TECHNICAL
DOCUMENTATION

Illumio Administration Guide 26.1 (SaaS)

"Show Me How" Videos and Animations

Administration Tasks		
Build Traffic Filters	Configure Access Restrictions	Export Events Using the Web Console
Configure Trusted Proxy IPs	Create API Keys	Modify LDAP Configuration
Enable SAML Request Signing	Enable LDAP Authentication	Verify LDAP Connectivity
Enable VEN Tampering Protection	Generate a VEN Maintenance Token	
Disable Health Check Forwarding	Generate a Support Report from the PCE	
Set Firewall Coexistence	Generate PCE Support Bundle	
Secure LDAP	Edit Expiration API Keys	

Table of Contents

PCE Administration 9
- Overview of PCE Administration 9
  - Before You Begin 9
  - PCE Architecture and Components 9
  - PCE Control Interface and Commands 13
  - PCE Organization and Users 14
- Manage PCE Nodes and Clusters 15
  - Recommendations for Data Management 16
- About the PCE Databases 16
  - Policy and Traffic Data Databases 16
  - Data Retention of Traffic Flow Summaries 17
  - Determine the Primary Database 17
  - Show Database Replication Information 18
  - Rotate Database Passwords and Other Secrets 18
  - Anonymize Database Export 18
  - View Events Using PCE Command Line 20
  - Database Migration, Failover, and Restore 21
  - Manage Multi-Node Traffic Database 24
  - PCE Database Backup 28
  - PCE Default Object Limits 33
- Role-based Access Control 41
  - Overview of Role-based Access Control 41
  - Use Cases 42
  - Features of Role-based Access Control 43
  - About Roles, Scopes, and Granted Access 44
  - Prerequisites and Limitations 52
  - Authentication 52
  - LDAP Authentication 56
  - Password Policy Configuration 62
  - Updates to Roles 67
  - OneLogin Single Sign-on 69
  - Okta Single Sign-on 71
  - Active Directory Single Sign-on 73
  - Setup for Role-based Access Control 80
  - Roles, Scopes, and Granted Access 90
  - Configure Access Restrictions and Trusted Proxy IPs 98
- Role-based Access for Application Owners 102
  - Overview 102
  - Updates to Roles 103
  - Configuration 105
  - Facet Searches for Scoped Roles 106
  - Ruleset Viewer 106
  - Scoped Roles and Permissions 107
  - Scoped Users and PCE 111

Labeled Objects 115
Rulesets and Rules 116
App Group Map 117
Policy Generator and Explorer 117
My Roles 118
Labeled Objects 118
Policies 119
Ruleset Viewer 120
Scoped Roles and Permissions 120
Scoped Users and PCE 125
Azure AD Single Sign-on 129
Prerequisites 129
STEP 1: Obtain URLs from the Illumio PCE Web Console 129
STEP 2: Configure SSO settings in Azure AD 130
STEP 3: Obtain SAML certificate and URLs from Azure AD 133
STEP 4: Configure SAML SSO settings in the Illumio PCE 133
STEP 5: Create App Roles in Azure AD 134
STEP 6: Assign users and groups to app roles in Azure AD 135
STEP 7: Add External Groups and assign roles in the PCE Web Console 136
STEP 8: Turn on SAML authentication in the PCE Web Console 138
STEP 9: Test SSO 138
Prerequisites 139
STEP 1: Obtain URLs from the Illumio PCE Web Console 139
STEP 2: Configure SSO settings in Azure AD 140
STEP 3: Obtain SAML certificate and URLs from Azure AD 143
STEP 4: Configure SAML SSO settings in the Illumio PCE 143
STEP 5: Create App Roles in Azure AD 144
STEP 6: Assign users and groups to app roles in Azure AD 145
STEP 7: Add External Groups and assign roles in the PCE Web Console 146
STEP 8: Turn on SAML authentication in the PCE Web Console 148
STEP 9: Test SSO 148
Monitor and Diagnose PCE Health 149
PCE Health Monitoring Techniques 149
Monitor PCE Health 149
Support Reports for PCE 161
PCE High Availability and Disaster Recovery Concepts 164
Goals for PCE High Availability 164
PCE High Availability and Disaster Recovery Requirements 166
PCE Replication and Failover 168
Core Node Failure 178
Data Node Failure 180
PCE Failures and Recoveries 185

Site Failure (Split Clusters) ... 209
Connectivity Settings ... 216
Private Data Centers ... 216
Offline Timers ... 216
Set the IP Version for Workloads ... 220
Manage Security Settings ... 221
SecureConnect Setup ... 223
AdminConnect Setup ... 231
PCE Administration Troubleshooting Scenarios ... 234
Transaction ID Wraparound in PostgreSQL Database ... 234
Best Practices for Handling Scanner Traffic for the Illumio PCE ... 236
Identify Exclusions Using the Core Services Detector ... 236
Allow and Filter Scanner Traffic from the Environment ... 236
Build Traffic Filters ... 237
Reduce Traffic Stored in Flow Logs ... 237
Review Detection Rules on a Schedule ... 237
Events Administration ... 239
About this guide ... 239
Before reading this guide ... 239
Notational conventions in this guide ... 239
Events Framework ... 240
Events Lifecycle for Resources ... 242
Event Types, Syntax, and Record Format ... 243
Types of Events ... 243
Anonymized Database Dumps ... 243
REST API Events Schema ... 244
Event Syntax ... 244
Events Record Information ... 244
Event Record Structure ... 245
Events Displayed in PCE Web Console ... 246
Cross-Site Request Forgery Protection ... 246
Events Monitoring Best Practices ... 247
Examples of Events ... 251
List of Event Types ... 266
View and Export Events ... 278
Events Settings ... 280
Events Are Always Enabled ... 280
Event Settings in PCE Web Console ... 280
Configure Events Settings in PCE Web Console ... 283
Requirements for Events Framework ... 284
SIEM Integration for Events ... 286
Syslog Forwarding ... 287
Showing Rule ID in Syslog ... 290
Traffic Flow Types and Properties ... 291
Visibility Settings ... 291
Event Types ... 291

Showing the Data Transfer Amount 293
Traffic Flow Summary Examples 294
Export Traffic Flow Summaries 300
Manage Traffic Flows Using REST API 303
VEN Administration Guide 310
Overview of VEN Administration 310
About This Administration Guide 310
VEN Architecture and Components 311
About VEN Administration on Workloads 318
Firewall Coexistence 333
`illumio-ven-ctl` General Syntax 342
Useful VEN and OS Commands 343
VEN State 351
VEN Startup and Shutdown 351
Disable and Enable VENs (Windows only) 354
VEN Suspension 355
Deactivate and Unpair VENs 361
Deactivate Using the VEN Command Line 361
Unpair Using the VEN Command Line 362
Unpair Using System Commands 364
Effects of Unpairing VENs 365
Safeguard or Remove Inactive VENs 366
Safeguard VENs from accidental unpairing 366
Remove inactive VENs 367
Monitor and Diagnose VEN Status 369
VEN-to-PCE Communication 370
VEN Status Command and Options 378
VEN and Workload States 382
VEN Logging 385
Tuning the IPFilter State Table (AIX/Solaris) 387
Manage Conntrack Table Size (Linux) 390
VEN Firewall Tampering Detection 392
VEN Tampering Protection 396
VEN Support Reports 402
VEN Troubleshooting 404
Illumio Core PCE CLI Tool Guide 1.4.3 409
Overview of the CLI Tool 409
Before You Begin Using the CLI Tool 409
CLI Tool Versioning 409
CLI Tool and PCE Resource Management 410
The `ilo` Command 411
HTTP Response Codes and Error Messages 412
Environment Variables 412
Installation and Authentication 414
Prerequisite Checklist 414
CLI Tool Installation Prerequisites 414

Install, Upgrade, and Uninstall the CLI Tool 416
Authenticate with the PCE 418
CLI Tool Commands for Resources 421
View Workload Rules 421
View Report of Workload Services or Processes 422
Use the `list` Option for Resources 423
List Draft or Active Version of Rulesets 429
Support for Proxy 430
Import and Export Security Policy 433
Upload Vulnerability Data 436
CLI Tool Tutorials 452
How to Import Traffic Flow Summaries 452
How to Create Kerberos-Authenticated Workloads 453
How to Work with Large Datasets 454
How to Upload Vulnerability Data 456
Illumio Core PCE CLI Tool Guide 1.4.2 458
Overview of the CLI Tool 458
About This Guide 458
CLI Tool and PCE Resource Management 460
The `ilo` Command 461
HTTP Response Codes and Error Messages 462
Environment Variables 462
Installation and Authentication 464
Prerequisite Checklist 464
Installation Prerequisites 464
Install, Upgrade, and Uninstall the CLI Tool 466
Authenticate with the PCE 468
CLI Tool Commands for Resources 471
View Workload Rules 471
View Report of Workload Services or Processes 473
Use the `list` Option for Resources 473
List Draft or Active Version of Rulesets 479
Import and Export Security Policy 480
Upload Vulnerability Data 483
CLI Tool Tutorials 499
How to Import Traffic Flow Summaries 499
How to Create Kerberos-Authenticated Workloads 500
How to Work with Large Datasets 501
How to Upload Vulnerability Data 503
Error Messages 505
Containers Error Messages 505
Containers 505
NEN, VEN Error Messages 505
NEN 506
VEN 507
PCE Error Messages 507

PCE Error Messages and Recommended Actions 508
Legal Notice 510

PCE Administration

Abstract

This guide describes maintaining and operating the Policy Compute Engine (PCE). It also includes other tasks to manage your PCE deployment and help you with ongoing PCE operations and 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

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:

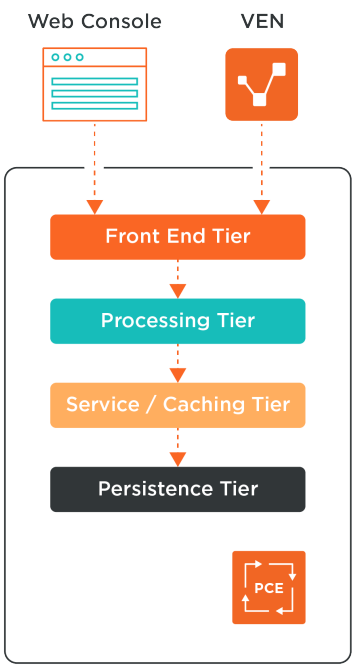

Each of these service tiers is responsible for various functions, as shown below and described further in the following table.

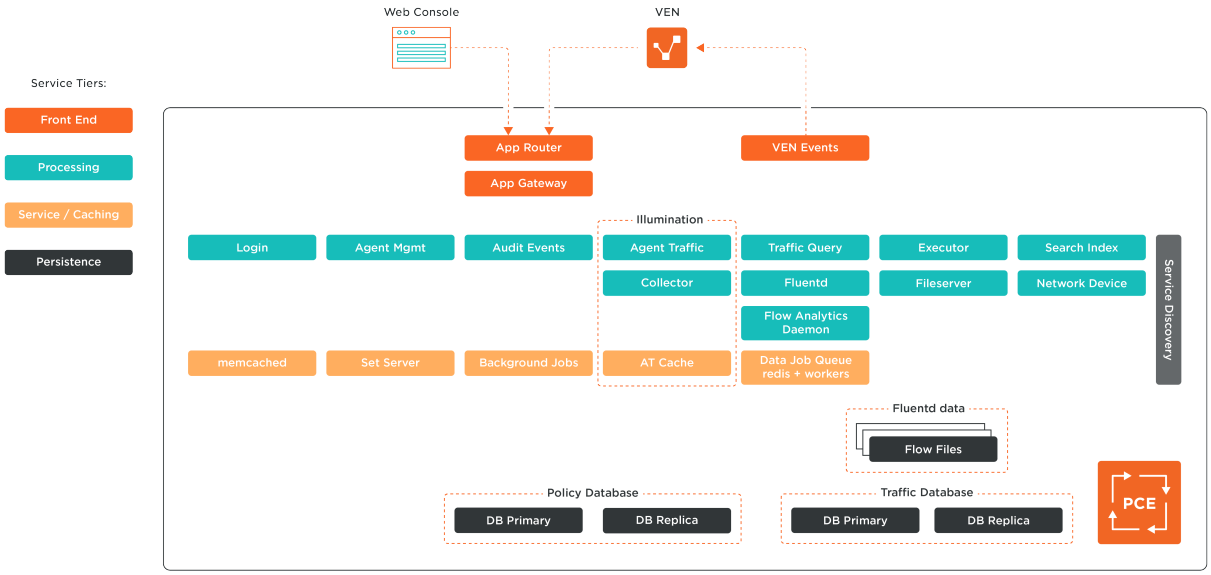

Description of PCE Components

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 the PCE manages.

Tier	PCE component	Description
Service	memcached	Open source component: in-memory cache.
	Background Jobs	The backbone for asynchronous job execution is report generation and background jobs.
	Set Server	In-memory cache to aid in policy calculations.
	Agent Traffic cache	It stores the traffic flow data and graphs for Illumination. See Agent Traffic. In the PCE architecture diagram, it is labeled “AT Cache.”
	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. This data store backs Traffic Explorer. 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.

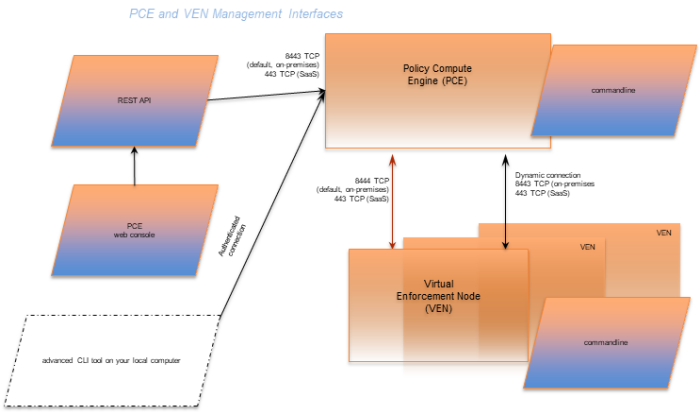

This guide focuses on using the `illumio-pce-ctl` control script and related administrative programs on the PCE itself.

Interface	Notes
PCE web console	With the PCE web console, you can perform many common tasks for managing the Illumio Segmentation for Data Centers.
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 Illumio Segmentation for Data Centers, including installing and updating the VEN.
REST API	With the 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 (policies, IP lists, services, and security settings), and labels.

Your Illumio administrator initially sets up organizations. When an organization is created, an email containing a user login is sent. The organization is created when this user logs in, and users can now be invited to join.

Invite Users to Your Organization

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

When you invite a user to your organization, they receive an email at the specified address containing 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 Segmentation for Data Centers 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 any anti-spam protection does not block its emails.
- Check your PCE's `runtime_env.yml` file to make sure that the `smtp_relay_address` value is correct.

Manage PCE Nodes and Clusters

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

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
- General PCE functionality issues

Recommendations for Data Management

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, follow these 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.

About the PCE Databases

This section outlines the key concepts you must understand to effectively 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 require backup or restoration.

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 to traffic data:

- The default storage limit on traffic data from 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 or 90-day limits are reached, traffic flow data is pruned. The pruning order is as follows: first, data from endpoints, then Kubelink, and lastly, Server VENs.

Determine the Primary Database

Policy Database

Run the following command to determine the primary policy database:

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

Traffic Database

Run the following command to determine the primary traffic database:

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

Show Database Replication Information

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

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

Rotate Database Passwords and Other Secrets

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

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

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

Anonymize Database Export

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

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

   Policy database

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

   Traffic database

   ```
   sudo -u ilo-pce /opt/illumio_pce/illumio-pce-db-management
   traffic dump --for-masking --file traffic_backup_filename
   ```
2. Anonymize the policy or traffic dump file by running one of the following commands.

Policy dump file

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

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

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

Optional `--tmpdir` parameter

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

Example command output

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

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

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

Illumio recommends using the same dictionary file consistently. This approach ensures that the exact 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 is based on the 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   |
+-----------------------------------+-------+------------+
```

Database Migration, Failover, and Restore

This section outlines the steps for performing database management tasks.

Migrate PCE Databases

These steps outline the process of migrating the database from a previous version to the 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 the 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 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 within 15 minutes, the PCE may only apply the 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. The `agent_traffic_redis_server` service is not running on runlevel 1.

   The following command is used to identify the Database Primary node to which to restore the database:

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

5. On the *primary data node* identified in the previous step, restore the policy database:

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

6. Copy the Illumination data file from the primary *data node* 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
   ```

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

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

   When prompted to bring the PCE to runlevel 5, reply "`yes`" if you want the PCE to automatically complete the traffic database migration and bring the PCE to fully operational status. Reply "`no`" if you don't want to migrate the traffic database.
8. 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
```

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

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

10. 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.

- A multi-node traffic database may be required 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).

Expand the Existing Traffic Database to Multiple Nodes

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

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

   ```
   sudo -u ilo-pc e illumio-pce-db-management traffic dump --
   file trafficdb-backup.tar.gz
   ```
2. On *any data node*, run the following command to back up the reporting database:

   ```
   sudo -u ilo-pc e illumio-pce-db-management report dump --
   file reportdb-backup.tar.gz
   ```
3. On *all new nodes*, run the following command to allow multi-node traffic, where the address is the IP address of each new node:

   ```
   illumio-pce-ctl cluster-nodes allow <address>
   ```
4. On *all nodes*, stop the PCE:

   ```
   sudo -u ilo-pce illumio-pce-ctl stop
   ```
5. Install the PCE software on the new coordinator and worker nodes, using the same version of the PCE on the existing cluster nodes. There must be exactly two (2) coordinator nodes and two (2) or more pairs of worker nodes.
6. Please update the runtime_env.yml configuration on every node (both the new ones you just added and those 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.tar.gz
```

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

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

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

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

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

If you do not choose to go to runlevel five (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

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

WARNING

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

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

   ```
   sudo -u ilo-pce  illumio-pce-db-management traffic dump --
   file trafficdb_backup.tar.gz
   ```
2. On *any node*, set the PCE to runlevel 1:

   ```
   sudo -u ilo-pce  illumio-pce-ctl set-runlevel  1
   ```
3. When removing a node, run the following command on the node you are removing:

   ```
   sudo -u ilo-pce  illumio-pce-ctl cluster-leave
   ```
4. On *all nodes*, stop the PCE cluster:

   ```
   sudo -u ilo-pce  illumio-pce-ctl cluster-stop
   ```
5. On *every PCE node,* update the value of traffic_datastore.num_worker_nodes in runtime_env.yml. The value should always be twice as large as the number of individual worker nodes, because the worker nodes are configured in pairs.
6. On *all nodes*, start the PCE at runlevel 1:

   ```
   sudo -u ilo-pce  illumio-pce-ctl start --runlevel 1
   ```
7. On the *data node where you took the backup*, restore the traffic database backup that you made in step 1:

   ```
   sudo -u ilo-pce  illumio-pce-db-management traffic restore
   --file trafficdb_backup.tar.gz
   ```
8. On *any node*, set the PCE to runlevel 5:

   ```
   sudo -u ilo-pce  illumio-pce-ctl set-runlevel 5
   ```
9. Verify the successful setup. Log in to the PCE web console and verify that the PCE's health 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

The following are some useful commands for getting information about a cluster in which 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 the 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 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 [16] 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 [32].

About PCE Database Backup

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

IMPORTANT

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

When to Back Up

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

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

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

Back Up the Policy Database

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

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

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

   You see the following output:

   ```
   SERVICES (runlevel: 5) NODES (Reachable: 1 of 1)
   ====================== ============================
   agent_background_worker_service 192.168.33.90
   agent_service NOT RUNNING
   agent_slony_service 192.168.33.90
   agent_traffic_redis_cache 192.168.33.90
   agent_traffic_redis_server 192.168.33.90          <=== run
   the dump command from this node
   agent_traffic_service NOT RUNNING
   ...
   ```

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

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

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

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

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

Back Up the Traffic Database

The traffic database dump can be very large, depending on the traffic data-store 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 [31].

