It is supported on inbound flows for Linux and	String	No	

Property	Description	Type	Required	Possible Values
	Windows VEN and on outbound flows for only Windows VEN. **NOTE** This information might not be available on short-lived processes, which are Linux-specific. Currently, flows are aggregated, so this value might represent only the first process detected across all aggregated flows. If network communication is done by an OS component (or a driver), no process is associated with it.			
`un`	The username is associated with the first flow of the summary. It is supported on inbound flows for Linux and Windows VEN and on outbound flows for only Linux VEN. On Windows, it can include the username of the user account that initiated the connection. **NOTE** This information might not be available on short-lived processes.	String	No	
`sn`	Service name associated with the first flow in the summary. It is supported only on inbound flows on Windows VEN.	String	No	
`src_hostname`	Hostname of the source workload that reported the flow.	String	No	
`src_href`	HREF of the source workload that reported the flow.	String	No	
`src_labels`	Labels applied to the source workload. **NOTE** The `src_hostname`, `src_href`, and `src_labels` values are not be included in a traffic summary if the source of the flow is not an Illumio-labeled workload. For example, Internet traffic or a managed workload without any labels applied.	Object	No	
`dst_hostname`	Hostname of the destination workload that reported the flow.	String	No	

Property	Description	Type	Required	Possible Values
`dst_href`	HREF of the destination workload that reported the flow.	String	No	
`dst_labels`	Labels applied to the destination workload. **NOTE** The `dst_hostname`, `dst_href`, and `dst_labels` values are not be included in a traffic summary if the destination of the flow is not an Illumio-labeled workload. For example, Internet traffic or a managed workload without any labels applied.	Object	No	
dst_vulns	Information about the vulnerabilities on the destination of the traffic flow with the specific port and protocol. **NOTE** • Vulnerabilities are defined by Common Vulnerabilities and Exposures (CVE), with identifiers and descriptive names from the U.S. Department of Homeland Security National Cybersecurity Center. • The vulnerability information is sent only when the Vulnerability Maps feature is turned on via a license and the information is imported into the PCE from a Vulnerability Scanner, such as Qualys.	Object	No	
`fqdn`	Fully qualified domain name	String	No	

The following table describes the sub-properties for the `dst_vulns` property:

Sub-property	Description	Type	Required
`count`	The total number of existing vulnerabilities on the destination port and protocol.	Integer	No
`max_score`	The maximum of all the scores for the vulnerabilities on the destination port and protocol.	Number	No
`cve_ids`	The list of CVE-IDs associated with the vulnerabilities that have the maximum score. Up to 100 displayed .	Array	No

Export Traffic Flow Summaries

Decide where to export the traffic flow summaries: syslog or Fluentd.

CAUTION

By default, from the 19.3.0 release on, the PCE generates all traffic flow summaries and sends them to syslog.

If you have not configured syslog, the syslog data by default is written to a local disk. For example, it is written to `/var/log/messages`.

Export to Syslog

To configure and export the traffic flow summaries to a remote syslog, follow these steps:

1. From the PCE web console menu, choose **Settings** > **Event Settings**.
2. Enable a remote syslog destination.
3. Select specific traffic flow summaries to be sent to remote syslog.
   This filters the selected traffic flow summaries and send those to the remote syslog.

To prevent the syslog data from being written to a local disk based on your preference, deselect the Events checkboxes on the **Settings** > **Event Settings** > Local page in the PCE web console. For more information, see Events Settings. [309]

NOTE

The generation of all traffic flow summaries is implemented to ensure that all of the traffic flow summaries are controlled from the PCE web console only.

This example shows the `runtime_env.yml` configuration to generate all types of flow summaries.

Export to Syslog

```
export_flow_summaries_to_syslog:
- accepted
- potentially_blocked
- blocked
```

This example shows the `runtime_env.yml` configuration if you do not want to generate any types of flow summaries.

Export to Syslog

```
export_flow_summaries_to_syslog:
- none
```

NOTE

Illumio does not currently support having a primary and secondary syslog configuration, with disaster recovery and failover.

You can configure it on a system syslog (local) and use the internal syslog configuration to send messages to local, which sends to system syslog.

Export to Fluentd

To generate and export the traffic flow summaries to Fluentd, follow these steps:

1. Set the `export_flow_summaries_to_fluentd` parameter in `runtime_env.yml`.
2. Set the `external_fluentd_aggregator_servers` parameter in `runtime_env.yml`.

This example shows the `runtime_env.yml` configuration to generate two types of flow summaries, out of the three possible types.

Export to Fluentd

```
external_fluentd_aggregator_servers:
- fluentd-server.domain.com:24224
export_flow_summaries_to_fluentd:
- accepted
- blocked
```

Flow Duration Attributes

The 20.2.0 VEN sends two new attributes to the syslog and fluentd output. The new attributes describe the flow duration and are appended to the flow data.

- **Delta flow duration in milliseconds (`ddms`)**: The duration of the aggregate within the current sampling interval. This field enables you to calculate the bandwidth between two applications in a given sampling interval. The formula is dbo (delta bytes out) / delta_duration_ms, or dbi / delta_duration_ms.
- **Total flow duration in milliseconds (`tdms`)**: The duration of the aggregate across all sampling intervals. This field enables you to calculate the average bandwidth of a connection between two applications. The formula is tbo (total bytes out) / total_duration_ms, or tbo / total_duration_ms. It also enables you to calculate the average volume of data in a connection between two applications. The formula is tbo (total bytes out) / count (number of flows in an aggregate), or tbi / count.

Traffic Flow Summary Examples

The following topic provides examples of traffic flow summaries in JSON, CEF, and LEEF, and messages that appear in syslog.

JSON

```
{
  "interval_sec": 600,
```

```
  "count": 1,
  "tbi": 73,
  "tbo": 0,
  "pn": "example-daemon",
  "un": "example",
  "src_ip": "xxx.xxx.xx.xxx",
  "dst_ip": "xxx.x.x.xxx",
  "timestamp": "2018-05-23T16:07:12-07:00",
  "dir": "I",
  "proto": 17,
  "dst_port": 5353,
  "state": "T",
  "src_labels": {
    "app": "AppLabel",
    "env": "Development",
    "loc": "Cloud",
    "role": "Web"
  },
  "src_hostname": "test-ubuntu-3",
  "src_href": "/orgs/1/workloads/xxxxxxxx-7741-4f71-899b-d6f495326b3f",
  "dst_labels": {
    "app": "AppLabel",
    "env": "Development",
    "loc": "AppLocation",
    "role": "Database"
  },
  "dst_hostname": "test-ubuntu-2",
  "dst_href": "/orgs/1/workloads/xxxxxxxx-012d-4651-b181-c6f2b269889e",
  "pd": 1,
  "dst_vulns": {
    "count": 8,
    "max_score": 8.5,
    "cve_ids": [
      "CVE-2016-2181",
      "CVE-2017-2241"
    ]
  },
  "fqdn" : "xxx.ubuntu.com",
  "version": 4
}
```


Syslog


```
2019-02-11T22:50:15.587390+00:00 level=info host=detest01 ip=100.1.0.1
program=illumio_pce/collector| sec=925415.586 sev=INFO pid=9944
 tid=30003240
rid=bb8ff798-1ef2-44b1-b74e-f13b89995520 {"interval_sec":1074,
"count":1,"tbi":3608,
"tbo":0,"pn":"company-daemon","un":"company","src_ip":"10.0.2.15",
"dst_ip":"211.0.0.232",
"class":"M","timestamp":"2019-02-11T14:48:09-08:00","dir":"I",
"proto":17,
"dst_port":5353,"state":"T","src_labels":{"app":"AppName",
"env":"Development","loc":"Cloud","role":"Web"},
"src_hostname":"dev-ubuntu-1",
"src_href":"/orgs/1/workloads/773f3e81-5779-4753-b879-35a1abe45838",
```

"dst_labels":{"app":"AppName","env":"Development","loc":"Cloud2",
"role":"Web"},
"dst_hostname":"dev-ubuntu-1","dst_href":"/orgs/1/workloads/
773f3e81-5779-4753-b879-35a1abe45838","pd":0,"dst_vulns":{"count":1,
"max_score":3.7,
"cve_ids":["CVE-2013-2566","CVE-2015-2808"]},"fqdn":"xxx.ubuntu.com",
"version":4}

Allowed Flow Summary (pd = 0)

