




Illumio Core 23.2 Visualization User Guide

Published: 2024



This guide describes the visualization tools in the Explorer category: Map, Traffic, Mesh, Reports, and App Groups. Visualization tools allow you to see the traffic flows in your network and help you configure policies to secure your applications.

Table of Contents

Legal Notice	4
Security Advisories	5
September 2024 Security Advisories	5
Ruby SAML gem component authentication bypass vulnerability	5
Severity	5
Affected Products and Patch Information	5
Resolution	5
References	6
Skipped Critical Patch Updates	6
Discovered By	6
Frequently Asked Questions	6
Modification History	7
September 2023 Security Advisories	7
Authenticated RCE due to unsafe JSON deserialization	7
Severity	7
Affected Products and Patch Information	7
Resolution	8
References	8
Skipped Critical Patch Updates	8
Discovered By	8
Frequently Asked Questions	8
Visualization	10
Visualization Tools	10
About the Visualization Tools	10
About the Map	25
Traffic Table	31
Mesh View	36
Work with the Visualization Tools	38
App Group Map	46
About the App Group Map	46
Work with the App Group Map	50
Vulnerability Map	52
About Vulnerability Map	52
Work with Vulnerability Maps	54
Reports and Statistics	58
Ransomware Protection Dashboard for Servers	58
VEN Dashboard	62
About Reports	63
Work with Reports in the PCE	67

Legal Notice

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

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

Resources

- [Legal information](#)
- [Trademarks statements](#)
- [Patent statements](#)
- [License statements](#)

Contact Information

- [Contact Illumio](#)
- [Contact Illumio Legal](#)
- [Contact Illumio Documentation](#)

Security Advisories

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

September 2024 Security Advisories

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

Ruby SAML gem component authentication bypass vulnerability

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

Severity

Critical: CVSS score is 9.9

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

Affected Products and Patch Information

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

Table 1. Products Affected by the Security Vulnerability

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

Resolution

Upgrade to the latest release for a given major version.

References

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

Skipped Critical Patch Updates

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

Discovered By

External Security Firm

Frequently Asked Questions

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

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

Modification History

- September, 2024: Initial Publication of CVE

September 2023 Security Advisories

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

Authenticated RCE due to unsafe JSON deserialization

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

Severity

Critical: CVSS score is 9.9

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

Affected Products and Patch Information

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

Table 2. Products Affected by the Security Vulnerability

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

Resolution

Upgrade to the latest release for a given major version.

References

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

Skipped Critical Patch Updates

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

Discovered By

External Security Firm

Frequently Asked Questions

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

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

Visualization

Visualization Tools

In Illumio Core 23.4, Illumio introduces a new user interface for Illumio Core Cloud customers. Illumio Core Cloud customers control the pace at which they adopt the New PCE UI. By default, current customers see the Classic PCE UI when they log in. The PCE Classic UI includes the Illumination Plus feature introduced in Illumio Core 22.5. This feature is an enhanced version of the data visualization features that combines and expands functionality and supports using the new 23.4 feature for custom label types.

In the Classic UI, you can still access the Illumination Classic feature by updating your setting in your My Profile. See [Configure Visibility Display](#) in the PCE Web Console for information.

In Illumio Core 23.4, the New UI includes all the Illumination Plus features under the **Explore** category of left navigation.

About the Visualization Tools

In the PCE UI, you can use the visualization tools to reveal the traffic flows in your network and to help you configure policies to secure your applications. These tools include the Map, Traffic table, and Mesh.



IMPORTANT

The visualization tools are available in the PCE Classic UI and the PCE New UI.