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

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

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

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

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

Using pgbackrest for Traffic Data Backups

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

Hardware Requirements

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

The NFS mount can be used to store other data in addition to the traffic data. For example, it could store the policy database and `runtime_env.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 backupLabel
```

Back Up the PCE Runtime Environment File

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

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

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

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

IMPORTANT

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

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 based on the PCE's tested performance and capacity.

Types of Object Limits

This section describes the difference between soft and hard limits.

Soft Limits

Soft limits serve as early warnings of 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 high volume of soft-limit alerts, consult Illumio Customer Support to adjust or suppress these alerts.

NOTE

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

Hard Limits

Hard limits protect the PCE from usage and performance overload, such as creating too many workloads or enforcing an overly restrictive security policy.

When you receive a hard limit warning, Illumio recommends investigating it immediately. When a hard limit is reached during a service outage, the PCE core capacity may become overloaded.

When a hard limit is reached, any attempt to create more objects of that type will fail, returning an error message in the PCE web console or an HTTP 406 response in the 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.

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 concurrently reach the hard object limit for that object during 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 transactions 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 key name 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
Custom Labels	Maximum number of labels	20	20
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, each containing 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
Label Dimensions (`total_label_dimensions`	Total number of label dimensions	none	20
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	Total number of network interfaces supported per unmanaged workload	102	128

Object	Description	Soft Limit	Hard Limit
(`interfaces_per_unmanaged_workload`)	An unmanaged workload does not have a VEN installed.		
Interfaces per VEN (`interfaces_per_agent`)	Total number of interfaces supported per managed workload A managed workload has a VEN installed.	32	None (-1)
Items per Rule (`total_actors_per_rule`)	The total number of items allowed per rule in the Providers and Destinations fields. A rule contains labels, workloads, and IP lists. When a rule has two Provider items and two source items, it has four 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. When you click **Stop Pairing**, the pairing key becomes inactive and is no longer counted in the object limit.	1200	5K
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: an RBAC user or user group, a role, and a 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	Total number of port ranges per service	50	None (-1)

Object	Description	Soft Limit	Hard Limit
(`port_ranges_per_policy_service`)			
Services per Rule (`total_services_per_rule`)	The total number of services that can be associated with a single rule	40	50
Ports per Rule (`total_service_ports_per_rule`)	The total number of ports that can be associated with a single rule. Each service has a fixed number of ports or a port range. 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 adding a distinct port or port range to a rule.	400	500
Rules (`total_rules`)	Total number of all rules in all rule sets	40K	50K
Scopes and Rules (`total_scopes_rules`)	Sum of the total number of rules times the total number of scopes in all rule sets For example, you have two rule sets: 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_org_auth_security_principals`)	Total number of all RBAC users and groups	1600	2000

Object	Description	Soft Limit	Hard Limit
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 services over the 200 limit.	160	200
Workloads (`total_workloads`)	Total number of managed and unmanaged workloads A managed workload has a VEN installed, whereas 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 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 managed by a Kubelink running in Cluster Local Actor Store (CLAS) mode.	8K	10K

Object	Description	Soft Limit	Hard Limit
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 (`total_active_sessions`)	The maximum number of user sessions on a single PCE cluster at the same time. This limit includes only actual logged-in user sessions and omits impersonated sessions, such as scheduled jobs that log in to access PCE data. When the limit is exceeded, anyone attempting to log in is denied access with an explanatory message.	100	125

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 Segmentation for Data Centers 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 se-

curity 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 Segmentation for Data Centers. 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 Segmentation for Data Centers.

  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 Segmentation for Data Centers.

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 Segmentation for Data Centers 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 source like AWS.
- To decentralize their security policy to specific application teams allowing them to act quickly when managing application security without waiting for the central security team.
- To bifurcate the security of their infrastructure in such a way that one user is responsible only for the West coast assets and another user is responsible for the East coast assets.

Features of Role-based Access Control

Built-in Roles

Illumio Segmentation for Data Centers includes several roles that grant users access to perform operations. Each role is matched with a scope. See About Roles, Scopes, and Granted Access [44] for information.

Granular Permissions

You can assign multiple roles to one user and by mixing and matching the different roles, you can achieve different levels of granularity of permissions.

You can grant different permissions to different users for different resources by defining scopes. For example, you might allow some users complete access to add rulesets for all workloads in your staging environment. For other users, you might grant access to all workloads in all environments. Users can be assigned exactly one role, representing their singular job function while other users can be assigned multiple roles, representing multiple job functions.

Identity Federation Using External Users and Groups

You can connect to external LDAP directories to manage users and user groups by configuring single sign-on (SSO) for the PCE.

Using this feature, you can create and manage users locally in PCE, or use an IdP to manage users and user groups from an existing directory. External user and user groups authenticate with the external IdPs.

Custom Role Assignments

You can customize access to suit your organization by specifying specific scopes for the Ruleset Manager and Ruleset Provisioner roles.

Audit Information

You can access an audit trail of user activity through the following reports:

- The User Activity page, which displays the authentication details for each user, when they logged in, and whether they are online.
- The Organization Events page, which displays when Organization Owners granted users access, when users logged in and out, and the actions they performed.

About Roles, Scopes, and Granted Access

Illumio Segmentation for Data Centers 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 Segmentation for Data Centers.

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. This user can change settings in UI to allow administrators to choose whether to show warning for editing labels used by enforced workloads or not. The default is do not show warning.
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 Segmentation for Data Centers 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 source matches the specified scope. The rule destination 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 source and destination 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 destinations, 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 source and destination 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 destination is not restricted. However, in extra-scope rules, the Environment label of the resource selected as the destination 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 destinations 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.

These users can write rules for their workloads and provision them when the rules do not have dependencies on global objects, such as services or IP lists.

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.

These users can create rulesets within their assigned scopes and write rules that are not dependent on global objects. However, they can provision any workloads, even those containing services, IP lists, and label groups.

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.

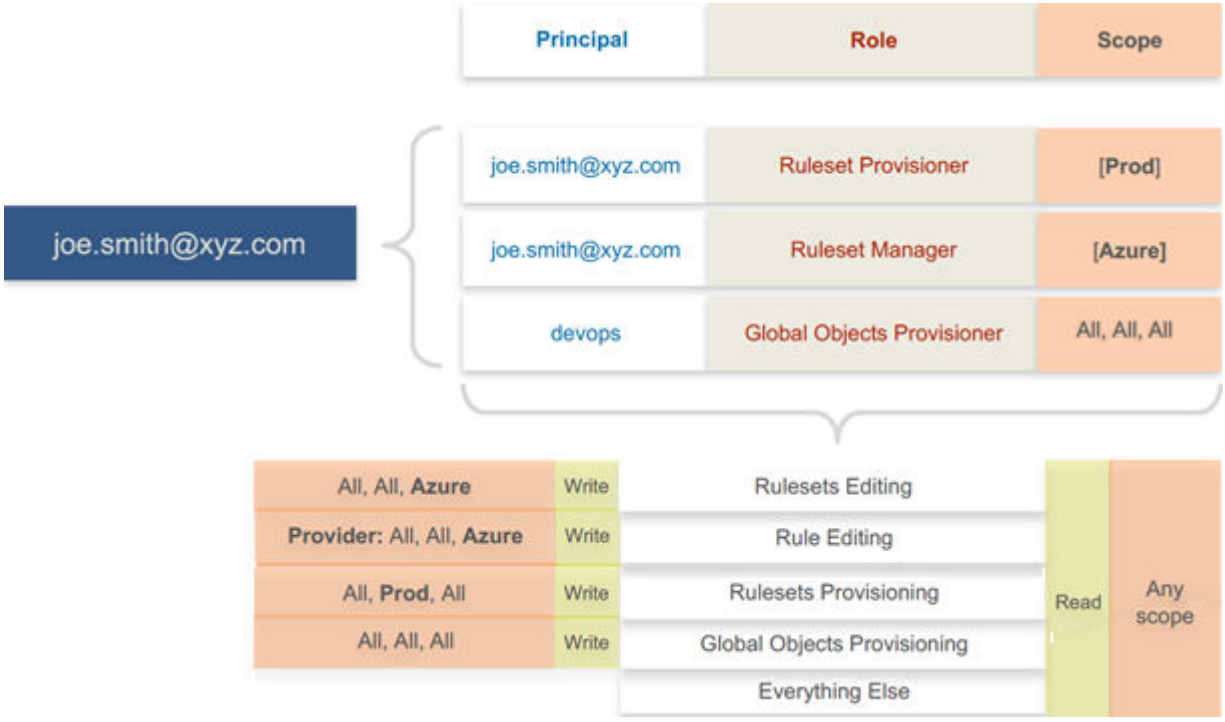

Exercise Caution when Combining Roles

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.

Authentication

The Illumio PCE supports either SAML SSO or LDAP as an external authentication method. Both SAML SSO and LDAP cannot be used simultaneously. When LDAP is turned on, the use of SAML SSO, if already configured, is disa-

bled. Similarly, enabling SAML SSO after LDAP is enabled will disable LDAP authentication.

SAML SSO Authentication

When you use a third-party SAML-based Identity source (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 Illumio Segmentation for Data Centers. It 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 configuring SSO in the PCE, you must configure SSO on your chosen IdP and obtain the required SSO information. After obtaining the IdP SSO information, log in to 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

The user's email address is the primary attribute used by the PCE to identify users uniquely.

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. Locate the SAML configuration panel on the Authentication Settings screen and click **Configure**.
3. Use the displayed information (as shown in the example below) while configuring your specific IdP.

Information for Identity Provider

Authentication Method	Unspecified
Force Re-authentication	No
SAML Version	2.0
Issuer	https://c[illegible]3/login
NameID Format	urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress
Assertion Consumer URL	https://[illegible]3/login/acs/a63e[illegible]49598e
Logout URL	https://[illegible]43/login/logout/a63[illegible]349598e

NOTE

Even though the SAML NameID format specifies an email address, 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**.

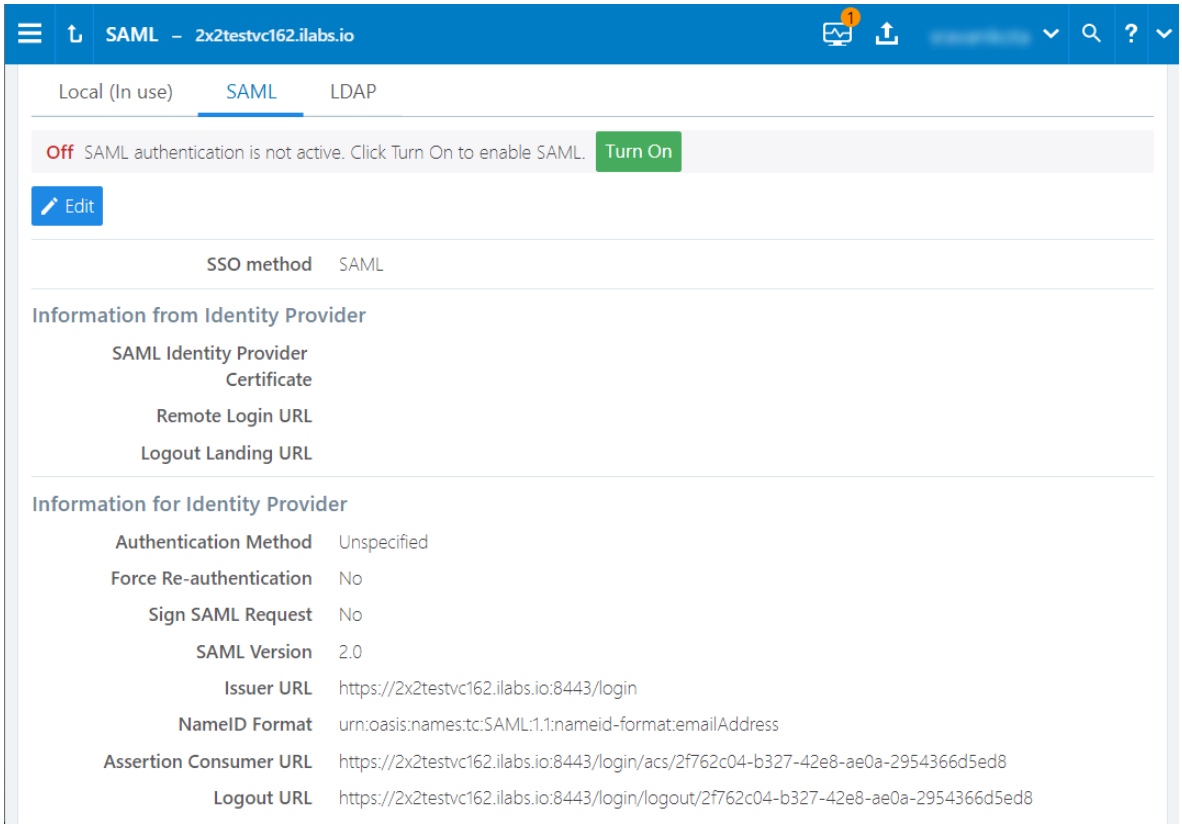

4. In the pop-up screen, click **Confirm**.

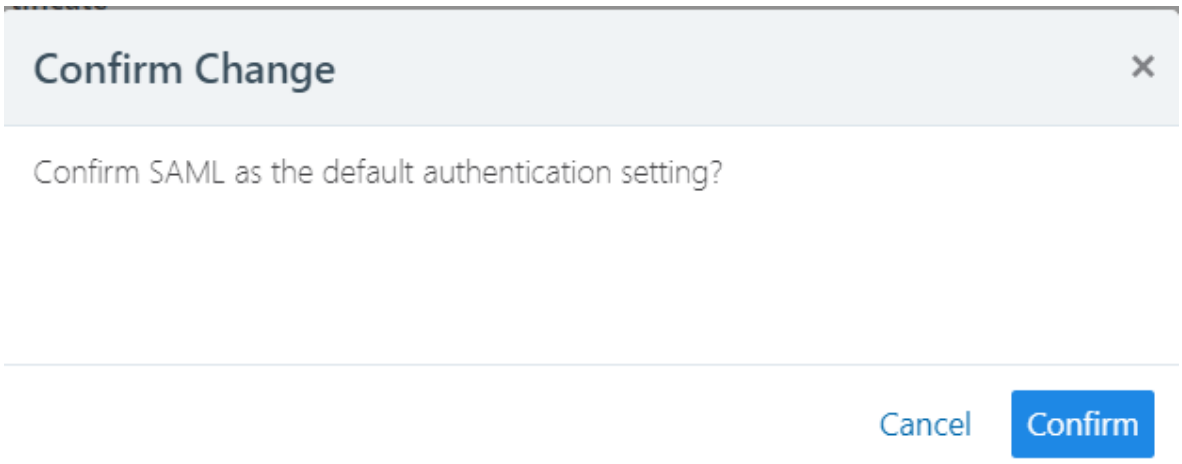

The updated SAML screen shows that SAML authentication is active.

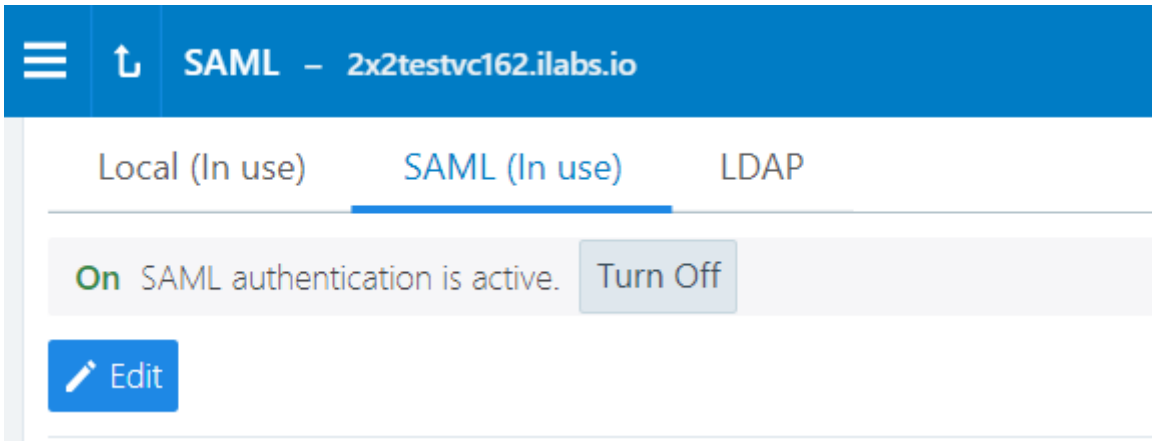

If necessary, you can disable it at any time.

Once configured using these steps, the lifetime of the SAML certificate is ten years.

LDAP Authentication

The Illumio PCE supports either SAML SSO or LDAP as an external authentication method. Both SAML SSO and LDAP cannot be used simultaneously. 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.

LDAP Authentication Support

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. Cloud does not support LDAP authentication.

To use LDAP authentication:

1. Review the Prerequisites and Limitations [57].
2. Enable the PCE to use LDAP authentication. See Enabling LDAP Authentication [58].
3. Set up an LDAP configuration. See Configuring LDAP Authentication [58].
4. Map your LDAP groups to one or more PCE roles. See Map LDAP Groups to User Roles [59].

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.

Contact your LDAP administrator for assistance if you cannot get this information.

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 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 the roles mapped to the local user will be applied to login. 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 use LDAP or SAML single sign-on to authenticate 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 Management** > **Authentication**.
3. In the Authentication Settings screen, locate the LDAP configuration panel and select **Configure**.
4. In the LDAP Authentication screen, select **Turn On**.

Configure LDAP Authentication

Follow these steps to configure LDAP authentication on the PCE. First, make sure you have followed the steps in Enable LDAP Authentication [58].

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access Management** > **Authentication**.
3. Locate the LDAP configuration panel on the Authentication Settings screen and click **Configure**.
4. Make sure LDAP is enabled on the LDAP authentication screen.
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: This is required only when a 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, but it can be configured to any value between 1 and 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: This flag is enabled by default and specifies whether to verify the server certificate when establishing an SSL connection to the LDAP server. Disabling it is not recommended.
- User Base DN: Base DN of the LDAP directory to search for users.
- User Search Filter: A search filter is used to query the LDAP tree for users.
- User Name Attribute: An attribute on a user object that contains the username, such as uid, sAMAccountName, or userPrincipalName.
- Full Name Attribute: Attribute of a user object that contains the full name. For example, cn, commonName, displayName.
- Group Membership Attribute: An attribute of a user object containing group membership information, such as memberOf or isMemberOf.

7. Click **Test Connection** to verify that the PCE can connect to the LDAP server successfully. If Test Connection fails, check your LDAP configuration and retry.

You can enter up to three LDAP server configurations for a PCE. See How the PCE Works with Multiple LDAP Servers for more information about using multiple LDAP servers.

Map LDAP Groups to User Roles

After configuring 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 an individual user's permissions, 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, add the user to that LDAP group on your LDAP server.
- Remove a user from a PCE role: Remove the user from the corresponding LDAP group on your LDAP server.

A user can be a member of several roles. In that case, the user can access all the capabilities available for any of those roles. For example, if a user is a member of both the docs and eng LDAP server groups, and the docs group

is mapped to the PCE user role "Ruleset Manager" and the eng group is mapped to "Ruleset Provisioner," the user obtains all permissions assigned to both the "Ruleset Manager" and "Ruleset Provisioner" roles.

NOTE

The PCE checks LDAP membership information when a user attempts to log in. You do not need to reload the authentication configuration when adding or removing users.

For details about how to map external groups to PCE user roles, see Setup for Role-based Access Control [80].

Modify LDAP Configuration

Follow these steps to update or delete an LDAP configuration in the PCE. You are assumed to have already followed the steps in Enable LDAP Authentication [58] and Configure LDAP Authentication [58].

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access Management** > **Authentication**.
3. Locate the LDAP configuration panel on the Authentication Settings screen and click **Configure**.
4. Make sure LDAP is enabled on the LDAP authentication screen.
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). You are assumed to have already followed the steps in Enable LDAP Authentication [58] and Configure LDAP Authentication [58].

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access Management** > **Authentication**.
3. Locate the LDAP configuration panel on the Authentication Settings screen and click **Configure**.
4. Make sure LDAP is enabled on the LDAP authentication screen.
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 works.

Secure LDAP with SSL/TLS Certificates

The PCE supports LDAPS and LDAP with STARTTLS. To use the PCE with secure LDAP, add the certificate chain to the local certificate store on the PCE. Follow these steps to configure secure LDAP. It is assumed you have already followed the steps in Enable LDAP Authentication [58] and Configure LDAP Authentication [58].

1. Log in to the PCE as a Global Organization Owner.
2. Choose **Access Management** > **Authentication**.
3. Locate the LDAP configuration panel on the Authentication Settings screen and click **Configure**.
4. Make sure LDAP is enabled on the LDAP authentication screen.
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**
> Self-signed certificates are not recommended for an LDAP server. Illumio recommends using 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 expires or fails, the PCE does not attempt to log in using LDAP authentication.
2. If the local user does not exist, the PCE attempts to log in to LDAP (if enabled).

How the PCE Works with Multiple LDAP Servers

You can configure up to three LDAP servers for each PCE. The Illumio Core platform can support up to three LDAP servers per region in a PCE super-cluster deployment.

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, and C. The PCE attempts to connect to the servers in that order. If it fails to connect to A, it attempts to connect to the remaining servers: first B, then C, after the connection timeout expires.

When the PCE successfully connects to an LDAP server, it searches for the user. 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, it 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.

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 (the required number and types of characters), and expiration, re-use, and history.

About Password Policy for the PCE

You must be a Global Organization Owner to view the Password Policy feature under the Settings > Authentication menu options.

The password length can be set to a maximum of 64 characters.

NOTE

The Password Policy feature does not apply to 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 must 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 eight 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 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 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 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 a password's similarity by not allowing a user to change it 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 1 to 24 passwords before reuse. The default setting for password reuse is five password changes before reuse is permitted.

Caveats

- When a Global Organization Owner increases the minimum password length policy or the password complexity requirements and enables password expiration (1-999 days), all existing users must reset their passwords based on the new policy.
- When a Global Organization Owner configures the password never to expire, all users who migrated from an older release to 18.2.0 must reset their passwords when they 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): The user will sign in to the PCE only with a local credential provided by the user's organization's password policy.

   SAML (IN USE): SAML users can 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 change
Similarity	1 character and position from the current password
Session timeout	You can configure the session timeout value using the PCE web console. 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 sessions frequently expiring. The timeout value depends 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; it does not affect existing browser sessions.
Timeout	30 minutes

SAML (in use)	
Information from identity source	
SAML Identity source certificate	-----BEGIN CERTIFICATE----- MIICpDCCAYwCCQD05WZzgx RugDANBgkqhkiG9w0BAQsFADAUMRIwEAYDVQQDDAlsb2NhbGhvc 3QwHhcNMTgxMTE0MjAyNzM2WhcNMjgxMTExMjAyNzM2WjAUMRIw EAYDVQQDDAlsb2NhbGhvc3QwggEiMA0GCSqGSIb3DQEBAQUAA4I BDwAwggEKAoIBAQDXs/OhH90IPQ8qBrUMqzQZb5MI72fu+Ay0s P8gI1v8RiUqSl+WJNo8s9L8GNI9hnQT+OXg99PNmoE41xiAlnx qx8T78Qxb9zX3uc4hec+9bMSF7iieUiFXWQQrIUVM3g8TWI6B5g Uapt0vZcxNok2eNhiFvVTLgPzB06vb2/yU68ilwQ8wz/MGO00Un/ lRw3LORynEA1uMeT6terWtX8JQGbvc1qYddnXD86Y5MOP1AXU+ 1w1w1JFxD0uKiuOHJvNYfJjkisEbDis9bO/EO0SyayVA7ABELaw QTfeWM6xLrNhZCTGeQiKb4XHMBgeliAloEvNDDofKbLDQrWUyIf7 TAgMBAAEwDQYJKoZIhvcNAQELBQADggEBANLhqsZsFUnq7kc+B5a vMmOXbCNJmSaASBULsX+akexhyJdMZUxmN6wfLjZ3FOwxvFuheTa Zpkp1UtC+2E9YlxY//FxOX/YyvNT/xfOBzqZ9SCsNxpCBsSRK5X4 DS+2jGQuz3fwbJDxTXP4sKNUZ/E9Z+dC9Npdq7xtcXr7pWhI2qe MO8E9LdvfWLcsqq8Z0VtxyHYYZYNh8KN0Q6ObfK1sPC4QZ/292B xm2ckxsWDTyONV8ytLQKwp93exxqmzzpbz6qi23y0B4u4af+/SW9 ukjzD/ atP34bY1YjeLBCsKEgy1nDTVgypAZSEy46kJ9mAu6t3r4/gEg XTkMYQDtr- PA= -----END CERTIFICATE-----
Remote login URL	https://hohoho.illumio.com
Logout landing URL	https://hohoho.illumios.com/1logout
Information for the identity source	
Authentication method	unspecified
Force re-authentication	no
Sign SAML request	no
SAML version	2.0
Issuer URL	https://2x2testlab360.ilabs.io:8443/login
NameID format	urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress
Assertion destination URL	https://2x2testlab360.mylabs.io:8443/login/acs/6b5243ef-2305-4ffd-bf81-4fa97fb91a5b

SAML (in use)	
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 to all the LDAP servers.
6. To create an LDAP server, click on **Create Server**.

   To continue with LDAP server configuration, see Enable LDAP Authentication [58].

Updates to Roles

Illumio Segmentation for Data Centers provides two types of user roles - Global and Scoped. It also provides the ability to stack multiple roles for the same user. A PCE owner can assign a combination of multiple roles to the same user. The resulting set of permissions is the summation of all permissions included with each of the stacked roles. With these updates:

- Existing scoped roles enhanced to restrict reads by scope.
- New scope based *read-only* role limits read access by labels.
- Scoped users get limited visibility into objects 1-hop away (this applies to Explorer, App Group Maps, Rule Search, and Traffic).
- Global read-only disabled by default for new PCE installations.
- PCE performance and scale enhanced to support concurrently active users.

Global Roles

Global roles provide the user with permissions to view everything and to perform operations globally. The four Global roles are :

- Global Organization Owner: Allowed to manage all aspects of the PCE, including user management.
- Global Administrator: Allowed to managed most aspects of the PCE, with the exception of user management.
- Global Viewer: Allowed to view everything within the PCE in a read-only capacity. This role was previously called "Global Read-only".
- Global Policy Object Provisioner: Allowed to provision global objects that require provisioning such as, Services and Label Groups.

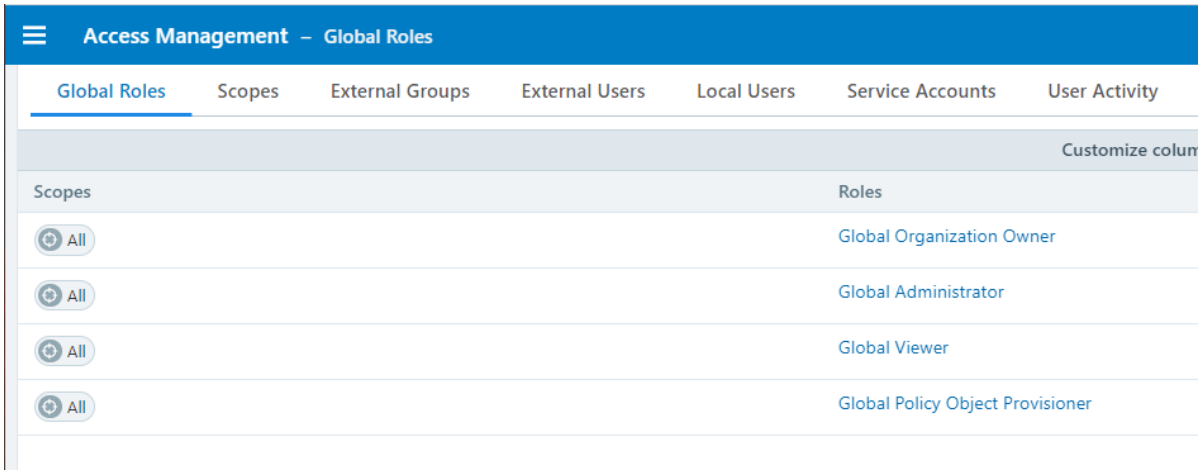

Scoped Roles

The Scoped roles are defined using labels. The permissions included with the assigned role apply only to the assigned scope where the scope is defined using a combination of as many label types you have defined (and with only one label value per type). To provide permissions to different applications for a user, each of the application scopes has to be added to the same user.

All the Scoped roles have been enhanced to restrict reads and writes by Scope. The Scoped roles are :

- Ruleset Viewer: A new scope-based read-only role. A user with this role has read-only permissions within the assigned scope. The user can view policy, application groups, incoming and outgoing traffic, and labeled objects such as, workloads, within the assigned scope.
- Ruleset Manager (Limited or Full): An existing scope-based read/write role. A user with this role can read/write policy within the assigned scope. The user can also view application groups, incoming and outgoing traffic, and labeled objects, within the assigned scope.
- Ruleset Provisioner: This role allows a user to provision changes to the scoped objects, provided the objects are inside the user's assigned scope. A user with this role can provision changes to policies within the assigned scope. The user can also view application groups, incoming and outgoing traffic, and labeled objects, within the assigned scope.
- Workload Manager: Allows a user to perform workload-specific operations such as pairing, unpairing, assignment of labels, and changing of policy state. A user with this role cannot view policies and traffic, and cannot provision changes.

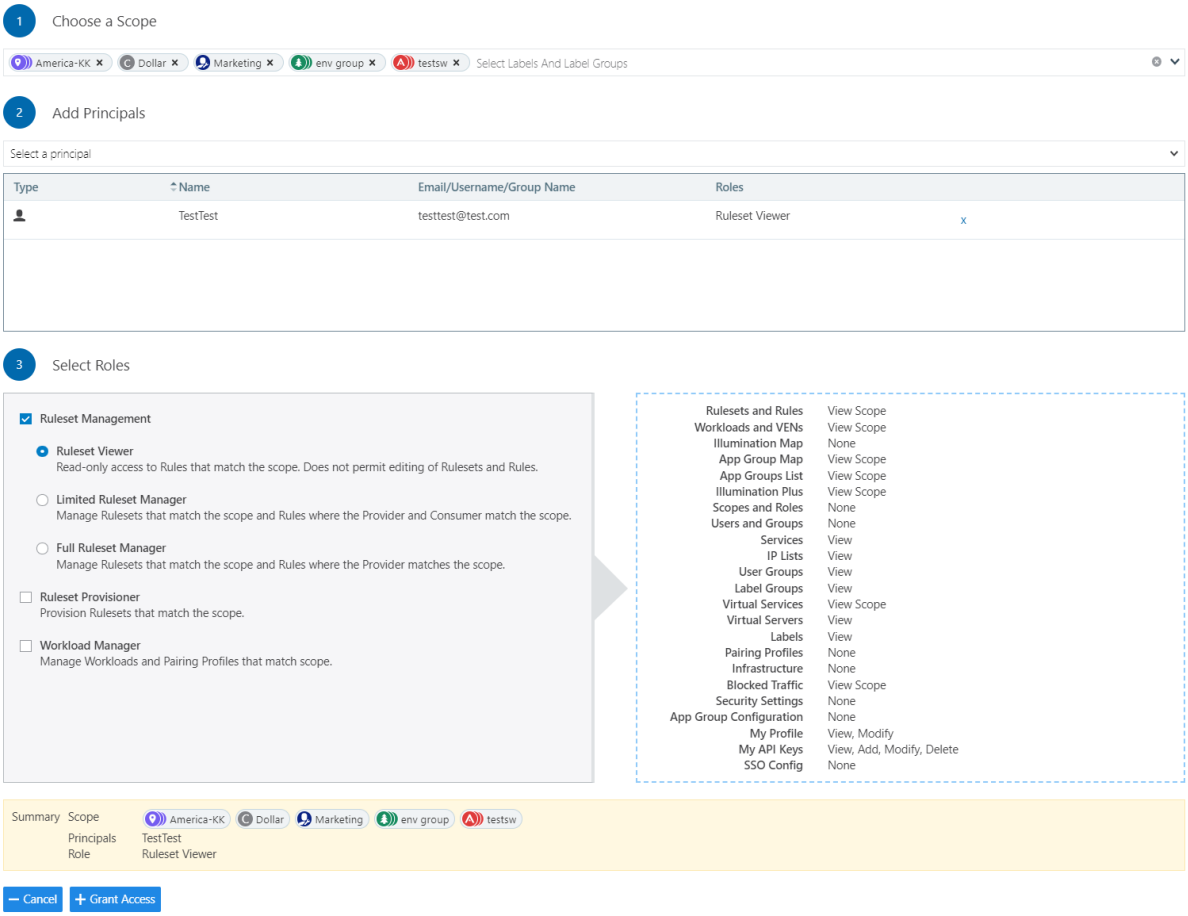

OneLogin Single Sign-on

This section describes how to configure SSO for OneLogin.

Configure SSO for OneLogin

This task shows you how to configure SSO to authenticate 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

To perform this task, your PCE user account must have Owner or Admin privileges.

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.
6. In the Information for Identity Provider section, make note of the following fields:
   - Issuer
   - Assertion source 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**: The user must 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

   The preferences are used when SSO is configured both in Illumio Segmentation for Data Centers and for the IdP. When SSO is not configured, the default IdP settings are used.

9. Click **Save**.
10 Log in to your OneLogin account.
.
11. Select the Illumio Segmentation for Data Centers 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.

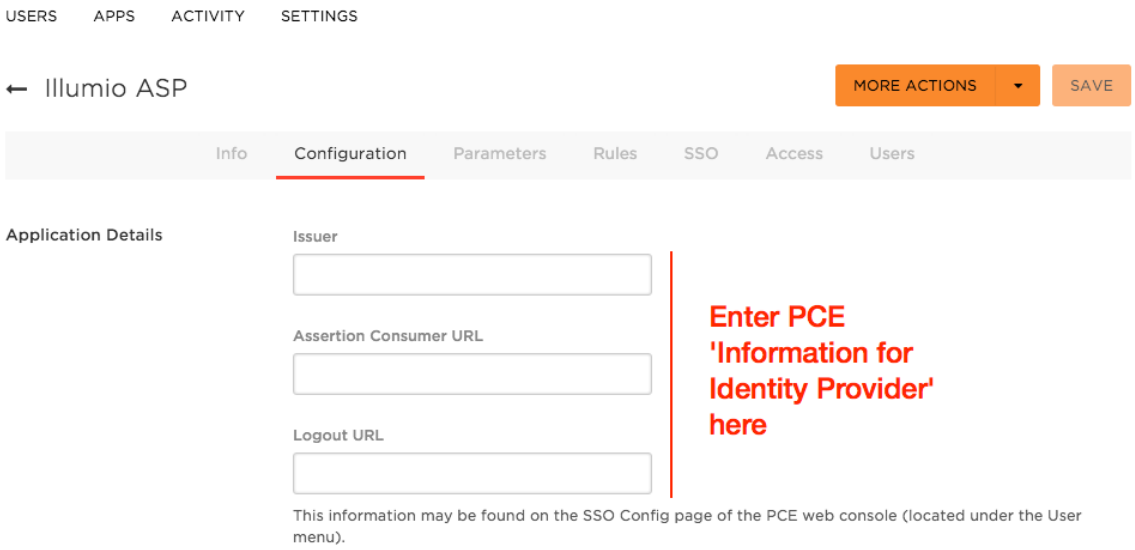

13. Click **Save**.
    Your PCE is now configured to use OneLogin SSO to authenticate users with the PCE.

Okta Single Sign-on

This section explains configuring 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

To perform this task, your PCE user account must have Owner or Admin privileges.

Configure the PCE for Okta SSO

1. From the PCE web console menu, choose **Access Management** > **Authentication**.
2. Locate the SAML configuration panel on the Authentication Settings screen 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 using Okta to authenticate with the PCE. When you select No Access, SSO users from your Okta account must be added manually before logging into the PCE.
5. In the Information for Identity Provider section, make note of the following fields:
   - Issuer
   - Assertion source URL
6. Select the authentication method from the drop-down list:
   - **Unspecified:** Uses the IdP default authentication mechanism.
   - **Password-Protected Transport:** The user must 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

   The preferences are used when SSO is configured both in Illumio Segmentation for Data Centers and for the IdP. When SSO is not configured, the default IdP settings are used.

8. Click **Save**.
9. Log in to your Okta account.

10. Select the Illumio Segmentation for Data Centers app, select the General tab, and click **Edit**.
11. Enter the values you copied from the PCE SSO Configuration page's Information for Identity Provider section.

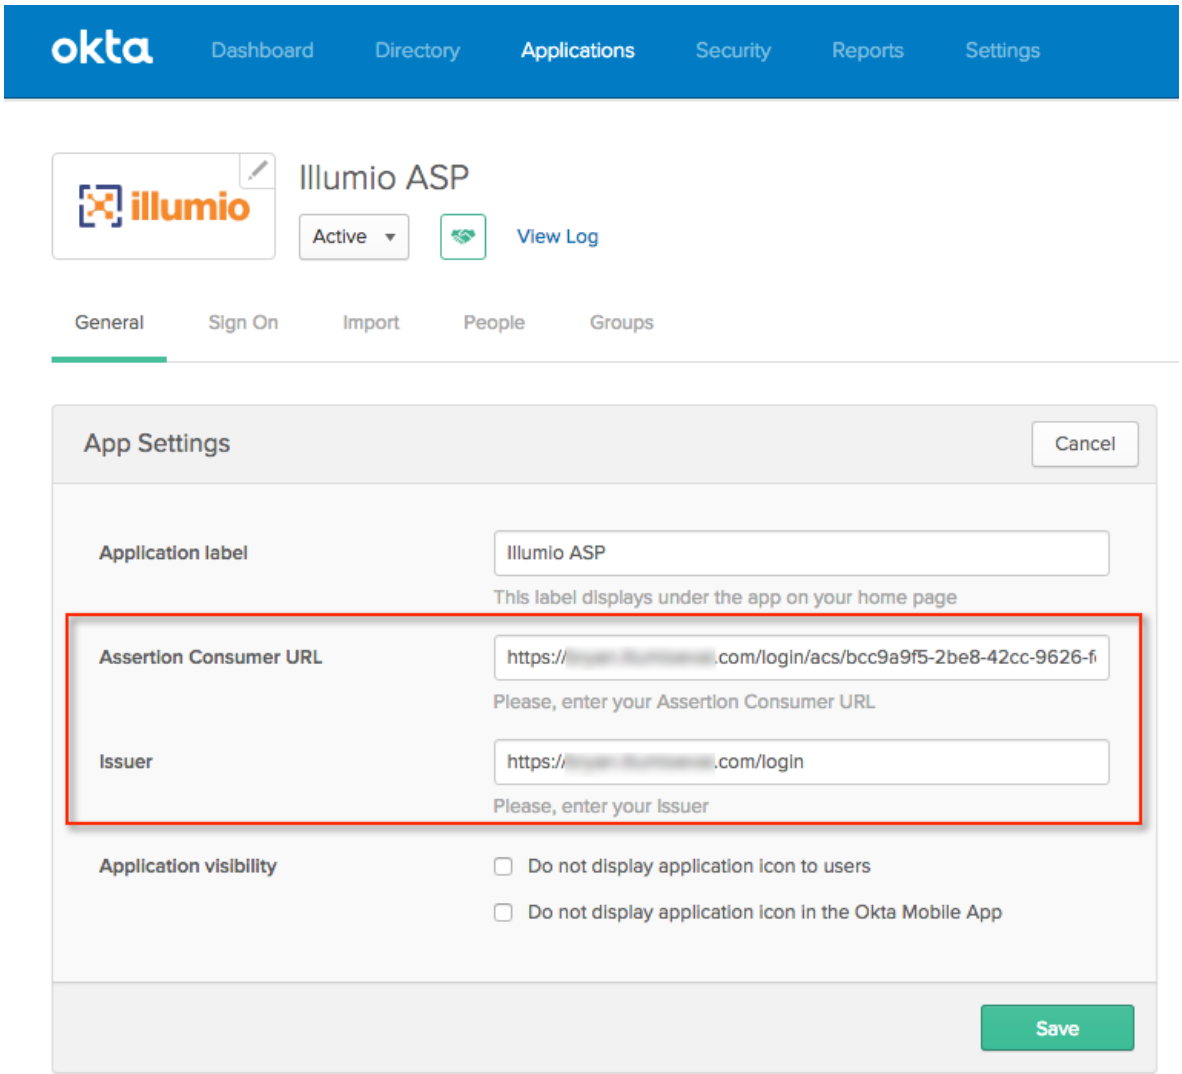

12. Click **Save**.

    Your PCE is now configured to use Okta SSO to authenticate users with the PCE.

Active Directory Single Sign-on

Learn 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@yourADDomainName. For example, a user named "test" in an Active Directory domain called "testing.com" would have a UPN of test@testing.com.

Configure AD Users to Use Different UPN Suffixes

To configure a different UPN suffix as the source for NameID:

1. To 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**.

   Figure 1. Alternative UPN suffixes

   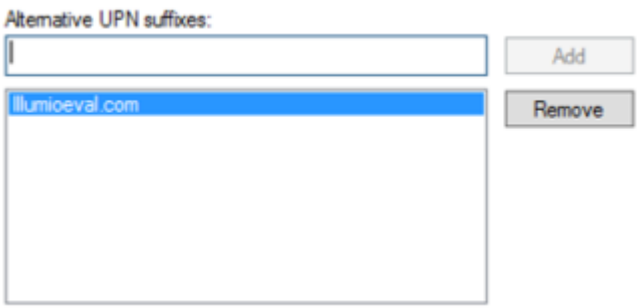

   You can now assign a custom UPN to an Active Directory user for the SAML response.
4. You can add multiple UPNs if needed, and select the UPN created in the previous steps.

   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 Segmentation for Data Centers.

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.

   Follow the Active Directory Federation Service Wizard to import a self-signed certificate.
5. Specify your Federation Service Name, enter a display name for this instance of AD FS, and click **Next**
6. Specify your service account and click **Next**.
7. Select "Create a database on this server using Windows Internal Database" or choose the SQL server option, and click **Next**.
8. Review your selected options and click **Next**.
9. Click Configure to complete the basic configuration of AD FS.
10 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 it, 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 source URL" (obtained from the PCE web console).

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

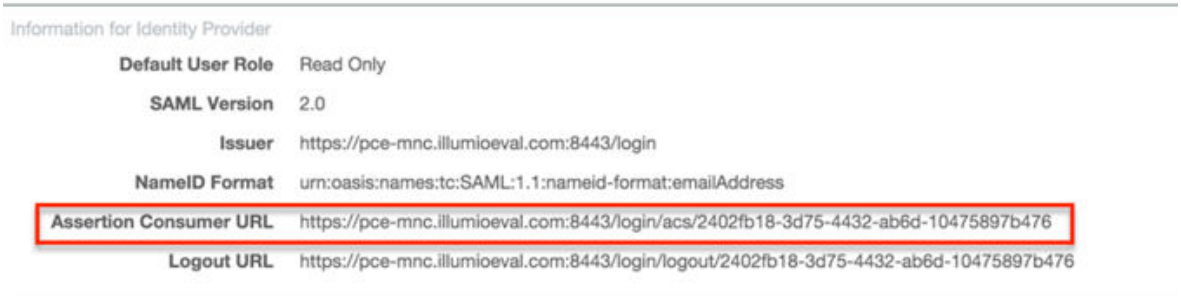

9. On the Configure Identifiers page, use the same URL for the Relying party trust identifier, without the `/acs/<randomNumbers>`.

   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 Group Claim Rule.

    To add groups to the 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 in to 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**.

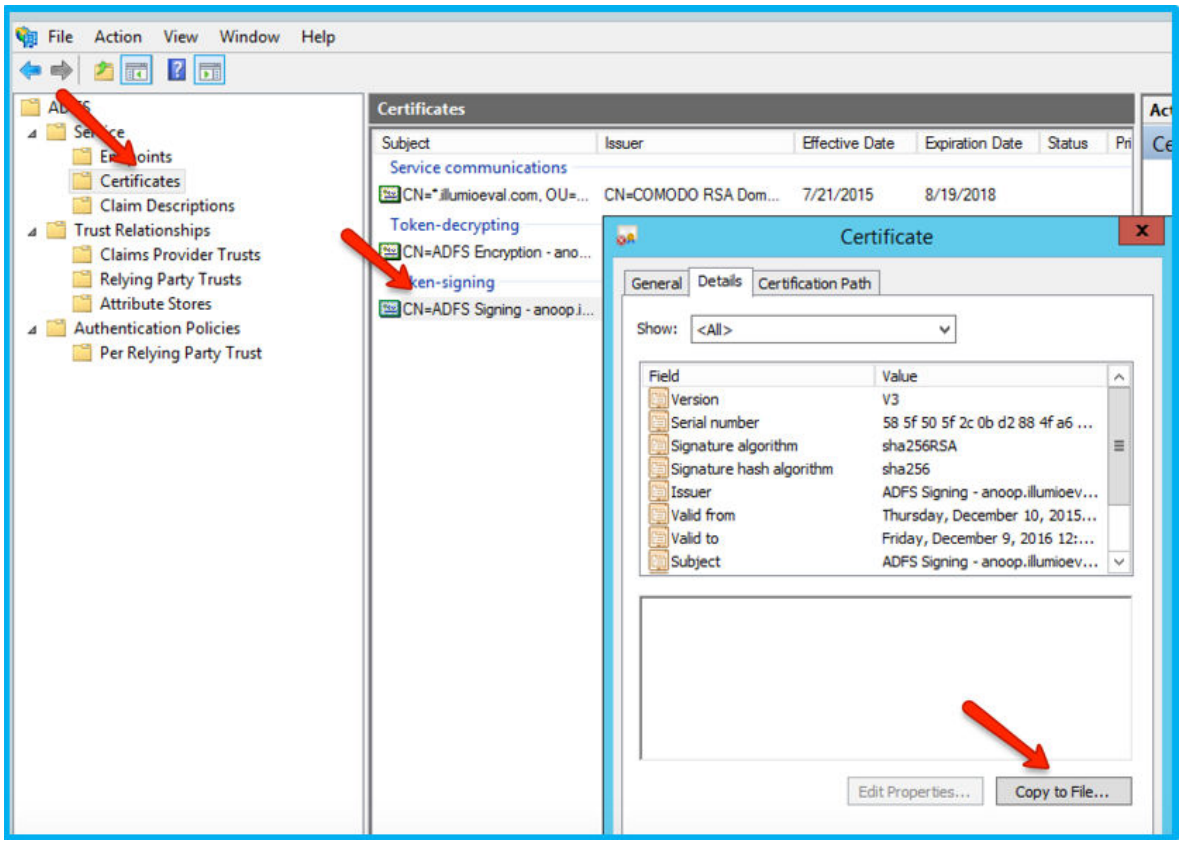

6. When the Certificate Export Wizard launches, click **Next**.
7. Verify that the “No - do not export the private key” option is selected and click **Next**.
8. Select Base 64 encoded binary X.509 (.cer) and click **Next**.
9. Select where you want to save the file, name the file, and click **Next**.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

```
format:persistent</NameIDFormat><NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid
-format:transient</NameIDFormat><SingleSignOnService

Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
Location="https://[redacted].illumio[redacted].com/adfs/ls/"/><SingleSignOnService
Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
Location="https://anoop.illumioeval.com/adfs/ls/"/><Attribute
Name="http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress"
```

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 single sign-on (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 [77].

NOTE

When SSO is configured in Illumio Segmentation for Data Centers and for the IdP, the preferences in Illumio Segmentation for Data Centers are used. When SSO is not configured in Illumio Segmentation for Data Centers, 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 single sign-on (SSO) authentication.

Setup for Role-based Access Control

This section describes configuring 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 complete access to add rulesets for all workloads in your staging environment and grant other users access to all workloads in all environments.

1. From the PCE web console menu, choose **Access Management** > **Scopes**.
2. Click **Add**.
   The Access Wizard appears.
3. Define the role's scope by selecting labels or label groups for Applications, Environment, and Location or for any other custom label types you have defined using Flexible Labels.
4. Add a principal -- a local user, external user, or user group -- to the role.
5. Select roles.
6. Click **Grant Access** > **Confirm**.
   The newly-added role is displayed on the Scopes page, and you can select it to edit or remove access.

Manage a Local User

Local users are created in the PCE (an IdP does not manage them). When they log into the PCE, they must enter their email addresses and passwords, which the Illumio PCE encrypts and stores.

When you install the PCE, the first user account it creates is a local user. You can create additional local users as a backup if your external IdP goes offline or the SAML server is inaccessible.

To add a local user:

1. Choose **Access Management** > **Local Users** tab.
2. Click **Add**.
3. Enter a name and an email address.
   The email address must be used in the format xxxx@yyyy.zzzz with 255 characters or fewer. You can add email addresses with an apostrophe (').
   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 it expires.

4. Select a role for the user:
   - None
   - Global Organization Owner
   - Global Administrator
   - Global Viewer
5. Choose an access restriction for the user (or None for no restrictions).

You can change a user's role membership after adding them by going to the user's details page or from a role details page. From the 20.1.0 release onwards, the "My Roles" feature lets you view the assigned permissions (roles) list.

To remove a local user:

1. From the PCE web console menu, choose **Access Management** > **Local Users**.
2. Select the user you want to remove.
3. Click **Remove User**.

When you remove a local user while the user is online, the PCE logs the user out as soon as the user is removed.

The user is removed from the Local Users tab; however, the user remains in the User Activity page and is designated 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, email address, or username.

To edit a local user:

1. From the PCE web console menu, choose **Access Management** > **Local Users**.
2. Click the name of the user you want to edit.
3. Click **Edit User**.
4. Change the user's name or access restriction and click **Save**.

You cannot edit a user's email address. You must remove and re-add the user with the new email address.

Changing a local user's name in the RBAC Roles and Users and Groups pages only changes it. The name is not changed in the user's profile or 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:

1. From the PCE web console menu, choose **Access Management** > **Local Users**.
2. Click the name of the user.
3. Click **Convert User**.

You can convert a local user to an external user so that your corporate IdP manages the user authentication credentials. When you convert a user to an external user, the user retains all their role memberships.

To invite a local user:

1. From the PCE web console menu, choose **Access Management** > **Local Users**.
2. Click the name of the user.
3. Click **Re-invite**.

You can send a new email to 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:

1. From the PCE web console menu, choose **Access Management** > **Local Users**.
2. Click the name of the user.
3. Click **Lock** or **Unlock**.

Manually locking out users with the above steps locks the user out indefinitely until you manually unlock their access.

Local users are 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 con-

sole reports that the username or password is invalid even when a user enters valid credentials. If automatically locked out for consecutive login failures, the user's account resets after 15 minutes and does not require an Illumio administrator to unlock it.

Manage a Service Account

The Global Organization Owner can create an API key without creating a new user account to associate with it. Instead, the API key can be associated with a service account. The service account is a security principal, just like a user.

- A service account can perform any API operation using its API key.
- Permissions for service accounts are specified with a combination of one or more PCE roles (Global Owner, Global Admin, etc.). You can include multiple roles for a single service account, just as you can for a user account.
- Access restrictions are supported. You can limit service account API keys by IP addresses, just as you can for user API keys.
- Service accounts support audit events. All audit events triggered by a service account indicate the name of the service account and the ID of its API key.

To create a service account:

1. Choose **Access Management** > **Service Accounts**.
2. Click **Add**, and give the account a unique name.
3. Enter an optional description for the account.
4. Specify the access restrictions for the account.
5. Optionally change the API key expiration duration from the default value. This duration cannot exceed your organization's setting.
6. Set the roles and scopes to determine the permissions granted to the service account. Click **Add** to assign an additional role or scope in the Roles and Scopes table.
7. To create an API key for the service account, click **Save** and **Download Credentials**.

   The new credentials are saved in the API Key section of the **Service Accounts**page.

Add or Remove an External User

Using RBAC, you can control access for users who are externally authenticated by a corporate IDP. Your corporate IDP manages authentication so that when these users log into the PCE, they are redirected to the IDP to authenticate. The PCE does not validate their usernames or passwords.

Using RBAC, you control external users' access to features and functionality. When you add an external user to the PCE, you specify that user's access by assigning the user to Illumio roles and scopes.

To add an external user:

1. Choose **Access Management** > **External Users** tab.
2. Click **Add**.
3. Enter a name and an email address or username.

   Whether you enter an email address or username for the user depends on how you have configured your IdP to identify corporate users.

   The username can contain up to 225 alphanumeric and special characters (. @ / _ % + -).

   In the PCE, you can have duplicate names for external users but you cannot have duplicates email addresses or usernames.

   When your IdP is configured to identify users using email addresses, the PCE emails the user at the address you specified an invitation to with a link to create their Illumio user account.

   If your IdP is configured to use usernames, you must provide the user with your Illumio PCE web console URL.
4. Select a role for the user:
   - None
   - Global Organization Owner
   - Global Administrator
   - Global Viewer
5. Specify an access restriction for the user, or leave it as None.

In general, a user must have a role to access the PCE. However, when default read-only is ON, any user accessing the org inherits the Global Read-Only role and can access the PCE. You can enable and disable Read-Only User access in the Global Read-Only role.

You can change a user's role membership after adding them by going to the user's details page or from a role details page.

To change an external user's name, click **Edit User** from the user's details page. You cannot edit the email address or username for an external user. You must remove and re-add the user with the new information.

To remove an external user:

1. From the PCE web console menu, choose **Access Management** > **External Users** tab.
2. Select the user you want to remove.
3. Click **Remove**.

Removing an external user removes the user from the **External Users** tab and all of the user's RBAC role memberships. However, the user's authentication is still managed by your corporate IDP.

If Read Only User access to the PCE is enabled for your organization, the user can still log into the PCE and view resources after you remove the user.

When you remove an external user while the user is online, the PCE logs the user out when they take the next action after being removed.

Add or Remove an External Group

The RBAC feature integrates with the user groups maintained in your corporate IDP so that you can manage user authentication centrally. 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 group, the PCE determines what roles and scopes are assigned. The user is granted access to the resources associated with the roles and scopes.

A user can belong to multiple external groups. When a user belongs to multiple groups, the user is granted access to Illumio resources based on the most permissive role and scopes defined for each group.

To add an external group:

1. From the PCE web console menu, choose **Access Management** > **External Groups** tab.
2. Click **Add**.
3. In the Name field, enter up to 225 alphanumeric or special characters.

4. In the External Group field, enter the group name as it is configured in your IdP.
5. 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 "User.MemberOf" attribute in the SAML assertion from your IdP.
6. The "User.MemberOf" attribute is a multiple-value attribute that contains the list of distinguished names for groups that contain the group as a member.
7. Click **Add**.

Click **Edit Group** from the group's details page to change an external group's name. You cannot edit the External Group field. You must remove and re-add the group with the new information.

To remove an external group:

1. Choose **Access Management** > **External Groups** tab.
2. Select the external group you want to remove.
3. Click **Remove**.

Removing an external group from the PCE removes all the group's RBAC role memberships and, therefore, removes access for all the group members. User authentication for the group members is still managed by your corporate IdP.

If Read Only User access to the PCE is enabled, the external group members can still log into the PCE and view resources after you remove the group. See About Roles, Scopes, and Granted Access [44] for more information.

Change Users and Groups Added to Roles

When you change a role's membership, the affected users must log out and log in again to access the new capabilities.

When you revoke a user's access to scopes or global objects while the user is online, the PCE logs the user out the next action they make after having their access revoked.

1. From the PCE web console menu, choose **Access Management** > **Global Roles**.
2. Click the name of the role you want to change users or groups in.
3. To remove one or more users or groups from the role, select them and click **Remove**.

4. To add a user or group to a role, click **Add**.
5. Click in the Add Principals field to select a user or group to add to the role. Continue to click in the field to select additional principals to add.
6. Click **Grant Access**.

Alternatively, you can select users or groups to add to roles from the various group and user details pages under **Access Management**, by clicking the desired user or group name, clicking **Add Role**, choosing **Add Global Role** or **Add Scoped Role**, and following the steps in the Access Wizard.

View User Activity

You can access a historical audit trail of user activity through the following reports:

- **User Activity:** Go to **Access Management** > **User Activity**
  - It 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 assigned to that user.

  The User Activity page also displays users who were removed and are designated as offline.

  > **NOTE**
  > The names that appears in the User Activity pages can be different from the various user and groups pages under **Access Management** > **Users and Groups** when users edit their profiles or an Organization Owner changes names at those pages.

- **Events:** Go to **Troubleshooting** > **Events**

  The Events page provides an ongoing log of all Organizational events that occur in the PCE. For example, it captures actions, such as users logging in and logging out and failed log-in attempts; when a system object is created, modified, deleted, or provisioned; and when a workload is paired or unpaired.

  Each of these events has a severity level and is exportable in JSON format. You can narrow the search for many events by event type, severity, status, timestamp, user-generated, or agent-generated filters.

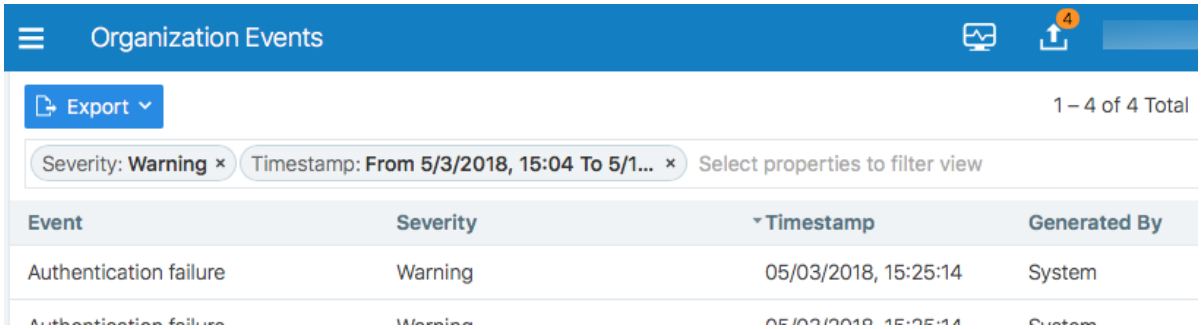

Change Your Profile Settings

If you want to change the password you use to access the PCE web console, you can do so from the User menu located at the top right corner of the console.

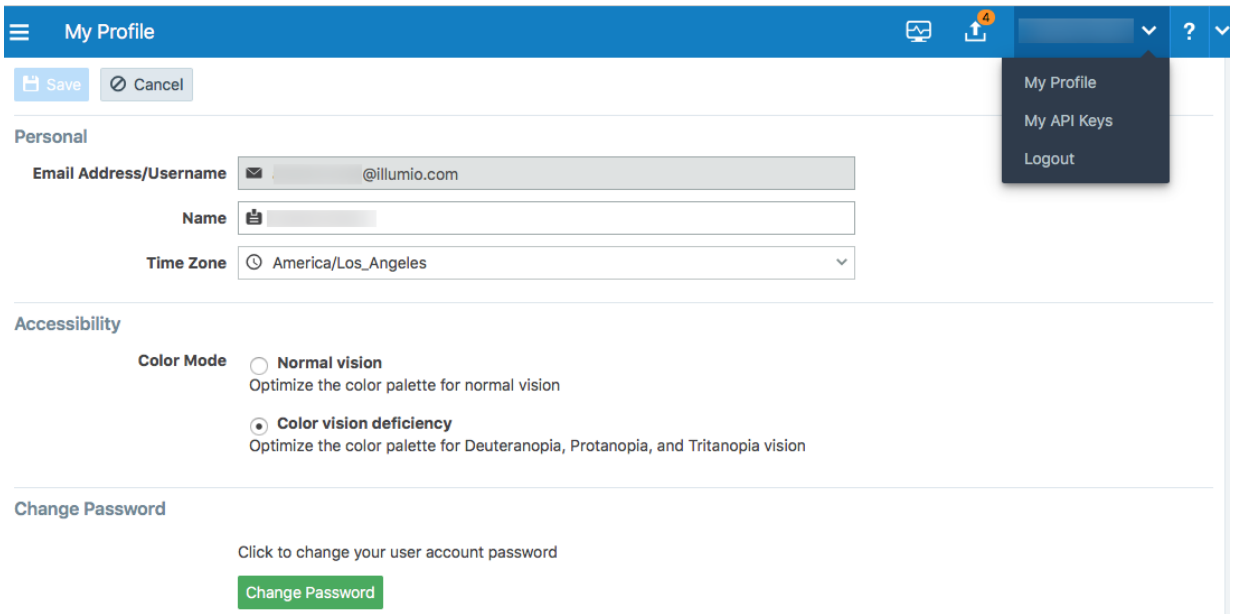

To change your password:

1. From the User menu in the PCE web console, select **My Profile**.
2. Click **Change Password**.
3. On the change password screen, enter your current password, and then your new password twice.
4. Click **Change Password**.

Color Vision Deficiency Mode

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

1. From the User menu in the PCE web console, select **My Profile**.
2. In the Accessibility section, select the **Color vision deficiency** radio button.

   NOTE
   To restore the default setting, select the **Normal vision** radio button.

3. Click **Save**.

Roles, Scopes, and Granted Access

This section describes the concepts of role-based access control (RBAC) and how it works with the PCE.

Illumio Segmentation for Data Centers 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.

Global Roles use the scope of All Applications, Environments, and Locations. You cannot change the scope for these roles. The roles have the following capabilities Illumio Segmentation for Data Centers.

Global Roles

The following table presents various roles and their granted access.

Table 1. Roles and Granted Access

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. This user can change settings in the UI to allow administrators to choose whether to show a warning for editing labels used by enforced workloads. The default is not to show a warning.
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, 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 Segmentation for Data Centers 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 using their corporate usernames and passwords. However, you haven't added all your external users to the PCE or assigned them 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 role. Users

have access to this role organization-wide because you 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. • You can add, edit, and delete rules when the source matches the specified scope. The rule destination can match any scope. **NOTE** With the Full Ruleset Manager role, you can choose the All Applications, All Environments, and All Locations scope.
Limited Ruleset Manager	• Add, edit, and delete all rulesets within the specified scope. • Add, edit, and delete rules when the source and destination 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 destinations, or have internet connectivity. **NOTE** With the Limited Ruleset Manager role, you cannot choose the All Applications, All Environments, and All Locations scope.
Ruleset Viewer	• Read-only access to rules that match the specified scope. • Ruleset Viewers cannot edit rules or rulesets.
Ruleset Provisioner	Provision rulesets within the specified scope. **NOTE** With the Ruleset Provisioner role, you can choose the All Applications, All Environments, and All Locations scope and custom scopes.
Workload Manager	Manage workloads and pairing profiles within the specified scope—read-only access provided to all other resources.

Workload Manager Role

Use Case 1

You want to use scripts in your development environment to spin up and bring down workloads programmatically; 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 workload security posture, such as changing 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 scoped; 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 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 outside their scopes, but can only modify those 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 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 source and destination 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 destination is not restricted. However, in extra-scope rules, the Environment label of the resource selected as the destination 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 destinations 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. Mixing and matching them allows you to effectively control your company's access needs.

Ruleset Only Roles

You can add users to the Full Ruleset Manager and Ruleset Provisioner roles to edit the security policies on the workloads within their assigned scopes

without affecting other entities, such as services, virtual services, or virtual servers.

These users can write rules for their workloads and provision them when the rules do not depend on global objects such as services or IP lists.

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.

These users can create rulesets within their assigned scopes and write rules not dependent on global objects. However, they can provision any workloads, including services, IP lists, and label groups.

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 policies for an enterprise.

To function as a security policy administrator, they can modify global objects such as services and labels, add workloads, pair workloads, and change workload modes.

Role Access is Additive

In the following example, Joe Smith is added to two user roles and one external group, each assigned a specific role and scope. Joe's ability to manage security for his company is a union of these roles and scopes.

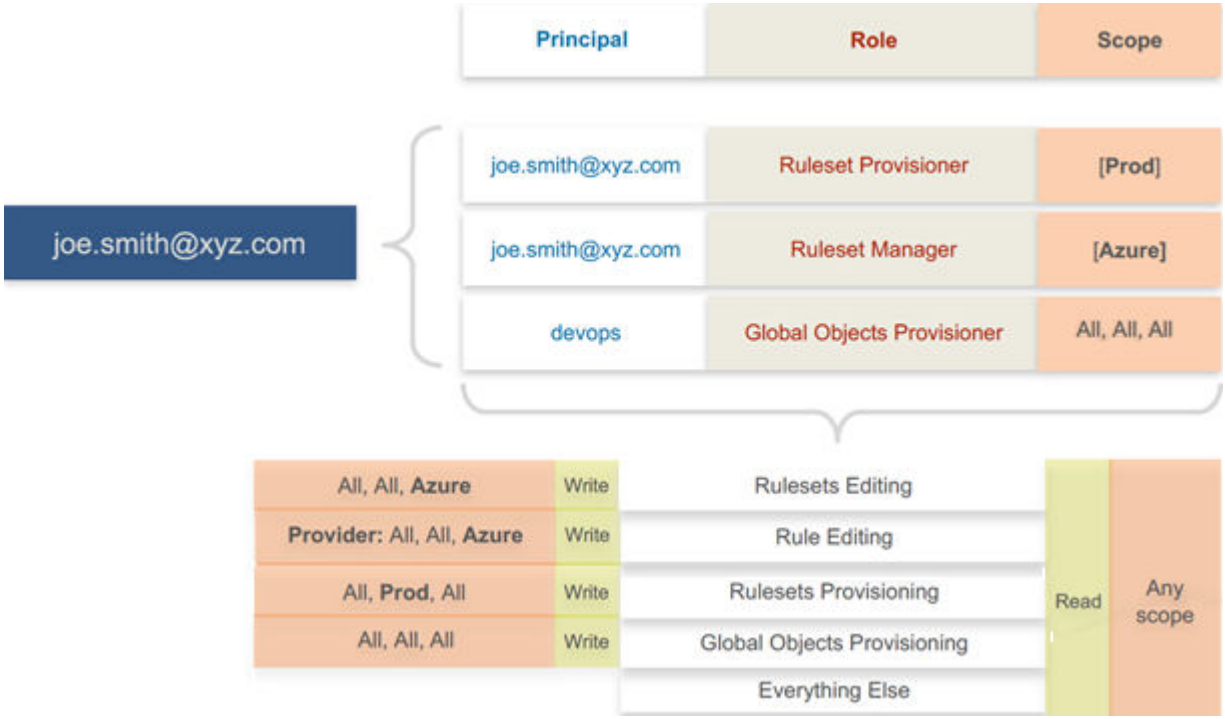

Exercise Caution when Combining Roles

Because role access is additive, 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 assign the Global Read Only role to the same user, the user will be able to view all workloads, including those outside the scope defined in the first role.

Example Role Workflows

The following example shows the handoffs between a user who is a member of the Global Organization Owner role and a member of the 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 those 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

- To manage users, roles, and scopes in the PCE, you must be a member of the Global Organization Owner role.
- Configuring SSO for an Illumio supported IdP is required for using RBAC with external users and groups. See PCE Information Needed to Configure SSO [53] for information.

  If you have not configured SSO, you can still add external users and external groups to the PCE; however, these users will not be able to log into the PCE because they cannot 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 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, delete the users or groups and re-add them to the PCE.
- An App Owner in charge of the application in both production and development environments does not have permission 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 five 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.

Configure Access Restrictions and Trusted Proxy IPs

To employ automation for managing the PCE environment, you can use API Keys created by an admin user and automate PCE management tasks. This section tells how you can restrict the use of API keys and the PCE web interface by IP address. In this way, you can block API requests and users coming in from non-allowed IP addresses.

Configure Access Restrictions

This section tells how to use the Illumio web console UI to configure access restrictions. You can also configure access restrictions programmatically using the REST API calls described in "Access Restrictions and Trusted Proxy IPs" in the REST API Developer Guide.

- You must have the global Org Owner role to view or change access restrictions.
- A maximum of 50 access restrictions can be defined.

To configure access restrictions:

1. Log in to the PCE web console as a user with the Global Org Owner role.
2. Open the menu and choose **Access Management** - **Access Restrictions**.
   The Access Restriction page opens with a list that shows which IP addresses are allowed and where the restrictions have been applied.
3. To add a new restriction, click **Add**.
   The Add Access Restriction page opens.
   Provide the required attributes:
   - Provide a name.
   - In **Restriction Applies To**, choose User Session, API Key, or Both. Access restrictions can be applied to these different types of user authentication.
   - List a maximum of eight IPv4 addresses or CIDR blocks.
4. Click **Edit** to edit the restriction.
5. View the access restrictions applied to local users. The default is blank, no restrictions.
6. You can assign access restrictions to local and external users. To add a local user:
   a. Click **Add**.
   b. In **Access Restriction**, choose the type of access restriction.
   c. Click **Add**.
7. View the local user's detail page. To modify the user settings, click **Edit User**.
8. Use the Edit User dialog to apply restrictions.
   If an Org Owner assigns an access restriction to any Org Owner, a warning is shown, because this can result in the Org Owner user losing access to the PCE.
9. View the list of API keys in the API Keys page and the Event page.

Configure Trusted Proxy IPs

This section tells how to use the Illumio web console UI to configure trusted proxy IPs. You can also configure trusted proxy IPs programmatically using the REST API calls as described in "Access Restrictions and Trusted Proxy IPs" in the REST API Developer Guide.

When a client is connected to the PCE's haproxy server, this connection can traverse one or more load balancers or proxies. Therefore, the source IP address of a client connection to haproxy might not be the actual public IP address of the client.

1. Log in to the PCE web console as a user with the Global Org Owner role.
2. 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 up to eight 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 existing addresses by hovering over the number in front of the address and then clicking the **Trash Can** icon that shows up.

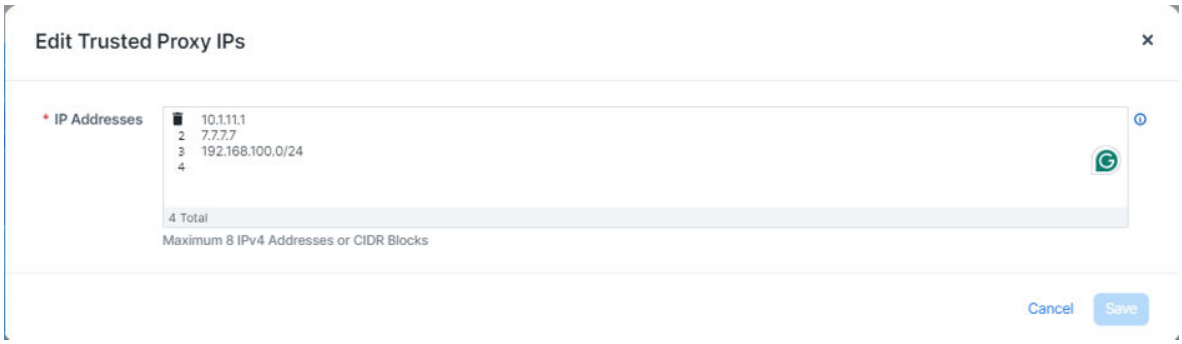

8. Once you have added or deleted the proxy addresses for your needs, click **Save**.

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.

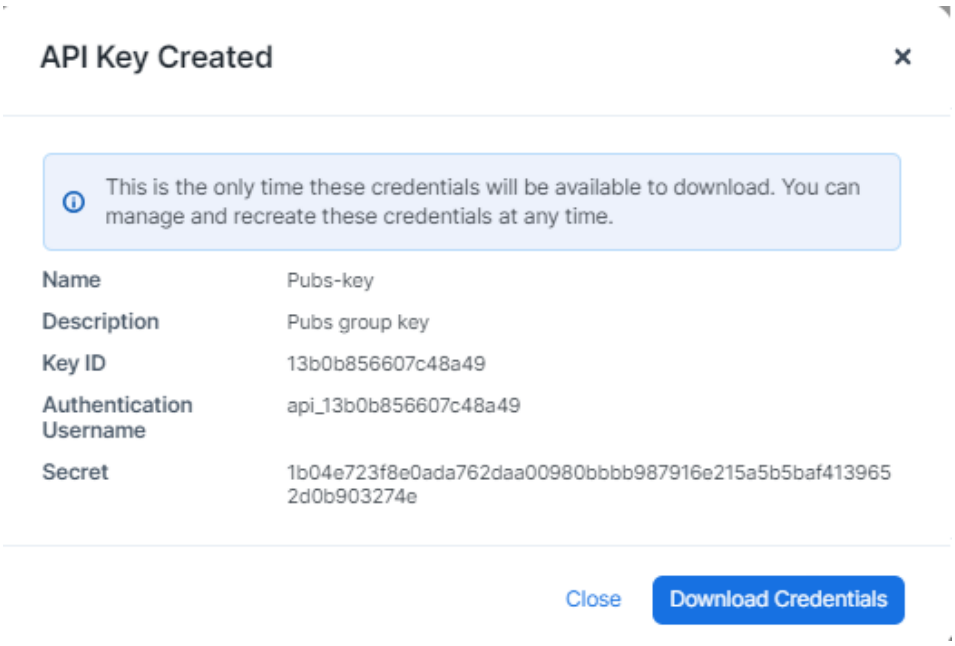

6. To download the credentials, click on **Download Credentials**.

   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":"1b04e723f8e0ada762daa00980bbbb987916e215a5b5baf4139652d0b903274e"}
```