2016-01-12T05:23:30+00:00 level=info host=myhost ip=127.0.0.1 program=illumio_pce/
collector| sec=576210.952 sev=INFO pid=25386 tid=16135120 rid=0
{"interval_sec":1244,"count":3,"dbi":180,"dbo":180,"pn":"sshd","un":"root",
"src_ip":"10.6.0.129","dst_ip":"10.6.0.129","timestamp":"2017-08-16T13:23:57-07:00",
"dir":"I","proto":6,"dst_port":22,"state":"A","dst_labels":{"app":"test_app_1","env":
"test_env_1","loc":"test_place_1","role":"test_access_1"},"dst_hostname":"corp-vm-2",
"dst_href":"/orgs/1/workloads/5ddcc33b-b6a4-4a15-b600-64f433e4ab33","pd":0,
"version":4}

Potentially Blocked Flow Summary (pd = 1)

2016-01-12T05:29:21+00:00 level=info host=myhost ip=127.0.0.1 program=illumio_pce/
collector| sec=576561.327 sev=INFO pid=25386 tid=16135120 rid=0 sec=920149.541
sev=INFO pid=1372 tid=30276700 rid=136019d0-f9d8-45f3-ac99-f43dd8015675
{"interval_sec":600,"count":1,"tbi":229,"tbo":0,"src_ip":"172.16.40.5",
"dst_ip":"172.16.40.255","timestamp":"2017-08-16T14:45:58-07:00","dir":"I",
"proto":17,"dst_port":138,"state":"T","dst_labels":{"app":"test_app_1",
"env":"test_env_1","loc":"test_place_1","role":"test_access_1"},"dst_hostname":
"corp-vm-2","dst_href":"/orgs/1/workloads/5ddcc33b-b6a4-4a15-b600-64f433e4ab33",
"pd":1,"version":4}

Blocked Flow Summary (pd = 2)

2016-01-12T05:23:30+00:00 level=info host=myhost ip=127.0.0.1 program=illumio_pce/
collector| sec=576210.831 sev=INFO pid=25386 tid=16135120 rid=0 sec=915000.311
sev=INFO pid=1372 tid=30302280 rid=90a01be5-a3c1-44f9-84fd-3c3a5eaec1f8
{"interval_sec":589,"count":1,"src_ip":"10.6.1.89","dst_ip":"10.6.255.255",
"timestamp":"2017-08-16T13:22:09-07:00","dir":"I","proto":17,"dst_port":138,
"dst_labels":{"app":"test_app_1","env":"test_env_1","loc":"test_place_1",
"role":"test_access_1"},"dst_hostname":"corp-vm-1","dst_href":"/orgs/1/workloads/
a83ba658-576b-4946-800a-b39ba2a2e81a","pd":2,"version":4}

Unknown Flow Summary (pd = 3)

2019-06-14T05:33:45.442561+00:00 level=info host=devtest0 ip=127.0.0.1
program=illumio_pce/collector| sec=490425.442 sev=INFO pid=12381 tid=32524120
rid=6ef5a6ac-8a9c-4f46-9180-c0c91ef94759 {"dst_port":1022,"proto":6,"count":20,
"interval_sec":600,"timestamp":"2019-06-06T21:03:57Z","src_ip":"10.23.2.7",
"dst_ip":"10.0.2.15","dir":"O","state":"S","pd":3,"src_href":"/orgs/1/workloads/
a0d735ce-c55f-4a38-965f-bf6e98173598","dst_hostname":"workload1",
"dst_href":"/orgs/1/workloads/a20eb1b5-10a4-419e-b216-8b35c795a01e","src_labels":
{"app":"app","env":"Development","loc":"Amazon","role":"Load Balancer"}
,"version":4}

CEF

CEF:0|Illumio|PCE|2015.9.0|flow_potentially_blocked|Flow Potentially Blocked|3|
act=potentially_blocked cat=flow_summary deviceDirection=0 dpt=137 src=someIPaddress

```
dst=someIPaddress proto=udp cnt=1 in=1638 out=0 rt=Jun 14 2018 01:50:14
cn1=120 cn1Label=interval_sec cs2=T cs2Label=state cs6=/orgs/1/workloads/
someID cs6Label=dst_href cs4={"app":"CRM","env":"Development","loc":"AppLocation",
"role":"Web"} cs4Label=dst_labels dhost=connectivity-check.someDomainName
cs1={"count":1,"max_score":3.7,"cve_ids": ["CVE-2013-2566","CVE-2015-2808"]}
cs1Label=dst_vulns dvchost=someDomainName
```

Unknown Flow Summary (pd = 3)

```
2019-06-14T21:02:55.146101+00:00 level=info host=devtest0 ip=127.0.0.1
program=illumio_pce/collector| sec=546175.145 sev=INFO pid=15416 tid=40627440
 rid=f051856d-b9ee-4ac8-85ea-4cb857eefa82 CEF:0|Illumio|PCE|19.3.0|flow_unknown|
Flow Unknown|1|act=unknown cat=flow_summary deviceDirection=0 dpt=22 src=10.0.2.2
 dst=10.0.2.15 proto=tcp cnt=6 in=6 out=6 rt=Jun 14 2019 21:02:25 duser=root
dproc=sshd cn1=31 cn1Label=interval_sec cs2=S cs2Label=state dhost=workload1
cs6=/orgs/1/workloads/a20eb1b5-10a4-419e-b216-8b35c795a01e cs6Label=dst_href
dvchost=devtest0.ilabs.io msg=
{"trafclass_code":"U"}
```

LEEF

```
LEEF:2.0|Illumio|PCE|2015.9.0|flow_blocked|cat=flow_summary
devTime=2018-06-14T10:38:53-07:00 devTimeFormat=yyyy-MM-dd'T'HH:mm:ssX
proto=udp sev=5 src=someIPaddress dst=someIPaddress dstPort=5353 count=15
dir=I intervalSec=56728 dstHostname=someHostName dstHref=/orgs/1/workloads/
someID dstLabels={"app":"CRM","env":"Development","loc":"Cloud","role":"Web"}
dstVulns={"count":2,"max_score":3.7} dstFqdn=someDomainName "cve_ids":
["CVE-2013-2566","CVE-2015-2808"]}
```

Unknown Flow Summary (pd = 3)

```
2019-06-14T19:25:53.524103+00:00 level=info host=devtest0 ip=127.0.0.1
program=illumio_pce/collector| sec=540353.474 sev=INFO pid=9960 tid=36072680
 rid=49626dfa-d539-4cff-8999-1540df1a1f61 LEEF:2.0|Illumio|PCE|19.3.0|
flow_unknown|cat=flow_summary devTime=2019-06-06T21:03:57Z
devTimeFormat=yyyy-MM-dd'T'HH:mm:ssX proto=tcp sev=1 src=10.23.2.7
dst=10.0.2.15 dstPort=1022 count=20 dir=O intervalSec=600 state=S
srcHref=/orgs/1/workloads/a0d735ce-c55f-4a38-965f-bf6e98173598 srcLabels=
{"app":"app","env":"Staging","loc":"Azure","role":"API"}
dstHostname=workload1 dstHref=/orgs/1/workloads/a20eb1b5-10a4-419e-b216-8b35c795a01e
```

Illumio Core PCE CLI Tool Guide 1.4.2

Overview of the CLI Tool

This topic provides an overview of the CLI Tool, describes the general syntax of the CLI Tool command, and lists the environment variables you can use to customize the CLI Tool.

IMPORTANT

See the *Illumio Core CLI Tool 1.4.0 Release Notes* and *Illumio Core CLI Tool 1.4.1 Release Notes* and *Illumio CORE CLI Tool 1.4.2* Release Notes in your respective Illumio Core Technical Documentation portal for the updates to the CLI Tool for these releases.

About This Guide

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

CLI Tool Versioning

Illumio Core CLI Tool version 1.4.2 is compatible with Illumio Core PCE versions:

PCE 19.3.6-H2 (LTS)

PCE 21.2.4 (LTS)

PCE 21.5.20 (LTS)

PCE 22.1.1 (Standard)

PCE 22.2.0 (Standard)

The CLI Tool version numbering is independent from the release and version numbering of Illumio Core PCE and VEN. The CLI Tool works with multiple versions of the PCE and the VEN and does not necessarily need software changes in parallel with releases of the PCE or the VEN.

IMPORTANT

See the *Illumio Core CLI Tool 1.4.0 Release Notes*, *Illumio Core CLI Tool 1.4.1 Release Notes* and *Illumio Core CLI Tool 1.4.2 Release Notes* in your respective Illumio Core Technical Documentation portal for the updates to the CLI Tool for these releases.