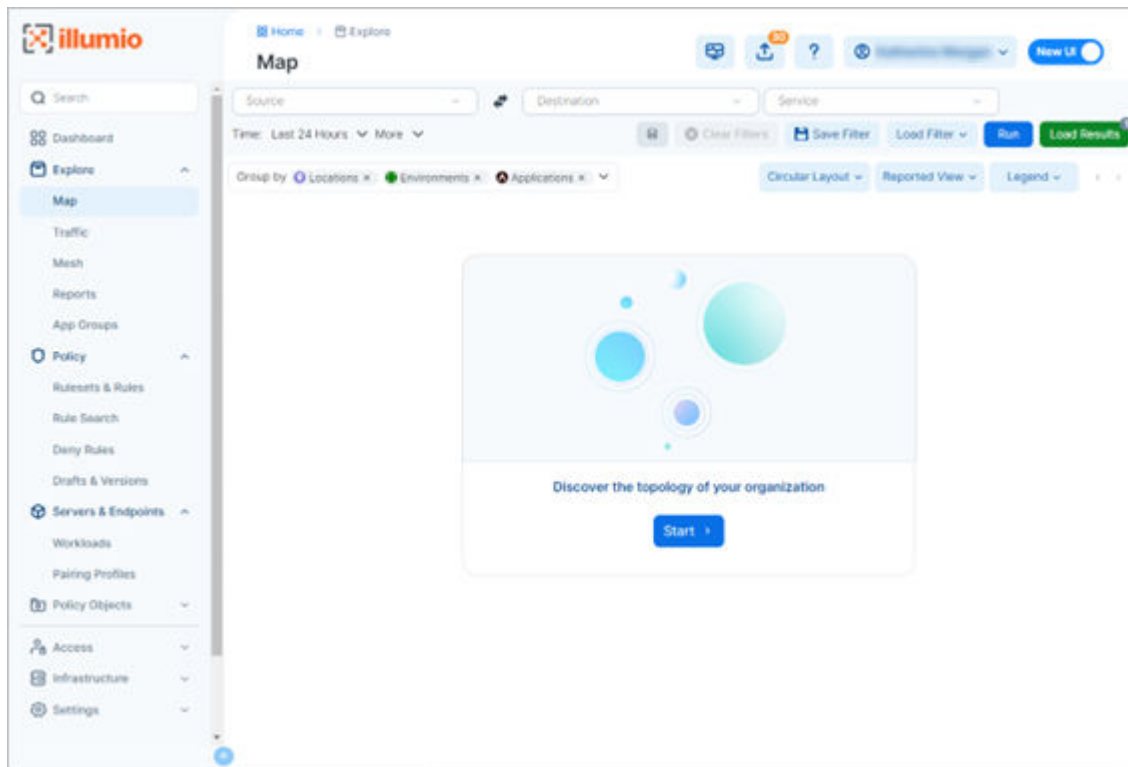
To access these features in each of the PCE UIs:

- In the Classic UI, choose **Illumination Plus** from the left navigation; select the type of visualization feature (**Map**, **Table**, or **Mesh**) from the left drop-down list on the page toolbar.
- In the New UI, choose **Map**, **Traffic**, or **Mesh** under the **Explore** category of the left navigation.

Other than the differences in the navigation, the functionality of the visualization tools is comparable across both PCE UIs.

When you open a visualization tool for the first time or the first time during a 24-hour period, the PCE UI displays a landing page with tiles to the different views and a message to run your first query.

The following image shows the start page that appears in the PCE New UI.



Types of Visualization Features

You can view detailed information about your environment by filtering your traffic flows in the following visualization tools:

- **Map**

Graphically visualizes workloads that form logical groups (based on labels attached to workloads) and provides an understanding of the traffic flows between workloads. You select groups in the Map view to view details about that group and develop policy for the workloads in the group.

- **Traffic**



NOTE

In the PCE Classic UI, this feature is referred to as the **Table** view.

Displays details about your traffic flows in columns and rows. Using this view, you query the PCE traffic database for historical data that can be used for compliance and audit, as well as policy development. With an easy-to-use interface, you enter your search parameters using plain-text language and filter results by a specific time period; specific ports, protocols, or processes; and actions that were taken on that traffic based on policies (for example, “allowed” vs. “potentially blocked” vs. “blocked”).