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 the 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`.

Role-based Access for Application Owners

The enhancements made to the Role-based Access Control (RBAC) framework in the Illumio Segmentation for Data Centers 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 Segmentation for Data Centers 20.1.0 restricted "writes" based on user role and scope, users had visibility into all aspects of the PCE irrespective of the role. With these new RBAC controls, application owners get visibility into the applications within their assigned scopes, specifically the PCE information relevant to their applications. Depending on the user's role, application owners can:

- Read/write policies to manage application segmentation.
- View inbound and outbound traffic flows as well as use Explorer.
- View labeled objects used in policies.
- View details of global objects such as, IP Lists and Services used by their applications.

Benefits

The key benefits of the RBAC framework in the PCE are as follows:

- Provides a label based approach to define user permissions.
- Provides roles based on application owner personas to manage application segmentation.
- Provides a building block based approach to stack permissions for users.
- Offers flexibility to delegate read-write and read-only privileges to same user for different sets of applications.
- Enables enforcement of least privilege by hiding information outside of an application scope.

- Allows application owners to effectively manage segmentation for their applications.

Updates to Roles

As described in "About Roles, Scopes, and Granted Access", Illumio Segmentation for Data Centers provides two types of user roles - Global and Scoped. It also provides the ability to stack multiple roles for the same user. A PCE owner can assign a combination of multiple roles to the same user. The resulting set of permissions is the summation of all permissions included with each of the stacked roles. With these updates:

- Existing scoped roles enhanced to restrict reads by scope.
- New scope based *read-only* role limits read access by labels.
- Scoped users get limited visibility into objects 1-hop away (this applies to Explorer, App Group Maps, Rule Search, and Traffic).
- Global read-only disabled by default for new PCE installations.
- PCE performance and scale enhanced to support concurrently active users.

Global Roles

Global roles provide the user with permissions to view everything and to perform operations globally. The four Global roles are :

- Global Organization Owner: Allowed to manage all aspects of the PCE, including user management.
- Global Administrator: Allowed to managed most aspects of the PCE, with the exception of user management.
- Global Viewer: Allowed to view everything within the PCE in a read-only capacity. This role was previously called "Global Read-only".
- Global Policy Object Provisioner: Allowed to provision global objects that require provisioning such as, Services and Label Groups.

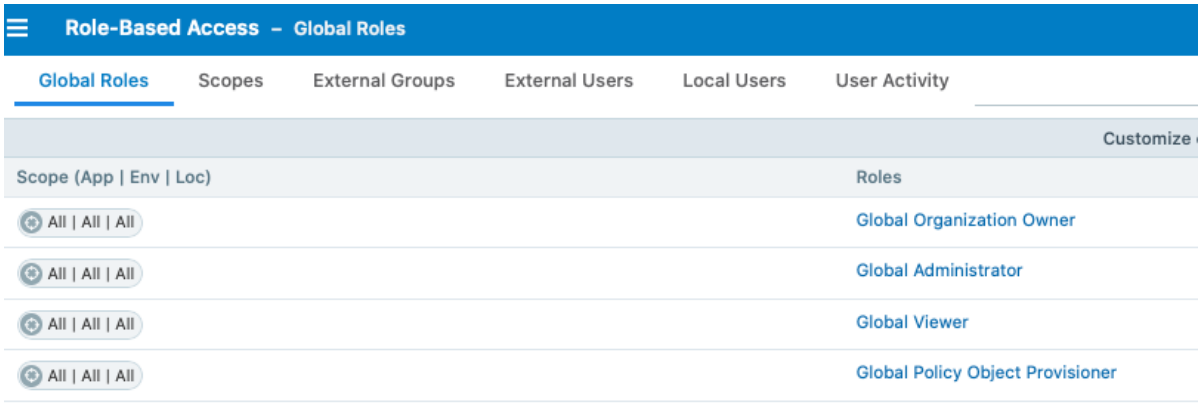

Scoped Roles

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 Application, Environment, and Location labels. 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 four 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: An existing scope-based read/write role. A user with this role can read/write policy within the assigned scope. The user can also view application groups, incoming and outgoing traffic, and labeled objects, within the assigned scope.
- Ruleset Provisioner: This role allows a user to provision changes to the scoped objects, provided the objects are inside the user's assigned scope. A user with this role can provision changes to policies within the assigned scope. The user can also view application groups, incoming and outgoing traffic, and labeled objects, within the assigned scope.
- Workload Manager: Allows a user to perform workload-specific operations such as, pairing, unpairing, assignment of labels, and changing of policy state. A user with this role cannot view policies and traffic, and cannot provision changes.

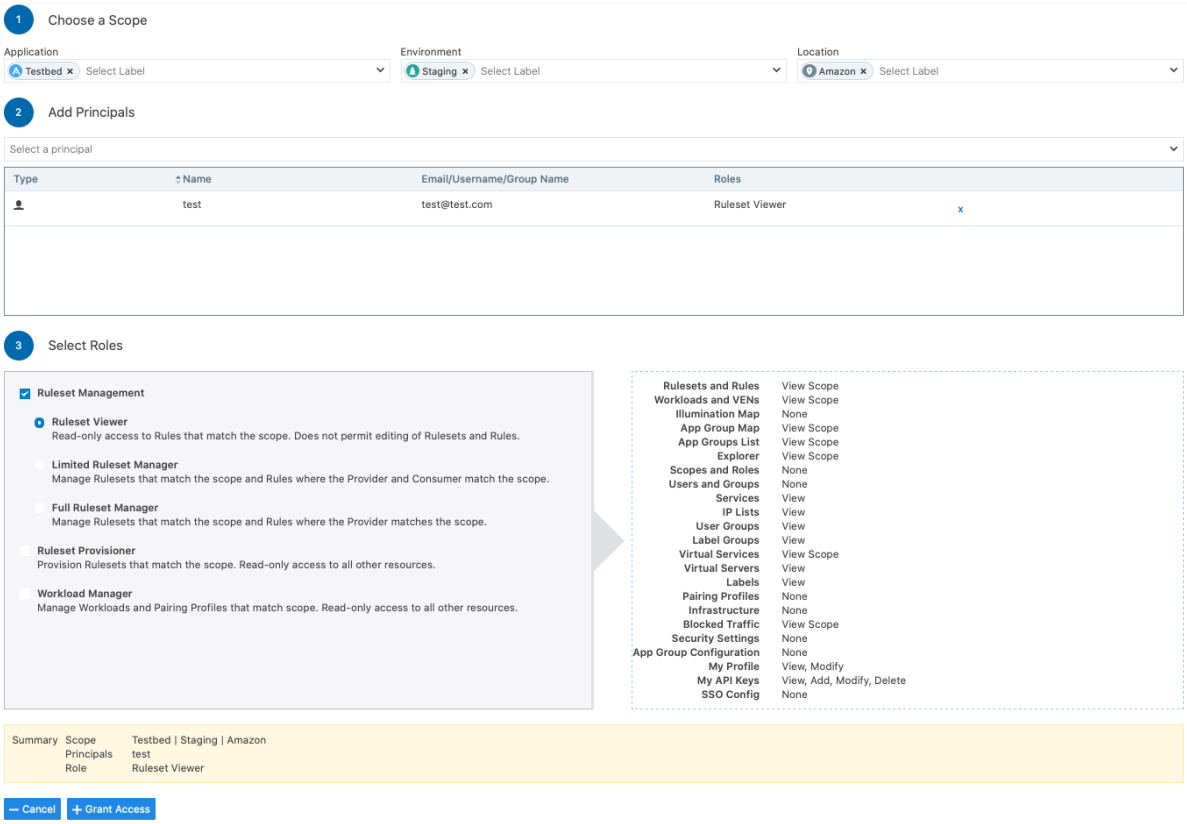

Configuration

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 Role-based Access > Global Roles > Global Viewer is set to **Off**.

> **NOTE**
> The Global Read-only user setting is disabled by default.

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

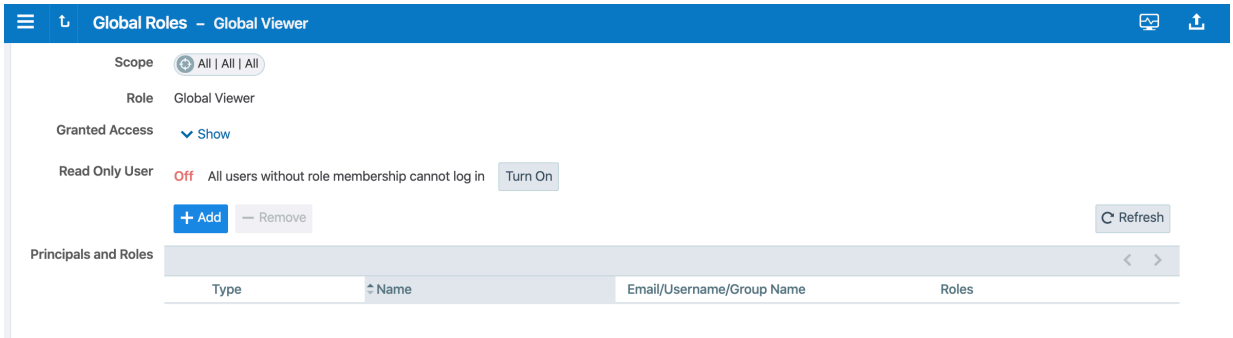

Facet Searches for Scoped Roles

The Scopes page now features a search bar with auto-complete and facets. This is restricted to users with a Global Organization Owner role. To use this feature, navigate to Role-based Access > Scopes. The search bar allows Organization Owners to query a list of users by a user's role. They can search by label type, such as Application, Environment, or Location, to get a list of users with the selected label(s) in their assigned scope(s). They can also select Principals to search for a specific user.

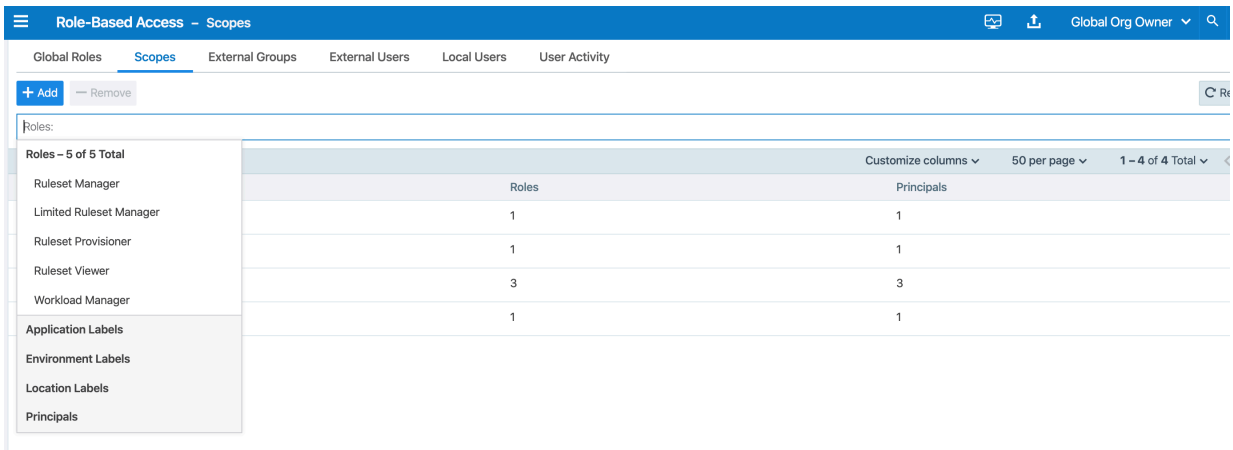

Ruleset Viewer

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 Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
Traffic - Illumination, App Group, Explorer					
Illumina-tion Lo-cation Map	---	---	---	---	---
App Group Policy Map	Read (R)	Read (R)	Read (R)	---	Read (R)
App Group Vulnera-bility 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 Genera-tor	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)
Provi-sioning Draft Changes	Read (R)	Read (R)	Read+Write (R)	---	Read+Write (R)
Policy Versions	Read (R)	Read (R)	Read (R)	---	Read (R)

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
Provi-sioning Status	Read (R)	Read (R)	Read (R)	---	Read (R)
Labeled Objects					
Work-loads	Read (R)	Read (R)	Read (R)	Read+Write (R)	Read+Write (R)
Contain-er Work-loads	Read (R)	Read (R)	Read (R)	Read (R)	Read (R)
Virtual Enforce-ment No-des	Read (R)	Read (R)	Read (R)	Read+Write (R)	Read+Write (R)
Pairing Profiles	---	---	---	Read+Write (R)	Read+Write (R)
Virtual Services	Read (R)	Read (R)	Read (R)	Read (R)	Read (R)
Virtual Servers	Read	Read	Read	Read	Read
Global Policy Objects					
Services	Read	Read	Read	Read	Read
IP Lists	Read	Read	Read	Read	Read
User Groups	Read	Read	Read	Read	Read
Labels	Read	Read	Read	Read	Read
Label Groups	Read	Read	Read	Read	Read
Settings					
Segmen-tation Tem-plates	---	---	---	---	---
Role-Based Access	---	---	---	---	---

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
Global Roles					
Role-Based Access Scoped Roles	---	---	---	---	---
Role-Based Access Users and Groups	---	---	---	---	---
Role-Based Access User Ac-tivity	---	---	---	---	---
Load Bal-ancers	---	---	---	---	---
Contain-er Clus-ters	---	---	---	---	---
Bi-direc-tional 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 Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
VEN Li-brary	---	---	---	Read	Read
My Pro-file	Read+Write	Read+Write	Read+Write	Read+Write	Read+Write
My API Keys	Read+Write	Read+Write	Read+Write	Read+Write	Read+Write
Other					
Support Reports	---	---	---	Read+Write (R)	Read+Write (R)
Events	---	---	---	---	---
Reports	Read (R, T)	Read (R, T)	Read (R, T)	Read (R, T)	Read (R)
Support	Read	Read	Read	Read	Read
PCE Health	---	---	---	---	---
Product Version	Read	Read	Read	Read	Read
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 Il-lumio	Read	Read	Read	Read	Read

Scoped Users and PCE

Each scoped role has different permissions that impact an application own-er'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 options available for Application Owner are limited. For example, a user 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 users of different scopes.

- 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 Provi-sioner	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 Pro-files	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 Re-ports	N/A	N/A	N/A	Y
Settings	N/A	N/A	N/A	See row be-low
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 Mem-bers are visi-ble)
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 > Con-tainer Workloads	Y	Y	Y	Y
Workload Management > Virtual Enforcement Nodes (Agents)	Y	Y	Y	Y

Page	Ruleset Viewer	Ruleset Manager	Ruleset Provi-sioner	Workload Manager
Provision > Draft Changes	Y	Y	Y	N/A
Provision > Policy Versions	Y	Y	Y	N/A
Policy Objects > IP Lists	Y	Y	Y	Y
Policy Objects > Services	Y	Y	Y	Y
Policy Objects > Labels	Y	Y	Y	Y
Policy Objects > User Groups	Y	Y	Y	Y
Policy Objects > Label Groups	Y	Y	Y	Y
Policy Objects > Virtual Serv-ices	Y	Y	Y	Y
Policy Objects > Virtual Serv-ers	Y	Y	Y	Y
Troubleshooting > Blocked Traffic	Y	Y	Y	N/A
Troubleshooting > Export Re-ports	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.

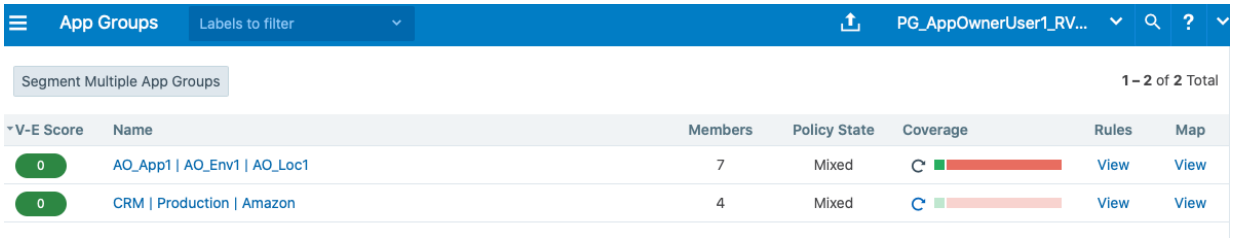

Labeled Objects

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 re-

striction 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.

App Group Map

The App Group Map provides complete visibility into applications and everything inside the application. Scoped users, with the exception of Workload Managers, can view App Group Maps. Scoped users get complete visibility into everything inside their application group. Scoped users can see workload objects, labels, traffic flows, and every other detail within their application group.

Scoped users get limited visibility for connected App Groups such as Providing App Groups and Consuming App Groups. Scoped users get limited information on endpoints with traffic flows to their application. For an endpoint in a connected App Group from which there is traffic flow, scoped users can get limited information such as labels, role names, and hostnames. The scoped user is not allowed to view any other endpoints in the connected App Group from which there are no traffic flows.

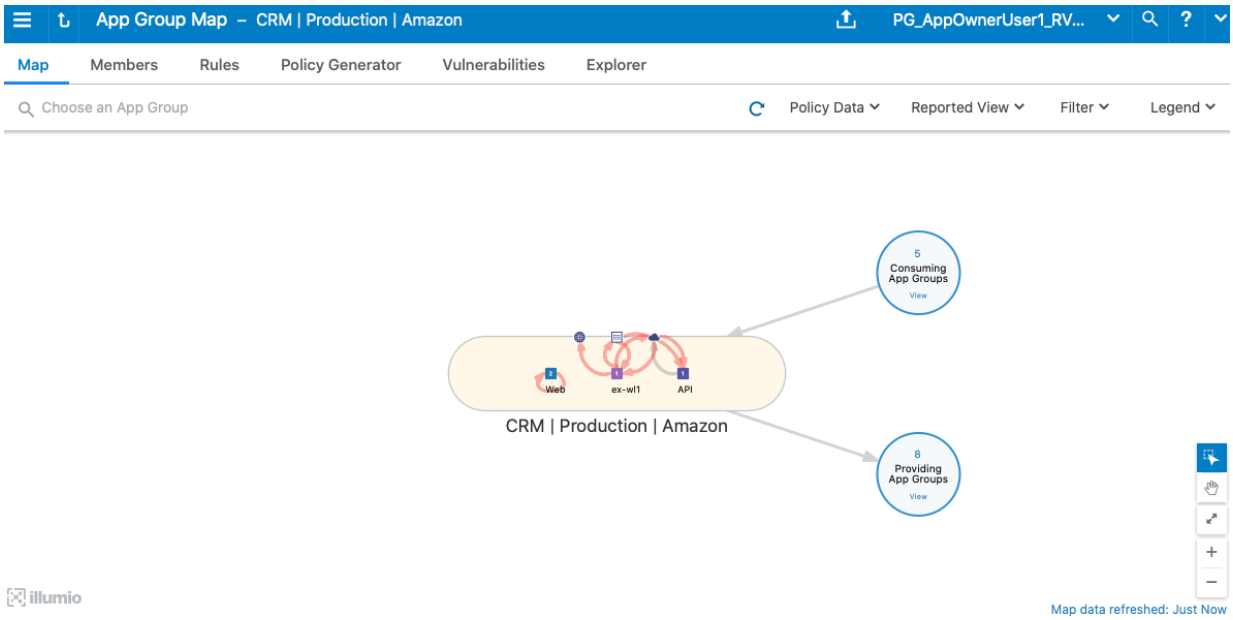

Policy Generator and Explorer

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).

Labeled Objects

Labeled objects, such as workloads are filtered by the scope of the user. On the Workloads page, you will only see the list of the workloads within the application scope. You cannot see any workloads that are outside the application scope. This applies to any labeled object, such as workloads, containers, Virtual Services, and Virtual Enforcement Nodes (VENs).

The menu functions and buttons change dynamically to reflect a user's permissions. If you are logged in as a Ruleset Manager, you are not allowed to manage workloads. So, all the workload-specific operations buttons are disabled. However, you are allowed to view the list of workloads within the scope and get details for individual workloads, except for Virtual Servers.

NOTE

While Virtual Servers are considered labeled objects, they are visible to all scoped users regardless of object scope.

Facet Searches and Auto-complete

The search bar with auto-complete and facets is scoped for labeled objects and policies. For example, you search for Application Labels, then you can only select the Application Labels under the assigned scope. This applies to other label types such as Environment labels and Location labels. However, Role labels are excluded since Role labels are not part of the user scope. The restriction of visibility by scope applies to facets such as hostname, IP address, and others. The search bar automatically filters the facets to the list of facets in the user's assigned scope.

Global Objects

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

NOTE

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

Policies

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

When label groups are used within the scope of a policy, a Ruleset Manager may not be allowed to edit the policy and its rules even if there is a scope match between the user's assigned scope and the underlying scope of the policy. The user will, however, be able to view the rules within such a policy.

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

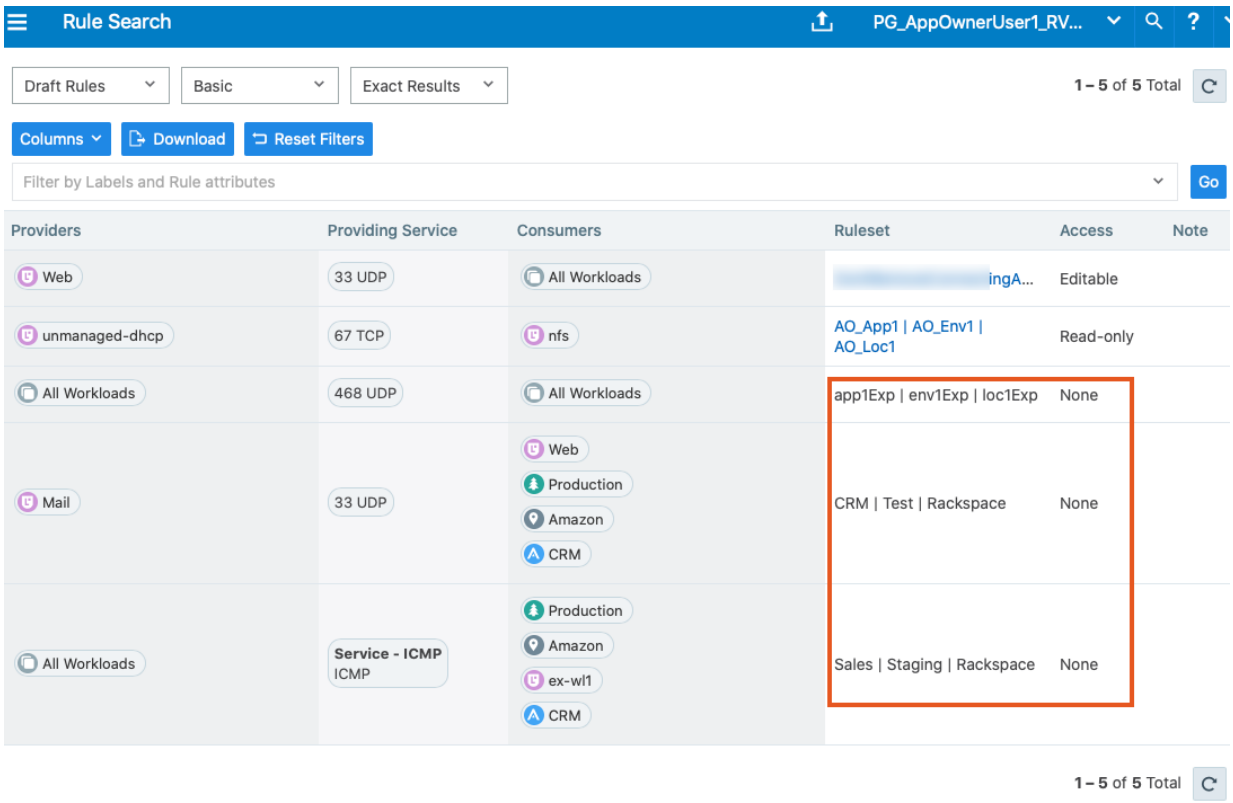

Ruleset Viewer

Ruleset Viewer is a new scope-based read-only role. When assigned, a user gets read-only visibility into the assigned application scope. As a Ruleset Viewer, you can view all the policies 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 Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
Traffic - Illumination, App Group, Explorer					
Illumina-tion Lo-cation Map	---	---	---	---	---
App Group Policy Map	Read (R)	Read (R)	Read (R)	---	Read (R)
App Group Vulnera-bility 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 Genera-tor	Read (R)	Read+Write (R)	Read (R)	---	Read+Write (R)
Policies	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)
Provi-sioning Draft Changes	Read (R)	Read (R)	Read+Write (R)	---	Read+Write (R)
Policy Versions	Read (R)	Read (R)	Read (R)	---	Read (R)

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
Provi-sioning Status	Read (R)	Read (R)	Read (R)	---	Read (R)
Labeled Objects					
Work-loads	Read (R)	Read (R)	Read (R)	Read+Write (R)	Read+Write (R)
Contain-er Work-loads	Read (R)	Read (R)	Read (R)	Read (R)	Read (R)
Virtual Enforce-ment No-des	Read (R)	Read (R)	Read (R)	Read+Write (R)	Read+Write (R)
Pairing Profiles	---	---	---	Read+Write (R)	Read+Write (R)
Virtual Services	Read (R)	Read (R)	Read (R)	Read (R)	Read (R)
Virtual Servers	Read	Read	Read	Read	Read
Global Policy Objects					
Services	Read	Read	Read	Read	Read
IP Lists	Read	Read	Read	Read	Read
User Groups	Read	Read	Read	Read	Read
Labels	Read	Read	Read	Read	Read
Label Groups	Read	Read	Read	Read	Read
Settings					
Segmen-tation Tem-plates	---	---	---	---	---
Role-Based Access	---	---	---	---	---

Page	Ruleset Viewer (Scoped Read-Only)	Ruleset Manager	Ruleset Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
Global Roles					
Role-Based Access Scoped Roles	---	---	---	---	---
Role-Based Access Users and Groups	---	---	---	---	---
Role-Based Access User Ac-tivity	---	---	---	---	---
Load Bal-ancers	---	---	---	---	---
Contain-er Clus-ters	---	---	---	---	---
Bi-direc-tional 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 Provision-er	Workload Manager	Application Owner (Combined Permis-sions)
VEN Li-brary	---	---	---	Read	Read
My Pro-file	Read+Write	Read+Write	Read+Write	Read+Write	Read+Write
My API Keys	Read+Write	Read+Write	Read+Write	Read+Write	Read+Write
Other					
Support Reports	---	---	---	Read+Write (R)	Read+Write (R)
Events	---	---	---	---	---
Reports	Read (R, T)	Read (R, T)	Read (R, T)	Read (R, T)	Read (R)
Support	Read	Read	Read	Read	Read
PCE Health	---	---	---	---	---
Product Version	Read	Read	Read	Read	Read
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 Il-lumio	Read	Read	Read	Read	Read

Scoped Users and PCE

Each scoped role has different permissions that impact an application own-er'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 Provi-sioner	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 Pro-files	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 be-low
Support > 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 Mem-bers are visi-ble)
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
Policies	Y	Y	Y	N/A
Rule Search	Y	Y	Y	N/A
Workload Management > Workloads	Y	Y	Y	Y
Workload Management > Con-tainer Workloads	Y	Y	Y	Y
Workload Management > Virtual Enforcement Nodes (Agents)	Y	Y	Y	Y

Page	Ruleset Viewer	Ruleset Manager	Ruleset Provi-sioner	Workload Manager
Provision > Draft Changes	Y	Y	Y	N/A
Provision > Policy Versions	Y	Y	Y	N/A
Policy Objects > IP Lists	Y	Y	Y	Y
Policy Objects > Services	Y	Y	Y	Y
Policy Objects > Labels	Y	Y	Y	Y
Policy Objects > User Groups	Y	Y	Y	Y
Policy Objects > Label Groups	Y	Y	Y	Y
Policy Objects > Virtual Serv-ices	Y	Y	Y	Y
Policy Objects > Virtual Serv-ers	Y	Y	Y	Y
Troubleshooting > Blocked Traffic	Y	Y	Y	N/A
Troubleshooting > Export Re-ports	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

Azure AD Single Sign-on

This topic describes how to configure Azure Active Directory (AD) to provide SSO authentication to the Illumio PCE.

> **TIP**
> Because you'll configure settings in both the Illumio PCE Web Console and in Azure AD, have both applications open in adjacent browser tabs.

Prerequisites

To perform this configuration, you need the following:

- An Azure AD subscription. If you don't have a subscription, you can get a free account.
- An Illumio single sign-on (SSO) enabled subscription.

STEP 1: Obtain URLs from the Illumio PCE Web Console

In this step you'll copy and preserve URLs from the Illumio PCE for use in Step 2.

1. Log in to the PCE as a Global Organization Owner.
2. Go to **Access Management > Authentication**.
3. On the **SAML** tile, click **Configure**.
4. Copy and preserve the following URLs needed to complete the Azure configuration in a later step:

   > **TIP**
   > Make sure to replace the x's in the URLs below with the actual values from your implementation.

   - **Issuer:** https://PCE.xxxx:8443/login
   - **NameID Format:** urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress

- **Assertion Source URL:** https://PCE.xxxx:8443/login/acs/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
- **Logout URL:** https://pce.xxxx:8443/login/logout/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx

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.

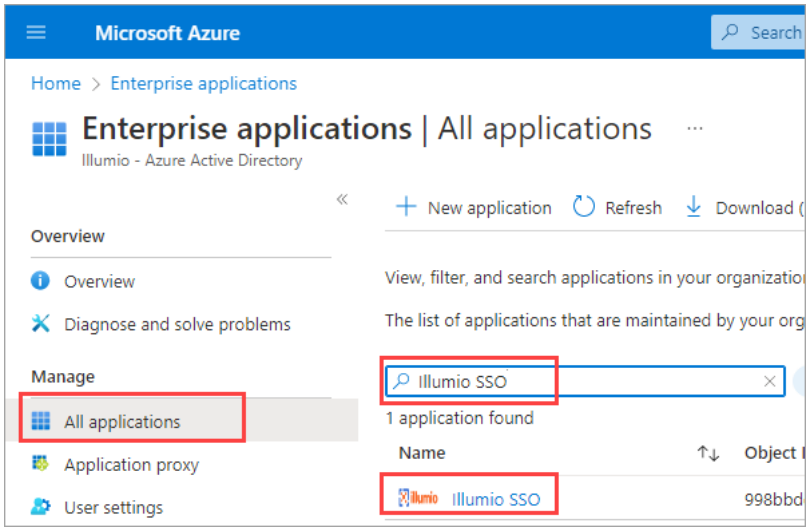

4. In the center of the page under **Getting Started**, click **Get started** on the **Set up single sign on** tile.

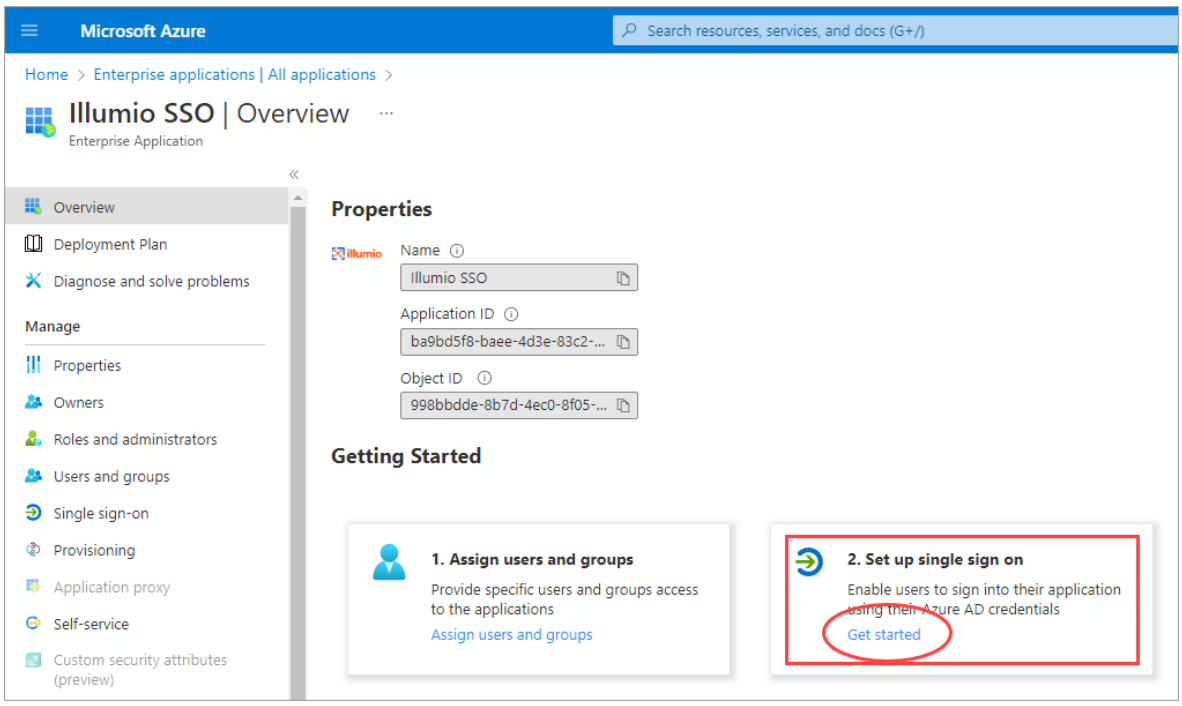

5. If prompted to select a single sign-on method, click **SAML**.
6. Configure Basic SAML:
   a. On the **Set up Single-Sign On with SAML** page **Basic SAML Configuration** tile, click **Edit**.