How to Use This Guide

This guide includes several major sections:

- Overview to the CLI Tool
- Installation
- Formal syntax of the `ilo` command
- Tutorials for various operations
- Uploading vulnerability data
- Security policy import and export

Before Reading This Guide

Before performing the procedures in this guide, be familiar with the following information:

- The CLI Tool interacts with the PCE; therefore, be familiar with PCE concepts such as core and data nodes, workloads, and traffic. See the PCE Administration Guide.
- The CLI Tool is often used to upload vulnerability data; therefore, understand how vulnerability data is used in the PCE web console. See the "Vulnerability Maps" topic in Visualization Guide.
- The CLI Tool can be used with workload data; therefore, you must understand what workloads are. See the "VEN Architecture and Components" topic in the VEN Administration Guide.
- The CLI Tool can be used with security policy rules, rulesets, labels, and similar resources; therefore, be familiar with these concepts. See "The Illumio Policy Model" in the Security Policy Guide.

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
...
```

CLI Tool and PCE Resource Management

With the Illumio CLI Tool, you can manage many of your PCE's resources directly from your local computer.

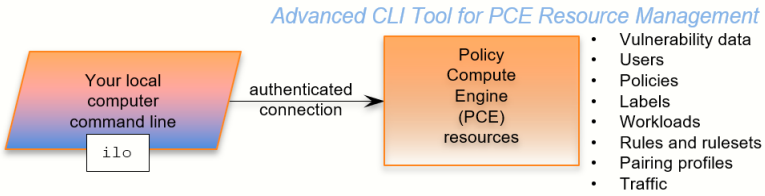

Some purposes of the CLI Tool include the following capabilities:

- Import vulnerability data for analysis with Illumination.
- Help with tasks such as directly importing workload information to create workloads in bulk.
- Create, view, and manage your organization's security policy rules, rulesets, labels, and other resources.

CAUTION

The CLI Tool is a powerful way to work with your PCE resources. Exercise caution to make sure that your use of the tool does not adversely affect your system. If possible, test your CLI Tool commands against a non-production system before using them on your production PCEs.

The CLI Tool is named `ilo`. It is a wrapper around the Illumio Core REST API. No knowledge of the REST API is required.

The `ilo` Command

This section describes the general syntax of the CLI Tool command, `ilo`, and tells how to use command-line help to get more specific syntax information.

Formal Syntax

The formal syntax for the `ilo` command is as follows:

```
ilo resource_or_specialCommand argument options
```

Where:

- `resource_or_specialCommand` represents either a resource managed by the PCE or a command that is not related to a particular resource.

  A resource is an object that the PCE manages, such as a workload, label, or pairing profile.

  Example resource command on Linux (create a workload):

  ```
  ilo workload create --name FriendlyWorkloadName --hostname
  myWorkload.BigCo.com
  ```

  A special command is a command that is not related to a specific resource. Special commands include `user`, `login`, `use_api_key`, and `node_available`.

  Example special command on Windows (log out of PCE):

  ```
  ilo user logout --id 6
  ```
- The `argument` represents an operation on the resource or special command.
- The `options` are allowed options for the `resource_or_specialCommand`. The specific option depends on the type of resource or special command.

CLI Tool Help

To get a complete list of all the available CLI Tool commands, use the `ilo` command without options. This command displays the high-level syntax of special commands, resources, and their allowable options.

For details about a resource's or special command's arguments, specify the name of the resource followed by the argument followed by the `--help` option. For example:

```
ilo workload create --help
```

HTTP Response Codes and Error Messages

This section describes the response codes and error messages that can be returned when you use CLI Tool commands.

REST API HTTP Response Codes

At the end of its output, the `ilo` command displays the REST API HTTP response code from the command. For example, a successful operation shows the following output:

```
...
200, OK
```

Error Messages

For many syntactical or other types of errors, the CLI Tool displays a general message encouraging you to verify your syntax with the CLI Tool help:

```
The ilo command has encountered an error. Check your syntax with either of the
following commands:

- ilo
- ilo <command> --help
```

In addition, in some circumstances, the CLI Tool writes a detailed log of errors:

```
For detailed error messages, see the file:
location-of-local-temp-directory/illumio-cli-error.log
```

Where `location-of-local-temp-directory` is as follows:

- Linux: `/tmp`
- Windows: `C:\Windows\Temp`

Environment Variables

Illumio provides Linux environment variables to allow users to customize operation of the CLI tool.

Environment Variable	Purpose
ILO_API_KEY_ID	API key for non-password based authentication and cookie less session with PCE. See "Authenticate with an API Key".
ILO_API_KEY_SECRET	API key secret for non-password based authentication and cookie less session with PCE. See "Authenticate with an API Key".
ILO_API_VERSION	API version to be used to execute CLI commands. Set this if you want to override the default API version. See "Set the Illumio ASP REST API Version." Default: v2. Example: $ export ILO_API_VERSION=v1
ILO_CA_DIR	Directory that contains certificates. See "TLS/SSL Certificate for Access to the PCE".
ILO_CA_FILE	Absolute path to certificate file. See "TLS/SSL Certificate for Access to the PCE".
ILO_DISPLAY_CONFIG	Absolute path to the display configuration file that is to be used with the list command. See "Linux Save Specific Fields to File For Reuse".
ILO_INSECURE_PASSWORD	Provide a password for login. If this variable is set, the login password prompt does not appear, and this password is used instead. Do not use in a production system when authentication security is desired. Example: `$ export ILO_INSECURE_PASSWORD=myInsecurePassword`
ILO_KERBEROS_SPN	Kerberos service principal name (SPN). Specify this variable when using Kerberos authentication.
ILO_LOGIN_SERVER	PCE login server FQDN. Use this variable when the login server FQDN is not the same as the PCE FQDN. See "Explicit Log into the PCE".
ILO_ORG_ID	Organization identifier for certificate-authenticated session with PCE. Value is always 1. Does not need to be explicitly set The environment variable is set by the system and should not be explicitly set. See "Authentication to PCE with API Key or Explicit Login".
ILO_PCE_VERSION	PCE version for the CLI to use. Default: 19.1.0 Example: `$ export ILO_PCE_VERSION=18.2.5`
ILO_PREVIEW	Enable any preview features that are included in this release. To disable preview features, remove this variable from the environment.
ILO_SERVER	FQDN of PCE for login and authentication with PCE. See "Authentication to PCE with API Key or Explicit Login".
TSC_ACCESS_KEY TSC_SECRET_KEY	These two ENV variables have been added in the release 1.4.2 to set up the Tenable SC API keys, which are used for authentication.
TSC_HOST	Variable that specifies the target host for Tenable
QAP_HOST	Variable that specifies the target host for Qualys

Installation and Authentication

This section describes how to install the CLI Tool. It also describes how to set up authentication, how to upgrade the tool, and how to uninstall it.

Installation Prerequisites

This section details prerequisites and the installation of the CLI Tool. Be sure you meet the prerequisites in the checklist.

Prerequisite Checklist

- ☐ License for vulnerability data upload
- ☐ Vulnerability data for upload
- ☐ Functional PCE
- ☐ Supported operating systems
- ☐ TLS/SSL certificate for authenticating to the PCE
- ☐ API version set in configuration
- ☐ The CLI Tool installation program

License for Vulnerability Data

The Illumio Core Vulnerability Maps license is required to import vulnerability data into the Illumio PCE. For information about obtaining a license, contact Illumio Customer Support. For information on activating the license, see Add the License for Vulnerability Data Upload [347].

Upload Vulnerability Data

When you plan on using the CLI Tool to upload vulnerability data, make sure you have the data to upload in advance. See Supported Vulnerability Data Sources [350] for information.

Install Functional PCE

Because the CLI Tool is for managing resources on your PCE, you need to have already installed a fully functional PCE.

Supported Computer Operating Systems

The CLI Tool is supported on the following operating systems.

Linux

- Ubuntu 18.04
- Ubuntu 20.04
- Centos/RHEL 7.9
- Centos/RHEL 8.4

Microsoft Windows

NOTE

The CLI Tool is not supported on Windows 32-bit CPU architecture. Ensure that you run it on Windows 64-bit CPU architecture.

- Windows 2012 64 bit
- Windows 2016 64 bit
- Windows 10 64 bit

NOTE

CLI 1.4.2 is no longer supported on Windows 2008 R2 (EOL). The CLI Tool should work and can be used at your own risk.

TLS/SSL Certificate for Access to the PCE

You need a TLS/SSL certificate to securely connect to the PCE. Requirements for this certificate are provided in the PCE Installation and Upgrade Guide.

Alternative Trusted Certificate Store

To secure the connection to the PCE, by default, the CLI Tool relies on your computer's trusted certificate store to verify the PCE's TLS certificate. You can specify a different trusted store. When you have installed a self-signed certificate on the PCE, the alternative trusted store might be necessary.

Example: Set envar for alternative trusted certificate store z