- **Mesh**

Using vertical axes, displays traffic flows as lists of destinations, sources, and the port being used in the traffic flows. The traffic flows between destinations and sources connect along parallel coordinates. You can sort the results based on port number or the number of traffic flows. Click any item in the results to focus on specific traffic flows.

**NOTE**

The PCE Classic UI uses the terms consumer (instead of destination) and provider (instead of source).

In the **PCE Classic UI**, you switch between views by selecting the view from the top-right corner of the **Illumination Plus** page:



In the **PCE New UI**, select the visualization tool you want to use from the **Explore** category in the left navigation.

Filters for the Visualization Tools

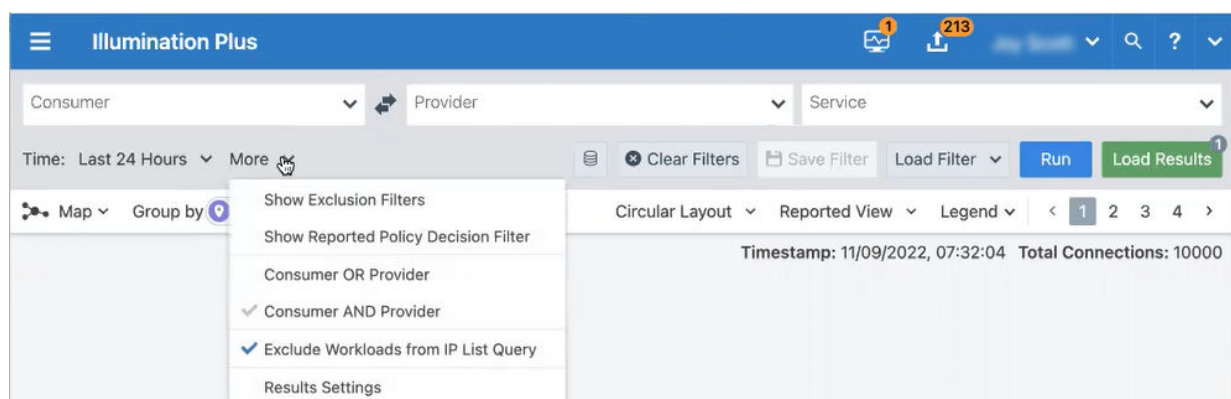
For each of the visualization tools, you can set one of several traffic filters to show or hide different elements of your data and focus on what is most important to you. All views allow you to filter your data by destination, source, and service. By default, you only see the Include filters to begin with.

**NOTE**

The PCE Classic UI uses the terms consumer (instead of source) and provider (instead of destination).

To modify the filters, open the **More** menu to select additional filter options.

Page from Illumination Plus in the Classic UI:



**NOTE**

The filters selected in previous sessions don't persist unless you've added values to them. For example, the Exclusion filters won't appear by default when you open the page unless you've explicitly excluded traffic in the past.

**TIP**

To search for traffic flows with a specific policy decision reported by the VENs, select the **Show Reported Policy Decision** option. This option controls the type of policy decision (allowed, potentially blocked, blocked, or unknown) that the Table and Map views display.