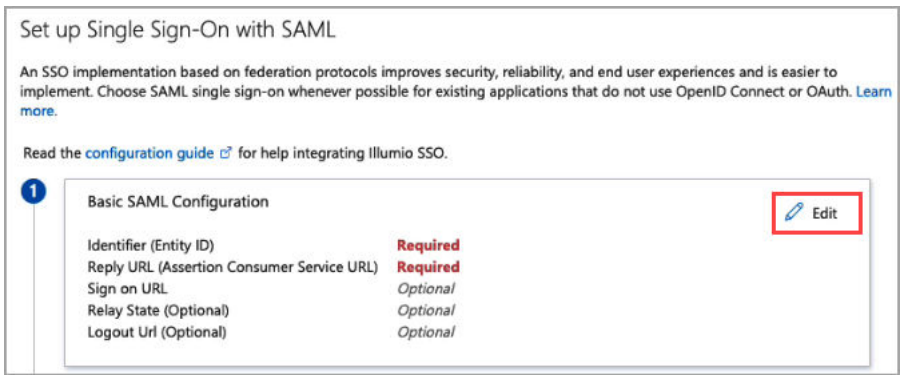

   b. On the **Basic SAML Configuration** panel that opens, populate the fields with the values you copied and preserved.
      - In the **Identifier (Entity ID)** field, paste the **Issuer URL** you copied from the Illumio PCE.
      - In the **Reply URL (Assertion source Service URL** field, click **Add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

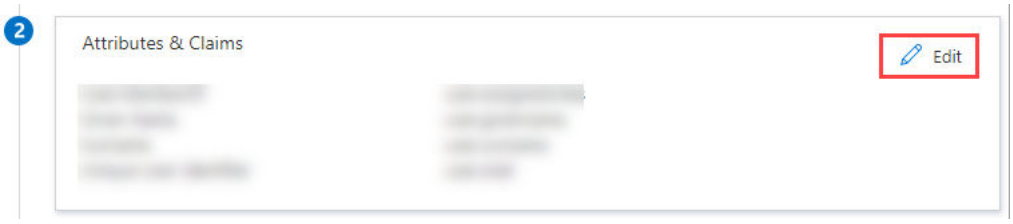

8. Under **Required claim**, update the **Claim name**:

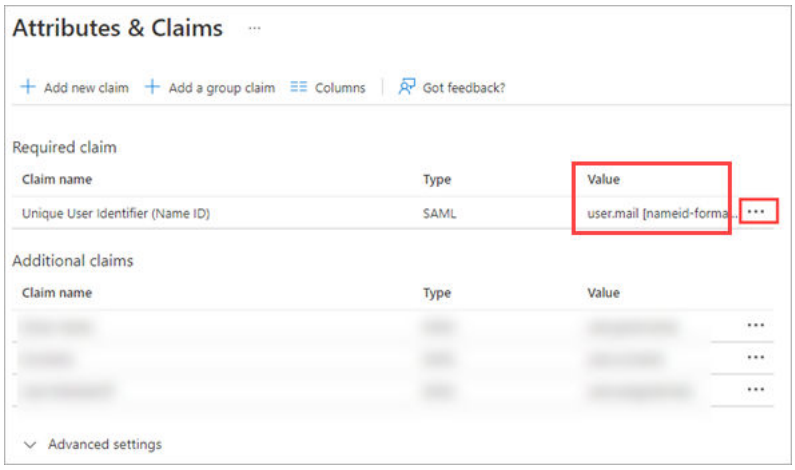

   a. Click the three dots.
   b. On the **Manage claim** page, click in the **Source attribute** field and select **user.mail** from the dropdown.
   c. Click **Save**.
9. Back on the **Attributes & Claims** page, delete **all** of the existing claims in the **Additional claims** section by clicking the three dots for each one and then clicking **Delete**.

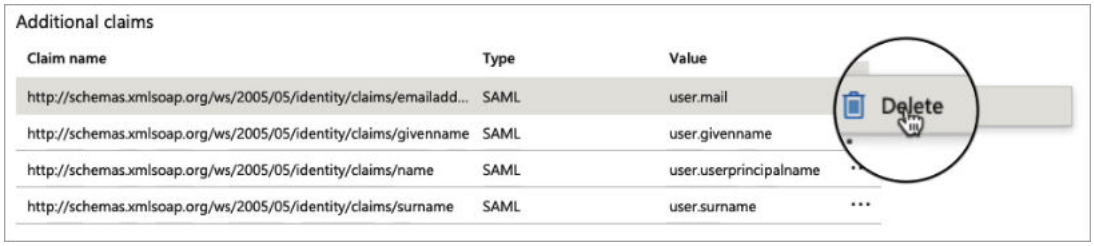

10 Click **Add new claim** and add three new claims:

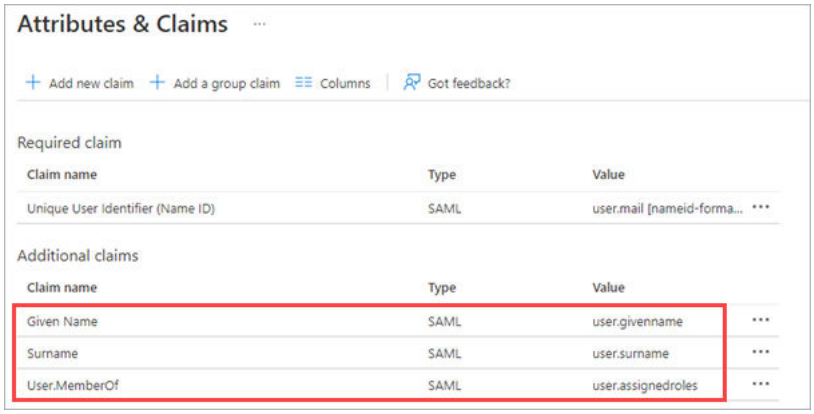

Given Name

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.

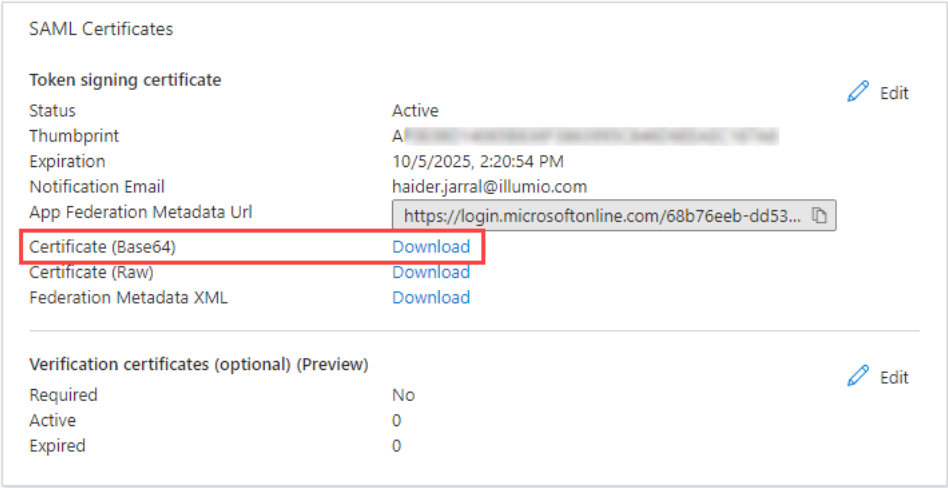

2. On the **Set up Illumio SSO** tile, copy and preserve the following URLs that you'll later paste into the Illumio PCE SAML setup.

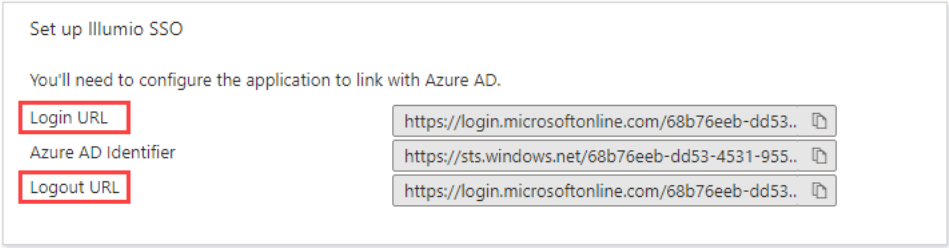

STEP 4: Configure SAML SSO settings in the Illumio PCE

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 you downloaded and copy and paste the contents.
   - **Remote Login URL**: Paste the Login URL you copied from Azure AD.
   - **Logout Landing URL**: Paste the Logout URL you copied from Azure AD.
5. In the **Information for Identity Destination** section:
   a. Choose an authentication method:
      - **Unspecified** uses the IdP default authentication mechanism.
      - **Password Protected Transport** requires the user to log in with a password in a protected session.
   b. If you want to require users to re-enter login credentials to access Illumio (even if the session is still valid), select **Force Re-authentication**. This allows users to log in to the PCE using login credentials different from their default computer login credentials.
6. Click **Save**.

STEP 5: Create App Roles in Azure AD

In this step you'll create app roles in Azure AD that you'll map to roles in the Illumio PCE Web Console in "Add External Groups and assign roles in the PCE Web Console".

For reference in this step, here's a list of the Global Roles available in the PCE Web Console:

- Global Organization Owner
- Global Administrator
- Global Viewer
- Globally Policy Object Provisioner

1. In Azure AD, go to **Users and Groups** and then click **application registration**.

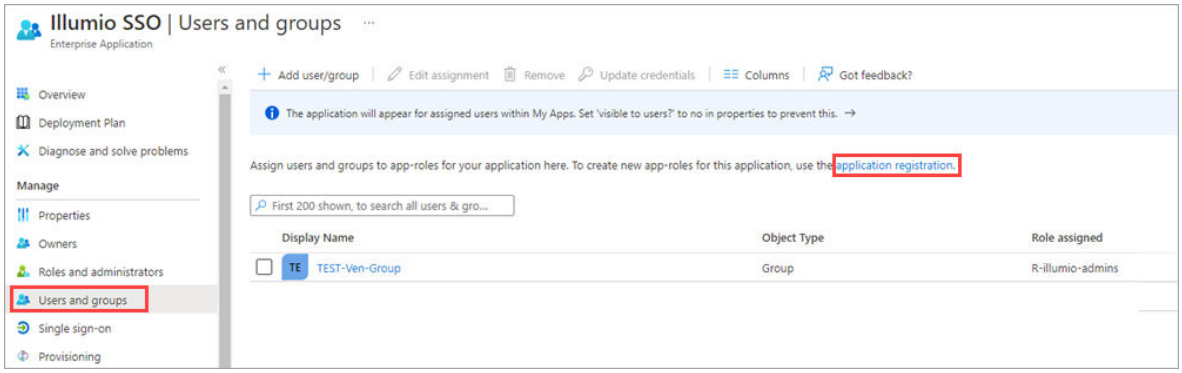

2. Create the roles you want by clicking **+ Create app role** and entering the required information for each role:
   - **Display name:** For example, enter one of the Global Roles that appear in the PCE Web Console.
   - **Value:** This must match the name you'll enter in the **Add External Groups** dialog box in "Add External Groups and assign roles in the PCE Web Console".
   - **Description:** The description will appear as help text in the app assignment and consent experiences.
3. Click **Apply** for each role that you create.
4. Delete the default app role **msiam_access**. Note that you must disable the default app role before you can delete it.
   a. Click **msiam_access** to open the **Edit app role** panel.
   b. Deselect **Do you want to enable the app role?**
   c. Click **Apply**. The side panel closes.
   d. Click **msiam_access** again to to open the **Edit app role** panel again.
   e. Click **Delete**.

When you're done creating roles in Azure AD, the **App roles** section should look similar to this:

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 in "Create App Roles in Azure AD".

1. In Azure AD, go to **Users and groups**.
2. Select the Illumio SSO app.
3. Click **Remove** to remove the current app assignments.
4. Click **Yes** to confirm removal.
5. Click **Add user/group**.

6. On the **Add Assignment** page, assign desired role(s) to users or groups:
   a. Under **User and groups**, click **None Selected**.
   b. In the **Users and groups** panel that opens, search for your desired user/group, click to select it, and then click **Select** at the bottom of the panel.
   c. Back on the **Add Assignment** page, under **Select a role***, click **None Selected**.
   d. In the **Select a role** panel that opens, find and click the role you want to assign, and then click **Select** at the bottom of the panel.
   e. Back on the **Add Assignment** page, click **Assign** at the bottom of the page.
   f. Repeat these sub-steps for each user and/or to which you want to assign app roles.

STEP 7: Add External Groups and assign roles in the PCE Web Console

In this step, you'll add external groups in the PCE Web Console and assign them the relevant global or scoped roles in Illumio RBAC.

TIP

Alternatively, you can add individual users by going to the **External Users** tab and following the onscreen prompts.

1. On the PCE Web Console, go to **Access Management > External Groups**.
2. Click **Add**.
3. In the **Add External Group** dialog box:
   - Enter a **Name**.
   - Enter an **External Group**.

     IMPORTANT

     This must match the **Value** that you specified for the app role inn"Create App Roles in Azure AD".

   - Click **Add**.
4. Repeat for additional groups.

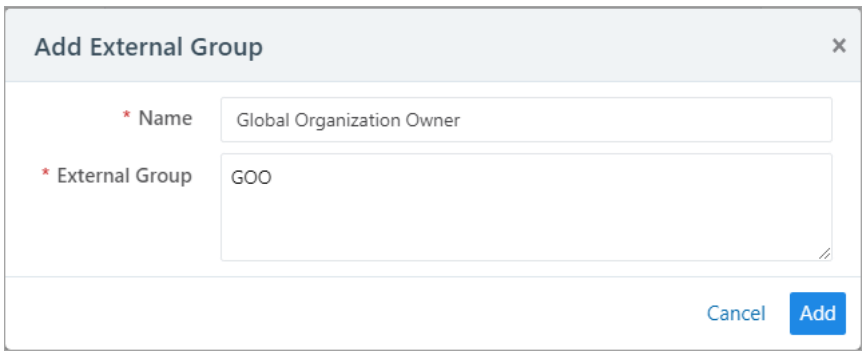

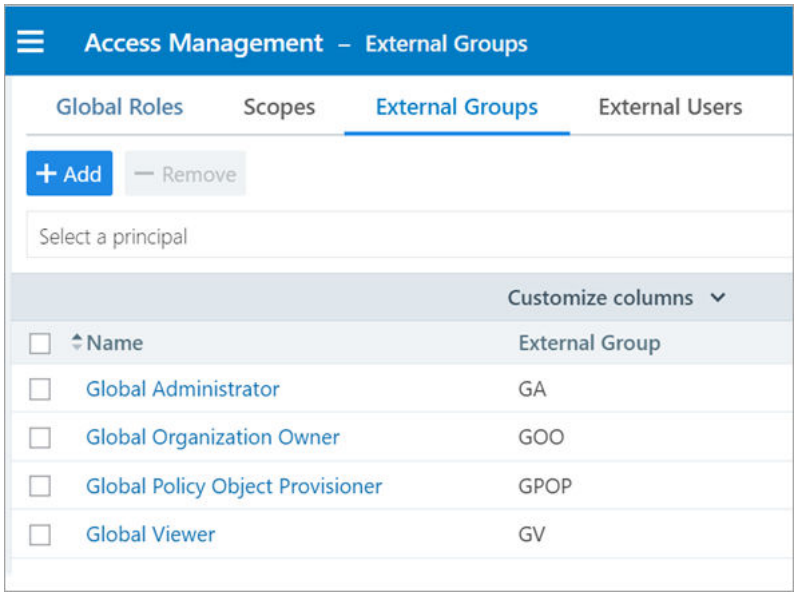

5. Click to open a group you created in the above step.
6. Click **Add Role > Add Global Role** or **Add Scoped Role**.
7. In the **Access Wizard**, select the appropriate **Role** and then click **Grant Access**.
8. Repeat for additional groups.

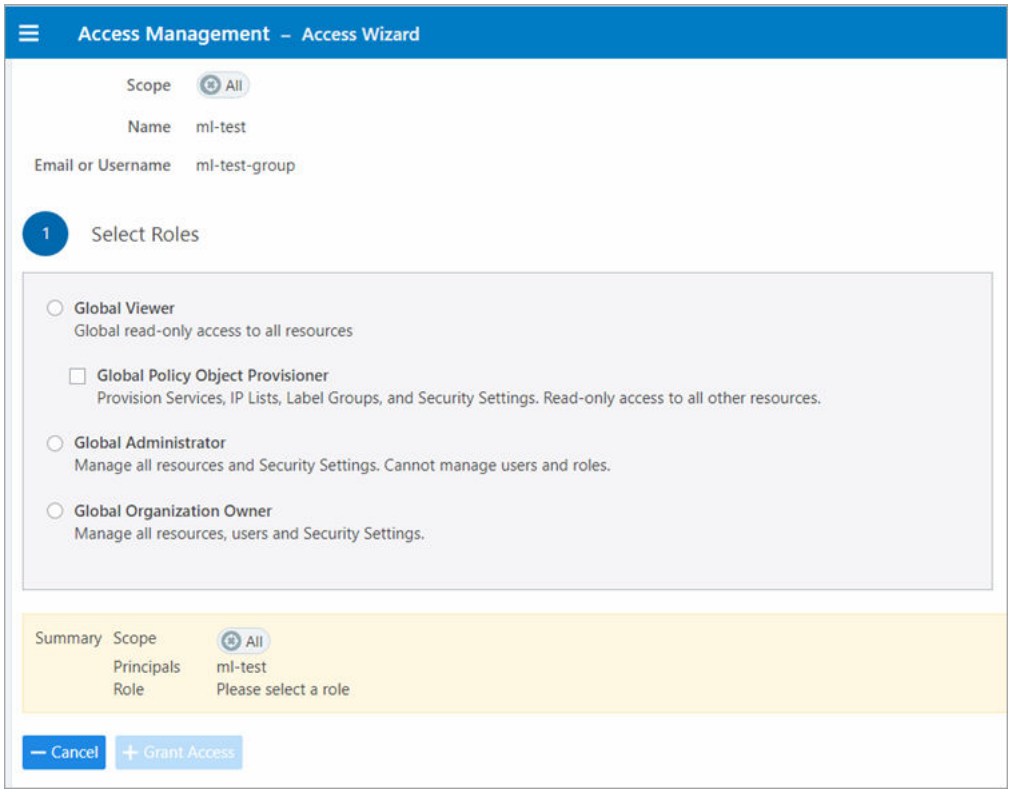

STEP 8: Turn on SAML authentication in the PCE Web Console

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**.

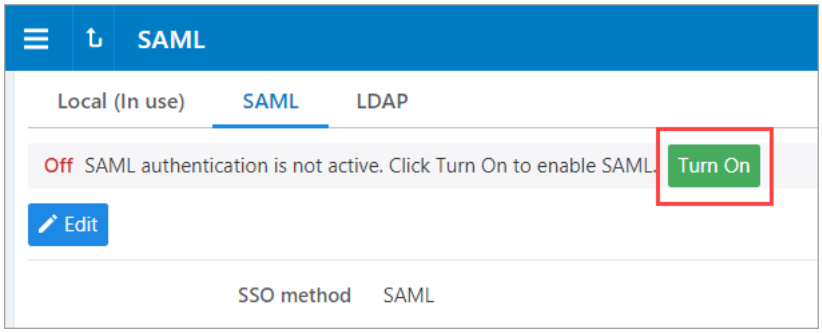

STEP 9: Test SSO

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**.

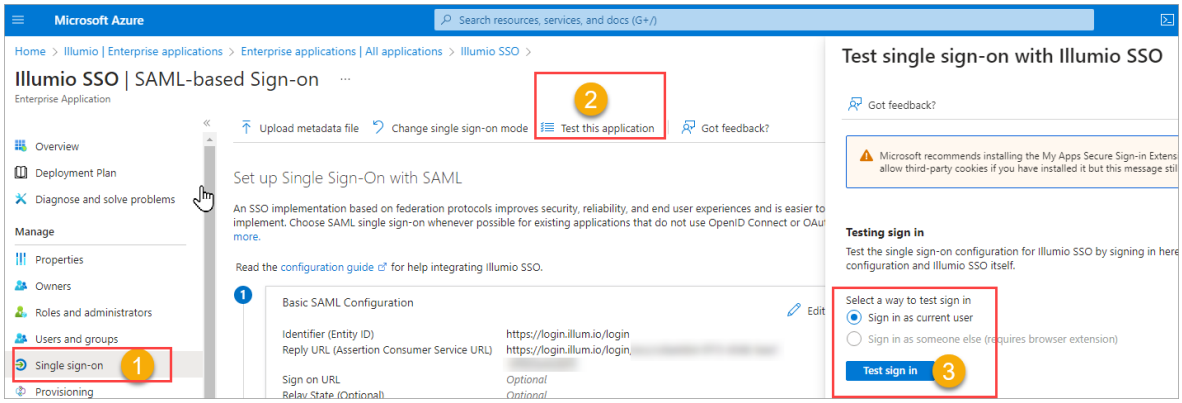

4. If the test is successful, the PCE will log you in to the **Welcome to Illumio** screen.

Prerequisites

To perform this configuration, you need the following:

- An Azure AD subscription. If you don't have a subscription, you can get a free account.
- An Illumio single sign-on (SSO) enabled subscription.

STEP 1: Obtain URLs from the Illumio PCE Web Console

In this step you'll copy and preserve URLs from the Illumio PCE for use in Step 2.

1. Log in to the PCE as a Global Organization Owner.
2. Go to **Access Management > Authentication**.
3. On the **SAML** tile, click **Configure**.
4. Copy and preserve the following URLs needed to complete the Azure configuration in a later step:

> **TIP**
> Make sure to replace the x's in the URLs below with the actual values from your implementation.

- **Issuer:** https://PCE.xxxx:8443/login

- **NameID Format:** urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress
- **Assertion Source URL:** https://PCE.xxxx:8443/login/acs/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
- **Logout URL:** https://pce.xxxx:8443/login/logout/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx

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.

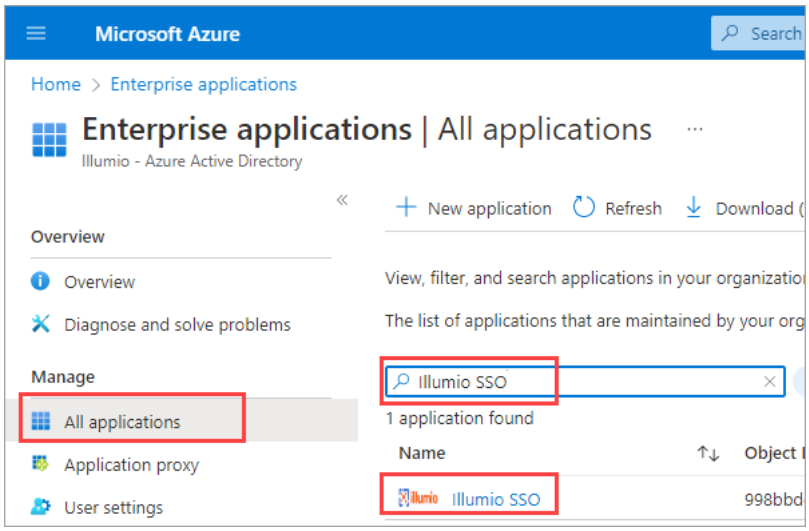

4. In the center of the page under **Getting Started**, click **Get started** on the **Set up single sign on** tile.

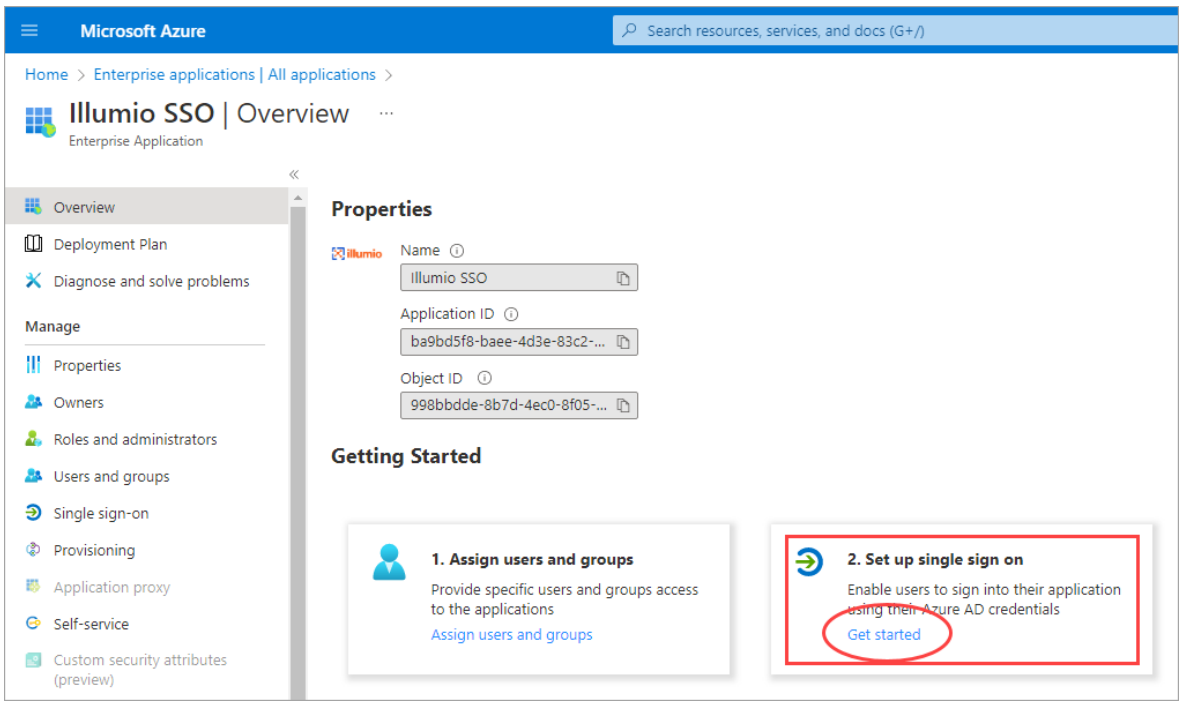

5. If prompted to select a single sign-on method, click **SAML**.
6. Configure Basic SAML:
   a. On the **Set up Single-Sign On with SAML** page **Basic SAML Configuration** tile, click **Edit**.

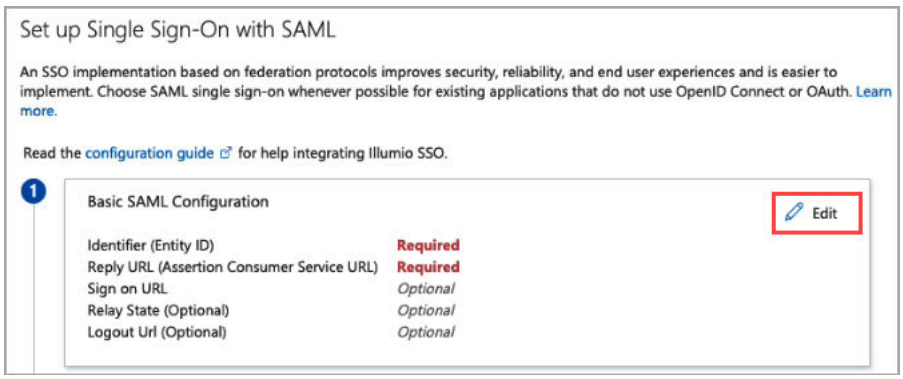

   b. On the **Basic SAML Configuration** panel that opens, populate the fields with the values you copied and preserved.
      - In the **Identifier (Entity ID)** field, paste the **Issuer URL** you copied from the Illumio PCE.
      - In the **Reply URL (Assertion source Service URL** field, click **Add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

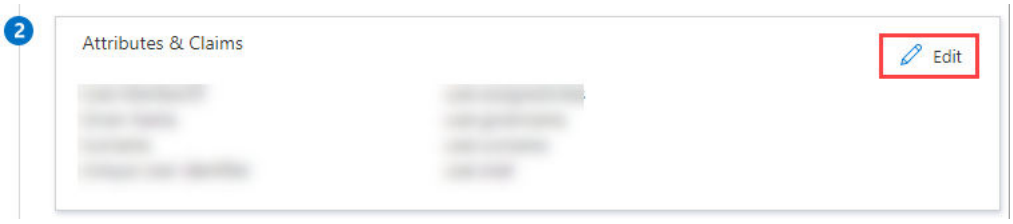

8. Under **Required claim**, update the **Claim name**:

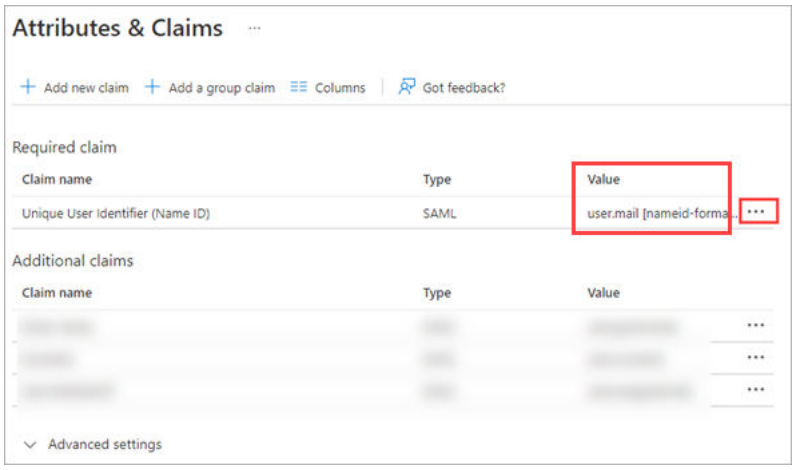

   a. Click the three dots.
   b. On the **Manage claim** page, click in the **Source attribute** field and select **user.mail** from the dropdown.
   c. Click **Save**.

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**.

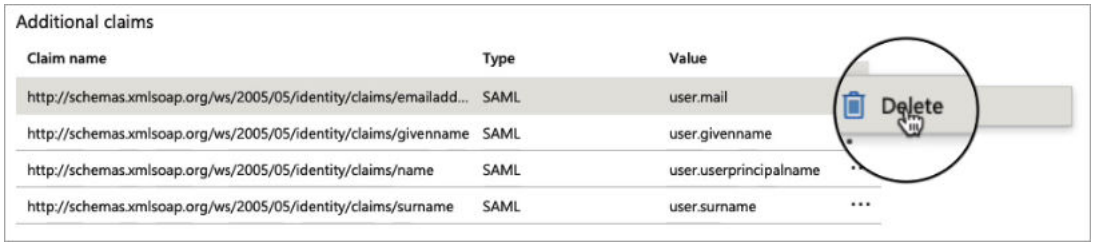

10 Click **Add new claim** and add three new claims:

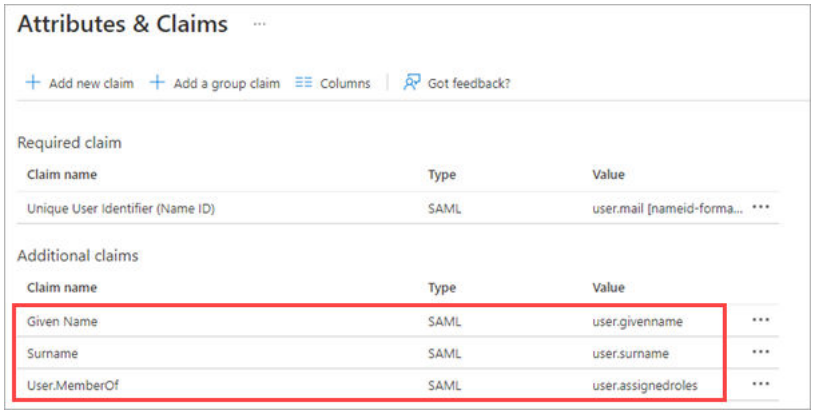

Given Name

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.

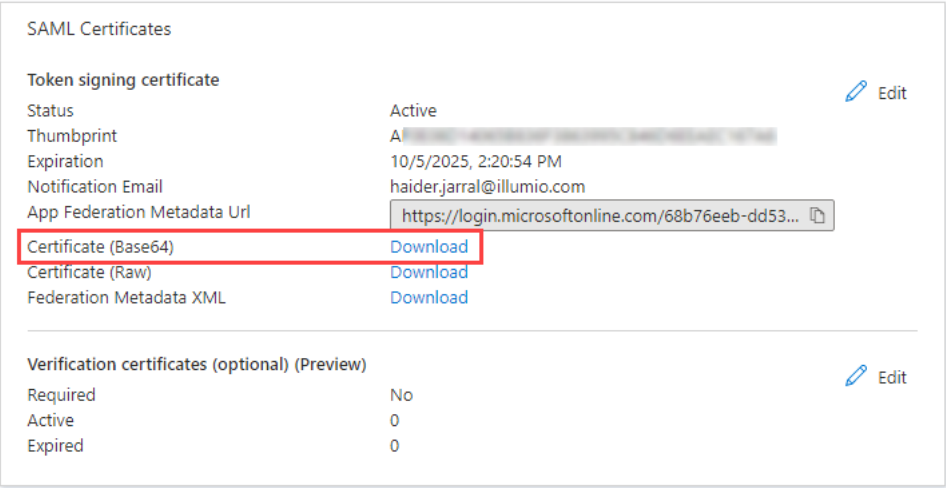

2. On the **Set up Illumio SSO** tile, copy and preserve the following URLs that you'll later paste into the Illumio PCE SAML setup.

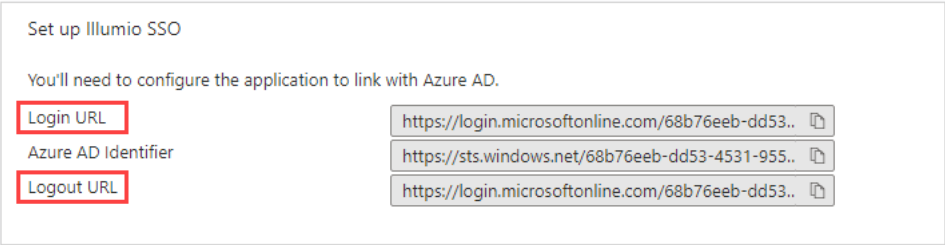

STEP 4: Configure SAML SSO settings in the Illumio PCE

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 you downloaded and copy and paste the contents.
   - **Remote Login URL**: Paste the Login URL you copied from Azure AD.
   - **Logout Landing URL**: Paste the Logout URL you copied from Azure AD.
5. In the **Information for Identity Destination** section:
   a. Choose an authentication method:
      - **Unspecified** uses the IdP default authentication mechanism.
      - **Password Protected Transport** requires the user to log in with a password in a protected session.

   b. If you want to require users to re-enter login credentials to access Illumio (even if the session is still valid), select **Force Re-authentication**. This allows users to log in to the PCE using login credentials different from their default computer login credentials.
6. Click **Save**.

STEP 5: Create App Roles in Azure AD

In this step you'll create app roles in Azure AD that you'll map to roles in the Illumio PCE Web Console in "Add External Groups and assign roles in the PCE Web Console".

For reference in this step, here's a list of the Global Roles available in the PCE Web Console:

- Global Organization Owner
- Global Administrator
- Global Viewer
- Globally Policy Object Provisioner

1. In Azure AD, go to **Users and Groups** and then click **application registration**.

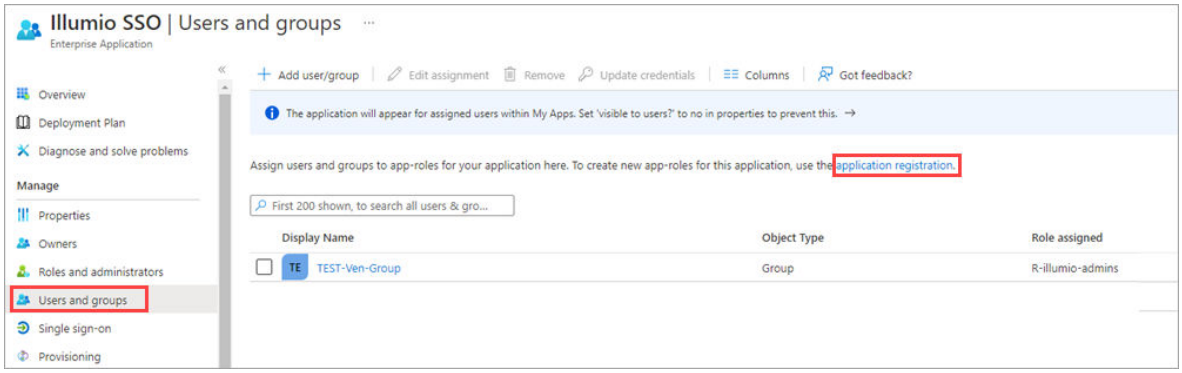

2. Create the roles you want by clicking **+ Create app role** and entering the required information for each role:
   - **Display name:** For example, enter one of the Global Roles that appear in the PCE Web Console.
   - **Value:** This must match the name you'll enter in the **Add External Groups** dialog box in "Add External Groups and assign roles in the PCE Web Console".
   - **Description:** The description will appear as help text in the app assignment and consent experiences.
3. Click **Apply** for each role that you create.
4. Delete the default app role **msiam_access**. Note that you must disable the default app role before you can delete it.
   a. Click **msiam_access** to open the **Edit app role** panel.
   b. Deselect **Do you want to enable the app role?**
   c. Click **Apply**. The side panel closes.
   d. Click **msiam_access** again to to open the **Edit app role** panel again.
   e. Click **Delete**.

When you're done creating roles in Azure AD, the **App roles** section should look similar to this:

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 in "Create App Roles in Azure AD".

1. In Azure AD, go to **Users and groups**.
2. Select the Illumio SSO app.
3. Click **Remove** to remove the current app assignments.
4. Click **Yes** to confirm removal.
5. Click **Add user/group**.

6. On the **Add Assignment** page, assign desired role(s) to users or groups:
   a. Under **User and groups**, click **None Selected**.
   b. In the **Users and groups** panel that opens, search for your desired user/group, click to select it, and then click **Select** at the bottom of the panel.
   c. Back on the **Add Assignment** page, under **Select a role***, click **None Selected**.
   d. In the **Select a role** panel that opens, find and click the role you want to assign, and then click **Select** at the bottom of the panel.
   e. Back on the **Add Assignment** page, click **Assign** at the bottom of the page.
   f. Repeat these sub-steps for each user and/or to which you want to assign app roles.

STEP 7: Add External Groups and assign roles in the PCE Web Console

In this step, you'll add external groups in the PCE Web Console and assign them the relevant global or scoped roles in Illumio RBAC.

TIP

Alternatively, you can add individual users by going to the **External Users** tab and following the onscreen prompts.

1. On the PCE Web Console, go to **Access Management > External Groups**.
2. Click **Add**.
3. In the **Add External Group** dialog box:
   - Enter a **Name**.
   - Enter an **External Group**.

     IMPORTANT

     This must match the **Value** that you specified for the app role inn"Create App Roles in Azure AD".

   - Click **Add**.
4. Repeat for additional groups.

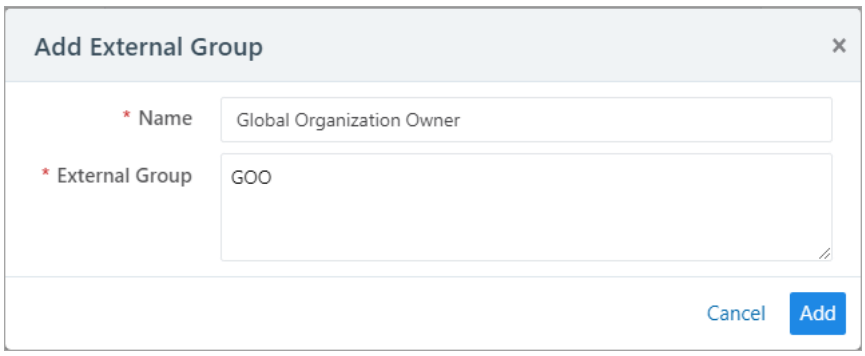

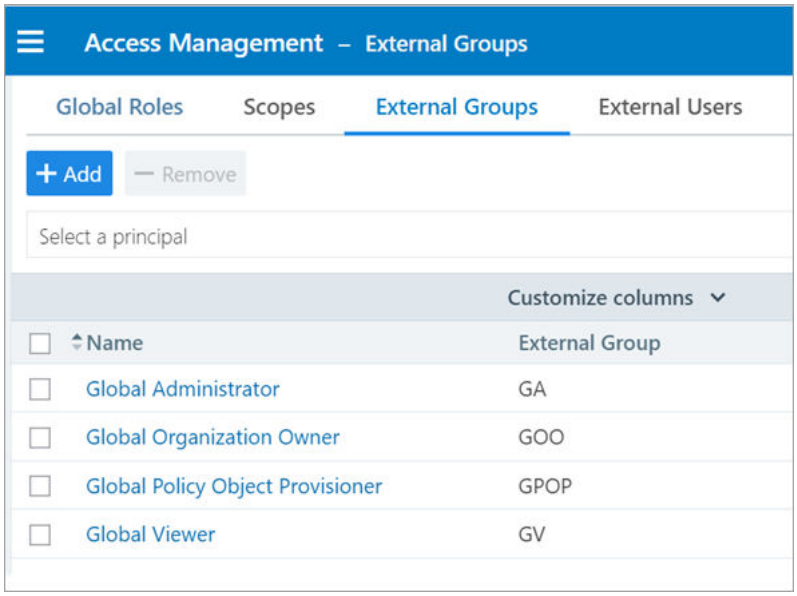

5. Click to open a group you created in the above step.
6. Click **Add Role > Add Global Role** or **Add Scoped Role**.
7. In the **Access Wizard**, select the appropriate **Role** and then click **Grant Access**.
8. Repeat for additional groups.