```
export ILO_CA_FILE=~/self-signed-cert.pem
```

Set the Illumio Core REST API Version

The CLI Tool uses v2 of the Illumio Core REST API by default.

Install, Upgrade, and Uninstall the CLI Tool

This section explains how to install, upgrade, or uninstall the CLI Tool on Linux or Windows.

Download the Installation Package

Download the CLI Tool installation package from the Tools Catalog page (login required) to a convenient location on your local computer.

Install Linux CLI Tool

The CLI Tool installer for Linux is delivered as an RPM for RedHat/CentOS and DEB for Debian/Ubuntu.

The CLI Tool is installed in the local binaries directory `/usr/local/bin`.

Log into your local Linux computer as a normal user and then use `sudo` to run one of the following commands.

RedHat/CentOS:

```
$ sudo rpm -ivh /path_to/nameOfCliRpmFile.rpm
```

Debian/Ubuntu:

```
$ sudo dpkg -i / path_to / nameOfCliDebFile .deb
```

Upgrade Linux CLI Tool

Log into your local Linux computer as a normal user and then use `sudo` to run one of the following commands.

RedHat/CentOS:

```
$ sudo rpm -Uvh /path_to/nameOfCliRpmFile.rpm
```

Debian/Ubuntu:

```
$ sudo dpkg -i / path_to / nameOfCliDebFile .deb
```

The same option, `-i`, is used for installation or upgrade.

Uninstall Linux CLI Tool

Log into your local Linux computer as a normal user and then use `sudo` to run one of the following commands.

RedHat/CentOS:

```
$ sudo rpm -e nameOfCliRpmFile
```

Debian/Ubuntu:

```
$ sudo dpkg -r nameOfCliDebFile
```

Install Windows CLI Tool

The CLI Tool installer for Windows is delivered as an .exe file.

Log into your local Windows computer as administrator and start the installation program in any of the following ways.

- In the Windows GUI, double-click the .exe file.
- In a cmd window, run the .exe.
- In a PowerShell window, run the .exe.

After starting the installation program, follow the leading prompts.

A successful installation ends with the "Installation Successfully Completed" message and the help text for the CLI Tool is displayed.

Upgrade Windows CLI Tool

The CLI Tool cannot be directly upgraded from an existing CLI Tool installation.

If you have already installed a previous version of the CLI Tool, manually uninstall it with the Windows Control Panel's Add/Remove Programs.

After uninstalling the previous version of the CLI Tool, install the new version of the CLI Tool as described in Install Windows CLI Tool [336].

Uninstall Windows CLI Tool

Log into your local Windows computer as an administrator, and from the Windows Control Panel, launch Add/Remove Programs.

Select Illumio CLI from the list and click the **Uninstall** button.

Authenticate with the PCE

When using the CLI Tool, you can authenticate to your PCE in the following ways:

- **With an API key and key secret:**
  This is the easiest way. Before you create the API key and secret, you need to log in to authenticate to the PCE. After creating and using the key, you do not have to specify your username and password again.
- **With the explicit command to log in:**
  This always requires a username and password.
  This method also requires you to log out with a user ID displayed at login. The explicit login times out after ten minutes of inactivity, after which you must log in again.

For both authentication mechanisms, on the command line, you always need to specify the FQDN and port of your PCE. The default port for the PCE is 8443. However, your system administrator can change this default. Check with your system administrator to verify the port you need.

Authenticate with an API Key

To authenticate to the PCE with an API key, you must first explicitly log into the PCE, create the API key, and then use the key to authenticate.

1. Authenticate via explicit login:

   ```
   ilo login --server yourPCEfqdn:itsPort
   ```
2. Create the API key:

   ```
   ilo api_key create --name someLabel
   ```

   `someLabel` is an identifier for the key.
3. Use the API key to authenticate:

   ```
   ilo use_api_key --server yourOwnPCEandPort --key-id yourOwnKeyId --org-id --key-secre
   ```

Create an API Key

On Linux, for later ease of use, with the `api_key --create-env-output` option, you can store the API key, API secret, and the PCE server name and port as environment variables in a file that you source in future Linux sessions.

Linux Example

This example creates the API key and secret and stores them as environment variables in a file named `ilo_key_MY_SESSION_KEY`.

```
# ilo api_key create --name MY_SESSION_KEY --create-env-output
# Created file ilo_key_MY_SESSION_KEY with the following contents:

export ILO_API_KEY_ID=14ea453b6f8b4d509
export ILO_API_KEY_SECRET=e1fa1262461ca2859fcf9d91a0546478d10a1bcc4c579d888
a4e1cace71f9787
export ILO_SERVER=myPCE.BigCo.com:8443
export ILO_ORG_ID=1

# To export these variables:
# $ source ilo_key_MY_SESSION_KEY
```

Log Into the PCE

Without an API key, you must explicitly log into the PCE.

For on-premises PCE deployments, the login syntax is the FQDN and port of the PCE:

`ilo login --server` *`yourPCEfqdn:itsPort`*

For `yourPCEfqdn:itsPort`, do not specify a URL instead of the PCE's FQDN and port. If you do, an error message is displayed.

For the Illumio Secure Cloud customers, the login syntax is:

`ilo login --server` *`URL_or_bare_PCEfqdn:itsPort`* `--login-server login.illum.io:443`

See the explanation above about the argument to the `--server` option.

- After login, the output of the command shows a user ID value. Make a note of this value. You need it when you log out.
- The session with the PCE remains in effect as long as you keep using the CLI Tool. After 10 minutes of inactivity, the session times out, and you must log in again.

Example

In this example, the user ID is 6.

```
C:\Users\marie.curie> ilo login --server myPCE.BigCo.com:8443
Enter User Name: albert.einstein@BigCo.com
Enter Password: Welcome Albert!
User ID = 6
Last Login Time 2018-08-10T-09:58:07.000Z from someIPaddress
Access to Orgs:
Albert: (2)
Roles: [3]
Capabilities: {"basic"=>["read", "write"], "org_user_roles"=>["read", "write"]}
```

```
User Time Zone: America/Los_Angeles
Server Time: 2018-08-12T17:58:07.522Z
Product Version: 16.09.0-1635
Internal Version: 48.0.0-255d6983962db54dc7ca627534b9f24b94429bd5
Fri Aug 6 16:11:50 2018 -0800
Done
```

Log Out of the PCE

To end a session with the PCE, use the following command:

```
ilo user logout --id valueOfUserIdFromLogin
```

Where:

- `valueOfUserIdFromLogin` is the user ID from your login. See Log Into the PCE [338] for information.

Example

In this example, the user ID is 6.

```
ilo user logout --id 6
```

CLI Tool Commands for Resources

This section describes how to use the CLI Tool with various PCE resources.

View Workload Rules

You can view a specific workload's rules with the following command:

```
ilo workload rule_view --workload-id UUID
```

Where:

- UUID is the workload's UUID. See About the Workload UUID [343] for information.

In the example below, the workload's UUID is as follows:

```
2ca0715a-b7e3-40e3-ade0-79f2c7adced0
```

Example View Workload Rules

```
ilo workload rule_view --workload-id 2ca0715a-b7e3-40e3-ade0-79f2c7adced0
+------------+-------+
| Attribute  | Value |
+------------+-------+
|  providing |  []   |
+------------+-------+
```

```
Using
+---------------------+-------------------------------------------------------------------
| Ports And Protocols | Rulesets
+---------------------+-------------------------------------------------------------------
| [[-1, -1, nil]]     | [{"href"=>"/api/v2/orgs/28/sec_policy/8/rule_sets/1909", "name"=
+---------------------+-------------------------------------------------------------------
200, OK
```

View Report of Workload Services or Processes

The following command lists all running services or processes on a workload:

```
ilo workload service_reports_latest --workload-id UUID
```

Where:

- UUID is the workload's UUID. See About the Workload UUID [343].

In the example, the workload's UUID is as follows:

```
2ca0715a-b7e3-40e3-ade0-79f2c7adced0
```

Example Workload Service Report

```
ilo workload service_reports_latest --workload-id 2ca0715a-b7e3-40e3-ade0-79f2c7adced0
+-----------------+---------------------------+
| Attribute       | Value                     |
+-----------------+---------------------------+
|  uptime_seconds |  1491                     |
|  created_at     |  2015-10-20T15:13:00.681Z |
+-----------------+---------------------------+
Open Service Ports
+----------+---------+-------+--------------+-----------------+---------+---------------
| Protocol | Address | Port  | Process Name | User            | Package | Win Service Na
+----------+---------+-------+--------------+-----------------+---------+---------------
| udp      | 0.0.0.0 | 5355  | svchost.exe  | NETWORK SERVICE |         | Dnscache
...
| tcp      | 0.0.0.0 | 135   | svchost.exe  | NETWORK SERVICE |         | RpcSs
+----------+---------+-------+--------------+-----------------+---------+---------------
200, OK
```

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
```

Use the `list` Option for Resources

Many resources take the `list` option. This section details some of its uses.