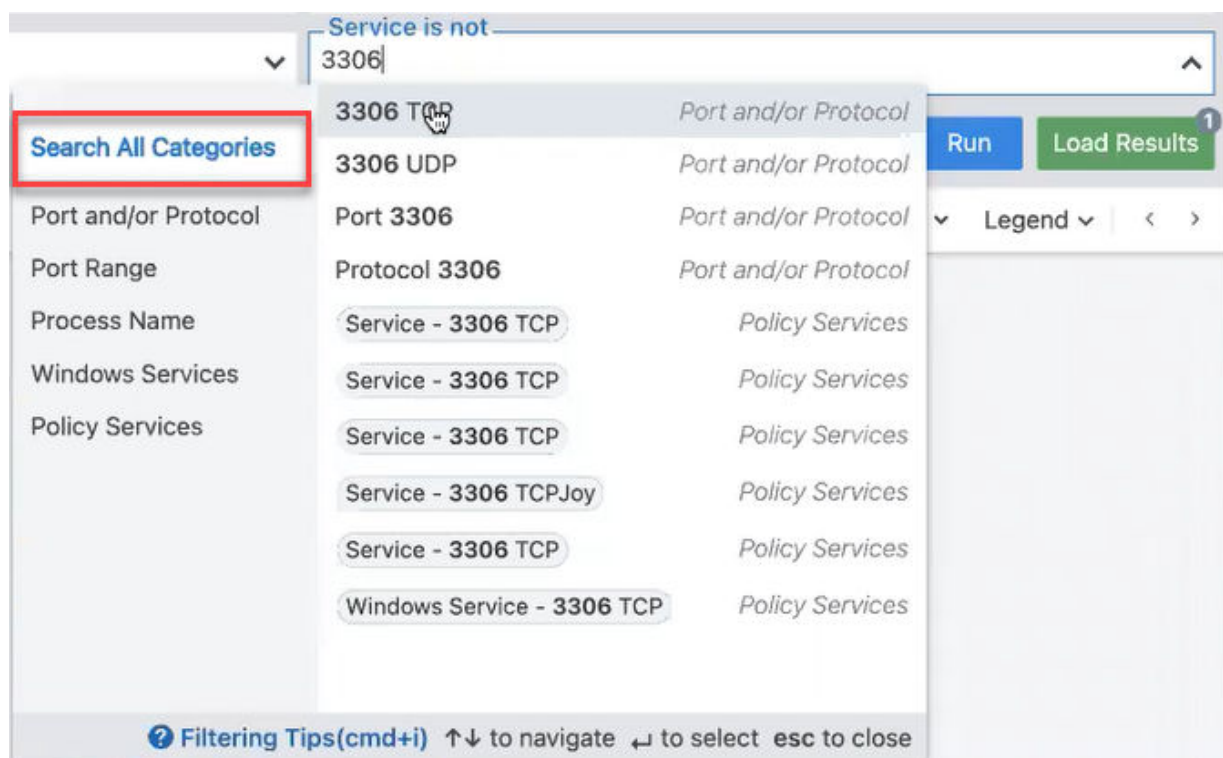
The Source and Destination filters include the following query options:

- Label and Label Groups
- App Groups
- Workloads
- IP Lists
- IP Address/CIDR Block
- FQDN
- Transmission

Using the **Search All Categories** feature, you don't have to enter a category first in the filters.

The Label and Label Groups category restricts the Map to only those entities that have the labels you enter in the filters. The filter does not filter the selected group. Only the connected groups are filtered.

From the **Service** drop-down list, search by port and protocol. You can select a specific protocol and the page allows you to search through all the services.



When you enter text in this filter, the PCE UI gives you the option to select whether that text is a process name or a service. Once selected, the UI specifies which option you chose; for example: `Process Name: dfsfjklsf x`

Example Search using Filters

Before you write policy rules to either allow or block traffic, you want to determine if there are any traffic flows between them. For example, you might want to find traffic between Development or Testing environments from your Production environments.

Using the visualization tools, you can run, for example, the following query:

Any traffic flows during the last week between my Development and Production environments, over any port except port 80, excluding any workloads that have a Role label named "Domain Controller"

The following steps show how you use the filters in the PCE New UI for this search to reveal certain traffic flows but not others.



NOTE

The PCE Classic UI uses the terms consumer (instead of source) and provider (instead of destination).

1. In the **PCE New UI**, choose **Explore > Map** or **Explore > Traffic**.

The page appears. To exclude criteria, go to **More > Show Exclusion Filters** if they don't already appear in the page.

2. Under Destination, enter or select the Environment label named "Development" from the Destination drop-down list.
3. Under Destination, enter or select the Role label named "Domain Controller" from the Destination is not drop-down list.
4. Under Source, enter or select the Environment label named "Production" from the Source drop-down list.
5. Under Source, enter or select the Role label named "Domain Controller" from the Source is not drop-down list.
6. Under Service, leave the Service field blank (which means "any") and under Service is not enter "80."
7. Under Time, select **Anytime**.
8. Click **Run**.