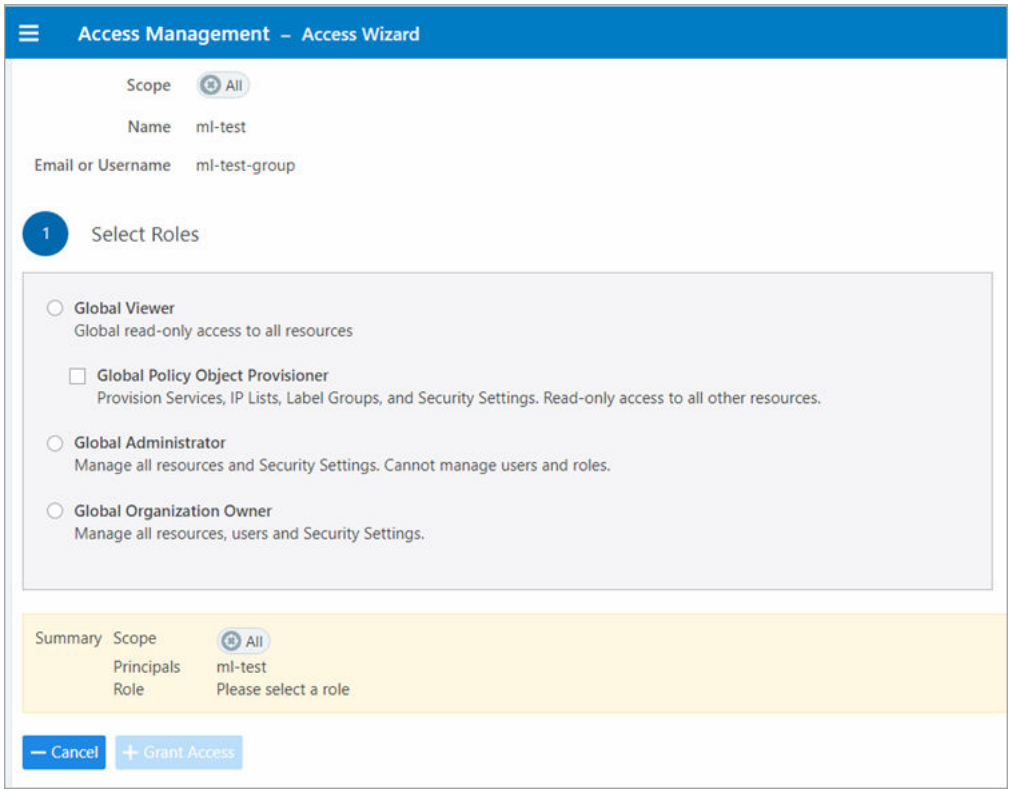

STEP 8: Turn on SAML authentication in the PCE Web Console

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**.

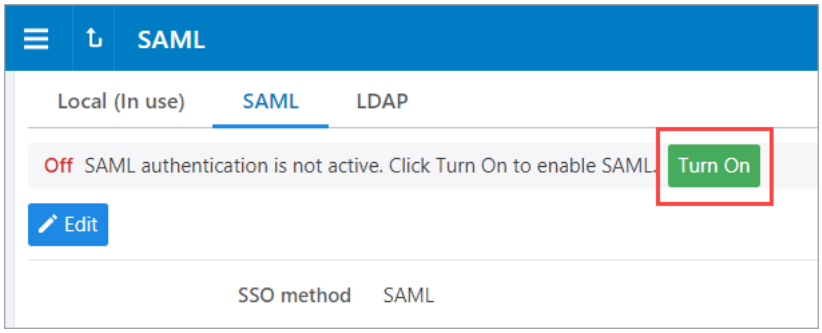

STEP 9: Test SSO

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**.

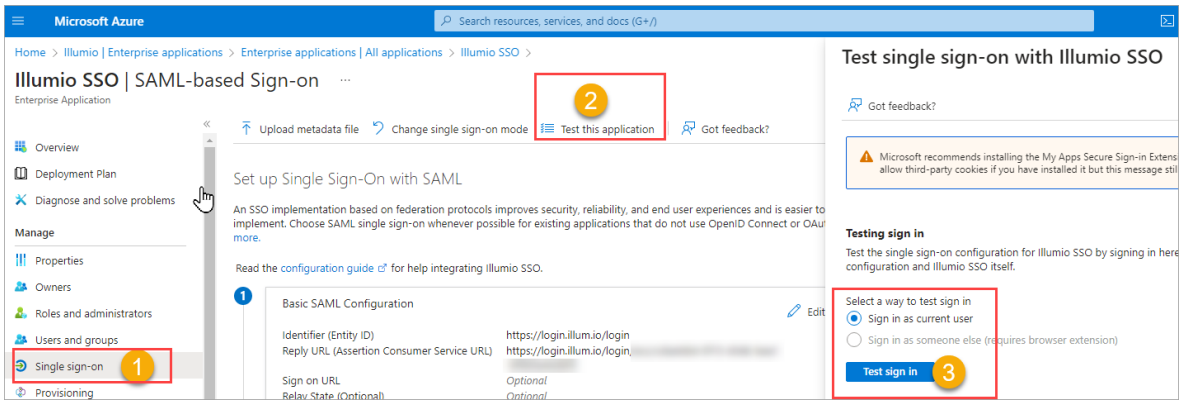

4. If the test is successful, the PCE will log you in to the **Welcome to Illumio** screen.

Monitor and Diagnose PCE Health

Learn about 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 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.

Monitor PCE Health

Learn how to monitor the health of the PCE.

For a consolidated description of the possible health-related states for VENs and workloads, see VEN and Workload States.

Minimum Required Monitoring

The PCE provides several different methods for monitoring PCE health.

No matter which technique you use, one main signal 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 [151].
- If you are using the API, monitor the status field similarly. For details, see Health Monitoring Using Health REST API [152].
- 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 following section details the meaning of the various PCE health metrics and what to do if you see a warning or error state.

PCE Health Status Codes

This 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. • The CPU, memory, 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, memory, 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 the PCE Web Console

Click the **Health** icon on 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 provision meets the requirements 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 set for every node. The hardware requirements check needs to know the cluster type to use the correct 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 the 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 issue in the tab.

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/system_health| sec=312831.757
sev=INFO pid=9231 tid=12334020 rid=0 leader=10.0.24.26
database_replication_lag=3.869344 cpu=2% disk=11% memory=19%
```

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

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

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

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

Health Monitoring Using PCE Command Line

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

Monitor a PCE Cluster

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

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

Return codes:

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

For example:

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

SERVICES (runlevel: 5)          NODES (Reachable: 4 of 4)
======================          =========================
agent_background_worker_service 10.0.26.49      10.0.6.171
agent_service                   10.0.26.49      10.0.6.171
agent_traffic_redis_cache       10.0.11.96      10.0.25.197
agent_traffic_redis_server      10.0.25.197
agent_traffic_service           10.0.26.49
10.0.26.49      10.0.6.171      10.0.6.171
auditable_events_service        10.0.26.49      10.0.6.171
collector_service               10.0.26.49
10.0.26.49      10.0.6.171      10.0.6.171
database_monitor                10.0.11.96      10.0.25.197
database_service                10.0.25.197
database_slave_service          10.0.11.96
ev_service                      10.0.26.49      10.0.6.171
executor_service                10.0.26.49      10.0.6.171
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      10.0.6.171
pg_listener_service             10.0.11.96
search_index_service            10.0.26.49      10.0.6.171
server_load_balancer            10.0.26.49      10.0.6.171
service_discovery_agent         10.0.25.197
service_discovery_server        10.0.11.96
10.0.26.49      10.0.6.171
set_server_redis_server         10.0.11.96
traffic_worker_service          10.0.26.49      10.0.6.171
web_server                      10.0.26.49      10.0.6.171
```

This command displays the members of the PCE cluster:

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

For example:

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

Monitor Database Replication

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

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

The PCE updates this information every two minutes.

IMPORTANT

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

For example:

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

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

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

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

PCE Health Troubleshooting

This section explains what to do 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 cannot diagnose and fix it yourself, contact Illumio Support.

The health metrics may occur in the PCE web console, API response status field, or 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 the following, here are the suggested actions.

Name	Troubleshoot
Disk Latency	Warning/Critical: Disk latency on data nodes indicates that the DB/Traffic service needs to be investigated further for possible performance issues. Typically, higher disk latency numbers indicate Disk I/O bottlenecks.
CPU	CPU usage increases when the PCE is under heavy load, 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	Memory usage increases when the PCE is under heavy load, and the Warning status is reported. Typically, the load should decrease without intervention in less than 20 minutes. Increasing the available memory if the Warning condition persists for 30 minutes or more.
Disk Space	The PCE manages disk space using log rotation, which 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 auto-vacuum's underperformance 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 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 VEN policy change rate, 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:

Name	Troubleshoot
	By default, a 4x2 PCE cluster is configured to handle approximately 10,000 flow summaries per second. If fewer posts are reported and you see many 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 hours. 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 and lose data. Adding additional `flow_analytics` daemons will distribute the work, but eventually, PostgreSQL could become the bottleneck, requiring DX. • Backlog size, node: If the backlog's size 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 cannot keep up with the posted data rate. The PCE collector flow summary rate and PCE traffic summary ingest rate need to be roughly equal, or this buffered backlog will grow.
Database Replication Lag	Warning: Check whether the PCE is running properly, and verify 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 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 defining each health metric's normal, warning, and critical status. 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 modify any Boolean, number, float, 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 pairs for lag in the policy and traffic databases are:

```
policy_database_replication_lag=x seconds
```

```
traffic_database_replication_lag=x seconds
```

Examples of when you might want to use this feature:

- The default memory threshold is set to 80% at a larger installation, but memory usage routinely spikes to 95%. The PCE Health page displays a warning whenever memory utilization exceeds the threshold. By configuring a higher threshold, you can reduce the frequency of warnings.
- Database replication lag can exceed a threshold briefly, 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, they 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 modification thresholds.

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 discusses 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
```

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.

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/PCE/Generate_PCE_Support_Bundle.mp4

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=09/14/2018T15:40:00
```

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.
`procs`	Includes process details in the Support Report.
`stats`	Includes statistics in the Support Report.

Run PCE Support or Inventory Report at Command Line

To run the PCE Support Report:

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

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

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

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

To run a PCE inventory report:

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

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

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

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

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

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

View Host and System Inventory

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

To show host inventory for the "local" node:

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

To show system inventory for the PCE:

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

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

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

PCE High Availability and Disaster Recovery Concepts

Learn how the PCE provides high availability (HA) and disaster recovery (DR).

Goals for PCE High Availability

The PCE is designed to handle system or network failures based on the following 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 failures without human intervention.
- **Reliable recovery: Recovery from failure is done with minimal loss of data.**

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

- Quorum [165]
- Service Availability [165]
- Cluster Capacity [166]

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

Quorum

A PCE cluster relies on *quorum*, a sufficient number of servers to ensure consistent operation. Quorum prevents the so-called "split-brain" case, where two cluster parts operate autonomously. Any node 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. Most 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, 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 that at least one instance of all required PCE services is 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 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)

Cluster Capacity

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

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 High Availability and Disaster Recovery Requirements

Learn how to make sure that your underlying systems are sufficient to 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

IMPORTANT

Make sure 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 consists of an *active* PCE and a *standby* PCE. A combination of continuous real-time replication and periodic synchronization keeps 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 is replicated:

- database_service
- citus_coordinator_service
- reporting_database_service
- agent_traffic_redis_server
- fileserver

Standby PCE Prerequisites

WARNING

Active Standby assumes the same certificate is used for all cluster nodes.

WARNING

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

Before designating a standby PCE, follow these 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:

- The hardware configuration and capacity of the two PCEs are as nearly identical as possible.
- The PCE software version is the same on both PCEs.
- Back up all `runtime.yml` files, typically located at `/etc/illumio-pce/runtime_env.yml`

Reset Any Repurposed PCE

If you are repurposing an existing PCE as the standby, ensure 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

Ensure that the required ports are open on both PCEs to allow network traffic between the active PCE and the standby PCE, enabling data replication. All the same service ports must be open on both the standby PCE and the active PCE. For a list of the required ports, see Port Ranges for Cluster Communication 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 Fully Qualified Domain Names (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 currently active PCE's Fully Qualified Domain Name (FQDN). The FQDN assigned to `active_pce_fqdn` should be resolved 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 necessary, to include 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, setting `active_pce_fqdn` to the same value is easier. Before you set up the standby PCE, change the `pce_fqdn` of the active PCE to something other than the `active_pce_fqdn`.

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 a new active PCE takes 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 described in "Set Up FQDNs".

Set Up VEN Library

The PCE is 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

Ensure the VEN versions in the library are supported by the installed PCE version.

Set Up a Standby PCE

Use the following steps to set up a standby PCE and associate it with its active PCE partner.

1. Complete the prerequisite steps in Standby PCE Prerequisites [168].
2. At this point, you should have two PCEs. Keep the active PCE in runlevel 5 (with UI available), ready to generate the API key. In step 3, you will bring the standby PCE to runlevel 2.
3. *On the active PCE*, which should be at Runlevel 5 with UI available, generate an API key. This API key is used only during the setup of the standby PCE.

   NOTE
   Users might get stuck at this step, trying to create an API key with both PCEs down.

4. 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 must remain at Runlevel 1 (if you switch runlevel, your replication may fail).
5. 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, provide the API key.

   ```
   $ export ILO_ACTIVE_PCE_USER_NAME=username
   $ export ILO_ACTIVE_PCE_USER_PASSWORD=secret
   ```
6. 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.

7. After replication is set up for the first time, the status of some services, such as the `citus_coordinator_service`, might be NOT RUNNING for a long time, and the cluster status is stuck in PARTIAL. This is usually because the service performs 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.

8. 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:

```
sudo -u ilo-pce ILO_ACTIVE_PCE_USER_NAME=api_17abrwerwe
ILO_ACTIVE_PCE_USER_PASSWORD=6efefeafe34ewroopll494934kdf
illumio-pce-ctl
setup-standby-pce --active-pce active.pce.com:8443$ export
```

Then short afterwards:

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

NOTICE

If you receive a 401 unauthorized error message/HTTP code, do not panic.

Go back to the active PCE, which is in runlevel 5 with UI open, and quickly create a new API Key with username and password.

The unauthorized error 401 might be a timeout, and you should use the active PCE to create a new API key.

The alternative for the above example in an environment with limited resources in terms of RAM and CPU is as follows:

```
sudo -u ilo-pce ILO_ACTIVE_PCE_USER_NAME=api_17abrwerwe
ILO_ACTIVE_PCE_USER_PASSWORD=6efefeafe34ewroopll494934kdf
illumio-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. Please verify that the PCE you are about to promote is indeed a standby PCE and is at runlevel 2.

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

   The output should say "standby."
2. Verify that the active PCE has failed and is offline. Data must not be replicated at 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 it 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" will be generated.

4. Ensure that DNS recognizes this as the new active PCE FQDN so that devices in your network can locate 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 [168].
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 monitor the output to ensure 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 may take some time, depending on the DNS Time To Live (TTL) value.

6. When ready, connect a new standby PCE to the new active PCE. Repeat the steps in Standby PCE Prerequisites [168] 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 the standby PCE. The Standby Replication tab displays the replication lag for the active and standby PCEs across the traffic database, policy database, reporting database, job queue Redis, and traffic data Redis. (The file server is not shown.)

Figure 2. PCE Health

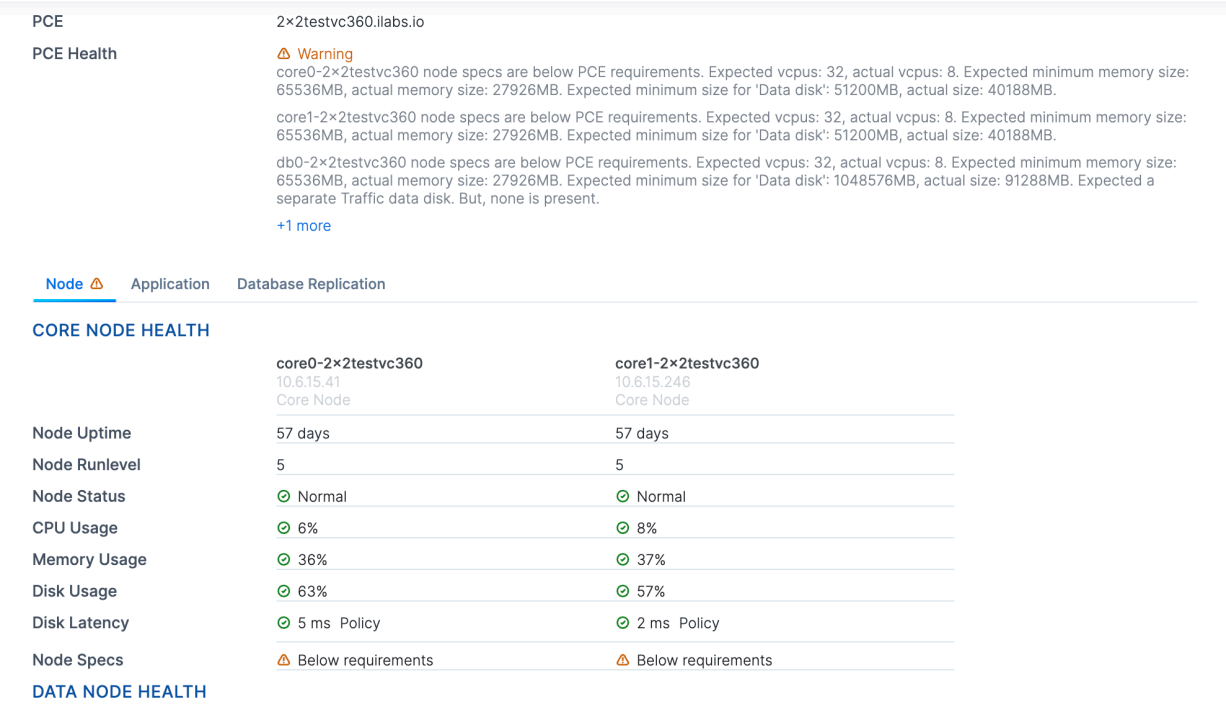

Another way that the PCE administrator can monitor replication is by watching the service discovery log for WAL segment missing errors. This error may occur when the standby traffic database service cannot keep up with the synchronization from the active traffic database service. When this error occurs, 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 a 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 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. Each WAL segment is 16 MB, so 5120 WAL segments would require approximately 82 GB of additional 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:

- File server 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 usually 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 the active PCE, and it already has an up-to-date copy of the data due to the ongoing replication between the two PCEs.

WARNING

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

Core Node Failure

In this failure case, one of the core nodes completely fails. This situation occurs anytime a node is not communicating with any of the other nodes in the cluster; for example, a node is destroyed, the node's SDS fails, or the node is powered off or disconnected from the cluster.

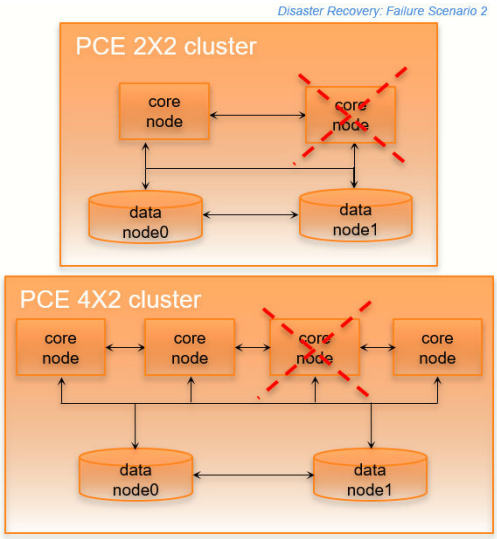
Disaster Recovery: Failure Scenario 2
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Precon-ditions	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 is always aware 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 Behav-ior	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 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.
Recov-ery	• **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.

Data Node Failure

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

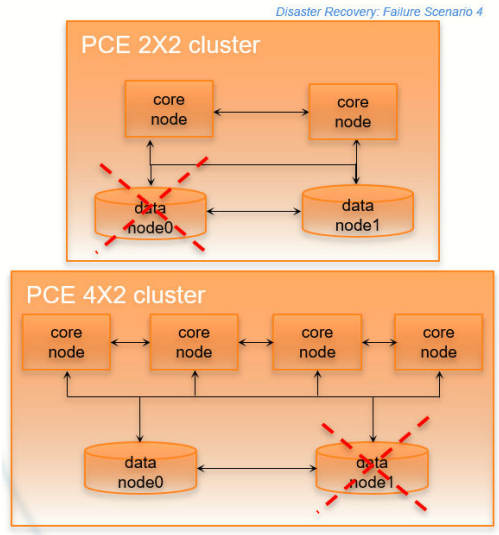
Disaster Recovery: Failure Scenario 4
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Precon-ditions	You should continually monitor the replica database's replication lag to ensure 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 Behav-ior	**PCE** • The PCE is temporarily unavailable. • Users may be unable to log in to 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 depend on the failed services may 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 the 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 the current policy continues to be enforced. • When a VEN misses a heartbeat to the PCE, it retries in 5 minutes.

Recov-ery	• **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`: All VENs must synchronize for an additional time, which varies depending on the number of VENs and the complexity of the policy.
	• **RPO:** Service-specific, based on the data services running on the failed data node.
	  • `database_service`: This implies that the failed data node was the primary database. All data committed to the primary database and not replicated to the replica is lost, typically in 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 Re-covery	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.

Primary Database Doesn't Start

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

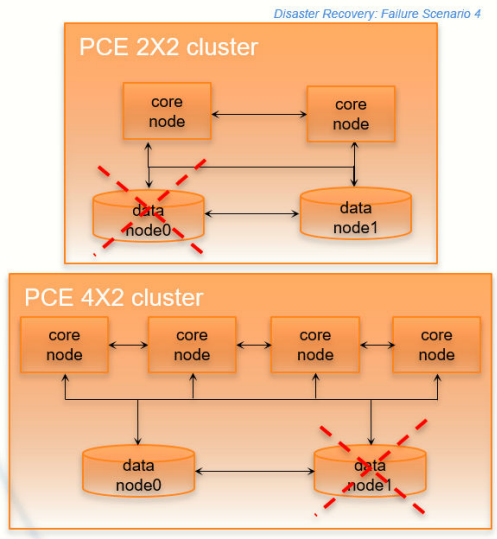

Stage	Details
Precondi-tions	The primary database node does not start.
Failure Behavior	The database cannot be started. Therefore, the entire PCE cluster cannot be started.
Full Re-covery	**Recovery type:** Manual. You have two recovery options: • Find and correct the root cause of the primary database failure. If necessary, contact Illumio Customer Support for assistance. • 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 this problem affects the PCE Supercluster, 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. If needed, contact Illumio Customer Support for help. • Promote the replica data node to the primary data node. **WARNING** **Promoting a replica to primary risks data loss** Consider this option as a last resort because, depending on the replication lag, it could result in data loss. 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.

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.

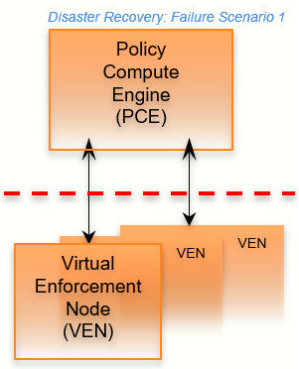

NOTICE

The VEN status command `illumio-ven-ctl status` doesn't report whether or not the VEN - PCE connection is degraded. If you suspect that the connection may be degraded, you can look in `agentmgr.log` and `platform.log` for messages that mention "degraded."

Stage	Details
Precon-ditions	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 for more 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 recalculates its current policy and excludes 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 this state, the VEN ignores all asynchronous commands received as lightning bolts from the PCE except commands that initiate software upgrades and Support Reports.

Recov-ery	• **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 a new policy from PCE is 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 emerges from 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.

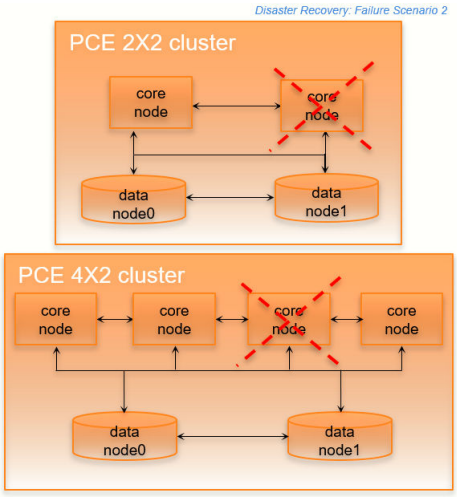

Stage	Details
Precondi-tions	None.
Failure Behavior	**PCE** • The PCE might be temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return an HTTP 502 response and the `/node_available` API request might return an HTTP 404 error. • Other services that are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

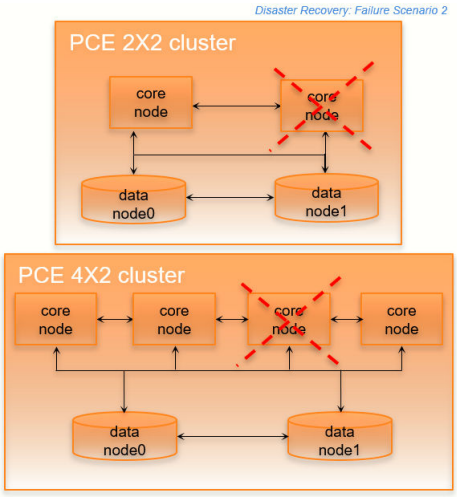
Disaster Recovery: Failure Scenario 2
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Precon-ditions	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 Behav-ior	**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.
Recov-ery	• **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.

Data Node Failure

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

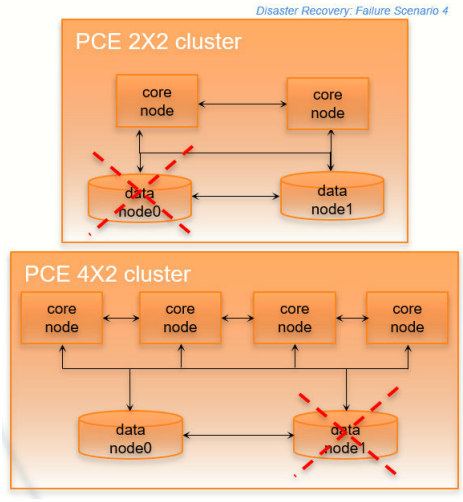
Disaster Recovery: Failure Scenario 4
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Precon-ditions	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 may be unable to log in to 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 depend on the failed services may 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 the 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.

Primary Database Doesn't Start

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

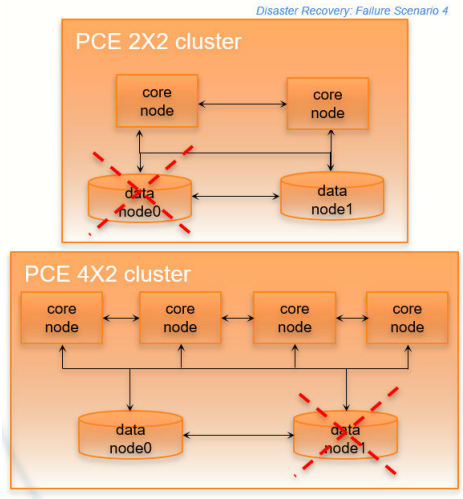

Stage	Details
Precondi-tions	The primary database node does not start.
Failure Behavior	The database cannot be started. Therefore, the entire PCE cluster cannot be started.
Full Re-covery	**Recovery type:** Manual. You have two recovery options: • Find the root cause of the primary database failure and 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 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 the primary data node. **WARNING** **Promoting a 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.

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.

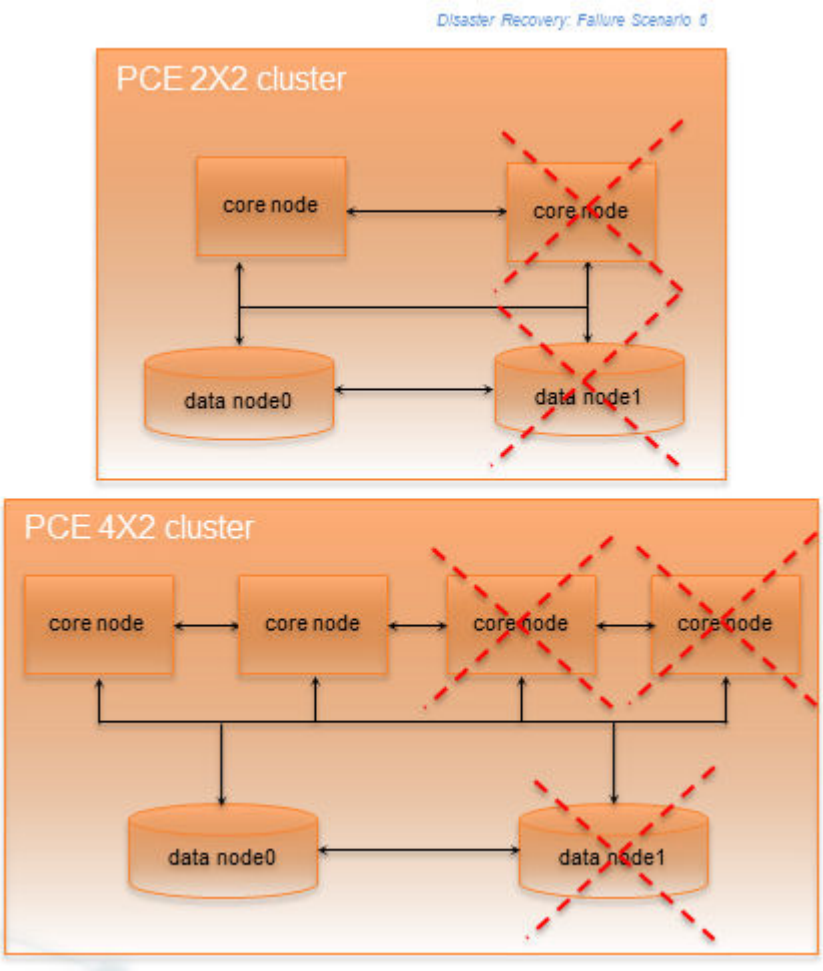

Stage	Details
Precon-ditions	None.
Failure Behavior	**PCE** • The PCE is temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return a HTTP 502 response and the `/node_available` API request might return am HTTP 404 error. • Other services that are dependent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

Split Cluster Failure Involving Data0

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

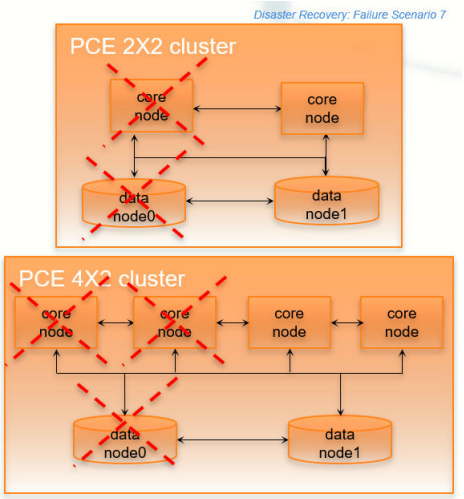
Disaster Recovery: Failure Scenario 7
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Precon-ditions	**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 Behav-ior	**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.

Recov-ery	• **Recovery type:** Manual intervention is required to recover from this failure case.
	• **Recovery procedure**
	• **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 Re-covery	See Revert Standalone Cluster Back to a Full Cluster [200] for information.

Configure Data1 and Core Nodes as Standalone Cluster

To enable the surviving data1 and core nodes to operate as a standalone 2x2 or 4x2 cluster, follow these steps in this exact order.

1. On the *surviving data1 node and all surviving core nodes*, stop the PCE software:

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

2. On *any surviving core node*, promote the core node to be a standalone cluster leader:

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

3. On the *surviving data1 node*, promote the data1 node to be the primary database for the new standalone cluster:

```
sudo -u ilo-pce illumio-pce-ctl promote-data-node <promoted-core-node-ip-address>
```

   For the IP address, enter the IP address of the promoted core node from step 2.

4. **(4x2 clusters only)** On the *other surviving core node*, join the surviving core node to the new standalone cluster:

```
sudo -u ilo-pce illumio-pce-ctl cluster-join <promoted-core-node-ip-address> --split-cluster
```

   For the IP address, enter the IP address of the promoted core node from step 2.

5. Back up the surviving data1 node.

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.

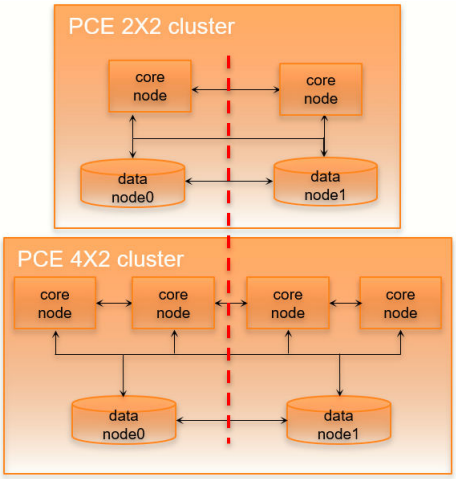

Stage	Details
Precon-ditions	None.
Failure behavior	**PCE** • The PCE is temporarily unavailable. • Users might be unable to log into the PCE web console. • The PCE might return an HTTP 502 response and the `/node_available` API request might return an HTTP 404 error. • Other services that are dependent 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.

Recov-ery	• **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 Re-covery	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.

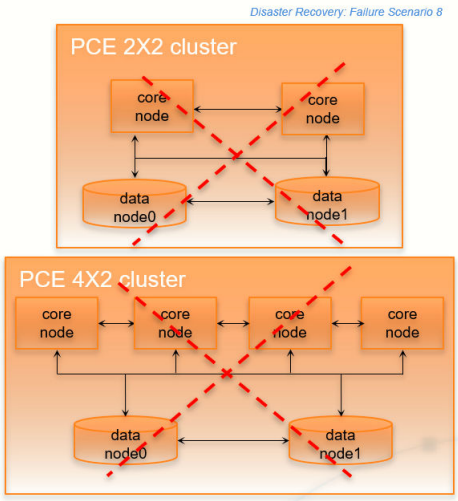

Stage	Details
Precon-ditions	Illumio assumes that you have met the following conditions before the failure occurs for this failure case. **IMPORTANT**: You must consistently and frequently back up the PCE primary database to an external storage system that can be used for restoring the primary database 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 behav-ior	**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.
Recov-ery	• **Recovery type:** Manual intervention is required to fully recover from this failure case. • **Recovery procedure:** See Complete Cluster Recovery [207] 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 Re-covery	See Complete Cluster Recovery [207] 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 file>
   ```
5. Migrate the database by running this command:

   ```
   sudo -u ilo-pce illumio-pce-db-management migrate
   ```
6. Copy the Illumination data file from the primary database to the other data node. The file is located in the following directory on both nodes:

   ```
   <persistent_data_root>/redis/redis_traffic_0_master.rdb
   ```
7. Bring the PCE cluster to runlevel 5:

   ```
   sudo -u ilo-pce illumio-pce-ctl set-runlevel 5
   ```
8. Verify that you can log into the PCE web console.

PCE-Based VEN Distribution Recovery

When you rely on the PCE-based distribution of VEN software, after you have recovered from a PCE cluster failure, you need to reload or redeploy 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 in-

formation 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.

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.

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.

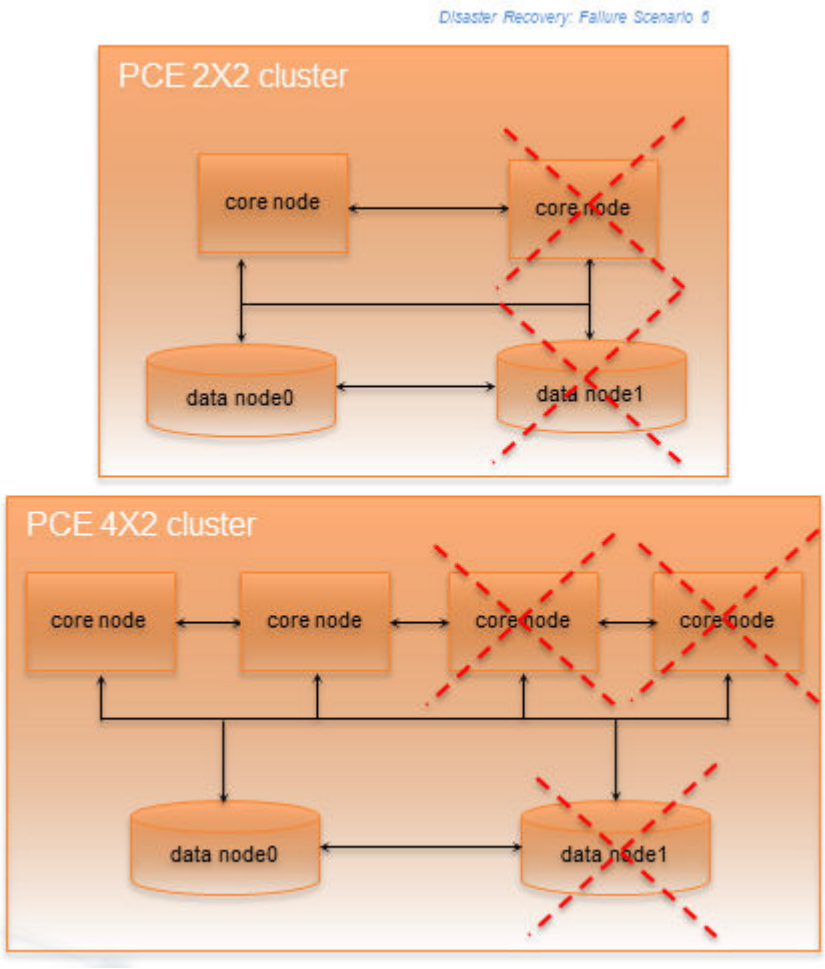
Disaster Recovery: Failure Scenario 6
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Precon-ditions	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 Re-covery	Either recover the failed nodes or provision new nodes and join them to the cluster. For recovery information, see Replace a Failed Node.

Split Cluster Failure Involving Data0

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

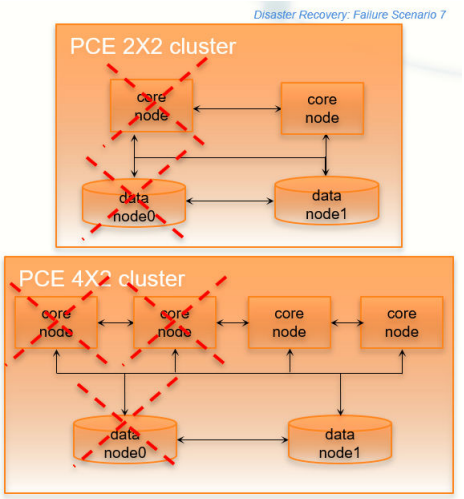
Disaster Recovery: Failure Scenario 7
PCE 2X2 cluster
core node
core node
data node0
data node1
PCE 4X2 cluster
core node
core node
core node
core node
data node0
data node1

Stage	Details
Precon-ditions	**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 Behav-ior	**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.

Recov-ery	• **Recovery type:** Manual intervention is required to recover from this failure case.
	• **Recovery procedure**
	• **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 Re-covery	See Revert Standalone Cluster Back to a Full Cluster [214] for information.

Configure Data1 and Core Nodes as Standalone Cluster

To enable the surviving data1 and core nodes to operate as a standalone 2x2 or 4x2 cluster, follow these steps in this exact order.

1. On the *surviving data1 node and all surviving core nodes*, stop the PCE software:

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

2. On *any surviving core node*, promote the core node to be a standalone cluster leader:

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

3. On the *surviving data1 node*, promote the data1 node to be the primary database for the new standalone cluster:

```
sudo -u ilo-pce illumio-pce-ctl promote-data-node <promoted-core-node-ip-address>
```

   For the IP address, enter the IP address of the promoted core node from step 2.

4. **(4x2 clusters only)** On the *other surviving core node*, join the surviving core node to the new standalone cluster:

```
sudo -u ilo-pce illumio-pce-ctl cluster-join <promoted-core-node-ip-address> --split-cluster
```

   For the IP address, enter the IP address of the promoted core node from step 2.

5. Back up the surviving data1 node.

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
> ```

Connectivity Settings

This section describes how to modify PCE settings that affect connectivity.

NOTE
Permission to edit these settings depends on your role.

Private Data Centers

The PCE uses connectivity settings to decide whether workloads can 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 hosted in private data centers that do not have duplicate IP addresses in the network. When your network environment hosts workloads in your private data center and in a public cloud, and you want to change this setting, contact Illumio Support.

Offline Timers

You can configure Offline Timers in **Settings > Offline Timers** 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 belonging to workloads that have been disconnected from those allowed to communicate with the disconnected workloads. You need to remove the disconnected workloads from the PCE to ensure 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 to the PCE every 5 minutes and a goodbye message when it is gracefully shut down. The PCE marks a workload offline when these conditions occur:

- The PCE hasn't received a heartbeat message from the VEN for the configured period (whether default or custom).
- The PCE receives a goodbye message from the VEN.