Default List of All Fields

The default `list` command displays all fields associated with the resource:

```
ilo resource list
```

List Only Specific Fields

With the `--field` option, specify the fields to display:

```
ilo resource list --field CSV_list_of_fieldnames
```

For example, to display a list of labels with only the href, key, and value fields, use the `--field` option with those fields as comma-separated arguments.

Example List with Selected Fields

```
ilo label list --fields href,key,value
+---------------------+------+-----------------+
| Href                | Key  | Value           |
+---------------------+------+-----------------+
| /api/v2/2/labels/1  | role | Web             |
| /api/v2/2/labels/2  | role | Database        |
...
| /api/v2/2/labels/48 | loc  | Asia            |
+---------------------+------+-----------------+
```

Nested Resource Fields and Wildcards

Some resources have hierarchical, nested fields. For example, the workload resource includes the following hierarchy for the agent field:

```
agent/config/log_traffic
```

- A field named agent
  - That has a field named `config`
    - That has a field named `log_traffic`

To list nested fields, separate the hierarchy of the field names with a slash to the depth of the desired field.

To see all nested fields of one of a resource's fields, use the asterisk (*) wildcard.

Examples

The following example displays all fields under the agent/config field.

Example of All Nested Fields with Wildcard (*)

```
ilo workload list --field agent/config/*
+-------------+------------------+-------------+
| Log Traffic | Visibility Level | Mode        |
+-------------+------------------+-------------+
| false       | flow_summary     | illuminated |
| false       | flow_summary     | idle        |
+-------------+------------------+-------------+
```

You can combine individual field names, nested field names, and the * wildcard.

Example Combination of Individual fields, Nested fields, and Wildcard

```
ilo workload list --fields href,hostname,agent/config/*,agent/status/uid,agent/status/st
+-----------------------------------------------------------+----------------------------
| Href                                                      | Hostname
+-----------------------------------------------------------+----------------------------
| /api/v2/1/workloads/527b8aca-97aa-43b9-82e1-29b17a947cdd | hrm-web.webscaleone.info
| /api/v2/1/workloads/4a8743a4-14ee-40d0-9ed2-990fe3f0ffb1 | hrm-db.webscaleone.info
+-----------------------------------------------------------+----------------------------
...
```

Linux: Save Fields for Reuse

On Linux, for ease of reuse of specific fields, create a display configuration file in YAML format and set the environment variable `ILO_DISPLAY_CONFIG` to point to that file. Thereafter, you no longer need to specify specific fields on the list command line.

Examples

Configure the workloads list command to display only the href, hostname, all agent configuration fields, and agent version:

Example Command to Save to List Configuration File

```
ilo workload list --fields href,hostname,agent/config/*,agent/status/agent_version
```

Add the field names to a display configuration file in the following YAML format:

Example YAML Layout of Display Configuration File

```
workload:
  fields:
    - href
    - hostname
```

```
agent:
  config:
    fields:
      - '*'
  status:
    fields:
      - agent_version
```

Set the Linux environment variable `ILO_DISPLAY_CONFIG` to the path to the YAML file:

Example ILO_DISPLAY_CONFIG environment variable

```
$ export ILO_DISPLAY_CONFIG=~/ilo_display/display_config.yaml
```

List of All Workloads

To view all details for all workloads, use the following command:

```
ilo workload list
```

About the Workload UUID

To view an individual workload, you need the workload's identifier, called the UUID, or Universal Unique Identifier.

The UUID is shown in the list of all workloads described in List of All Workloads [343]. The UUID is the last word of the value of the workload's href field, as shown in bold in the following example:

```
/api/v2/orgs/28/workloads/2ca0715a-b7e3-40e3-ade0-79f2c7adced0
```

View Individual Workload

To see the details about an individual workload, use the following command:

```
ilo workload read -workload-id UUID
```

Where:

- UUID is the workload's UUID. See About the Workload UUID [343] for information.

The details of an individual workload are grouped under major headings:

- Workload > Interfaces
- Workload > Labels
- Workload > Services
- Services > Open Service Ports
- Agent > Status

Example List of Individual Workload

```
ilo workload read --workload-id 2ca0715a-b7e3-40e3-ade0-79f2c7adced0
+-------------------------+-------------------------------------------------------------
```

```
| Attribute                | Value
+--------------------------+-------------------------------------------------------------
|  href                    |  /orgs/1/workloads/2ca0715a-b7e3-40e3-ade0-79f2c7adced0
|  deleted                 |  false
...
Workload -> Interfaces
+------+------------------+------------+-------------------------+------------+---------
| Name | Address          | Cidr Block | Default Gateway Address | Link State | Network
+------+------------------+------------+-------------------------+------------+---------
| eth0 | 10.0.0.16        | 8          | 10.0.0.1                | up         | 1
...
Workload -> Labels
+--------------------+
| Href               |
+--------------------+
| /orgs/1/labels/37 |
...
Workload -> Services
+-----------------+---------------------------+
| Attribute       | Value                     |
+-----------------+---------------------------+
|  uptime_seconds |  69016553                 |
...
Services -> Open Service Ports
+----------+---------+------+--------------+------+---------+------------------+
| Protocol | Address | Port | Process Name | User | Package | Win Service Name |
+----------+---------+------+--------------+------+---------+------------------+
| 17       | 0.0.0.0 | 123  | ntpd         | root |         |                  |
...
Workload -> Agent
+-----------+-------------------------------------------------------------------------------
| Attribute | Value
+-----------+-------------------------------------------------------------------------------
|  config   |  {"log_traffic"=>true, "visibility_level"=>"flow_summary", "mode"=>"enforc
|  href     |  /orgs/1/agents/16
...
Agent -> Status
+---------------------------+--------------------------------------+
| Attribute                 | Value                                |
+---------------------------+--------------------------------------+
|  uid                      |  db482b06-41c6-4297-a60c-396de13576ad |
|  last_heartbeat_on        |  2016-12-07T04:07:03.756Z            |
...
200, OK
```

List Draft or Active Version of Rulesets

A security policy item consists of ruleset, IP lists, label groups, services, and security settings. Before changes to these items take effect, the policy must be provisioned on the managed workload by setting its state to active with the CLI Tool or provisioning it with the PCE web console.

To view a ruleset and provisioning state use the following command:

```
ilo rule_set list --pversion state
```

Where `state` is one of the following values:

- Draft: Any policy item that has not yet been provisioned.
- Active: All policy items that have been provisioned and are enabled on workloads.

The provisioning states are listed in the Enabled column:

- True: The policy is provisioned.
- Empty: The policy is a draft.

Example Draft Versions of Rulesets

```
ilo rule_set list --pversion draft
+-------------------------------------------------+------------------------------+------
| Href                                            | Created By                   | Name
+-------------------------------------------------+------------------------------+------
| /api/v2/orgs/28/sec_policy/draft/rule_sets/2387 | {"href"=>"/api/v2/users/74"} | foo1
| /api/v2/orgs/28/sec_policy/draft/rule_sets/1909 | {"href"=>"/api/v2/users/0"}  | Defau
200, OK
```

The state of the policy is stored in the agent/status/status field. See Nested Resource Fields and Wildcards [341] for information.

Import and Export Security Policy

Using the CLI Tool, you can export and import security policy to and from the PCE. Importing and exporting security policy is particularly useful for moving policy from one PCE to another so you can avoid recreating policy from scratch on the target PCE. For example:

- You can test policy on a staging PCE and then move it to your production PCE.
- You can move policy from a proof-of-concept PCE deployment to your production PCE.

Export and Import Policy Objects

You can use the CLI Tool to export or import the following objects in the PCE:

- Labels: `labels`
- Label groups: `label_groups`
- Pairing profiles: `pairing_profiles`
- IP lists: `ip_lists`
- Services: `services`
- Rulesets and rules: `rule_sets`

About Exporting Rules

You can export rules for workloads, virtual services, or virtual servers.

For flexibility, Illumio recommends that you base your security policy rules on labels. Do not tie the rules to specific individual workloads, virtual services, or virtual servers.

Virtual servers and virtual services are not exported.

The CLI Tool policy export does not include such references. When you have rules that are tied to individual workloads, virtual services, or virtual servers, a warning is displayed on export. Attempts to import such rules fail and display the reason for the failure.