Under the following conditions, you can change the default Offline Timer settings before putting your workloads in enforcement:

- 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 and increasing the window of time during which 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 how much time must elapse before a managed workload is marked "offline" after it sends a goodbye message. By default, the High Security setting is:

- Server VENs: 15 minutes
- Endpoint VENs: 24 hours

Wait 1 hour/1 day - High Security (Default)

The PCE performs the following actions:

1. Listens for Goodbye messages from the VEN.
2. Pushes an updated policy to the peer workloads previously allowed to communicate with the removed workloads.
3. Immediately cleans up those workloads' IP addresses from its active policy.

Never remove IP addresses - Highest Availability.

This setting has the following effect on the PCE:

- Ignores Goodbye messages from workloads.
- Keeps all IP addresses in policy and never automatically removes unused IP addresses.
- Requires the removal of those unused IP addresses.

Custom Timeout

Enter a time period (minimum: 0 seconds).

The PCE performs the following actions:

- Listens for Goodbye messages from the VEN.
- Waits for the specified period before the cleanup of those workloads' IP addresses from its active policy.
- Pushes an updated policy to the peer workloads previously allowed to communicate with the removed workloads.

Disconnect and Quarantine Timer

Sets how much time must elapse before a managed workload is marked "offline" after the PCE has received no heartbeat from the VEN. By default, the High Security setting is:

- Server VENs: 60 minutes
- Endpoint VENs: 24 hours

Wait 1 hour/1 day - High Security (Default)

The PCE performs the following actions:

1. Waits for the configured time to receive a heartbeat from the disconnected workloads and then quarantines workloads that do not respond within that period.
2. Removes the quarantined workloads' IP addresses from its active policy.
3. Pushes an updated policy to the peer workloads previously allowed to communicate with the quarantined workloads.

Never remove IP addresses - Highest Availability.

This setting has the following effect 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 the removal of those unused IP addresses.

Custom Timeout

Enter a time period (minimum: 300 seconds).

The PCE performs the following actions:

1. Waits for the specified time period for the VEN to heartbeat.
2. Quarantine those workloads that do not respond within that time period.
3. Removes the quarantined workloads' IP addresses from its active policy.
4. Pushes an updated policy to the peer workloads previously allowed to communicate with the quarantined workloads.

Disconnect and Quarantine Warning

Sets how much time must elapse before the PCE emits a warning event to indicate that the VEN missed heartbeats. The server VEN will appear warning on the VEN pages.

The default settings are:

- Server VENs: Wait one-quarter of the Disconnect and Quarantine Timer.

- Endpoint VENs: Disabled.

Wait one-quarter of the Disconnect and Quarantine Timer - (Default) (applies to Server VENs only)

The PCE performs the following actions:

1. Wait one-quarter of the *Disconnect and Quarantine Timer* setting for the server VEN to heartbeat before emitting a warning event indicating that it has missed heartbeats. The server VEN appears warning on the VEN pages.
2. If the *Disconnect and Quarantine Timer* is set Never to remove IP addresses - Highest Availability, the PCE emits a warning event 15 minutes after receiving the previous VEN heartbeat.
3. If you set a custom time of 20 minutes or less for the *Disconnect and Quarantine Timer* and the PCE receives no heartbeat from the VEN at least 5 minutes after receiving the previous heartbeat, the PCE emits a warning event to indicate the missed heartbeat. The endpoint VEN will appear with a warning on the VEN pages.

Custom Timeout (applies to Server and Endpoint VENs)

Enter a period greater than 5 minutes (300 seconds) and less than the value specified for the Disconnect and Quarantine Timer.

1. Waits for the specified time period for the VEN to heartbeat.
2. VENs appear with a warning on the VEN pages.

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 allow only IPv4 traffic for your PCE organization, the VENs on your workloads drop IPv6 traffic when they are in Enforced mode. This decision can lead to delays and communication failures in applications because applications will wait for IPv6 connection attempts to time out before attempting to connect over IPv4.

The problem occurs because, by default, the Windows OS prefers IPv6 over IPv4 and will attempt to connect over IPv6 before IPv4. As a workaround, you can change the order of connection attempts so that IPv4 is preferred over IPv6. With this change, applications will connect over IPv4 first and succeed or fail as governed by the workload's firewall policies.

For information about changing the connection order to prefer IPv4 over IPv6, see the Microsoft KB article Guidance for configuring IPv6 in Windows for advanced users.

As explained in the KB article, run the following command and reboot the Windows workload:

```
reg add
hklm\system\currentcontrolset\services\tcpip6\parameters /v
DisabledComponents /t REG_DWORD /d 0x20
```

To avoid rebooting the Windows workload, run the following commands:

```
netsh interface ipv6 delete prefixpolicy ::ffff:0:0/96
netsh interface ipv6 add prefixpolicy ::ffff:0:0/96 60 4
```

Manage Security Settings

This section describes how to modify PCE settings that affect connectivity.

How to 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 the Enforcement state. Always allowed on AIX and Solaris workloads.
VENS (Versions lower than 20.2.0)	IPv6 traffic	Allow IPv6 traffic	All IPv6 traffic is allowed.
		Block IPv6 traffic	Blocked only in the 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 workload's public IP address as seen by the PCE is distributed along with the IP addresses of the workload's interfaces. Use this setting only if no shared SNAT IP addresses exist 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 workload's public IP address in policy computation. This setting is used in environments where workloads in a known public cloud (e.g., AWS or Azure) 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 not specific to any workloads. The policy distributes only the IP addresses of the network interfaces on the workload (usually the private IP addresses).

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 Kubelink reports 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**. If the feature is enabled, the IP Forwarding tab will appear.

   NOTE
   Enable this feature using the API call to the PCE so it appears as an option in the Security menu.

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 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, including traffic forwarded through the host, without visibility.

SecureConnect Setup

Enterprises are required 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 data centers. 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, Windows workloads, and Linux and Windows workloads.

Supported Encryption Protocols

These are the encryption protocols/ciphers enabled by SecureConnect when configuring IPsec between servers:

Encapsulating Security Payload (ESP)

- `sha1-aes256`
- `sha256-aes256`
- `aes256`
- `sha256-null`
- `sha1-null!`

Internet Key Exchange (IKE)

- `aes256-sha256-modp2048`
- `aes256-sha1-modp2048`
- `aes256-sha1-modp1024`
- `aes256-sha384-prfsha384-ecp384`

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. In this 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 used protocol will be changed to reflect the transport mode (ESP).

This change does not create underlying interfaces or change any other underlying networking infrastructure. 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 is 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 the 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 that relate to IPsec thereby tracking all IPsec activity.

Performance

The CPU processing power that a workload uses determines the capacity of the encryption. The packet size and throughput determine the amount of power that is required to process the encrypted traffic using this feature.

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

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 that are in both Idle and non-Idle policy states, you can impact the traffic between these workloads.
- SecureConnect is not supported on AIX and Solaris platforms.

SecureConnect 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.

Using Keys with SecureConnect

Use either pre-shared keys or PKI certificates with SecureConnect.

Use Pre-Shared Keys with SecureConnect

SecureConnect supports the use of 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 the PCE automatically generates. 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:

```
c4aeb6230c508063db3e3e1fac185bea9c4d17b4642a87e091d11c9564fbd07
5
```

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 regenerate 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.

asfasdfasdf

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. Se-

cureConnect 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 allows you to use client-side PKI certificates for IKE authentication and IPsec communication between managed workloads. If you have a certificate management infrastructure in place, you can leverage it for IKE authentication between workloads because it provides higher security compared to using pre-shared keys (PSKs).

Certificate-based SecureConnect works for connections between Linux workloads, between Windows workloads, and between Linux and Windows workloads.

The IPsec configuration uses the certificate with the distinguished name from the issuer field that you specify during PCE configuration for IKE peer authentication.

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 them will fail, and the workloads will not be able to communicate with each other.

Leaf certificate X.509 field requirement

- Version 3

- Subject Name DN must contain Common Name (example: OU=VEN, CN=centos6.ilabs.io)
- SubjectAltName (required for better compatibility) must contain an email address field that is identical to the Common Name of DN (example: DNS:centos6, email:centos6@ilabs.io)
- 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 macOS 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, Applications, Environments, and Locations.

Configuring SecureConnect to use PKI certificates in the global Security Settings page does not manage certificates for your organization or deliver them to your workloads.

NOTE

You must set up certificates on your Windows and Linux workloads independently. For information, see Requirements for Certificate Setup on Workloads [229].

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

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 a client's certificates-based identity to verify its authenticity before allowing it to connect.

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 a certificate for SecureConnect, see SecureConnect Setup [223].

The exact prerequisites and limitations for certificate setup apply to 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 usages 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. [231]
2. Add the laptop to the PCE by creating an unmanaged workload and assign the appropriate labels to it to be used for rule writing
3. Create rules using those labels to grant access to the managed workloads. For information, see "Enable AdminConnect for a Rule" in the Security Policy Guide.
4. Configure IPsec on a laptop.

To add a laptop to the PCE by creating an unmanaged workload:

To manage a laptop with AdminConnect, add the laptop to the PCE as an unmanaged workload.

1. Choose **Workloads** > **Add** > **Add Unmanaged Workload**.
2. Complete the fields in the General, Labels, Attributes, and Processes sections.
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:

You must configure IPsec so these clients can use the AdminConnect feature with laptops in your organization.

For information about configuring IPsec using netsh, see the Microsoft Technet article Netsh Commands for Internet Protocol Security (IPsec) .

See also the following examples for information about the IPsec settings required to manage laptops with the AdminConnect feature.

```
PS C:\WINDOWS\system32> netsh advfirewall show global

Global Settings:
----------------------------------------------------------------
-------
IPsec:
StrongCRLCheck                          0:Disabled
SAIdleTimeMin                           5min
DefaultExemptions                       NeighborDiscovery,DHCP
IPsecThroughNAT                         Server and client behind
NAT
AuthzUserGrp                            None
AuthzComputerGrp                        None
AuthzUserGrpTransport                   None
AuthzComputerGrpTransport               None

StatefulFTP                             Enable
StatefulPPTP                            Enable

Main Mode:
KeyLifetime                             60min,0sess
SecMethods                              ECDHP384-AES256-SHA384
ForceDH                                 Yes

Categories:
BootTimeRuleCategory                    Windows Firewall
FirewallRuleCategory                    Windows Firewall
StealthRuleCategory                     Windows Firewall
ConSecRuleCategory                      Windows Firewall

Ok.


PS C:\WINDOWS\system32> netsh advfirewall consec show  rule
name=all

Rule Name:                              telnet
----------------------------------------------------------------
-------
Enabled:                                Yes
Profiles:                               Domain,Private,Public
Type:                                   Static
Mode:                                   Transport
Endpoint1:                              Any
Endpoint2:
10.6.3.189/32,10.6.4.35/32,192.168.41.163/32
```

```
Port1:                                  Any
Port2:                                  23
Protocol:                               TCP
Action:                                 RequireInRequireOut
Auth1:                                  ComputerKerb,ComputerCert
Auth1CAName:                            CN=MACA, O=Company,
OU=engineering, S=CA, C=US, L=Sunnyvale, E=user@sample.com
Auth1CertMapping:                       No
Auth1ExcludeCAName:                     No
Auth1CertType:                          Intermediate
Auth1HealthCert:                        No
MainModeSecMethods:                     ECDHP384-AES256-SHA384
QuickModeSecMethods:                    ESP:SHA1-
AES256+60min+100256kb
ApplyAuthorization:                     No
Ok.
```

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 stores data in PostgreSQL databases. 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. Illumio Support requires immediate remediation.

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 "<database_name>"
```

```
Stop the postmaster and vacuum that database in single-user
mode.
```

Cause:

In a PostgreSQL database, transaction wraparound (also known as transaction ID exhaustion) can occur if many 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. Instilling commands manually is the only way to recover from a transaction ID wraparound.

To avoid this situation, PostgreSQL provides an autovacuum feature that recovers disk space by 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:

- The vacuum did not run on the tables.
- Temporary tables remained in the database rather than being dropped as they should have been. 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 a transaction ID wraparound failure.

WARNING

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

Best Practices for Handling Scanner Traffic for the Illumio PCE

Use these best practices to manage scanner traffic on the Illumio PCE while monitoring network security.

Modern data centers require scanners, but the traffic they generate is not representative of actual network traffic. Scanners can add noise, complexity, and overhead to an Illumio PCE environment without maintaining effective policy and monitoring.

Illumio recommends filtering scanner traffic to prevent the PCE from ingesting it. Filtering scanner flows enables longer retention of workload traffic and a more accurate view of the environment. This enables you to apply a more precise security policy while allowing scanners to perform their roles within the security stack.

Identify Exclusions Using the Core Services Detector

The Core Services Detector identifies well-known applications that assist with labeling and detecting scanner-generated flows. With the Core Services Detector, you can define security policies to allow, block, or filter known scanner traffic. See Core Services Detector.

Allow and Filter Scanner Traffic from the Environment

You should implement a policy that allows scanner traffic without retaining the flows. If you do not account for scanners in your policy, they will be blocked from scanning the environment when the policy is enforced, rendering them ineffective. Build a Flow Collection filter to discard the flows that the scanner generates. The PCE still applies policies for these flows, but it does not store them in the database because there is limited value in retaining them within the Illumio platform.

You can create Unmanaged Workloads to represent scanners in your Visualization Map. If you allow scanner flows, the vulnerability scanner can generate reports that you can upload into Illumio using the Vulnerability Maps feature. See Upload Vulnerability Data.

Build Traffic Filters

Use the UI or an API to identify and filter scanner sources so that traffic is filtered by source port, IP range, or protocols before it reaches the Flow Collector.

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/PCE/Build_Traffic_Filters.mp4

1. Navigate to **Settings** > **Flow Collection**, and click **Add**.
2. In the **When Traffic Matches the Following Conditions** section, set these values:
   - Transmission: Unicast
   - Enforcement Node Type: Any
   - Network: Any
   - Protocol: UDP
   - Source IP Address: <Scanner IP>
3. In the **Take the Following Action** section, select **Drop**.
4. Create another Flow Collection to drop TCP traffic for the same sources.
5. Create additional Flow Collections for scanner IPs and protocols.

Reduce Traffic Stored in Flow Logs

Note that when you create a Flow Collection filter, you can also filter transmission types like broadcast and multicast traffic. Create filters for any traffic you will not use to build a security policy, or for traffic that will not provide value in the Visualization Map.

You can aggregate or filter broadcast and multicast traffic to improve PCE performance by reducing traffic that is irrelevant to the security policy. Filter the flows based on transmission type, enforcement-node type, network, protocol, and source IP/port or destination IP/port.

See Reduce Traffic Stored in Flow Logs. (Log in to the Illumio Support portal.)

Review Detection Rules on a Schedule

Audit your filtering rules regularly to ensure they remain accurate, as scanner IPs change. Review scanner traffic trends regularly on a schedule that fits your organization's needs and update your exclusion policies based on new scanning methods.

- For large organizations, the best practice is to review rules monthly.
- The best practice for small and mid-size organizations is to review rules quarterly.

Events Administration

Abstract

Events Administration guide provides information on how to administer your PCE deployment: an overview of events and SIEM integration, event setup, event record formats, and event types by resource.

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

Before reading this guide

Illumio recommends that you be familiar with the following technology:

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

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 Segmentation for Data Centers

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 VENs are recorded, including workload firewall tampering.

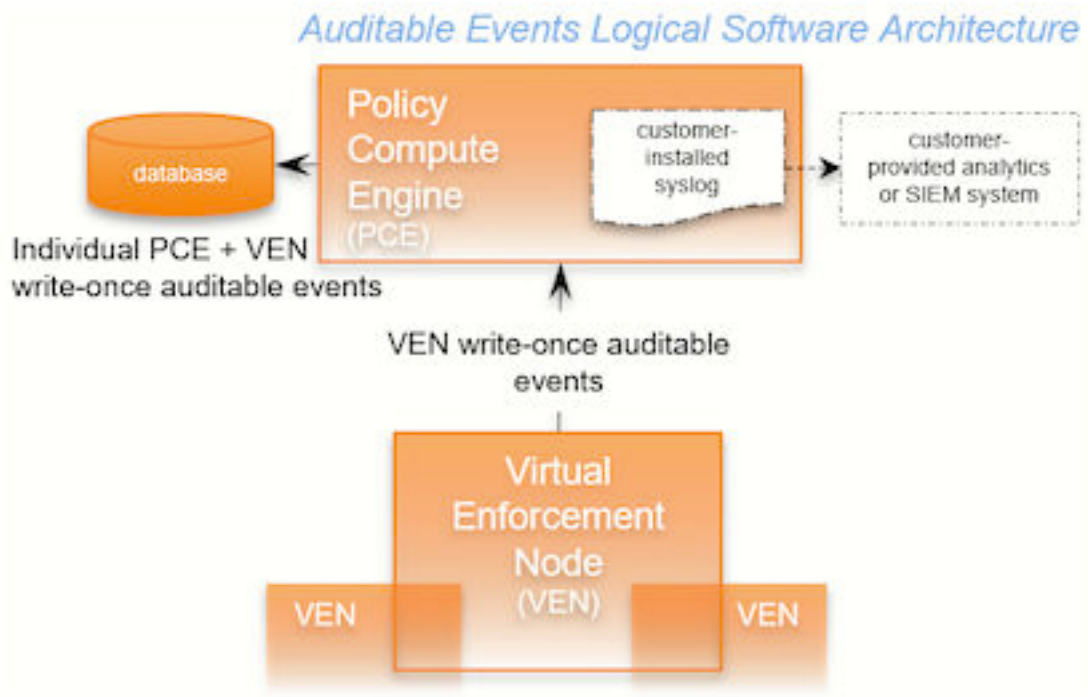

As required by auditing standards, every recorded change includes a reference to the program that made the change, the change's timestamp, and other fields. After recording, the auditable events are read-only.

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

Auditing Needs Satisfied by Framework

Need	Description
Audit and Compliance	Evidence to show that resources are managed according to rules and regulatory standards.
Resource Lifecycle Tracking	All information is necessary to track a resource through creation, modification, and deletion.
Operations	Trace of recent changes to resources.
Security	Evidence to show which changes failed, such as incorrect user permissions or failed authentication.

Benefits of Events Framework

The events framework 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 integration
  - 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 Illumio Segmentation for Data Centers

Events Lifecycle for Resources

Illumio resources progress through the lifecycle stages (creation, updating, deletion) and the Illumio Segmentation for Data Centers 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

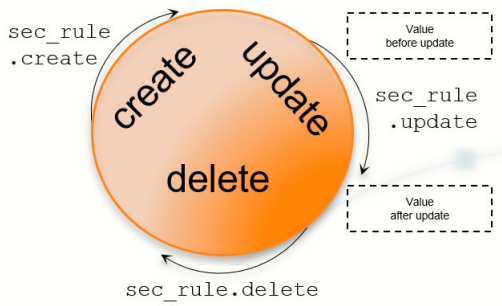

Other Resource Lifecycles

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`

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

Illumio Segmentation for Data Centers includes the following general categories of auditable events:

- Organizational events: Organizational events are further grouped by their source:
  - API-related events: Events occurring from the 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 REST API schemas' zipfile. Go to the Develop category > REST API Public Schemas (Archive File)from the documentation portal's Home page.

Event Syntax

The names of recorded auditable events 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 an event record, which answers the who, what, where, how, and when:

Type of information	Description
Who	• VEN identified by hostname and agent 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, the exception type and the exception message.
	  • All failures related to security, such as authentication and authorization.
	  • Severity as INFO, WARNING, or 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`	The exact time that the event occurred was in RFC 3339 format, which included fractional seconds.
`pce_fqdn`	The PCE's fully qualified domain name is especially useful for Supercluster deployments or if multiple PCEs send data to the SIEM server.
`created_by`	Identifies the event's creator, a user, the system, or a workload.
`event_type`	Name the event; see the List of Event Types table for more information.
`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 PCE Organizational events. For example, Organizational events capture actions such as users logging in and logging out, 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 to act 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 a website's trust for a user.

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

Illumio Segmentation for Data Centers 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 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 software.

IMPORTANT

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

Events Monitoring Best Practices

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 Segmentation for Data Centers events are structured and actionable. Using the event data, you can identify the severity, affected systems, and what triggered the event.

Illumio Segmentation for Data Centers 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 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 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 recommends 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"`	This document contains the information necessary to identify workloads with lost agents. A lost agent occurs when the PCE deletes a workload from its database but 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 (default: total of 15 minutes). Typically, this event precedes the PCE taking the VENs offline to perform internal maintenance. For Server VENs, 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 Server VEN has not sent a heartbeat for 15 minutes; if the offline timer is configured to 4 hours, an alert is sent after the Server VEN hasn't sent a heartbeat for 1 hour. Alerts are disabled by default for Endpoint VENs. **Recommendation:** Monitor these events for high-value workloads. The PCE can take these workloads offline when the VENs miss 12 heartbeats (usually 60 minutes).
`event_type="system_task.agent_offline_check"`	This event 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 a change in the affected workloads' security posture.
`event_type="agent.suspend"`	This event indicates that the VEN is suspended and no longer protecting the workload. If you did not intention-

Events	Description
	ally 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"`	This event indicates tampering with 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 event's cause. **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 **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 Segmentation for Data Centers 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"`
	These events contain labels indicating when all workloads are affected, all services, or a label/label-group is used as a rule source or destination.
	
	Illumio Segmentation for Data Centers generates these events when you create, update, or delete a draft ruleset. The removed or added labels could represent high-value applications or environments.

Events	Description
	Recommendation: Monitor these events for high-value labels.
`event_type="sec_policy.create"`	Illumio Segmentation for Data Centers contains the `workloads_affected` field, including the number of workloads a policy affects. This event is generated when you provision a draft policy, updating the policy for 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 affected workloads. If the number exceeds an acceptable threshold, investigate the associated policy.
`event_type="agent.clone_detected"`	The PCE detects cloned VENs based on clone token mismatch. The volume of these events makes the severity level important, 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.

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": [{
"source": true,
"actors": "ams"
}]
},
"destinations": {
"created": [{
"source": false,
"actors": "ams"
}, {
"source": 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 the 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
```


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 [276].

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.goodbye`	Agent disconnected
`agent.machine_identifier`	Agent machine identifiers updated
`agent.refresh_token`	Agent refreshed token
`agent.refresh_policy`	Success or failure to apply policy on VEN
`agent.request_upgrade`	VEN upgrade requested
`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_cont ainers`	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

JSON Event Type	Description
`api_key.create`	API key created
`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_services`	Container cluster services updated as 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_autoclean-up_started`	DB temp table cleanup started.
`database.temp_table_autocleanup_com-pleted`	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

JSON Event Type	Description
`enforcement_boundary.update`	Enforcement boundary updated
`event_settings.update`	Event settings updated
`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 created or updated.

JSON Event Type	Description
`login_proxy_ldap_config.create`	Interservice call to the login service to create the LDAP config.
`login_proxy_ldap_config.delete`	Interservice call to the login service to delete the LDAP config.
`login_proxy_ldap_config.update`	Interservice call to the login service to update the LDAP config
`login_proxy_ldap_config.verify_connection`	Interservice call to the login service to verify the connection to the LDAP server.
`logout_from_jwt`	User logged out
`lost_agent.found`	Lost agent found
`network.create`	Network created
`network.delete`	Network delete
`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.network_devices_network_endpoints_workloads`	Workload added to the network endpoint.

JSON Event Type	Description
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
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 the 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
request.authentication_failed	API request authentication failed

JSON Event Type	Description
`request.authorization_failed`	API request authorization failed
`request.internal_server_error`	API request failed due to an internal server error
`request.service_unavailable`	API request failed due to an unavailable service
`request.unknown_server_error`	API request failed due to an 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 destination services updated
`saml_config.create`	SAML configuration created
`saml_config.delete`	SAML configuration deleted
`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
`secure_connect_gateway.create`	SecureConnect gateway created

JSON Event Type	Description
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_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
slb.create	Server load balancer created
slb.delete	Server load balancer deleted.
slb.update	Server load balancer updated.
support_report_request.create	Support report requested
support_report_request.delete	Deleted a request for a support report
support_reports	Support report added
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	The agent missed heartbeats.
system_task.agent_offline_check	Agents marked offline
system_task.prune_old_log_events	Event pruning completed

JSON Event Type	Description
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
user.login_session_terminated	User login session terminated.
user.logout	User logged out
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 an expired password.
user_local_profile.create	User local profile created.
user_local_profile.delete	User's local profile deleted.
user_local_profile.reinvite	Invitation email resent for the local user.
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

JSON Event Type	Description
`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
`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 created or updated.
`workload.create`	Workload created
`workload.delete`	Workload deleted
`workload.online`	Workload online
`workload.recalc_rules`	Workload policy recalculated
`workload.redetect_network`	Workload network re-detected
`workload.undelete`	Workload undeleted
`workload.update`	Workload settings updated
`workload.upgrade`	Workload upgraded
`workload_interface.create`	Workload interface created

JSON Event Type	Description
workload_interface.delete	Workload interface deleted
workload_interface.update	Workload interface updated
workload_interfaces.update	Workload interfaces updated For example, IP address changes, a new interface is added, and an interface shuts down.
workload_service_report.update	Workload service report updated
workload_settings.update	Workload settings updated
workloads.apply_policy	Workload 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.
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

View and Export Events

You can view events by default in the PCE web console or 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 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 those that interest you most. You can search for Organization events by their universally unique identifier (UUID) and filter events by 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. You can filter events by criteria such as severity, status, timestamp, or who generated them.

> **NOTE**
> 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	

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 the PCE Web Console

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

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Export+Events+using+PCE+Console.mp4

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 events, filter the list and then click **Export** > **Export Filtered** to export only the filtered view.
4. Use the search filter for events based on event type, severity, status, timestamp, and who generated them.

Events Settings

The following section describes configuring the Events Settings in the PCE web console.

NOTE
Information about Event Settings applies only to the on-premises PCE.

Events Are Always Enabled

Events are enabled by default in the PCE and cannot be disabled by 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, ranging from 1 day to 200 days, with a default period of 30 days.
- **Event Pruning:** The system automatically prunes events based on disk usage and their age; events older than the retention period are pruned. When pruning is complete, the `system_task.prune_old_log_events` event is recorded.
- **Event Format:** This setting sets the message output to one of three formats. The selected message output format only applies to messages sent over Syslog to an 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	This should be corrected immediately.
Critical	Critical conditions
Error	Error conditions
Warning	Might indicate that an error will occur if action is not taken
Notice	Unusual events, 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 remain recorded in JSON.

Set Event Retention Values

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

If you use an 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. When setting it to 7 days, you can use the PCE Troubleshooting or Events Viewer to troubleshoot and diagnose events quickly. The benefit of setting it to 7 days is that if an issue occurs on a Friday, it can still be diagnosed the following Monday. Many events are generated in a dynamic environment, increasing the data stored (disk space used), backup size, etc. 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, such as the usage of APIs, ETL, Containers, and so on.

If you use a SIEM in a non-dynamic environment, consider setting the event retention period to 30 days. In a non-dynamic environment, fewer events are generated, and less disk space is used.

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 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 the primary storage of events is not on the PCE	7 days (PCE troubleshooting) 1 day (minimum)
No: Not primary storage for events	If events are stored primarily on the PCE, consider the organization's record retention policy and 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 outputs 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 the 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

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 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

  When the soft limit is reached, aggressive pruning is triggered. 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" 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.

Requirements for Events Framework

To use the events framework, 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 of 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 information about the default events data retention period, database dumps with and without events data, disk compacting, and more, see Manage Data and Disk Capacity in the PCE Administration Guide.

Events Preview Runtime Setting

If you participated in the preview of Events in 18.1.0, you enabled it 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 will not succeed.

Removing this preview parameter does not affect the ongoing recording of "organization events" records.

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 other previews, you can also remove the line `enable_preview_features`.

2. Save your changes.

SIEM Integration for Events

Event data can be sent using syslog to your own analytics or SIEM systems for analysis or other needs.

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) is needed for SIEM applications like Splunk®.
- Common Event Format (CEF) is needed for Micro Focus ArcSight®.
- Log Event Extended Format (LEEF) is needed for IBM QRadar®.

Illumio Tools for SIEM Integration

Illumio offers other tools for SIEM integration.

Illumio App for Splunk:

- Software: Technical Add-on for Illumio and Illumio App for Splunk
- Documentation: Illumio App for Splunk Guide 4.x

Illumio App for QRadar:

- Software: Illumio App for QRadar
- Documentation: Illumio App for QRadar Guide 1.4.0

Illumio App for ServiceNow:

- Software: Illumio App for CMDB
- Documentation: Illumio App for ServiceNow 2.1.0

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 are 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**.

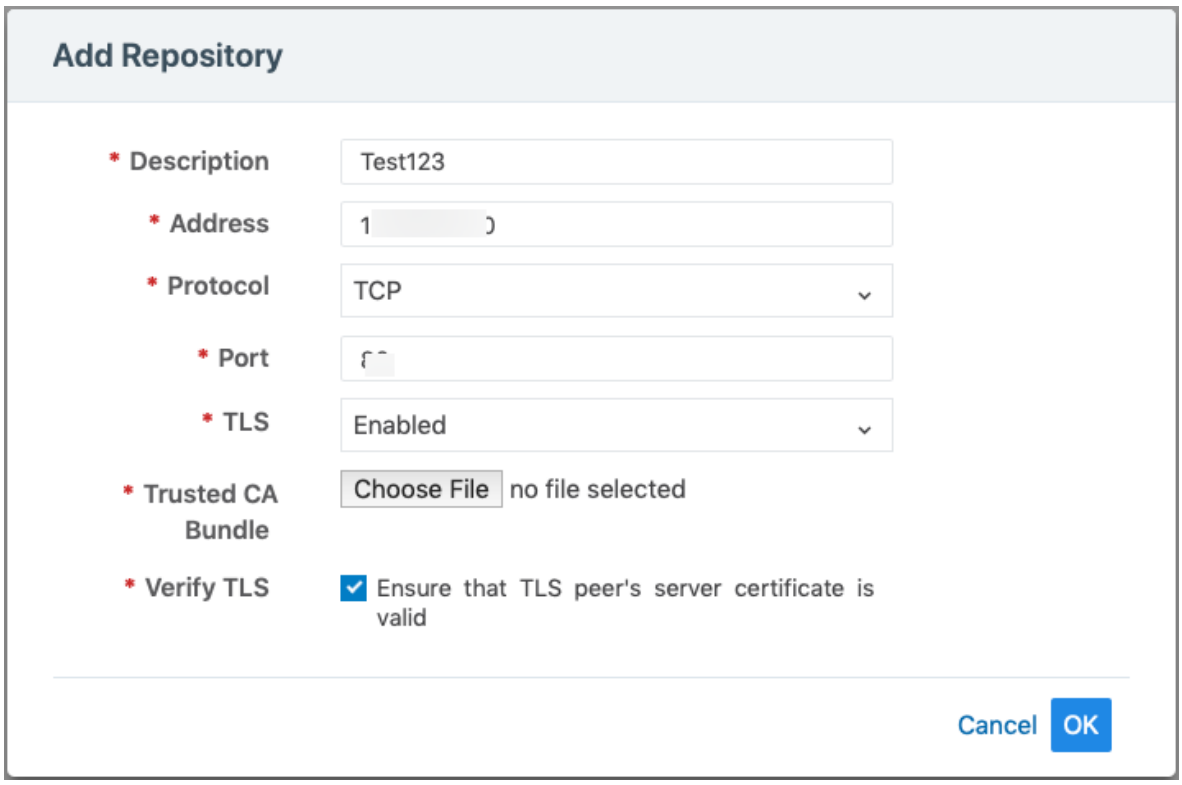

4. In the Add Repository dialog:
   - Description: Enter the name of the Syslog server.
   - Address: Enter the IP address for the Syslog server.
   - Protocol: Select TCP or UDP. If you select UDP, you only need to enter the port number and click **OK** to save the configuration.
   - Port: Enter the 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 where it is stored.

     The Trusted CA Bundle contains all the certificates 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 the TLS peer's server certificate is valid.
5. Click **OK**.

After ensuring that the events are being forwarded as configured to the correct external Syslog servers, you can stop using the "Local" server by editing the local server setting and deselecting all message types.

NOTE

You cannot delete the “Local” server.

Disable Health Check Forwarding

PCE system health messages are helpful for PCE operations and monitoring. If they are needed at the remote destination, you can choose to forward them.

For example, IBM QRadar is usually used by security personnel who might not need to monitor the health of the PCE system. The Illumio App for QRadar does not process the PCE system health messages.

The PCE system health messages are only 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.

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Disable+Health+Check+Forwarding.mp4

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 health messages in the system. If you are using either of these for SIEM, do not select the System Health Messages checkbox.

Showing Rule ID in Syslog

For large customers handling 10K+ messages per second, including rule IDs in the Syslog events will substantially increase the volume of recorded data.

In release 25.1.0, an organization-level feature flag `rule_info_exposure_to_syslog` (disabled by default) was added. This flag controls whether rule ID information is included in the syslog messages:

`rule_info_exposure_to_syslog`

To add the rule IDs to the syslog events, the API `optional_features_put` was changed by adding the new property `rule_info_exposure_to_syslog`.

To provision the firewall settings via the PCE console, follow these steps:

1. In the PCE console, go to **Settings > Event Settings** .
2. In the Event Settings dialog, click Add next to Event Forwarding.
3. Select **Local**.
4. Select check boxes for all events: Organizational Events, System Events, Allowed, Potentially Blocked, Blocked, and System Health Messages.
5. Click **Save**.

Enabling the Rule Data via API

To set the flag `enable_all_rule_hit_count_enabled` via API, use the following CURL command:

```
curl -u api_${ILO_API_KEY_ID}:${ILO_API_KEY_SECRET}
-H "Content-Type: application/json" -X
PUT -d '{"rule_hit_count_enabled_scopes":
[[]]}' https://${ILO_SERVER}/api/v2/orgs/${ILO_ORG_ID}/
sec_policy/draft/firewall_settings
```

For more details about using the Rule ID feature using the API, see Showing Rule ID in Syslog.

Traffic Flow Types and Properties

Visibility Settings

The table below indicates whether or not a traffic summary is logged as Allowed, Blocked, or Potentially Blocked according to a workload's visibility setting.

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 that the attribute might not appear in the summary.

Event Attributes	Allowed (pd=0)	Potentially Blocked (pd=1)	Blocked (pd=2)
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*	✓	✓	✓

Showing the Data Transfer Amount

The JSON, CEF, and LEEF for the accurate byte count work events are related to the 'Show Amount of Data Transfer' preview feature, which is available with the 20.2.0 release.

The PCE now reports the amount of data transferred into and out of workloads and applications in a data center. The number of bytes sent and received by an application's source is provided separately. These values can be seen in traffic flow summaries streamed from the PCE. You can enable this capability on a per-workload basis in the Workload page. You can also enable it in the pairing profile to directly pair workloads into this mode.

The direction reported in the flow summary is from the viewpoint of the source of flow:

Destination Total Bytes Out (

```
dst_tbo
```

): Number of bytes transferred out of source.

Destination Total Bytes In (

```
dst_tbi
```

): Number of bytes transferred to source.