Example Failed Attempt to Export Rules for Workload

```
WARNING: rule /orgs/1/sec_policy/active/rule_sets/3/sec_rules/39 contains non-transferra
Unable to proceed, please verify input
```

Workflow for Security Policy Export/Import

- Authenticate to the source PCE. See Authenticate with the PCE [337] for information.
- Export the policy to a file. Syntax summary:

  `ilo sec_policy export --file` *`someExportFilename`*
- Authenticate to the target PCE. See Authenticate with the PCE [337] for information.
- Import the saved policy. Syntax summary:

  `ilo sec_policy import --file` *`someImportFilename`*

Output Options, Format, and Contents

All exported policy is written to standard output. To write to a file, use the `--file` option.

Exported policy is in JSON format.

By default, all supported policy objects are exported. You can export a subset of policy by specifying one or more resource types with the `–resource` option (`labels`, `label_groups`, `pairing_profiles`, `ip_lists`, `services`, or `rule_sets`).

When a subset of policy items is exported (such as only labels), all referenced resources are also exported.

See also About Exporting Rules [345] for information.

Exported Rulesets

With the `-- rule_set` option, you can export multiple rulesets.

By default, only the most recently provisioned, active policy is exported. To export the current draft policy or a previous policy, use the `--pversion` state option. See List Draft or Active Version of Rulesets [344]for information.

For a single ruleset, make sure the `--pversion` state you specify matches the provisioned state of the ruleset. In the following example, the state is draft:

```
ilo sec_policy export --pversion draft --rule_set /orgs/1/sec_policy/draft/rule_sets/1
```

Effects of Policy Import

All imported policy is read from standard input, unless you import from a file with the `--file` option.

You can import policy file multiple times. Each import affects only a single copy of a resource.

All imported policy is set to the draft provisioned state. After the import, you must explicitly provision the active state.

Non-transferrable policy rules (that is, rules tied to specific workloads, virtual servers,and bound services), the import aborts with a warning. See About Exporting Rules [345] for information.

Policy items already on the target PCE are updated by imported resources whose names match the already existing resources' names. Services do not have to have the same names. Services match if they have the same set of ports and protocols.

Resources are not deleted by an import. For example, if you export policy from PCE-1 to PCE-2, delete a resource “R” from PCE 1, and then export and import again, resource “R” is still present on PCE 2. You must explicitly delete resource “R” from PCE2.

Upload Vulnerability Data

This section describes how to use the `ilo` commands to upload vulnerability data to the PCE for analysis in Illumination.

After uploading the data, you can use Vulnerability Maps in the PCE web console to gain insights into the exposure of vulnerabilities and attack paths across your applications running in data centers and clouds. See the "Vulnerability Maps" topic in the Visualization Guide for information.

Add the License for Vulnerability Data Upload

An Illumio Core Vulnerability Maps license is required to upload vulnerability data into the Illumio PCE. For information about obtaining the license, contact Illumio Customer Support.

You are provided with a license file named `license.json`. After you have obtained your license key, store it in a secure location.

NOTE

Before adding the license, you must first authenticate to the PCE. See Authenticate with the PCE [337] for information.

To add the license, you must be the organization owner or a be a user who has owner privileges.

Use the following command to inform the PCE of your valid license:

```
ilo license create --license-file "path_to_license_file/license.json" --feature "feature
```

Where:

What	Required?	Description
`"path_to_license_file/license.json"`	Yes	The quoted path to the `license.json` file from Illumio Example: `"~/secretDir/license.json"`
`"feature_name"`	Yes	The quoted string `"vulnerability_maps"`, which specifies the feature name the license enables
`debug`	No	Enable debugging
`v | verbose`	No	For verbose logging
`trace`	No	Enable API trace

Vulnerability Data Upload Process

On upload, the CLI Tool associates a workload's IP addresses with corresponding vulnerabilities identified for that workload.

Using API to Download Vulnerability Data

In release CLI 1.3, Tenable IO and tenable SC have been supporting both manual and API download of vulnerability data while Qualys tool was only available for manual download.

IMPORTANT

Starting from the release CLI 1.4, Qualys supports also API download, with some minor differences in options.

For the release CLI 1.4.1, it is suggested that users use an API key instead of a login session while using Qualys API download.

For the release CLI 1.4.2 for Tenable, the most reliable way to provide authentication is through API keys instead of username/password. If customers observe any authentication issues while using Tenable SC API upload, they are advised to use API keys for authentication.

There are 2 ENV variables to set up the Tenable SC API keys which are used for authentication:

`TSC_ACCESS_KEY`

TSC_SECRET_KEY

The API connects directly to the cloud instance of Tenable or Qualys and the vulnerability tool then scans new vulnerabilities and downloads them into the PCE.

Users can also set up cron jobs that run in the desired intervals and check the state of the vulnerability scanner.

Qualys and Tenable scanners work in a similar way, using the username and password and similar options.

Automating Vulnerability Imports from Tenable-SC

Users of Illumio vulnerability maps can automate the import of vulnerabilities from tenable-sc using a script.

Illumio CLI supports the API username and password as environment variables or a cmd line switch (such as `--api-password`).

The ILO-CLI tool was updated to add a switch for `--api-user`.

Kinds of Vulnerability Data Uploads

There are two kinds of upload: non-authoritative and authoritative.

- **Non-authoritative:** This is the default. A non-authoritative upload:
  - Appends incoming data to any previously loaded records
  - Accumulates records for the same workloads without regard to duplicates.

  You can repeat the non-authoritative upload as many times as you like until you are satisfied with the results.

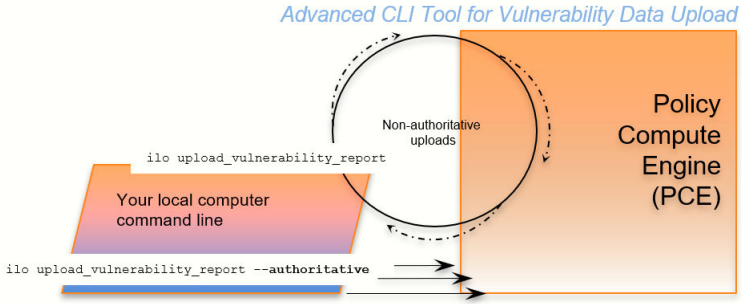

- **Authoritative:** You indicate authoritative data with the -authoritative option. An authoritative upload:
  - Overwrites any previously uploaded records for workloads matched to the incoming records.
  - Eliminates duplicate records.
  - Adds new records not previously written by other uploads.

  You can repeat the authoritative upload as many times as you like until you are satisfied with the results.

After either kind of upload, you can examine the uploaded data with the CLI Tool or the PCE web console. See “Vulnerability Maps” in the Visualization Guide for information.

Supported Vulnerability Data Sources

The CLI Tool works with vulnerability data from the following sources.

- Nessus Professional™
- Qualys®
- Tenable Security Center
- Tenable.io
- Rapid7©

NOTE

Before uploading Rapid7 data to the PCE, export the data from Rapid7 to Qualys format with Qualys XML Export.

Vulnerability Data Formats

In the CLI 1.4.0, 1.4.1 and 1.4.2 releases, Illumio supports the following report formats:

- For tenable-io: API, CSV
- For tenable-sc: API, CSV
- For nessus-pro: XML
- For qualys: API, XML

Common Vulnerabilities and Exposures (CVE)

Vulnerabilities are defined by Common Vulnerabilities and Exposures (CVE), with identifiers and descriptive names from the U.S. Department of Homeland Security National Cybersecurity Center.

Vulnerability Scores

Illumio computes a vulnerability score, which is a measure of the vulnerability of your entire organization. The score is displayed by the ilo vulnerability list command for all vulnerabilities or individual vulnerability via the vulnerability identifier.

Vulnerability Identifier

A uploaded vulnerability has an identifier as shown in the example below. The vulnerability identifier is tied to a specific CVE. You use this identifier with `--reference-id` option to examine specific uploaded vulnerabilities. See Example – List Single Uploaded Vulnerability [355] for information.

The following are examples of vulnerability identifiers.

- Nessus Professional: nessus-65432
- Qualys: qualys-23456
- Rapid7: qualys-98765. Because Rapid7 data is first exported from Rapid7 in Qualys format, it is given a Qualys identifier when uploaded to the PCE.

Vulnerabilities for Unmanaged Workloads

You can upload vulnerabilities for unmanaged workloads. However, unmanaged workloads do not have any vulnerability score or associated CVE. If the unmanaged workload is later changed to managed, this information becomes available.

Prerequisites for Vulnerability Data Upload

Before uploading vulnerability data, ensure that you are ready with the following requirements.

- An Illumio Vulnerability Maps license is required to upload vulnerability data to the PCE. See Add the License for Vulnerability Data Upload [347] for information.
- XML-formatted vulnerability data files from one of the supported sources.
- Authenticated CLI-tool access to the target PCE. See Authenticate with the PCE [337] for information.
- Authenticated access and necessary permissions in the PCE web console for working with vulnerability maps. See Authenticate with the PCE [337] for information.