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`

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":"2017-0
8-16T13:23:57-07:00",
"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":"172.1
6.40.5",
"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,"d
st_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
```

Export Traffic Flow Summaries

Decide where to export the traffic flow summaries: Syslog or Fluentd.

CAUTION

By default, the PCE generates all traffic flow summaries and sends them to Syslog.

If you have not configured Syslog, the Syslog data is written to a local disk by default. 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** > **Events Settings**.
2. Enable a remote Syslog destination.
3. Select specific traffic flow summaries to be sent to the remote syslog.

This filters the selected traffic flow summaries and sends 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.

NOTE

The generation of all traffic flow summaries is implemented to ensure that they can be controlled only from the PCE web console.

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 the local, which sends to the 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 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

VENs send two attributes to the Syslog and fluentd output. These 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 lets you 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.

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. The default is approximately 600 seconds (10 minutes), depending on the VEN's ability to report traffic and the 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 is associated with the first flow in the summary. This value exists only if the protocol is ICMP (1). Example: 3 for "Destination Unreachable."	Integer	No	Minimum=0 Maximum=255
`code`	The ICMP message code (subtype) is associated with the first flow in the summary. This value exists only if the protocol is ICMP (1). Example: 1 for "Destination host unreachable."	Integer	No	Minimum=0 Maximum=255
`dst_port`	Destination port. This value exists only if the protocol is not TCP (6) or UDP (17).	Integer	Yes	Minimum=0 Maximum=65535

Property	Description	Type	Required	Possible Values
`pd`	This is the 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, specifically those 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`	The session state for the traffic flows is in the flow summaries. Possible values: • **Active (A):** The connection was still open when the flow summary was logged. This applies to allowed and potentially blocked flows. • **Closed (C):** (Linux only) The connection was closed when the flow summary was logged. Applies to allowed and potentially blocked flows.	String	No	A, C, T, S, N

Property	Description	Type	Required	Possible Values
	• **Timed out (T):** The Connection timed out when 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 when 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 associated with a new connection. Applies to blocked TCP flows. The value is empty for blocked UDP flows.			
pn	The program name is associated with the first flow of the summary. It is supported on inbound flows for Linux and Windows VEN and 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 that this value might represent only the first process detected across all aggregated flows. No process is associated if network communication is done by an OS component (or a driver).	String	No	
un	The username is associated with the first flow of the summary. It is supported on inbound flows for Linux and Windows VEN and outbound flows for only Linux VEN. On Windows, it can include the username of the user account that initiated the connection.	String	No	

Property	Description	Type	Required	Possible Values
	NOTE This information might not be available on short-lived processes.			
`sn`	The service name associated with the first flow in the summary 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.	Object	No	
	NOTE The `src_hostname`, `src_href`, and `src_labels` values are not included in a traffic summary if the source of the flow is not an Illumio-labeled workload, such as Internet traffic or a managed workload without any labels applied.			
`dst_hostname`	Hostname of the destination workload that reported the flow.	String	No	
`dst_href`	HREF of the destination workload that reported the flow.	String	No	
`dst_labels`	Labels applied to the destination workload.	Object	No	
	NOTE The `dst_hostname`, `dst_href`, and `dst_labels` values are not included in a traffic summary if the flow's destination is not an Illu-			

Property	Description	Type	Required	Possible Values
	mio-labeled workload, such as Internet traffic or a managed workload without any labels applied.			
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 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 subproperties 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 vulnerability scores 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

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 Segmentation for Data Centers.

About This Administration Guide

This guide shows you how use `illumio-ven-ctl` (for Linux, AIX, and Solaris) and `illumio-ven-ctl.exe` (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 Segmentation for Data Centers 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.

Pairing or Activating the VEN

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*.

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 VEN is unpaired from the PCE, 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.

Mounted VEN Guidelines

If a VEN is installed on a mounted directory, the mount point must remain stable throughout its lifecycle, from VEN activation to unpairing. Follow these guidelines to avoid operational issues.

- **Mount Before Installing**

  Ensure the file system designated for VEN software and data is mounted before installing the VEN package.
- **Unpair Before Unmounting**

  If the VEN is installed or activated from a mounted directory, the mount point must remain mounted until after the VEN is unpaired. Unmounting the drive or directory before unpairing the VEN can lead to operational issues or data loss.
- **Avoid Symbolic Links**

  The absolute path to the VEN software and data install locations must not traverse any symbolic links.

- **Do Not Relocate Post-Install**
  After installation, the VEN software and data locations are fixed. Do not re-mount the file system to a different path or use symbolic links to move the VEN directories. Follow documented upgrade procedures when relocating VENs.
- **Mount at Boot**
  When the OS is booted, the VEN software will start automatically at the same time as OS networking services. This means that the file system that stores the VEN software and data should already be mounted at this time; that is, the VEN software and data should not be stored on a remote/networked file system.
- **Keep Mounted During Operation**
  The file system that stores the VEN software and data must remain mounted while the VEN is operational and while any VEN administrative actions are performed (including unpair and uninstall).

VEN Architectural Diagram

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

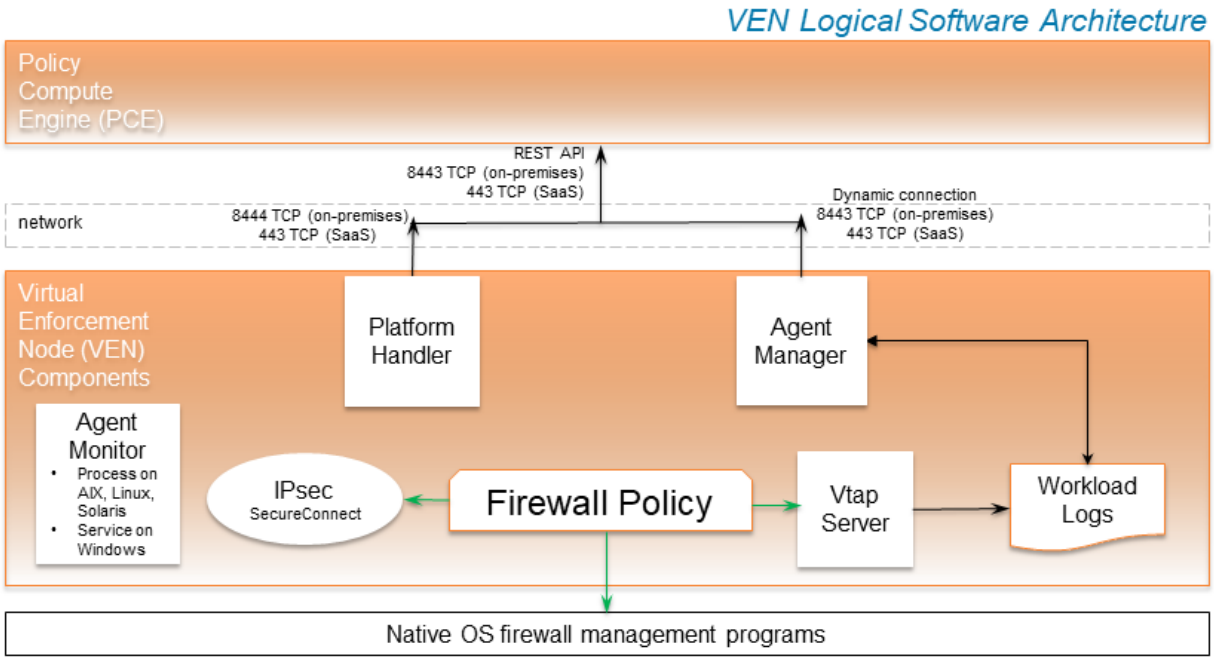

1. The VEN reports to the PCE the status of the workloads.
2. The PCE computes a unique security policy for each managed workload and transmits it to the VEN.
3. The VEN receives the policy and it programs a firewall by using the firewall platform of the OS. The VEN supports the following firewall platforms:
   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
Agent-Manager	• 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
Platform-Handler	• 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.

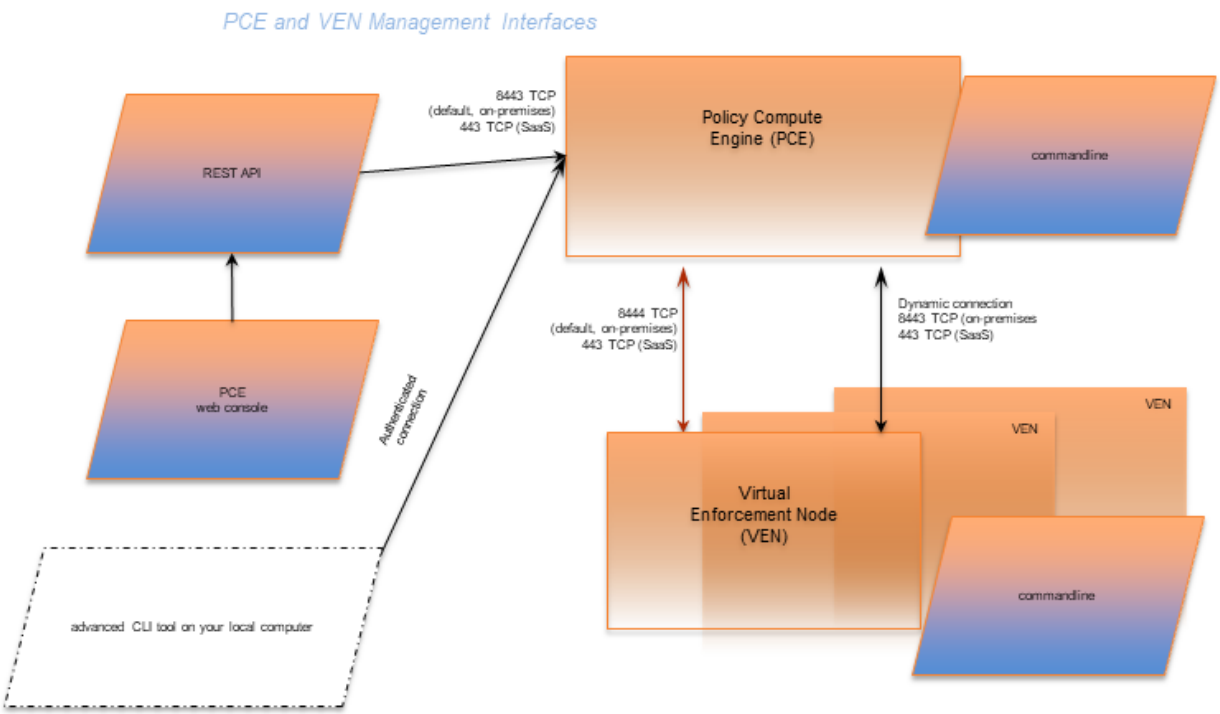

Interface	Notes	See...
PCE web console	With the PCE web console, you can perform many common tasks for managing Illumio Segmentation for Data Centers.	Security Policy Guide
PCE command line	Use of the command line directly on the PCE. A primary management tool on the PCE is the command line `illumio-pce-ctl` control script. You can perform many common tasks for managing the Illumio Segmentation for Data Centers 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 Segmentation for Data Centers 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

Support for Windows 10 Enterprise Multi-session

The VEN supports Windows 10 Enterprise Multi-session, versions 1809 and 1903, which allow you to have multiple concurrent interactive sessions. The number of concurrent sessions depends on your user activity and hardware resources such as vCPU, memory, disk, and vGPU. Windows 10 Enterprise multi-session platform can be Azure AD-joined to a production network. It is not supported for on-premises environments. The software image is downloadable from the Azure Gallery. For more information on its features, see Microsoft documentation.

About VEN Administration on Workloads

This topic describes the workload policy states and VEN enforcement characteristics you should 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	
		Enhanced Data Collection	VEN logs connection information for allowed, blocked, and potentially blocked traffic
			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	
		Enhanced Data Collection	VEN logs connection information for allowed, blocked, and potentially blocked traffic
			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	
		Enhanced Data Collection	VEN logs connection information for allowed, blocked, and potentially blocked traffic
			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 Segmentation for Data Centers 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

Feature	Initial Release
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	IP-Filter	IP-Filter/PF	PF
Firewall coexistence	✓	✓	✓	✓	✓	✓	-	-	✓
Container support	-	-	✓	✓	✓	✓	-	-	-
IPv6	✓	✓	✓	-	✓	✓	-	✓	✓
PCE repo upgrade	✓	✓	✓	✓	-	✓	-	-	✓
Aggressive Tampering Detection	✓	✓	✓	✓	-	-	-	-	-
Process-based policies	✓	✓	-	-	-	-	-	-	-
Extended process path/args (vtap)	✓	✓	✓	✓	✓	✓	✓	✓	✓
Flow-byte counting	✓	✓	✓	-	-	-	-	-	-

Feature	Windows	Windows Edge	Linux	RHEL 5	C-VEN	CentOS8	AIX	Solaris	MacOS (Endpoint)
Kerberos	✓	✓	✓	✓	✓	✓	✓	✓	✓
FIPS	✓	✓	✓	✓	✓	✓	✓	✓	-
FQDN Policies	✓	✓	✓	-	✓	✓	-	-	✓
FQDN Traffic reporting	✓	✓	✓	✓	✓	✓	-	-	-
IPSec (SecureConnect)	✓	✓	✓	✓	✓	✓	-	-	-
Installer	MSI; EXE (from 21.2.1)	MSI; EXE (from 21.2.1)	pkg	pkg	apk; rpm (from 19.3.2)	pkg	bff	pkg	dmg
Pairing script (one-liner from PCE UI)	✓	✓	✓	✓	✓	✓	✓	-	✓
Process-based policies	✓	✓	o e-bpf	o e-bpf	-	-	-	-	o (P1) network-extension

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 [327].

  > **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) [327].
- **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 Segmentation for Data Centers 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 Segmentation for Data Centers 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.

Firewall Coexistence

To provide additional security, you can supplement Illumio's firewall with your organization's firewalls using Firewall Coexistence. You can set the Illumio firewall 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.

NOTE

Using Firewall Coexistence requires careful consideration

Illumio cannot prevent any non-Illumio processes from programming the firewall, so interference from non-Illumio processes is always possible. The server VEN is able to detect many cases of such interference and will report them as firewall tampering. Although Illumio expects to have exclusive control over the firewall by default (preferred), it is possible to coexist with non-Illumio processes depending on exactly how the non-Illumio processes are programming the firewall.

Because the VEN has no way to know about the actions of non-Illumio processes, coexistence necessarily results in the loss of some visibility and clarity in traffic reporting. For example, server VENs are usually able to coexist with:

- Many versions of stand-alone Docker and other simple containers
- Manually programmed rules, depending on the precise details
- Windows GPO, depending on precise details
- Many anti-virus solutions

Server VENs aren't able to coexist with complex containers such as Kubernetes. For such cases, consider using the C-VEN.

Firewall Coexistence and Endpoint VENs

Endpoint VENs are in Firewall Coexistence by default. This cannot be changed.

IMPORTANT

The Firewall Coexistence feature deprecates these features:

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

Coexistence Mode by Operating System

Coexistence Mode is a key architectural feature of Illumio's enforcement model, designed to support interoperability with native and third-party firewalls. By evaluating rule injection behavior and maintaining policy integrity, Illumio ensures consistent security outcomes across diverse platforms. This section provides an OS-specific breakdown of how Coexistence Mode is implemented and what visibility and control mechanisms are in place.

Windows Server VEN

Coexistence Mode	Windows Firewall [1]	3rd Party WFP Firewalls	Tampering Behavior	Visibility
Exclusive	Illumio **disables** Windows Advanced Firewall/Defender Policy.	Illumio coexists with all 3rd party non-Illumio WFP filters. WFP is fundamentally a coexisting firewall and Illumio **coexists with all non-Microsoft Windows firewalls**.	Standard [2]	The VEN reports all traffic flows passing through the Illumio WFP filter. The reported policy decision is based solely on the Illumio policy.
Coexistence, Primary and Non-Primary	Illumio **coexists with** Windows Advanced Firewall/Defender Policy.	Illumio coexists with all 3rd party non-Illumio WFP filters. WFP is fundamentally a coexisting firewall and Illumio **coexists with all firewalls**.		

Windows Endpoint VEN

Coexistence Mode	Windows Firewall [1]	3rd Party WFP Firewalls	Tampering Detection Behavior	Visibility
Coexistence Primary [3]	Illumio coexists with Windows Advanced Firewall/Defender Policy.	Illumio coexists with all 3rd party non-Illumio WFP filters. WFP is fundamentally a coexisting firewall and Illumio coexists with **all** firewalls.	Standard [2]	The VEN reports all traffic flows passing through the Illumio WFP filter. The reported policy decision is based solely on the Illumio policy.

Linux IPTables

Coexistence Mode	IPtables Rules	3rd Party IPtables Firewalls	Tampering Detection Behavior	Visibility
Exclusive	Illumio rules are inserted at the top of the filter table and all **policy decisions are final**. Any rules appearing after the Illumio firewall **are not** evaluated or processed. Illumio doesn't alter existing services such as `firewalld`[4] and nftables (not to be confused with the kernel module of the same name that implements the firewall itself).	Overwrites or ignores 3rd Party Firewall Rules	Standard [2]	The VEN reports all traffic flows passing through the Illumio IPTables filter. The reported policy decision is based solely on the Illumio policy.
Coexistence Primary	Illumio rules are inserted at the top of the filter table and all **Blocks are final**. Any rules appearing after the Illumio firewall **are** evaluated or processed if the Illumio policy decision was to Allow. Illumio doesn't alter existing services such as `firewalld`[4] and nftables (not to be confused with the kernel module of the same name that implements the firewall itself).	Coexists with 3rd party IPTables firewalls	Will tolerate only rules appearing **after** Illumio in the filter table.	
Coexistence Non-Primary			Will tolerate rules appearing **before or after** Illumio in the filter table.	The VEN does **not** report traffic blocked from rules before Illumio because the Illumio IPTables filter never evaluates that traffic. The VEN reports all traffic flows passing through the Illumio IPTables filter.

Linux NFTables

Coexistence Mode	NFTables Rules	3rd Party nftables Firewalls	Tampering Detection Behavior	Visibility
All Modes: • Exclusive • Coexistence Primary • Coexistence Non-Primary	Illumio doesn't alter existing services such as `firewalld`[4] and nftables (not to be confused with the kernel module by the same name that implements the firewall itself). Illumio rules are inserted into Illumio-specific filters. Non-Illumio filters operating at higher priority may pre-empt Illumio filters, otherwise all **policy decisions are final**. Non-Illumio filters operating at lower priority **are not** evaluated or processed after Illumio filters.	Illumio coexists with all 3rd party non-Illumio nftable spaces. NFT is fundamentally a coexisting firewall, and Illumio coexists with all non-Illumio firewalls.	Illumio doesn't tolerate any rule changes within Illumio nftable spaces.	The VEN does **not** report traffic blocked from rules with a higher priority than Illumio because nftable never evaluates that traffic. The VEN reports all traffic flows passing through the Illumio nftable space.

Typical Firewall State on New Linux Installations

The following table details the typical state of the firewall on new installations of select Linux distributions.

OS version	Firewall Stack	Firewall Service Name	Typical State of Firewall on New OS Installation
CentOS 5 – 6	IP Tables	n/a	n/a
CentOS 7	IP Tables	`firewalld`	Disabled; no default rules; allows all traffic by default
CentOS 8 – 10	NetFilter Tables (aka NFT)		Enabled; active with default rules in "public" zone; restricts inbound traffic to essentially just SSH and DHCP by default
RedHat	see CentOS	`firewalld` (+ nftables)	See CentOS
SUSE 11 – 14	IP Tables	SuSEfirewall2	Enabled; active with default rules in "external" zone; restricts inbound traffic to essentially just SSH and DHCP by default.
SUSE 15 < SP3	IP Tables	`firewalld`	Enabled; active with default rules in "public" zone; restricts inbound traffic to essentially just SSH and DHCP by default.
SUSE 15 >= SP3	NetFilter Tables (aka NFT)		

Solaris/AIX

Coexistence Mode	Behavior
Coexistence is enabled by default and can't be disabled.	Illumio coexists with all 3rd party Packet Filter firewalls. [5]

macOS

Coexistence Mode	Behavior
Coexistence is enabled by default and can't be disabled.	Illumio coexists with all 3rd party Packet Filter firewalls.

Notes

[1] For example, Group Policy, not the Windows Filtering Platform (WFP).

[2] Standard tampering detection behavior means that Illumio reviews external rule injections for policy impact. Illumio permits such injections only if it determines that no conflict with Illumio's defined policy is possible. Non-disruptive injections are tolerated.

[3] On Endpoint VENs, Coexistence mode is enabled by default and can't be disabled.

[4] Services such as `firewalld` are often configured out-of-the-box to block most inbound ports by default. To avoid confusion around potentially conflicting rules, Illumio recommends (but doesn't require) that services like `firewalld` be disabled.

[5] Don't disable these services in Solaris:

- Packet Filter (PF)
  - Service: Managed by the SMF service `svc:/network/firewall`
  - Configuration: Rules are configured in the `/etc/firewall/pf.conf` file.
- IP Filter (IPF)
  - Service: Managed by the SMF service `svc:/network/ipfilter`
  - Configuration: Rules are configured in the `/etc/ipf/ipf.conf` file

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-Illlumio 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.

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

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/PCE/Set_Firewall_Coexistence.mp4

WARNING

- Endpoint VENs are in Firewall Coexistence by default. This cannot be changed.
- 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**: Ilumio Core is the exclusive firewall by default. You can configure firewall coexistence as needed for all workloads, for specific labels, or for both.
- **Firewall Coexistence:**

To add the scope for firewall coexistence:

1. Click **Add**.

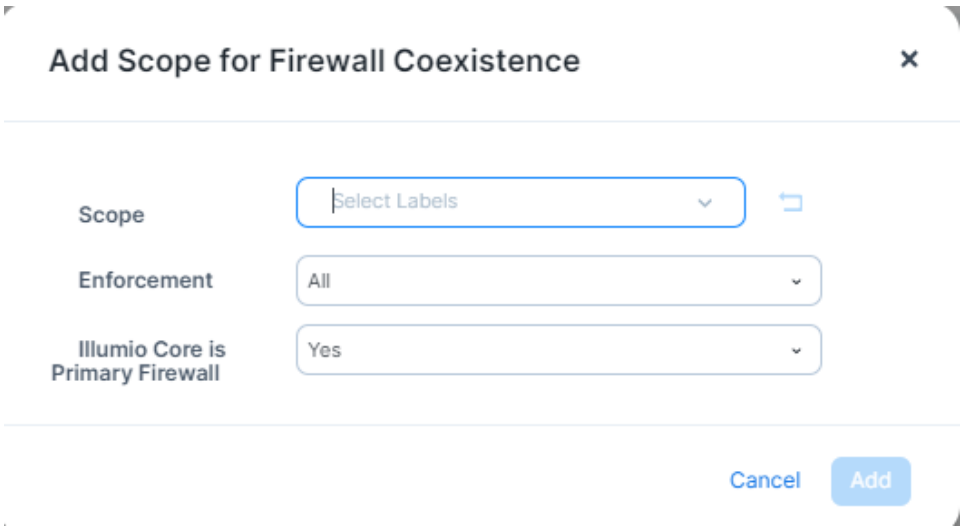

2. From the **Scope** drop-down list, select the labels.

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. When finished, click **Add**.

`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-supportreport|prepare|restart|restore|start|
status|stop|suspend|unpair|unsuspend|version|workloads}
```

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-ctl version
  ```
- 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 Segmentation for Data Centers REST API. With the REST API, the `agent-version` key and value are returned in the payload of every response.

Commonly Used VEN Commands

NOTE

The VEN's runtime_env.yml file is a private configuration file. Illumio advises that you not modify this file directly. To customize the VEN, use environment variables on Linux/Unix hosts or MSI variables on Windows hosts. For more information, see the topics "Linux: Install and Upgrade with CLI and VEN CTL" or "Windows: Install and Upgrade with CLI and VEN CTL" in the VEN Installation and Upgrade Guide.

Platform	Command	Description
Linux & macOS	`/opt/illumio_ven/illumio-ven-ctl`	VEN Linux shell control script to control VEN settings and functions

	`/opt/illumio_ven/illumio-ven-ctl status`	Returns VEN status.
		`Checking Runtime Environment.....`
		`Status for illumio-control:`
		`- Environment Illumio VEN Environment is setup`
		`- venAgentMgr venAgentMgr is running`
		`- IPSec IPSec feature not enabled`
		`- venPlatformHandler venPlatformHandler is running`
		`- venVtapServer venVtapServer is running`
		`- venAgentMonitor venAgentMonitor is running`
		`Agent state: idle`
		`#`

	`ps`	Native OS command to list all system processes
	`chkconfig`	Native OS command to update and query runlevel information for system services
Windows	`C:\Program Files\Illumio\illumio-ven-ctl.exe`	VEN CLI to control VEN settings and functions

	VEN releases 23.5 and earlier:
	`C:\Program Files\Illumio\illumio-ven-ctl.ps1 status`
	VEN releases 24.2.10 and later:
	`C:\Program Files\Illumio\illumio-ven-ctl.exe status`
		Returns VEN and server status.
		`Service venAgentMonitorSvc:  Running`
		`Service venAgentMgrSvc:  Running`
		`Service venPlatformHandlerSvc:  Running`
		`Service venVtapServerSvc:  Running`
		`Service venAgentMonitorSvc:  Enabled`
		`Service venAgentMgrSvc:  Enabled`
		`Service venPlatformHandlerSvc:  Enabled`
		`Service venVtapServerSvc:  Enabled`
		
		`Agent State: enforced`
		`Agent Type: server`

	`Get-Service`	Native OS PowerShell command to display system services
	`tasklist /svc`	Native OS command to display system services
	`wf.msc`	Native OS command to manage the Windows firewall

Platform	Command	Description
AIX/ Solaris	`/opt/illumio_ven/illumio-ven-ctl`	VEN AIX/Solaris shell control script to control VEN settings and functions
	`/opt/illumio_ven/illumio-ven-ctl status`	Returns VEN status. `Checking Runtime Environment..... Status for illumio-control: - Environment Illumio VEN Environment is setup - venAgentMgr venAgentMgr is running - IPSec IPSec feature not enabled - venPlatformHandler venPlatformHandler is running - venVtapServer venVtapServer is running - venAgentMonitor venAgentMonitor is running Agent state: idle #`
	`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 2. Usage

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

/opt/illumio-ven/illumio-ven-ctl <command> [command-options] <subcommand> [sub-
command-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 `scvs` 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 3. Commands by Operating System

Command	Description	Windows	AIX	CentOS	Debian	RHEL & macOS	Solaris	SUSE
activate	Activate VEN	Y	Y	Y	Y	Y	Y	Y
check-env	Check VEN **runtime_env.yml** settings	Y	Y	Y	Y	Y	Y	Y
conncheck	Query VEN policy	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
deactivate [--maintenance-token <token>] [--notify-pce <true | false>]	Deactivate VEN without uninstalling	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
prepare	Prepare VEN image	Y	Y	Y	Y	Y	Y	Y
restart [--maintenance-token <token>]	Restart VEN services	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:	Manage VEN proxy settings	Y	Y	Y	Y	Y for RHEL No for macoS	Y	Y

Command	Description	Windows	AIX	CentOS	Debian	RHEL & macOS	Solaris	SUSE
• IP address of the proxy (for example, `10.10.10.10:8080`) • FQDN of the proxy (for example, `proxy.example.com:8080`) • HTTP or HTTPS schema (for example, `https://proxy.example.com:8080`								
start	Start VEN services	Y	Y	Y	Y	Y	Y	Y
status [-v] [-x | --stdexit]								
status connectivity	Report VEN status	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
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
unpair [--maintenance-token <token>] <saved | open | recommended> [noreport]	Unpair VEN	Y	Y	Y	Y	Y	Y	Y
Subcommands:								

Command	Description	Windows	AIX	CentOS	Debian	RHEL & macOS	Solaris	SUSE
<saved | open | recommended>								
Subcommand arguments:								
[noreport]								
unsuspend [--maintenance-token <token>] [-y]	Unsuspend VEN (exit emergency state)	Y	Y	Y	Y	Y	Y	Y
version	Display VEN version	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.

For a consolidated description of the possible health-related states for VENs and workloads, see VEN and Workload States [382].

VEN Startup and Shutdown

This topic provides information on starting and stopping VENs.

Start Up VENs

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

Automatic VEN 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 VEN 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	• VEN releases 23.2.x and earlier
	`C:\Program Files\Illumio\illumio-ven-ctl.ps1 start`
	• VEN releases 24.2.10 and later
	`<VEN Installation Directory>\Illumio\illumio-ven-ctl.exe start`

Remote VEN Restart

Beginning with VEN release 25.2.10, you can restart a VEN directly from the PCE without physical access to the workload. Remote Restart is similar to other VEN operations that you can initiate from the PCE, such as unpairing and upgrading.

How VEN Remote Restart Works

- After you click **Restart**, the PCE waits for a heartbeat from the VEN before it sends a restart request to the VEN.
- After receiving the restart request from the PCE, the VEN restarts and then sends the Last VEN Service Restart Performed time on the next heartbeat.
- When the PCE receives the Last VEN Service Restart Performed time on the heartbeat, it marks the restart operation completed and displays Last VEN Service Restart Performed on the VEN details page. The reported time remains on the page for 1 hour.
- If you click **Restart** again before the previous restart operation has concluded, a message displays letting you know.
- The restart operation is typically unnoticeable from the PCE UI. The VEN sends a "goodbye" message while stopping and resumes heartbeating after it has restarted. Depending on how long the restart operation takes, the PCE may or may not report the VEN as "inactive."
- The restart operation will not trigger the VEN's offline timer.
- An audit event is logged when the VEN is restarted (see VEN Restart Audit Event [353]).
- As the VEN doesn't flush its policy during this restart operation, policy remains in the kernel.

VEN Restart Audit Event

VEN Remote Restart logs an audit event that captures event details.

Go to **Troubleshoot > Events**.

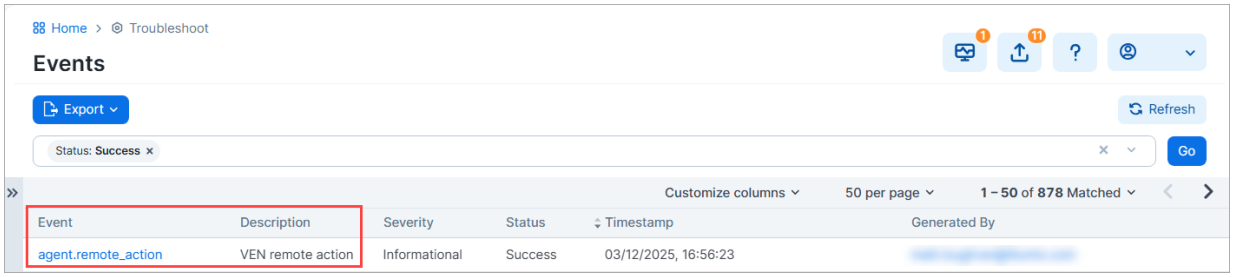

To restart a VEN from the PCE

NOTE

The Restart button is grayed out if the VEN is Suspended or Offline.

1. Go to **Servers & Endpoints > Workoads**.
2. Click the **VEN** tab.
3. In the VEN list page, click the VEN you want to restart.
4. On the VEN's detail page, click **Restart**.

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:	`<VEN Installation Directory>\illumio-ven-ctl.exe enable` `<VEN Installation Directory>\illumio-ven-ctl.exe disable`

VEN Suspension

Suspending a VEN allows you to isolate it 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. This means that the VEN is open and all traffic is allowed.

About VEN Suspension

When a VEN is suspended, the following are 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 Segmentation for Data Centers:

- 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 [358] in this topic.

Suspend and Unsuspend Commands

NOTE

Be aware that suspending the VEN removes all custom iptable filters and NAT rules from the workload. If you want to restore those rules and filters to the workload after you unsuspend the VEN, make sure to back them up **before you suspend a Linux VEN**.

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.`
			Before suspending a Linux VEN, back up the workload's custom iptables filter or NAT rules.
			
			After suspending a workload, restore the rules on the workload because all custom iptables filter or NAT rules are removed from the workload when it is suspended.

Windows	• Suspend
	• Unsuspend
		For Windows VEN releases **24.2 and later**, use `.exe` commands:
		
		`<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.`
		
		For Windows VEN releases **23.5 and earlier**, use PS commands:
		
		`PS C:\Program Files\Illumio> .\illumio-ven-ctl.ps1 suspend`
		`Suspending the VEN... The VEN has been suspended. PCE was notified.`
		
		`PS C:\Program Files\Illumio> .\illumio-ven-ctl.ps1 unsuspend`
		`Unsuspending the VEN... The VEN has been unsuspended. PCE was notified.`	

C-VEN	• Suspend
	• Unsuspend
		`/opt/illumio_ven/illumio-ven-ctl suspend`
		
		`/opt/illumio_ven/illumio-ven-ctl unsuspend`
			For C-VENs, these operations are supported only through `kubectl`, not through the PCE web console.

Mark VEN as Suspended Using the PCE Web Console

In addition to using the commands explained in the previous section, you can mark a workload as **Suspended** throughthe PCE web console (different for C-VENs; see note below).

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.

For C-VENs, suspend and unsuspend are supported only through `kubectl`, not through the PCE web console.

To mark a VEN Suspended:

1. Go to **Workloads and VENs > VENs** (in later versions, **Servers & End-points > Workloads > VENs tab)**.
2. Click the name of the VEN you want to mark as suspended.
3. On the VEN's detail page, click **Mark as Suspended**.
4. 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 **Workloads and VENs > VENs** (in later versions, **Servers & End-points > Workloads > VENs tab**).
2. Click the name of a VEN marked as suspended that you want to mark as unsuspended.
3. On the VEN's detail page, click **Clear Suspension**.
4. Click **Clear** to confirm.

Disable VEN Suspension on Workloads

You can disable the VEN suspension feature by modifying the following VEN environment variable. 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 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:

Examples:

For Windows VEN releases **24.2 and later**, use `.exe` commands:

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

For Windows VEN releases **23.5 and earlier**, use PS commands:

```
PS C:\> Set-ExecutionPolicy -ExecutionPolicy RemoteSigned
PS C:\> msiexec /i  ven_installation_filename.msi
VEN_NO_SUSPEND=1
```

Available options include:

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

AIX VENs

Before installing or upgrading the AIX VEN, enter the following 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 - default/no-action
```

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

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.

ABOUT VENS INSTALLED ON A MOUNTED DIRECTORY

If a VEN is installed on a mounted directory, the mount point must remain stable throughout its lifecycle, from VEN activation to unpairing. For important information about mounted VENs, see Mounted VEN Guidelines [312].

For more information, see Effects of Unpairing VENs

Deactivate Using the VEN Command Line

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

Using `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 the 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 may fail because it will no longer accept any connections 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

Issue `illumio-ven-ctl.exe unpair` to specify the post-deactivation state for the VEN:

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

Unpair Options for Windows VENs

NOTE

On Windows VENs, issuing the `unpair` command without specifying an option simply uninstalls the VEN and removes the Illumio policy from the workload. (It has the same effect as specifying the `saved` command).

- `recommended`: Temporarily allows 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`: Uninstalls the VEN and removes the Illumio policy from the workload. It has the same effect as not specifying any option.

  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`: Use this option when removing a Windows VEN that has never been paired to a PCE; it will leave the firewall configuration unchanged.

Unpair Using System Commands

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

> **IMPORTANT**
> While it is possible to use the `system uninstall` command to unpair the VEN, it is not recommended. You should use that command only if you're unable to unpair with `illumio-ven-ctl` or `illumio-ven-ctl.exe`.

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`

Effects of Unpairing VENs

> ⚠ **ABOUT VENS INSTALLED ON A MOUNTED DIRECTORY**
>
> If a VEN is installed on a mounted directory, the mount point must remain stable throughout its lifecycle, from VEN activation to unpairing. For important information about mounted VENs, see Mounted VEN Guidelines [312].

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 source.

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.

Safeguard or Remove Inactive VENs

This section describes how to safeguard VENs associated with inactive Windows Golden Images from accidental unpairing or deactivation. It also covers how to automatically identify and remove inactive VENs that match easily configurable rules.

Safeguard VENs from accidental unpairing

NOTE

The acronym **GIDV** is used in this topic to mean **Golden Image-designated VEN**.

This feature primarily protects Golden Image VENs from accidental unpairing or deactivation. In an Illumio context, a Golden Image is a standardized template for cloning multiple pre-configured workloads with a specific operating system, VEN version, labels, enforcement state, security patches, applications, settings, and/or hardware specifications.

After a Golden Image is cloned, it's typically inactive for extended periods, and the associated VEN doesn't send heartbeats to the PCE. This inactivity may make the VEN seem unnecessary and lead some users to unpair or deactivate it mistakenly. This feature is designed to prevent that.

Key benefits

- Protects **Golden Image-designated VENs** (**GIDVs**) from accidental unpairing or deactivation. A GIDV must first be undesignated before it can be unpaired or deactivated.
- Applies a unique icon to GIDVs so they're easily identifiable in the PCE.
- Suppresses events that are generated when a GIDV is cloned. This prevents clogging up the PCE's event stream with events unnecessarily. (However, an event is generated when a VEN is designated as a Golden Image.)
- Allows you to filter the VEN List page for GIDVs.

Keep in mind

- Designed primarily for Windows domain-joined VDI images, but works for non-VDI images too.
- Works with VEN releases 24.2.20 or later
- If you select multiple VENs on the VEN List page to perform a bulk unpair and your selection includes a mixture of GIDVs and non-GIDVs, only the non-GIDVs are unpaired.
- Designating a VEN as a Golden Image only prevents it from being unpaired through the PCE UI. It doesn't prevent it from being unpaired through a command line initiated from the VEN.
- Beginning with PCE release 25.4, the prepare script is no longer necessary to install VENs on a Windows Golden Image workload if the VEN is a GIDV.

Designate a VEN as a Golden Image

1. Go to the Servers & Endpoints > Workloads > VENs tab in the PCE Web Console.
2. On the VEN List page, click the VEN you want to designate as a Golden Image.
3. In the VEN's Detail page, move the Golden Image toggle to the right (**Yes**).
4. Click **Save**.
   - On the **List page**, a yellow icon appears next to the Golden Image-designated VEN.
   - On the **Details page > Status section**, **Yes** appears next to Golden Image.
5. To remove the Golden Image designation from a VEN, go to the VEN's Details page, click **Edit**, toggle the setting to **No**, and then click **Save**.

Remove inactive VENs

Some PCE Segmentation environments may have a number of inactive VENs that have stopped sending heartbeats to the PCE. This commonly occurs

in organizations that deliver Illumio VEN-protected Virtual Desktop Infrastructure (VDI) to end-users who may only use the VDI for a short time. Eventually, the VDI is turned off, abandoned, or destroyed. This can also occur when an organization decommissions a server or when a user's VEN-protected laptop is retired. Although the PCE in these cases can't receive heartbeats from the installed VEN (assuming it still exists), a VEN object representing the actual VEN still resides in the PCE database, unnecessarily consuming VEN licenses and possibly leading to unnecessary costs to the customer.

This feature allows you to automate the cleanup of inactive VENs based on easily configurable rules. Matching criteria includes:

- Label(s) assigned to VENs
- Inactivity Threshold (the time since the PCE last received a heartbeat from the inactive VENs you want to remove (minimum 1 day; maximum 90 days))

Specifications

- Maximum daily VEN cleanup limit: 5k VENs
- Performs a query for VENs that match Inactive VEN Cleanup rules every four hours
- You can create a maximum of 10 Inactive VEN Cleanup rules

Events and Notifications

The Inactive VEN Cleanup feature generates the following events and notification:

- `inactive_ven_cleanup_schedule.create`
- `inactive_ven_cleanup_schedule.delete`
- `system_task.inactive_ven_cleanup` lists all the resources that were affected by the cleanup
- Settings > Events > Notifications: **Object creation soft limit exceeded**

Create an Inactive VEN cleanup rule

WARNING

Inactive VEN cleanup rules are executed from the shortest to the longest inactivity threshold defined in your rules.

To avoid unintended data removal, Illumio strongly recommends that you make sure your scheduled VEN cleanup rules are up-to-date.

1. Go to **Settings > VEN Operations**.
2. Click **Add** under **Inactive VENs Cleanup Schedules**.
3. Define up to 10 rules:
   a. Enter a name.
   b. Specify the labels assigned to the inactive VENs you want to cleanup.
   c. **Inactivity Threshold**: Time since the PCE last received a heartbeat from the inactive VENs you want to cleanup. (Minimum 1 day, maximum 90 days.)
   d. Toggle the schedule to **On**.
   e. Click **Save**.

About Overlapping VEN Cleanup Rules

If multiple rules specify one or more of the same labels and the query finds one or more VENs that match the rules, which rule to apply is resolved using the following logic:

1. Shortest Inactivity Threshold wins.
   The rule with the shortest threshold is applied.
2. Tie-breaker: Evaluation order.
   If Inactivity Thresholds are equal, the rule evaluated first is applied.

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.

For a consolidated description of the possible health-related states for VENs and workloads, see VEN and Workload States [382].

VEN-to-PCE Communication

This topic discusses how the VEN communicates with the PCE for both Illumio Segmentation for Data Centers Cloud customers and Illumio Segmentation for Data Centers 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

By specifying a custom time through the PCE User Interface, you can change how long the VEN can go without heartbeating before it enters the Warning state:

1. Go to **Settings** > **Offline Timers**.
2. Select the **Server** or **Endpoint** tab.
3. Click **Edit**.
4. In the **Disconnect and Quarantine** section, select **Custom Timeout**.
5. Specify a wait time.
6. 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 [373].
Dead-peer interval	Configurable	Defaults are: • Server VENs: 60 minutes, or 12 heartbeats • Endpoint VENs: 24 hours See also VEN Offline Timers and Isolation [374].
VEN tampering detection	Within a few seconds on Windows and Linux.	For more information, see Monitor and Diagnose VEN Status [369].

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" (default: 15 minutes for Server VENs, 1 day for Endpoint VENs), 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 defaults are:

- Server VENs: 60 minutes, or 12 heartbeats
- Endpoint VENs: 24 hours

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 source of an application are provided separately. You can see these values in traffic flow summaries streamed out of the PCE. You can enable this capability on a per-workload basis in the Workload page. You can also enable it 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 source of the flow.

- Destination Total Bytes Out (`dst_tbo`): Number of bytes transferred out of source (Connection Responder)
- Destination Total Bytes In (`dst_tbi`): Number of bytes transferred in to source (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.
inter-val_sec T	**Time Interval in Seconds** Duration of latest sampled interval over which the above metrics are valid.

Connec-tion 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, it's `pid`
`IPSec`	Whether or not the `IPSec` feature is enabled
`venPlatformHandler`	`venPlatformHandler` status, and if running, it's `pid`
`venVtapServer`	`venVtapServer status`, and if running, it's `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

VEN releases 23.5 and earlier:

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

VEN releases 24.2.10 and later:

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

Return parameters

Windows

```
{
   "timestamp" : "2019-06-14T00:41:41Z",
   "id" : "xxxxxxxx940d0f4c2531b0d44400523dae055674-
xxxxxxxx7a6796c210fb846b0321847bc22d701e",
   "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 and Workload States

This topic consolidates information about VEN and Workload states and identifies where they appear in the PCE. You can also find much of the same information in other topics throughout Illumio documentation.

Workload Connectivity

Possible states	Definition	PCE UI Locations
Online	The workload is connected to the network and can communicate with the PCE.	• Workload List page • Workload Details page > Summary tab
Offline	The workload is not connected to the network and cannot communicate with the PCE.	
Unmanaged	No VEN is installed on the workload.	

Workload Policy Sync

Possible states	Definition	PCE UI Locations
Active	The most recent policy provisioning was successful, no unwanted changes to the workload's firewall have been reported, and all VEN processes are running correctly.	• Workload List page • Workload Details page > Summary tab
Active (Syncing)	Policy is being applied to the workload currently. Appears if the VEN is not currently heartbeating but the PCE has not received a goodbye event from the VEN, 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. 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	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.	
Error	One of the following errors has been reported by the VEN: • The most recent policy provisioning 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	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.	
Staged (PCE)	The PCE has successfully sent policy to the VEN and it is staged and scheduled to be applied by the user at a later time. Staged appears only if the Policy Update Mode is configured to use Static Policy. For more information, see Policy Update Mode.	

Possible states	Definition	PCE UI Locations
Staged (VEN)	The VEN has received the latest OS-level firewall rules from the PCE but has not applied them.	

VEN Health

NOTE

VEN health is independent from VEN status.

Possible states	Definition	PCE UI Location
Healthy	No specific error or warning conditions related to the VEN and its operation are currently present.	VEN details page > Summary tab
Warning	The VEN has missed 1 or more heartbeats.	
Error	• The VEN has missed heartbeats following an upgrade • The VEN is reported too many interfaces • A cloned VEN is detected	

VEN Status

NOTE

VEN status is independent from VEN health.

Possible states	Definition	UI Location
Active	The PCE is expecting the VEN to heartbeat.	VEN details page > Summary tab
Suspended	Either the VEN was suspended from the CLI and reported it to the PCE, or the user marked the VEN as suspended in the PCE UI. For more information, see VEN Suspension.	
Stopped	The VEN has sent a goodbye message to the PCE and the time specified in the Offline Timer has elapsed. The VEN's IP address is removed from policy. On the Workload list page, the "Connectivity" column is changed to "Status." On the Workload details pages, "VEN Connectivity" is changed to "VEN Status."	

See Also

Monitor and Diagnose PCE Health

VEN State

Monitor and Diagnose VEN Status

PCE Health

List of Event Types

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_view"`
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_flow_view"`

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)

NOTE

Illumio recommends that you upgrade the AIX ip-filter to the latest available version.

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
`IPFILTER_STATE_MAX=<value>`	`fr_statemax`	Maximum number of network connections stored in the state table. You must also set `IPFILTER_STATE_SIZE`.
`IPFILTER_STATE_SIZE=<value>`	`fr_statesize`	Size of the hash table. Must be a prime number. You must also set `IPFILTER_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_MAXBUCKET=<value>`	`fr_state_maxbucket`	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.
`IPFILTER_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.

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) [387].

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.

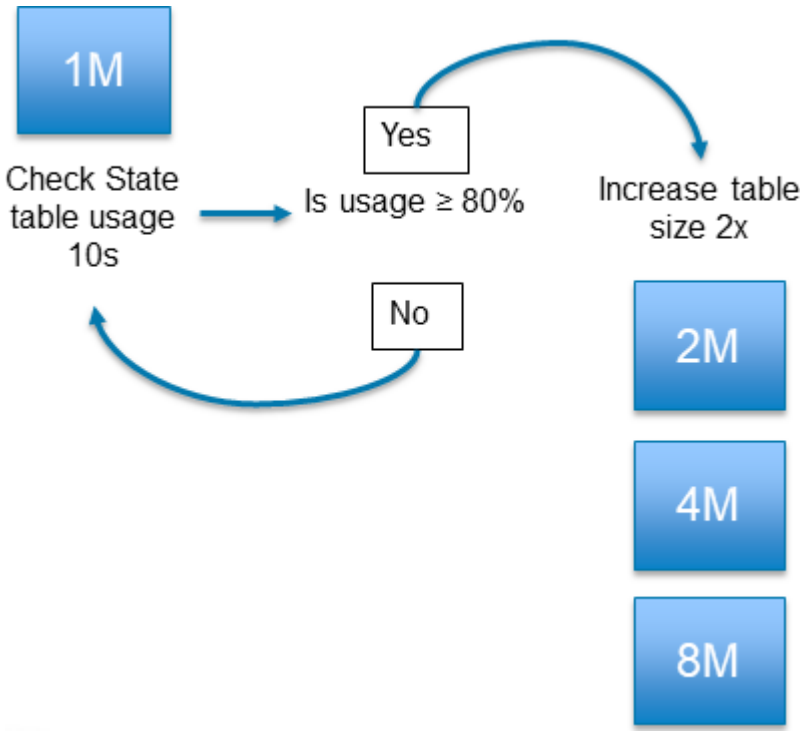

Customizing the VEN Adjustment Behavior

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 minutes), 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.

Work-load 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 Segmentation for Data Centers REST API.
- In the return from a query in Splunk or other SIEM software.

View Tampering Alerts on VEN Host

Workload 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)`

Windows	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 Segmentation for Data Centers 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.tampering -H "Accept: application/
json" -u $KEY:$TOKEN
```

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 "Accept: application/json" -u
$KEY:$TOKEN
```

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 Segmentation for Data Centers 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 Protection.

About Tampering Protection

NOTE

Before using this feature, complete the tasks in Requirements for Using VEN Tampering Protection [397].

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. See VEN Startup and Shutdown [351].
- Restarting the VEN. See To restart a VEN from the PCE [353].
- Suspending the VEN. See VEN Suspension [355].
- Deactivating the VEN. See Deactivate and Unpair VENs [361].

- Unpairing the VEN from the PCE. See Unpair Using the VEN Command Line [362].
- 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 [398].
2. Generate a maintenance token for all VENs or for specific VENs that you want protected. See Generate VEN Maintenance Token. [398]

   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 [399].

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.

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Enable+VEN+Tampering+Protection.mp4

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.

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Generate+a+VEN+Maintenance+Token.mp4

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 [398].

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 when 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 Is 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 [396] 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+L8m6VUV
I9FHRzT31sn37F+
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
eyJhY3Rpb25zIjpudWxsLCJleHBpcmVzX2F0IjpudWxsLCJhZ2VudF9pZHMiOm5
1bGwsIm9yZ19pZCI
6MX0=.MGUCMHSfLNS8yGHgFY0D3CuFvi+L8m6VUVI9FHRzT31sn37F+GsKecpSn
bR8abYuSoz2wgIxAL
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 Support Report from the PCE

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Generate+a+Report+Summary+from+the+PCE.mp4

1. In the PCE web console, go to **Workloads and VENs** or **Workloads** (depending on the PCE release).

2. Click the **VENs** tab.
3. Click the name of a VEN to go to its details page.
4. Click **Generate Support Bundle**. This process can take up to 10 minutes.
5. To view the status of the report, click the **Support Reports** link, which opens the **Support Reports** page and displays the 50 most recent reports that you have generated.
6. Click **Download** to download a report.

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 a Windows Support Report Using a 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.

> **NOTE**
> This applies to VEN releases 24.2.10 and later. For VEN releases 23.5 and earlier, see this guide.

To generate a VEN support report

```
illumio-ve-ctl.exe gen-supportreport
```

To bypass the confirmation prompt

```
illumio-ve-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.