Vulnerability Data Upload CLI Tool Syntax

The key argument and option for uploading vulnerability data are as follows. For readability, this syntax is broken across several lines.

```
ilo upload_vulnerability_report
--input-file path_to_datafile.xml [path_to_datafile.xml]...
--source-scanner [nessus-pro|qualys|tenable-sc|tenable-io]
--format xml
[--authoritative]
[ --api-user ApiServerUserName --api-server SourceApiServer:port ]
```

Where:

What	Required	Description
`--input-file path_to_datafile.xml [path_to_data-file.xml]...`	Yes	Location of one or more data files to upload. The path to the data file can be either an absolute path or a relative path. If more than one data file is listed (bulk upload), separate the file names with space characters.
`--debug`	No	Enable debugging
`--authoritative`	No	For uploading authoritative vulnerability data. The default command is without the `--authoritative` option. See Kinds of Vulnerability Data Uploads [349] for information.
`--workload-cache FILE`	No	DEBUGGING ONLY: Workload Cache file - use this if available
`--source-scanner [nessus-pro| qualys| tenable-sc]`	Yes	Indicates the source of the scan. Note for rapid data: • Vulnerability data from Rapid must have been exported from Rapid in Qualys XML format. • To load the Rapid data, use the 'qualys' argument

`--format`
 	
`REPORT_FORMAT`
	Yes	Report format. Allowed values are:
		
		`xml`
		
		• `--source-scanner nessus-pro`
		• `--source-scanner qualys`
		
		`csv`
		
		• `--source-scanner tenable-sc`
		• `--source-scanner tenable-io`
		
		`api`
		
		• `--source-scanner tenable-sc`
		• `--source-scanner qualys`
		• `--source-scanner nessus-pro`
		
		See also `--api-server` and `--api-user`.

`--api-server SourceApiServer:port`
 	
`SERVER_FQDN`
	Yes for
	
	Tenable with `--format api`
		API server FQDN. Allowed formats are `HOST` or `HOST:PORT`

`--api-user ApiServerUserName`
 	
`USERNAME`
	Yes for source API server authentication	The user name for authenticating to the SourceApiServer.
		
		You are always prompted to enter your password.

`--api-page-size`
 	
`PAGE_SIZE`
	Yes for Qualys and Tenable	Appropriate page size if API supports pagination. The default page is 1000.

What	Required	Description
`--skip-cert-verification`	Yes for Qualys and Tenable	Disable certificate verification for API.
`--on-premise`	Yes only for Tenable io	Tenable IO deployment is on premise.
`--mitigated`	Yes only for Tenable sc	Tenable SC input is exported from the mitigated vulnerabilities analysis view.
`--scanned-after SCANNED_AFTER`	Yes for Qualys	Qualys users can select scan data to process after a certain date, in ISO 8601 format. When the optional `scanned-after` option is not provided, the system will pull all the historical vulnerability records from your Qualys account. If your account has historical records, it may take a very long time for the first time. With the `scanned-after` option, vulnerability data scanned after a certain date will be extracted and uploaded. It is recommended to include a certain scanned-after time if you use Qualys API upload option for the first time.
`--severities SEVERITIES`	No	Qualys API users can select vulnerabilities with defined severity levels to include in their report. Users can filter based on severity and avoid severity levels 1 and 2, which are often very informational and noisy. Example: `--only-include-severity=3,4,5` For Windows, be sure to include quotes around the severity levels: Example: `--only-include-severity="3,4,5"` NOTE: This option was added in Release 1.4.1
`-v, --verbose`	No	Verbose logging mode
`--trace`	No	Enable API trace mode

Using the ILO Command with Windows Systems

Windows systems take a maximum of four options with the ILO command for the vulnerability data upload. Users who choose to use more optional parameters need to set `api-server`, `username`, and `password` as the environmental variables to use other options in the command.

Work with Vulnerability Maps in Illumination

See "Vulnerability Maps" in the Visualization Guide for information.

Vulnerability Data Examples

Example – Upload Non-Authoritative Vulnerability Data

In this example, the `--source-scanner nessus-pro` option indicates that the data comes from Nessus Professional. On Windows, provide the absolute path to the data file. This Windows example is broken across several lines with the PowerShell line continuation character (`).

```
C:\Users\donald.knuth> ilo upload_vulnerability_report `
--input-file C:\Users\donald.knuth\Desktop\vuln_reports\nessus3.xml `
--source-scanner nessus-pro --format xml

Elapsed Time [0.05 (total : 0.05)] - Data parsing is done.
Elapsed Time [1.08 (total : 1.13)] - Got workloads. Workload count: 5.
Elapsed Time [0.0 (total : 1.13)] - Built workload interface mapping. Total interfaces :
Elapsed Time [4.57 (total : 5.7)] - Imported Vulnerabilities..
Elapsed Time [0.0 (total : 5.7)] - Detected Vulnerabilities are associated with vulnerab
Elapsed Time [0.83 (total : 6.53)] - Report Imported.

Summary:
Processed the report with the following details :
Report meta data =>
Name          : Generic
Report Type   : nessus
Authoritative : false
Scanned IPs   : ["10.1.0.74", "10.1.0.223", "10.1.0.232", "10.1.0.221", "10.1.0.11", "10

Stats :
   Number of vulnerabilities           => 19
   Number of detected vulnerabilities  => 31

Done.
```

Example – Upload of Rapid7 Vulnerability Data

The syntax for uploading vulnerability data from Rapid7 is identical to the syntax for uploading vulnerability data from Qualys. On Windows, you use the `--format qualys` option and the absolute path to the data file. This Windows example is broken across several lines with the PowerShell line continuation character (`).

Rapid7 data exported in Qualys format

Before uploading to the PCE, Rapid7 vulnerability data must have been exported in Qualys format from Rapid7 with Qualys XML Export.

```
C:\Users\edward.teller> ilo upload_vulnerability_report `
--input-file C:\Users\edward.teller\Desktop\vuln_reports\rapid7.xml `
--source-scanner qualys --format xml
...
Done.
```

Example – Upload Authoritative Vulnerability Data

In this example, the prompt shows this is an authoritative upload.

To proceed, you must enter the word YES in all capital letters.

```
C:\Users\jrobert.oppenheimer> ilo upload_vulnerability_report --input-file dataDir/autho

Using /home/centos/.rvm/gems/ruby-2.4.1
Authoritative scan overwites the previous entries for all the ips within this scan. Ther
Are you sure this is an authoritative scan? (YES | NO)
YES
```

```
Elapsed Time [11.86 (total : 11.86] - Data parsing is done.
Elapsed Time [0.27 (total : 12.13] - Got workloads. Workload count: 3.
Elapsed Time [0.0 (total : 12.13] - Built workload interface mapping. Total interfaces :
Elapsed Time [3.02 (total : 15.15] - Imported Vulnerabilities..
Elapsed Time [0.0 (total : 15.15] - Detected Vulnerabilities are associated with vulnera
Elapsed Time [0.84 (total : 16.0] - Report Imported.
Summary:
Processed the report with the following stats -
    Number of vulnerabilities           => 14
    Number of detected vulnerabilities  => 48
Done.
```

Example – List Single Uploaded Vulnerability

This example uses a single Qualys vulnerability identifier to show the associated vulnerability. The value passed to the `--reference-id` option is shown as qualys-38173. See Vulnerability Identifier [350] for information.

```
$ ilo vulnerability read --xorg-id=1 --reference-id=qualys-38173
...

| Attribute | Value |
+-------------+-------------------------------------------------------------------+
| href | /orgs/1/vulnerabilities/qualys-38173 |
| name | SSL Certificate - Signature Verification Failed Vulnerability
| score | 39 |
| cve_ids | [] |
| created_at | 2018-11-05T18:16:56.846Z |
...
```

Example – List All Uploaded Vulnerabilities

This example highlights the vulnerability identifier, the CVE identifiers, and the description of the CVE. See Common Vulnerabilities and Exposures (CVE) [350] and Vulnerability Identifier [350] for information. The layout of the output is the same for all supported vulnerability data sources.

Nessus Professional

```
C:\Users\werner.heisenberg> ilo vulnerability list --xorg-id=1
...
| Href | Name | Score | Description | Cve Ids | Created At | Updated At | Created By | U
---------------------+-------------------------+---------------------+----------------
| /orgs/1/vulnerabilities/nessus-18405 | Microsoft Windows Remote Desktop Protocol Serve
...
```

Qualys

```
C:\Users\isaac.newton> ilo vulnerability list --xorg-id=1
...
| Href | Name | Score | Description | Cve Ids | Created At | Updated At | Created By | U
---------------------+-------------------------+---------------------+----------------
| /orgs/1/vulnerabilities/qualys-38657 | Birthday attacks against TLS ciphers with 64bit
...
```

Rapid7

Because Rapid7 vulnerability data must be in Qualys format before upload, the output is the same as for Qualys data, including the vulnerability identifier (qualys-38657 in the example above) and CVE. See Common Vulnerabilities and Exposures (CVE) [350] and Vulnerability Identifier [350] for information.