> **NOTE**
> The VEN's runtime_env.yml file is a private configuration file. Illumio advises that you not modify this file directly. To customize the VEN, use environment variables on Linux/Unix hosts or MSI variables on Windows hosts. For more information, see the topics "Linux: Install and Upgrade with CLI and VEN CTL" or "Windows: Install and Upgrade with CLI and VEN CTL" in the VEN Installation and Upgrade Guide.

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 queue
  ```

```
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 --ctstate NEW -j NFLOG --nflog-prefix
"0x800000000000025f " --nflog-threshold 1
    -A tcp_allow_out -d 54.185.43.60/32 -p tcp -m multiport
--dports 443 -m conntrack --ctstate NEW -j ACCEPT
    -A tcp_allow_out -d 204.51.153.0/27 -p tcp -m multiport --
dports 443 -m conntrack --ctstate NEW -j NFLOG --nflog-prefix
"0x8000000000000265 " --nflog-threshold 1
    -A tcp_allow_out -d 204.51.153.0/27 -p tcp -m multiport
--dports 443 -m conntrack --ctstate NEW -j ACCEPT
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 --ctstate NEW -j NFLOG --nflog-prefix
"0x8000000000000266 " --nflog-threshold 1
    -A tcp_allow_out -d 54.185.43.60/32 -p tcp -m multiport
--dports 444 -m conntrack --ctstate NEW -j ACCEPT
```

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**.

Illumio Core PCE CLI Tool Guide 1.4.3

Overview of the CLI Tool

Learn about the CLI Tool, become familiar with the general syntax of the CLI Tool commands, and learn the environment variables you can use to customize the CLI Tool.

IMPORTANT

See the *Illumio Core CLI Tool 1.4.3 What's New document*.

Before You Begin Using the CLI Tool

Before you start using the CLI Tool, review these prerequisites:

- The CLI Tool interacts with the PCEs. Become 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. Learn how vulnerability data is used in the PCE web console. See "Vulnerability Maps" in the Visualization Guide.
- The CLI Tool can be used with workload data. Learn about workloads. 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. Become familiar with these concepts. See "The Illumio Policy Model" in the Security Policy Guide.

CLI Tool Versioning

Illumio Segmentation for Data Centers CLI Tool 1.4.3 is compatible with Illumio Segmentation for Data Centers PCE versions:

- PCE 24.5
- PCE 24.2.10

- PCE 23.5
- PCE 23.2
- PCE 22.5
- PCE 22.2.0 (Standard)
- PCE 21.5.20 (LTS)
- PCE 21.2.4 (LTS)
- PCE 22.1.1 (Standard)

The CLI Tool version numbering is independent of the release and version numbering of Illumio Segmentation for Data Centers PCE and VEN. The CLI Tool works with multiple versions of the PCE and VEN and does not necessarily need software changes in parallel with releases of the PCE or the VEN.

CLI Tool and PCE Resource Management

The Illumio CLI Tool allows you to manage many of your PCE resources directly from your local computer.

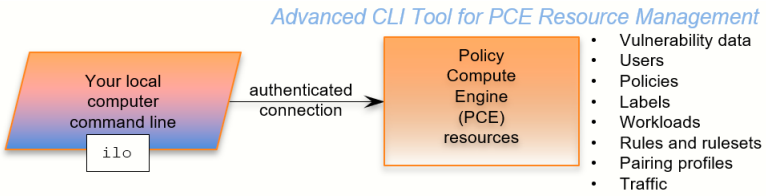

Use the CLI Tool to:

- 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 tool that you can use to work with your PCE resources. 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 Segmentation for Data Centers REST API. No knowledge of the REST API is required.

The `ilo` Command

Learn about the general syntax of the CLI Tool command, `ilo`, and how to use the command-line help to get more specific syntax information.

CLI Tool Formal Syntax

The formal syntax for the `ilo` command is:

```
ilo resource_or_specialCommand argument options
```

Where:

- `resource_or_specialCommand` represents either a resource managed by the PCE or a command unrelated 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 resource's name followed by the argument followed by the `--help` option. For example:

```
ilo workload create --help
```

HTTP Response Codes and Error Messages

Learn about the response codes and error messages that are returned using 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 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:

- Linux: `/tmp`
- Windows: `C:\Windows\Temp`

Environment Variables

Illumio provides Linux environment variables to allow you to customize the 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 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 the certificate file. See "TLS/SSL Certificate for Access to the PCE".
`ILO_DISPLAY_CONFIG`	An absolute path to the display configuration file 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`	The variable that specifies the target host for Tenable

Environment Variable	Purpose
`QAP_HOST`	The variable that specifies the target host for Qualys

Installation and Authentication

Learn how to install the CLI Tool, set up authentication, upgrade the tool, and uninstall it.

Review the prerequisites before you install the PCE CLI Tool.

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

CLI Tool Installation Prerequisites

Review the prerequisites before you install the CLI Tool.

License for Vulnerability Data

The Illumio Segmentation for Data Centers 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 [436].

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 [439].

Install Functional PCE

Because the CLI Tool is for managing resources on your PCE, you must already have installed a fully functional PCE.

Supported Computer Operating Systems

The following operating systems support the CLI Tool:

Linux

- Ubuntu 18.04
- Ubuntu 20.04
- Centos/RHEL 7.9
- Centos/RHEKL 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

TLS/SSL Certificate for Access to the PCE

You need a TLS/SSL certificate to connect to the PCE securely. 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, an 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 Segmentation for Data Centers REST API Version

The CLI Tool uses v2 of the Illumio Segmentation for Data Centers 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 (log-in 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 the 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 an administrator and start the installation program in 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 using 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 [417].

Uninstall the 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-secret yourOwnKeySecret
```

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=e1fa1262461ca2859fcf9d91a0546478d10a1bcc4c57
9d888
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 associated with your login. See Log Into the PCE [419] 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 [426] 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



    | Href                                          |
Name         |
+---------------------
+-------------------------------------------------------------------
--------------------------
| [[-1, -1, nil]]     | [{"href"=>"/api/v2/orgs/28/
sec_policy/8/rule_sets/1909", "name"=>"Default",
"secure_connect"=>false, "peers"=>[{"type"=>"ip_list",
"href"=>"/api/v2/orgs/28/sec_policy/8/ip_lists/188",
"name"=>"Any (0.0.0.0/0)",
"ip_ranges"=>[{"from_ip"=>"0.0.0.0/0"}]}]}] | /api/v2/orgs/28/
sec_policy/8/services/1153 | All Services |
+---------------------
+-------------------------------------------------------------------
--------------------------
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 [426].

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 Name |
+----------+---------+-------+--------------+-----------------
+---------+------------------+
| udp      | 0.0.0.0 | 5355  | svchost.exe  | NETWORK
SERVICE |         | Dnscache         |
...
| tcp      | 0.0.0.0 | 135   | svchost.exe  | NETWORK
SERVICE |         | RpcSs            |
+----------+---------+-------+--------------+-----------------
+---------+------------------+
200, OK
```

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/status
+-------------------------------------------------------------
+-----------------------------+-------------+-------
| Href
| Hostname                    | Log Traffic | Visibility Level
| Mode        | Uid                                  | Status |
+-------------------------------------------------------------
+-----------------------------+-------------+-------
| /api/v2/1/workloads/527b8aca-97aa-43b9-82e1-29b17a947cdd
| hrm-web.webscaleone.info    | false       | flow_summary
| illuminated | 0ffd2290-e26a-4ec6-b241-9e2205c0b730 | active |
| /api/v2/1/workloads/4a8743a4-14ee-40d0-9ed2-990fe3f0ffb1
| hrm-db.webscaleone.info     | false       | flow_summary
| illuminated | 145a3cc8-01a8-4a52-97b8-74264ad690e4 | active |
+-------------------------------------------------------------
+-----------------------------+-------------+----
...
```

Linux: Save Fields for Reuse

On Linux, to easily reuse specific fields, create a display configuration file in YAML format and set the environment variable `ILO_DISPLAY_CONFIG` to point to that file. 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 [426]. 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 [426] 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 Id | Network Detection Mode
+------+------------------+------------
+-------------------------+------------+------------
+--------------------
| eth0 | 10.0.0.16        | 8          |
10.0.0.1                | up         | 1          |
single_private_brn
...
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"=>"enforced"} |
|  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 includes a 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    | Description | Enabled |
+-------------------------------------------------
+------------------------------+---------+-------------+-----
| /api/v2/orgs/28/sec_policy/draft/rule_sets/2387 |
{"href"=>"/api/v2/users/74"} | foo1    |             | true
| /api/v2/orgs/28/sec_policy/draft/rule_sets/1909
| {"href"=>"/api/v2/users/0"}  | Default
|             | true    ...
200, OK
```

The state of the policy is stored in the agent/status/status field. See Nested Resource Fields and Wildcards [424] for information.

Support for Proxy

Release CLI 1.4.3 includes support for authenticated and unauthenticated proxies.

Type the `ilo login --help` command to see proxy-related options.

Table 4. ilo login --help

Command Options	Description
`-v, --verbose`	Verbose logging mode
`--trace`	Enable API Trace Mode
`--server SERVER_NAME`	Illumio API Access gateway server name
`--login-server LOGIN_SERVER`	Illumio login server name
`--kerberos-spn KERBEROS_SPN`	Illumio Kerberos SPN Kerberos authentication is only applicable to --login-server option
`--proxy-server PROXY_SERVER`	proxy server
`--proxy-port PROXY_PORT`	proxy port
`--proxy-server-username PROXY_SERVER_USERNAME`	proxy server username
`--proxy-server-password PROXY_SERVER_PASSWORD`	proxy server password
`--logout`	Logout
`--username USER`	User Name
`--username USER`	User Name
`--auth-token AUTH_TOKEN`	authorization token

Connecting via a Proxy

The command for connecting via an unauthenticated proxy:

```
ilo login --server <fqdn:port> --proxy-server <proxy_ip> --proxy-port <proxy_port> --user-name selfserve@illumio.com
```

An example of connecting via an unauthenticated proxy:

```
ilo login --server 2x2testvc308.ilabs.io:8443 --proxy-server 10.2.184.62 --proxy-port 3128 --user-name selfserve@illumio.com
```

An example of connecting via an authenticated proxy:

```
 ilo login --server 2x2testvc308.ilabs.io:8443 --proxy-
server devtest30.ilabs.io --proxy-port 3128 --user-name
selfserve@illumio.com --proxy-server-username proxy_user --
proxy-server-password proxy_124
```

After the command is executed, users are prompted to enter the PCE user's password, and then a session will be created in the context of the proxy server.

From this point on, all connections/traffic will use the proxy to send traffic.

Using API Keys and Secrets with a Proxy Server

With the command `ilo use_api_key` , you can use an API Key and a secret with a proxy server:

Table 5. ilo use_api_key --help

Command options	Description
`--key-id`	API Key ID
`--key-secret`	API Key Secret
`--org-id`	Illumio Org ID
`--user-id Illumio`	User ID
`-v, --verbose`	Verbose logging mode
`--trace`	Enable API Trace Mode
`--server SERVER_NAME`	Illumio API Access gateway server name
`--login-server LOGIN_SERVER`	Illumio login server name
`--kerberos-spn KERBEROS_SPN`	proxy server
`--proxy-port PROXY_PORT`	proxy port
`--proxy-server-username PROXY_SERVER_USER-NAME`	proxy server username
`--proxy-server-password PROXY_SERVER_PASS-WORD`	proxy server password

The command for using an API Key with an unauthenticated proxy:

```
 ilo use_api_key --key-id <key_id> --key-secret <secret>
--server <pce_fqdn> --org-id <orgid> --proxy-server
<proxy_server> --proxy-port <proxy_port>
```

The command for using an API Key with an authenticated proxy:

```
 ilo use_api_key --key-id <key_id> --key-secret
<secret> --server <pce_fqdn> --org-id <orgid> --proxy-
server <proxy_server> --proxy-port <proxy_port>  --proxy-
server-username <proxy_username> --proxy-server-password
<proxy_password>
```

After a command is executed, all connections/traffic from this point on will use the proxy.

Import and Export Security Policy

You can export and import security policy to and from the PCE using the CLI Tool. Importing and exporting security policy is particularly useful for moving policy from one PCE to another to avoid recreating policy from scratch on the target PCE. For example:

- You can test the policy on a staging PCE and then move it to your production PCE.
- You can move the 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.

Illumio recommends that you base your security policy rules on labels for flexibility. 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. A warning is displayed on export when you have rules tied to individual workloads, virtual services, or virtual servers. Attempts to import such rules fail, and the reason for the failure is displayed.

Example: Failed Attempt to Export Rules for Workload

```
WARNING: rule /orgs/1/sec_policy/active/rule_sets/3/
sec_rules/39 contains non-transferrable providers: workload /
orgs/1/workloads/a51ae67d-472a-44c3-984e-d518a8e95aee
Unable to proceed, please verify input
```

Workflow for Security Policy Export/Import

- Authenticate to the source PCE.
- Export the policy to a file. Syntax summary:

  ```
  ilo sec_policy export --file someExportFilename
  ```
- Authenticate to the target PCE.
- 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.

The 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 [433] 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 [429]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 policies are read from standard input unless you import from a file with the `--file` option.

You can import policy files multiple times. Each import affects only a single copy of a resource.

All imported policies are set in 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 [433] for information.

Policy items already on the target PCE are updated by imported resources whose names match existing resources' names. Services do not have to have the same names. Services match if they have the same set of ports and protocols.

An import does not delete resources. 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 Segmentation for Data Centers 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.

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_name" [debug [v | verbose]
trace]
```

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

Starting from the release of CLI 1.4.0, Qualys supports API downloads 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.

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.

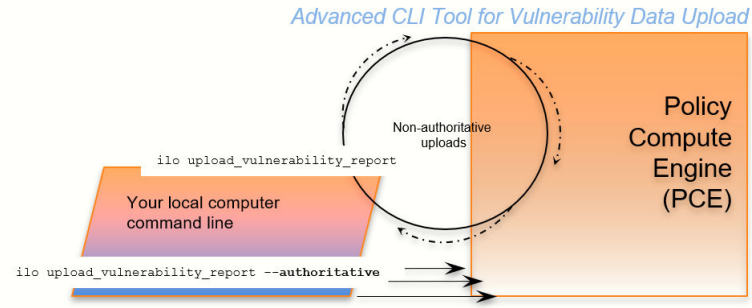

- **Authoritative:** You indicate authoritative data with the -authoritative option. An authoritative upload:
  - Overwrites any previously uploaded records for workloads matched to the incoming records.
  - Eliminates duplicate records.
  - Adds new records not previously written by other uploads.

  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 measures the vulnerability of your entire organization. The score is displayed by the ilo vulnerability list command for all vulnerabilities or individual vulnerabilities via the vulnerability identifier.

Vulnerability Identifier

An 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 [449] 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. This information becomes available if the unmanaged workload is later changed to managed.

Prerequisites for Vulnerability Data Upload

Before uploading vulnerability data, ensure 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 [436] for information.
- XML-formatted vulnerability data files from one of the supported sources.
- Authenticated CLI-tool access to the target PCE.
- Authenticated access and necessary permissions in the PCE web console for working with vulnerability maps.

Vulnerability Data Upload CLI Tool Syntax

The key argument and options 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

`--enable-proxy`
 	
NOTE This is available in CLI Tool 1.4.4.
	No	Use this to enable the proxy between tenable and CLI.
		
		Use this command to enable the proxy:
		
		`ilo upload_vulnerability_report --source-scanner tenable-sc --format api --severities=3  --enable-proxy -v --debug`
		
		Use this command if you do not want to enable the proxy:
		
		`ilo upload_vulnerability_report --source-scanner tenable-sc --format api --severities=3  -v --debug`

`--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.

What	Required	Description
		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 [438] 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 SourceApi-Server:port`
 	
`SERVER_FQDN`
	Yes for	API server FQDN. Allowed formats are `HOST` or `HOST:PORT`

What	Required	Description
	Tenable with `--format api`	

`--api-user ApiServerUserName`
 	
`USERNAME`
	Yes for SourceAPI server authent	The user name for authenticating to the SourceApiServer.
		
		You are always prompted to enter your password.

What	Required	Description
	ication	

`--api-page-size`
 	
`PAGE_SIZE`
	Yes for Qualys and Tenable	Appropriate page size if API supports pagination. The default page is 1000.

`--skip-cert-verification`	Yes for Qualys and Tenab	Disable certificate verification for API.

What	Required	Description
	le	
`--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 Qu	Qualys users can select scan data to process after a specific 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

What	Required	Description
	alys	data scanned after a specific date will be extracted and uploaded. Including a particular scanned-after time is recommended 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 reports. 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 up to four options with the ILO command for the vulnerability data upload. Users who choose to use more optional parameters must 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 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 : 11.
Elapsed Time [4.57 (total : 5.7)] - Imported Vulnerabilities..
Elapsed Time [0.0 (total : 5.7)] - Detected Vulnerabilities
are associated with vulnerability and workload data..
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", "10.1.0.82", "10.1.0.43",
"10.1.0.91", "10.1.0.8", "10.1.1.250"]

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/authoritativedata.xml --authoritative --
source-scanner qualys --format xml

Using /home/centos/.rvm/gems/ruby-2.4.1
Authoritative scan overwites the previous entries for all the
ips within this scan. There is no ROLLBACK
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 : 6.
Elapsed Time [3.02 (total : 15.15] - Imported
Vulnerabilities..
Elapsed Time [0.0 (total : 15.15] - Detected Vulnerabilities
are associated with vulnerability and workload data..
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 [440] 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) [439] and Vulnerability Identifier [440] 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 | Updated By |
----------------------+-------------------------
+----------------------+-----------------------+
| /orgs/1/vulnerabilities/nessus-18405 | Microsoft Windows
Remote Desktop Protocol Server Man-in-the-Middle Weakness
| 51    |              | ["CVE-2005-1794"]
| 2018-11-07T03:15:39.410Z | 2018-11-07T03:15:39.410Z
| {"href"=>"/users/1"} | {"href"=>"/users/1"} |
...
```

Qualys

```
C:\Users\isaac.newton> ilo vulnerability list --xorg-id=1
...
| Href | Name | Score | Description | Cve Ids | Created At |
Updated At | Created By | Updated By |
---------------------+-------------------------
+----------------------+----------------------+
| /orgs/1/vulnerabilities/qualys-38657 |
Birthday attacks against TLS ciphers with 64bit
block size vulnerability (Sweet32) | 69 | | ["CVE-2016-2183"]
| 2018-07-27T18:16:57.166Z | 2018-08-08T22:30:32.421Z
| {"href"=>"/users/1"} | {"href"=>"/users/16"} |
...
```

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) [439] and Vulnerability Identifier [440] 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 | Updated By |
+-------------------------------------------------------
+-------------+----------------------
+--------------------------+----------------------
| /orgs/1/vulnerability_reports/scan_1502310096_09344 |
qualys  | NewAuthoritativeScan | 2018-08-08T22:30:34.877Z |
2018-08-08T22:30:34.877Z | 62 | {"href"=>"/users/16"} |
{"href"=>"/users/16"} |
...
```

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 authenticated to the PCE and run two `ilo bulk_upload_csv` commands.

1. Authenticate to the PCE via API key or explicit login.
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 Kerberos requires, the Kerberos realm name is shown in all capital letters as `MYREALM`.
- 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.

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.
2. Create Kerberos-authenticated `myWorkload1` and set its `mode` to `idle`:

```
ilo workload create --hostname myPCE.BigCo.com
--name myWorkload1 --service-principal-name host/
myKerberosTicketGrantingServer@MYREALM --agent/config/mode
idle
```

For information about how the mode is a nested field, see Nested Resource Fields and Wildcards [424].

3. Create Kerberos-authenticated `myWorkload2` and set its `mode` to `illuminated`:

```
ilo workload create --hostname myPCE.BigCo.com
--name myWorkload2 --service-principal-name host/
myKerberosTicketGrantingServer@MYREALM --agent/config/mode
illuminated
```

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 data sets without waiting 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 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 | Description | Hostname
| Service Principal Name | Public
Ip              | Distinguished Name |
External Data Set | External Data Reference
| 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                  | Updated By
+-------------------------------------------------------------
+---------+------+-------------+----------------
... More lines not shown
-------------------------------------------------------------+
| /orgs/1/workloads/50ce441e-75ac-4be8-9201-96169545019c
| false   |      |             | 10.14.59.8.netstat
...
... 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 [436]. 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-scanner qualys --format xml
```

2. Examine a single uploaded vulnerability record identified by its vulnerability identifier, `qualys-38173`. See Vulnerability Identifier [440] 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-scanner qualys --format xml
   ```

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 --authoritative --source-
   scanner qualys --format xml
   ```

Results

The authoritative vulnerability data has been uploaded and is ready for use in Illumination.

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 1.4.2 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 Segmentation for Data Centers CLI Tool version 1.4.2 is compatible with Illumio Segmentation for Data Centers 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 of PCE and VEN's release and version numbering. 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 of the CLI Tool
- Installation
- The 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 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 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

The Illumio CLI Tool allows you to manage many of your PCE resources directly from your local computer.

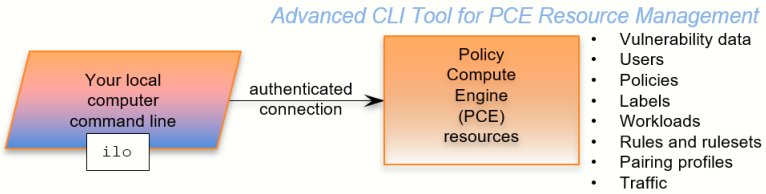

Use the CLI Tool to:

- 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 tool that you can use to work with your PCE resources. 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 Segmentation for Data Centers REST API. No knowledge of the REST API is required.

The `ilo` Command

Learn about the general syntax of the CLI Tool command, `ilo`, and how to use the command-line help to get more specific syntax information.

CLI Tool Formal Syntax

The formal syntax for the `ilo` command is:

```
ilo resource_or_specialCommand argument options
```

Where:

- `resource_or_specialCommand` represents either a resource managed by the PCE or a command unrelated 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 resource's name followed by the argument followed by the `--help` option. For example:

```
ilo workload create --help
```

HTTP Response Codes and Error Messages

Learn about the response codes and error messages that are returned using 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 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:

- Linux: `/tmp`
- Windows: `C:\Windows\Temp`

Environment Variables

Illumio provides Linux environment variables to allow you to customize the 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 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 the certificate file. See "TLS/SSL Certificate for Access to the PCE".
`ILO_DISPLAY_CONFIG`	An absolute path to the display configuration file 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`	The variable that specifies the target host for Tenable

Environment Variable	Purpose
QAP_HOST	The variable that specifies the target host for Qualys

Installation and Authentication

Learn how to install the CLI Tool, set up authentication, upgrade the tool, and uninstall it.

Review the prerequisites before you install the PCE CLI Tool.

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

Installation Prerequisites

This section details the prerequisites for installing 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 Segmentation for Data Centers 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 [436].

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 [439] for information.

Install Functional PCE

Because the CLI Tool is for managing resources on your PCE, you must already have installed a fully functional PCE.

Supported Computer Operating Systems

The CLI Tool is supported by 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

TLS/SSL Certificate for Access to the PCE

You need a TLS/SSL certificate to connect to the PCE securely. 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, an 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 Segmentation for Data Centers REST API Version

The CLI Tool uses v2 of the Illumio Segmentation for Data Centers 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 (log-in 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 the 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 an administrator and start the installation program in 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 using 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 [467].

Uninstall the 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-secret yourOwnKeySecret
   ```

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=e1fa1262461ca2859fcf9d91a0546478d10a1bcc4c57
9d888
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 associated with your login. See Log Into the PCE [470] 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 [426] 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



    | Href                                              |
Name         |
+---------------------
+-------------------------------------------------------------------
--------------------------
| [[-1, -1, nil]]     | [{"href"=>"/api/v2/orgs/28/
sec_policy/8/rule_sets/1909", "name"=>"Default",
"secure_connect"=>false, "peers"=>[{"type"=>"ip_list",
"href"=>"/api/v2/orgs/28/sec_policy/8/ip_lists/188",
"name"=>"Any (0.0.0.0/0)",
"ip_ranges"=>[{"from_ip"=>"0.0.0.0/0"}]}]}] | /api/v2/orgs/28/
sec_policy/8/services/1153 | All Services |
+---------------------
+-------------------------------------------------------------------
--------------------------
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 [426].

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 Name |
+----------+---------+-------+--------------+-----------------
+---------+------------------+
| udp      | 0.0.0.0 | 5355  | svchost.exe  | NETWORK
SERVICE |         | Dnscache         |
...
| tcp      | 0.0.0.0 | 135   | svchost.exe  | NETWORK
SERVICE |         | RpcSs            |
+----------+---------+-------+--------------+-----------------
+---------+------------------+
200, OK
```

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/status
+-------------------------------------------------------------
+-----------------------------+-------------+-------
| Href
| Hostname                    | Log Traffic | Visibility Level
| Mode        | Uid                                  | Status |
+-------------------------------------------------------------
+-----------------------------+-------------+-------
| /api/v2/1/workloads/527b8aca-97aa-43b9-82e1-29b17a947cdd
| hrm-web.webscaleone.info    | false       | flow_summary
| illuminated | 0ffd2290-e26a-4ec6-b241-9e2205c0b730 | active |
| /api/v2/1/workloads/4a8743a4-14ee-40d0-9ed2-990fe3f0ffb1
| hrm-db.webscaleone.info     | false       | flow_summary
| illuminated | 145a3cc8-01a8-4a52-97b8-74264ad690e4 | active |
+-------------------------------------------------------------
+-----------------------------+-------------+----
...
```

Linux: Save Fields for Reuse

On Linux, to easily reuse specific fields, create a display configuration file in YAML format and set the environment variable `ILO_DISPLAY_CONFIG` to point to that file. 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 [476]. 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 [476] 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 Id | Network Detection Mode
+------+------------------+------------
+-------------------------+------------+------------
+--------------------
| eth0 | 10.0.0.16        | 8          |
10.0.0.1                | up         | 1          |
single_private_brn
...
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"=>"enforced"} |
|  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 includes a 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    | Description | Enabled |
+-------------------------------------------------
+------------------------------+---------+-------------+-----
| /api/v2/orgs/28/sec_policy/draft/rule_sets/2387 |
{"href"=>"/api/v2/users/74"} | foo1    |             | true
| /api/v2/orgs/28/sec_policy/draft/rule_sets/1909
| {"href"=>"/api/v2/users/0"}  | Default
|             | true    ...
200, OK
```

The state of the policy is stored in the agent/status/status field. See Nested Resource Fields and Wildcards [424] for information.

Import and Export Security Policy

You can export and import security policy to and from the PCE using the CLI Tool. Importing and exporting security policy is particularly useful for moving policy from one PCE to another to avoid recreating policy from scratch on the target PCE. For example:

- You can test the policy on a staging PCE and then move it to your production PCE.
- You can move the 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.

Illumio recommends that you base your security policy rules on labels for flexibility. 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. A warning is displayed on export when you have rules tied to individual workloads, virtual services, or virtual servers. Attempts to import such rules fail, and the reason for the failure is displayed.

Example: Failed Attempt to Export Rules for Workload

```
WARNING: rule /orgs/1/sec_policy/active/rule_sets/3/
sec_rules/39 contains non-transferrable providers: workload /
orgs/1/workloads/a51ae67d-472a-44c3-984e-d518a8e95aee
Unable to proceed, please verify input
```

Workflow for Security Policy Export/Import

- Authenticate to the source PCE.
- Export the policy to a file. Syntax summary:

  ```
  ilo sec_policy export --file someExportFilename
  ```
- Authenticate to the target PCE.
- 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.

The 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 [481] 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 [429]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 policies are read from standard input unless you import from a file with the `--file` option.

You can import policy files multiple times. Each import affects only a single copy of a resource.

All imported policies are set in 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 [481] for information.

Policy items already on the target PCE are updated by imported resources whose names match existing resources' names. Services do not have to have the same names. Services match if they have the same set of ports and protocols.

An import does not delete resources. 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 Segmentation for Data Centers 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.

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_name" [debug [v | verbose]
trace]
```

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

Starting from the release of CLI 1.4.0, Qualys supports API downloads 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.

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.

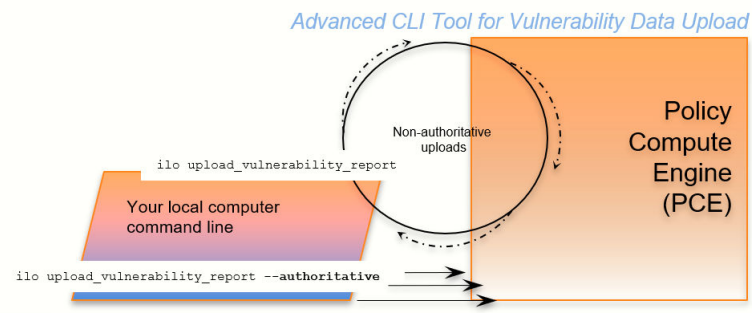

- **Authoritative:** You indicate authoritative data with the -authoritative option. An authoritative upload:
  - Overwrites any previously uploaded records for workloads matched to the incoming records.
  - Eliminates duplicate records.
  - Adds new records not previously written by other uploads.

  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 measures the vulnerability of your entire organization. The score is displayed by the ilo vulnerability list command for all vulnerabilities or individual vulnerabilities via the vulnerability identifier.

Vulnerability Identifier

An 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 [496] 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. This information becomes available if the unmanaged workload is later changed to managed.

Prerequisites for Vulnerability Data Upload

Before uploading vulnerability data, ensure 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 [483] for information.
- XML-formatted vulnerability data files from one of the supported sources.
- Authenticated CLI-tool access to the target PCE.
- Authenticated access and necessary permissions in the PCE web console for working with vulnerability maps.

Vulnerability Data Upload CLI Tool Syntax

The key argument and options 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

`--enable-proxy`
 	
NOTE This is available in CLI Tool 1.4.4.
	No	Use this to enable the proxy between tenable and CLI.
		
		Use this command to enable the proxy:
		
		`ilo upload_vulnerability_report --source-scanner tenable-sc --format api --severities=3  --enable-proxy -v --debug`
		
		Use this command if you do not want to enable the proxy:
		
		`ilo upload_vulnerability_report --source-scanner tenable-sc --format api --severities=3  -v --debug`

`--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.

What	Required	Description
		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 [485] 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 SourceApi-Server:port`
 	
`SERVER_FQDN`
	Yes for	API server FQDN. Allowed formats are `HOST` or `HOST:PORT`

What	Required	Description
	Tenable with `--format api`	

`--api-user ApiServerUserName`
 	
`USERNAME`
	Yes for SourceAPI server authent	The user name for authenticating to the SourceApiServer.
		
		You are always prompted to enter your password.

What	Required	Description
	ication	

`--api-page-size`
 	
`PAGE_SIZE`
	Yes for Qualys and Tenable	Appropriate page size if API supports pagination. The default page is 1000.

`--skip-cert-verification`	Yes for Qualys and Tenab	Disable certificate verification for API.

What	Required	Description
	le	
`--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 Qu	Qualys users can select scan data to process after a specific 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

What	Required	Description
	alys	data scanned after a specific date will be extracted and uploaded. Including a particular scanned-after time is recommended 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 reports. 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 up to four options with the ILO command for the vulnerability data upload. Users who choose to use more optional parameters must 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 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 : 11.
Elapsed Time [4.57 (total : 5.7)] - Imported Vulnerabilities..
Elapsed Time [0.0 (total : 5.7)] - Detected Vulnerabilities
are associated with vulnerability and workload data..
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", "10.1.0.82", "10.1.0.43",
"10.1.0.91", "10.1.0.8", "10.1.1.250"]

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/authoritativedata.xml --authoritative --
source-scanner qualys --format xml

Using /home/centos/.rvm/gems/ruby-2.4.1
Authoritative scan overwites the previous entries for all the
ips within this scan. There is no ROLLBACK
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 : 6.
Elapsed Time [3.02 (total : 15.15] - Imported
Vulnerabilities..
Elapsed Time [0.0 (total : 15.15] - Detected Vulnerabilities
are associated with vulnerability and workload data..
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 [487] 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) [486] and Vulnerability Identifier [487] 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 | Updated By |
----------------------+-------------------------
+----------------------+----------------------+
| /orgs/1/vulnerabilities/nessus-18405 | Microsoft Windows
Remote Desktop Protocol Server Man-in-the-Middle Weakness
| 51    |              | ["CVE-2005-1794"]
| 2018-11-07T03:15:39.410Z | 2018-11-07T03:15:39.410Z
| {"href"=>"/users/1"} | {"href"=>"/users/1"} |
...
```

Qualys

```
C:\Users\isaac.newton> ilo vulnerability list --xorg-id=1
...
| Href | Name | Score | Description | Cve Ids | Created At |
Updated At | Created By | Updated By |
---------------------+-------------------------
+----------------------+----------------------+
| /orgs/1/vulnerabilities/qualys-38657 |
Birthday attacks against TLS ciphers with 64bit
block size vulnerability (Sweet32) | 69 | | ["CVE-2016-2183"]
| 2018-07-27T18:16:57.166Z | 2018-08-08T22:30:32.421Z
| {"href"=>"/users/1"} | {"href"=>"/users/16"} |
...
```

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) [486] and Vulnerability Identifier [487] 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 | Updated By |
+-------------------------------------------------------
+-------------+----------------------
+--------------------------+----------------------
| /orgs/1/vulnerability_reports/scan_1502310096_09344 |
qualys  | NewAuthoritativeScan | 2018-08-08T22:30:34.877Z |
2018-08-08T22:30:34.877Z | 62 | {"href"=>"/users/16"} |
{"href"=>"/users/16"} |
...
```

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 authenticated to the PCE and run two `ilo bulk_upload_csv` commands.

1. Authenticate to the PCE via API key or explicit login.
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 Kerberos requires, the Kerberos realm name is shown in all capital letters as `MYREALM`.
- 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.

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.
2. Create Kerberos-authenticated `myWorkload1` and set its `mode` to `idle`:

```
ilo workload create --hostname myPCE.BigCo.com
--name myWorkload1 --service-principal-name host/
myKerberosTicketGrantingServer@MYREALM --agent/config/mode
idle
```

For information about how the mode is a nested field, see Nested Resource Fields and Wildcards [424].

3. Create Kerberos-authenticated `myWorkload2` and set its `mode` to `illuminated`:

```
ilo workload create --hostname myPCE.BigCo.com
--name myWorkload2 --service-principal-name host/
myKerberosTicketGrantingServer@MYREALM --agent/config/mode
illuminated
```

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 data sets without waiting 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 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 | Description | Hostname
| Service Principal Name | Public
Ip              | Distinguished Name |
External Data Set | External Data Reference
| 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                  | Updated By
+-------------------------------------------------------------
+---------+------+-------------+----------------
... More lines not shown
-------------------------------------------------------------+
| /orgs/1/workloads/50ce441e-75ac-4be8-9201-96169545019c
| false   |      |             | 10.14.59.8.netstat
...
... 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 [436]. 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-scanner qualys --format xml
```

2. Examine a single uploaded vulnerability record identified by its vulnerability identifier, `qualys-38173`. See Vulnerability Identifier [440] 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-scanner qualys --format xml
   ```

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 --authoritative --source-
   scanner qualys --format xml
   ```

Results

The authoritative vulnerability data has been uploaded and is ready for use in Illumination.

Error Messages

Learn about common error messages and ways you can resolve these errors on your own by performing specific recommended actions.

Containers Error Messages

Here are some common error messages and ways you can resolve them.

Containers

Table 6. Common Containers Error Messages

Action	Error	Description	Recommended Action
Kubelink	pkg.pcelink PCE request failed {"url": "https://<PCE>:443/api/v2/orgs/<OrgID>/container_clusters/<UUID>/put_from_cluster", "error": Put \"https://<PCE>:443/api/v2/orgs/<OrgID>/container_clusters/<UUID>/put_from_cluster\": dial tcp: lookup <PCE FQDN> on <IP>:53: server misbehaving"}	DNS Error	Check the DNS records for the PCE FQDN

NEN, VEN Error Messages

Here are some recommended actions you can take to help you troubleshoot common NEN and VEN errors.

NEN

Table 7. Common NEN Error Messages

Action	Error	Description	Recommended Action
Add NEN	ERROR -- : The property 'service_discovery_certificate' value='/var/lib/illumio-nen/cert/server.crt' file doesn't exist	NEN requires a separate certificate from the PCE	Create and add a certificate specific to the NEN
Add NEN	ERROR -- : The property 'service_discovery_private_key' value='/var/lib/illumio-nen/cert/server.key' is not a valid private key	NEN certificates require a valid certificate chain as part of the .crt file	Check the syntax of the certificate chain, including the order of the certificates

VEN

Table 8. Common VEN Error Messages

Action	Error	Description	Recommended Action
Operation after upgrade	agentmgr.log: ERROR:: [ApiHelper] Failed to PUT https://<PCE>:8443/api/v26/orgs/1/agents/<AgentID>/heartbeat (internal error 56)platform.log: ERROR:: SSL peer certificate or SSH remote key was not OK , SSL certificate problem: unable to get local issuer certificate	VEN cannot communicate with the PCE after an upgrade	Edit the PCE certificates to be up to date and include all intermediate certificate components. KB article
Install with Pairing Profile	{"error": "profile_id is invalid"}	The profile_id associated with the pairing profile is missing	Create a new pairing profile. KB article
Policy Sync	PCE Reports policy sync errors on VEN, you see two copies of Illumio VEN agent installation path	Multiple VEN binaries are installed in different paths after the VEN upgrade	Remove VEN binaries and rpm link to both directories. Install current VEN binaries in default path. KB article
SUSE Linux 15	failed to create symbolic link '/etc/init.d/rc2.d/S50illumio-ven': No such file or directory	Missing dependency for the VEN to run	insserv-compat is a prerequesit to have deployed before install/upgrade of the VEN KB article

PCE Error Messages

Here are some common error messages and ways you can resolve them.

PCE Error Messages and Recommended Actions

Table 9. Common PCE Error Messages

Action	Error	Description	Recommended Action
Daily operation	ERROR -- : Certificate validation issue with web_service_certificate : host: <PCE>	Expired certificate on the PCE	Run this command: illumio-pce-env setup --test 5 --list
Install compatibility matrix	There was an error (unsupported_ven_release): VEN release <X> is not supported by PCE <Y>	Incorrect/ outdated compatibility matrix installed	Get the latest compatibility matrix from the Support website.
Error while generating an inventory report	No such file or directory - sysctl		Get the path where sysctl is located i.e. /sbin or /usr/sbin and add it to the path. See KB article
Log In	DB connection error detected; could not fork new process for connection: Resource temporarily unavailable	File limit settings	Make sure to set the correct file limit and configuration settings of the PCE. See KB article
Log In	500 Internal Server Error / PCE UI Missing Error	This means the UI is not installed or the directory containing /opt/illumio-pce is incorrect.	Check that the UI is installed with: yum list installed | grep -i illumio-pce-ui Make sure that you set the correct owner recursively for the folder: /opt/illumio-pce with: chown -R root:ilo-pce /opt/illumio-pce See KB article
Pairing Profile Access	500 Internal Server Error	500 Internal Server Error	VEN Bundle missing from PCE with a Pairing Profile using that VEN version. See KB article
Enable feature for API to SaaS PCE	API call returned error code 403. Errors: forbidden_error: Access denied.		Check the permissions of the API user. Require elevated access to perform API actions.

Action	Error	Description	Recommended Action
			If this is a SaaS PCE, the feature flag may need to be enabled.

Legal Notice

Copyright © 2026 Illumio 920 De Guigne Drive, Sunnyvale, CA 94085. All rights reserved.

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

Resources

- Legal information
- Trademarks statements
- Patent statements
- License statements

Contact Information

- Contact Illumio
- Contact Illumio Legal
- Contact Illumio Documentation