Example – View Vulnerability Report

The Report Type column identifies the source of the scan; in this example, Qualys.

```
C:\Users\gracemurry.hopper> ilo vulnerability_report list --xorg-id=1
...
| Href | Report Type | Name | Created At | Updated At | Num Vulnerabilities | Created By
+-----------------------------------------------------+-------------+------------------
| /orgs/1/vulnerability_reports/scan_1502310096_09344 | qualys  | NewAuthoritativeScan |
...
```

Example - Upload a Qualys Report Using API

```
upload_vulnerability_report --source-scanner  qualys --format api
--api-server qualysguard.qg3.apps.qualys.com  --api-user um3sg
--scanned-after 2021-09-20
```

CLI Tool Tutorials

This section provides several hands-on exercises that demonstrate step-by-step how to perform common tasks using the CLI Tool.

How to Import Traffic Flow Summaries

Static Illumination provides “moment-in-time” visibility of inter-workload traffic. This visibility is useful to model policies, to look for specious traffic flows, and to ensure that metadata for labels is accurate.

Goal

Load workload and traffic data needed for analysis with static Illumination.

Setup

This tutorial relies on the following data to import.

- 1,000 workloads defined in the file `bulkworkloads-1000.csv`, which has the following columns:

  ```
  hostname,ips,os_type
  10.14.59.8.netstat,10.14.59.8,linux
  10.4.78.178.netstat,10.4.78.178,linux
  10.37.134.179.netstat,10.37.134.179,linux
  ...
  ```
- 1,000,000 traffic flows defined in the CSV file `traffic.clean-1m.csv`, which has the following columns:

```
src_ip,dst_ip,dst_port,proto
10.40.113.86,10.14.59.8,10050,6
10.14.59.8,10.8.251.138,8080,6
10.40.113.124,10.14.59.8,22,6
...
```

Steps

The workflow is authenticate to the PCE and run two `ilo bulk_upload_csv` commands.

1. Authenticate to the PCE via API key or explicit login. See Authenticate with the PCE [337] for information.
2. Load the workload data:

   ```
   ilo workload bulk_upload_csv --file bulkworkloads-1000.csv
   ```
3. Load the traffic flow data:

   ```
   ilo traffic bulk_upload_csv --file traffic.clean-1m.csv
   ```

Results

The data from the CSV files are uploaded.

How to Create Kerberos-Authenticated Workloads

This tutorial describes how to create workloads that use Kerberos for authentication. The tutorial makes the following assumptions:

- This tutorial assumes that you already have your Kerberos implementation in place.
- As required by Kerberos, the Kerberos realm name is shown in all capital letters as `MYR-EALM`.
- VEN environment variables must be set *before* VEN installation. Environment variables for Linux are detailed in the VEN Installation and Upgrade Guide.

Goals

- Create two workloads on Linux that are authenticated by Kerberos.
- Set the workloads' modes to idle and illuminated.
- Run the kinit command to get Kerberos tickets for the workloads.

Setup

The key data for using the `ilo` command to create these workloads are the name of the Kerberos realm and the Service Principle Name (SPN).

Steps

The workflow is authenticate, run two `workload create` commands that set the workloads' modes, set the VEN environment variables, install the VEN, and run two Kerberos `kinit` commands to get Kerberos tickets for the workloads.

1. Authenticate to the PCE via API key or explicit login. See Authenticate with the PCE [337] for information.
2. Create Kerberos-authenticated `myWorkload1` and set its `mode` to `idle`:

```
ilo workload create --hostname myPCE.BigCo.com --name myWorkload1 --service-principal
```

For information about how the mode is a nested field, see Nested Resource Fields and Wildcards [341].

3. Create Kerberos-authenticated `myWorkload2` and set its `mode` to `illuminated`:

```
ilo workload create --hostname myPCE.BigCo.com --name myWorkload2 --service-principal
```

4. Before installation, set VEN environment variables:

```
# Activate on installation
VEN_INSTALL_ACTION=activate
# FQDN and port PCE to pair with
VEN_MANAGEMENT_SERVER=myPCE.BigCo.com:8443
# Kerberos Service Principal Name
VEN_KERBEROS_MANAGEMENT_SERVER_SPN=host/myKerberosTicketGrantingServer
# Path to Kerberos shared object library
VEN_KERBEROS_LIBRARY_PATH=/usr/lib/libgssapi_krb5.so
```

5. Install the Linux VEN:

```
rpm -ivh illumio-ven*.rpm
```

6. Run `kinit` to get a Kerberos ticket for `myWorkload1`:

```
kinit -k -t /etc/krb5.keytab host/myWorkload1.BigCo.com@MYREALM
```

7. Run `kinit` to get a Kerberos ticket for `myWorkload2`:

```
kinit -k -t /etc/krb5.keytab host/myWorkload2.BigCo.com@MYREALM
```

Results

The Kerberos-authenticated workloads are created, set in the desired modes, and given a Kerberos ticket.

How to Work with Large Datasets

The `--async` option is for working with large sets of data without having to wait for the results. The option works like “batch job.”

The option can be used with any resource. The workflow is as follows:

1. You issue the desired `ilo` command with the `--async` option, which displays a job ID.
2. You take note of the job ID.
3. Your session is freed up while the job runs.
4. The job creates a data file, which you then view with `datafile --read --job-id jobID`.

Goal

Get a report of a large workload data set.

Steps

1. Issue the `--async` request for a workload list. Take note of job ID which is the final word of the href displayed on the Location line.

```
[kurt.goedel~]$ ilo workload list --async
Using /home/kurt.goedel/.rvm/gems/ruby-2.2.1
```

```
Location: /orgs/1/jobs/fe8a1c2b-1674-4b83-8967-eb56c4ffa1e3
202, Accepted
```

2. Check to see if the job completed. Use the job ID from the `Location` output in previous command:

```
[sigmund.freud~]$ ilo job read --job-id fe8a1c2b-1674-4b83-8967-eb56c4ffa1e
Using /home/sigmund.freud/.rvm/gems/ruby-2.2.1
```

3. Download the resulting data file, specifying the job ID with `-uuid jobID`:

```
[bill.gates ~]$ ilo datafile read --uuid 1e1c1540-8a01-0136-ec14-02f4d6c1190c
Using /home/ bill.gates /.rvm/gems/ruby-2.2.1
+-----------------------------------------------------------+---------+------+--
... Many lines not shown
+----------------------------+---------------------+----------------------------+-
| Href                                                    | Deleted | Name | Descripti
| Service Principal Name | Public Ip             | Distinguished Name | External Data
| Interfaces | Ignored Interface Names | Service Provider | Data Center
| Data Center Zone | Os Id | Os Detail | Online | Labels  | Services | Agent
 | Created At                  | Created By           | Updated At                  |
+-----------------------------------------------------------+---------+------+----------
... More lines not shown
-----------------------------------------------------------+
| /orgs/1/workloads/50ce441e-75ac-4be8-9201-96169545019c | false   |      |
...
... Many lines not shown
...
```

How to Upload Vulnerability Data

This example tutorial shows how to upload vulnerability data to the PCE. For more information, see Upload Vulnerability Data [347]. The source of the vulnerability data in this example comes from Qualys®.

Goal

Upload authoritative vulnerability data for analysis in Illumination.

Steps

1. Do a non-authoritative upload of vulnerability data for examination:

```
ilo upload_vulnerability_report --input-file C:\Users\albert-einstein0.xml --source-s
```

2. Examine a single uploaded vulnerability record identified by its vulnerability identifier, `qualys-38173`. See Vulnerability Identifier [350] for information.

```
ilo vulnerability read --xorg-id=1 --reference-id=qualys-38173
```

3. Do another non-authoritative upload of vulnerability data.

```
ilo upload_vulnerability_report --input-file C:\Users\albert-einstein99.xml --source-
```

4. Do an authoritative upload of vulnerability data, overwriting any previously uploaded records and adding any new vulnerability records.

```
ilo upload_vulnerability_report --input-file C:\Users\albert.einstein_FINAL.xml --aut
```

Results

The authoritative vulnerability data has been uploaded and is ready for use in Illumination.

Legal Notice

Resources

- Legal information
- Trademarks statements
- Patent statements
- License statements

Contact Information

- Contact Illumio
- Contact Illumio Legal
- Contact Illumio Documentation