Query Results in the Visualization Tools

In all views, the PCE limits the number of connections you can load per page in the PCE UI to 10,000. You can't load your total number of connections in a single page. To handle this limitation, the PCE UI displays your connections in paginated results. To view all connections, you can paginate through your query results. For example, when you run a query that returns 200,000 traffic flows, you can paginate through your data to see all traffic flows.



To configure the maximum number of connections per page:

1. From the PCE left navigation, choose **Map**.
2. Choose **More > Results Settings**. The Results Settings dialog box appears.
3. Specify the maximum number of connections to display per page:
 In the **Displayed In Traffic** field, configure the maximum number of results that can be retrieved from the PCE database and displayed per page in all views.
 In the **Returned from Database** field, configure the results when the PCE is part of a Supercluster.



IMPORTANT

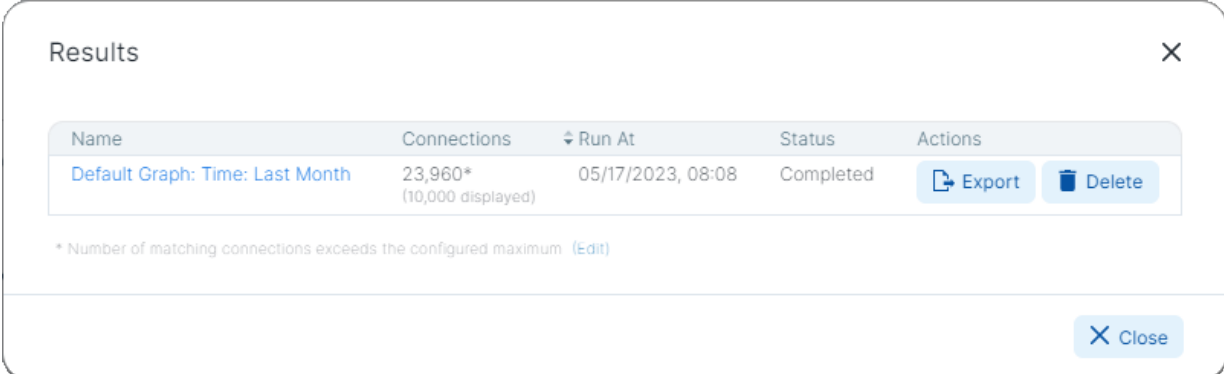
Configuration for a Supercluster deployment does not apply to Illumio Core Cloud customers; you must be an Illumio Core On-Premises customer to configure your Illumio deployment as a Supercluster.

In a Supercluster, a query run on the leader PCE can return 200,000 results for each PCE in the Supercluster, including the leader. For example, in a Supercluster with four regions, the maximum results is 800,000, and in a standalone PCE, it is 200,000. When logged into a member PCE on a Supercluster, the limits are the same as for any SNC or MNC. In every case, the maximum number of results that can be shown in the PCE UI is 100,000 results. If more than 100,000 results are retrieved, the full results are available as a downloaded CSV file, and the first 100,000 are available in the PCE UI.

For more information about PCEs in a Supercluster configuration, see the *PCE Supercluster Deployment Guide*.

Load Results in the Map or Traffic Table

As you run searches, the PCE caches your queries and saves them for a 24-hour period. Caching your query results is beneficial because the PCE displays pages quickly. To view and access your cached queries, click **Load Results** at the top-right corner of the page. The Results page appears.



The screenshot shows a modal window titled "Results" with a close button (X) in the top right corner. Inside the modal is a table with the following columns: Name, Connections, Run At, Status, and Actions. The table contains one row with the following data:

Name	Connections	Run At	Status	Actions
Default Graph: Time: Last Month	23,960* (10,000 displayed)	05/17/2023, 08:08	Completed	Export Delete

Below the table, there is a note: "* Number of matching connections exceeds the configured maximum [\(Edit\)](#)". At the bottom right of the modal is a "Close" button.

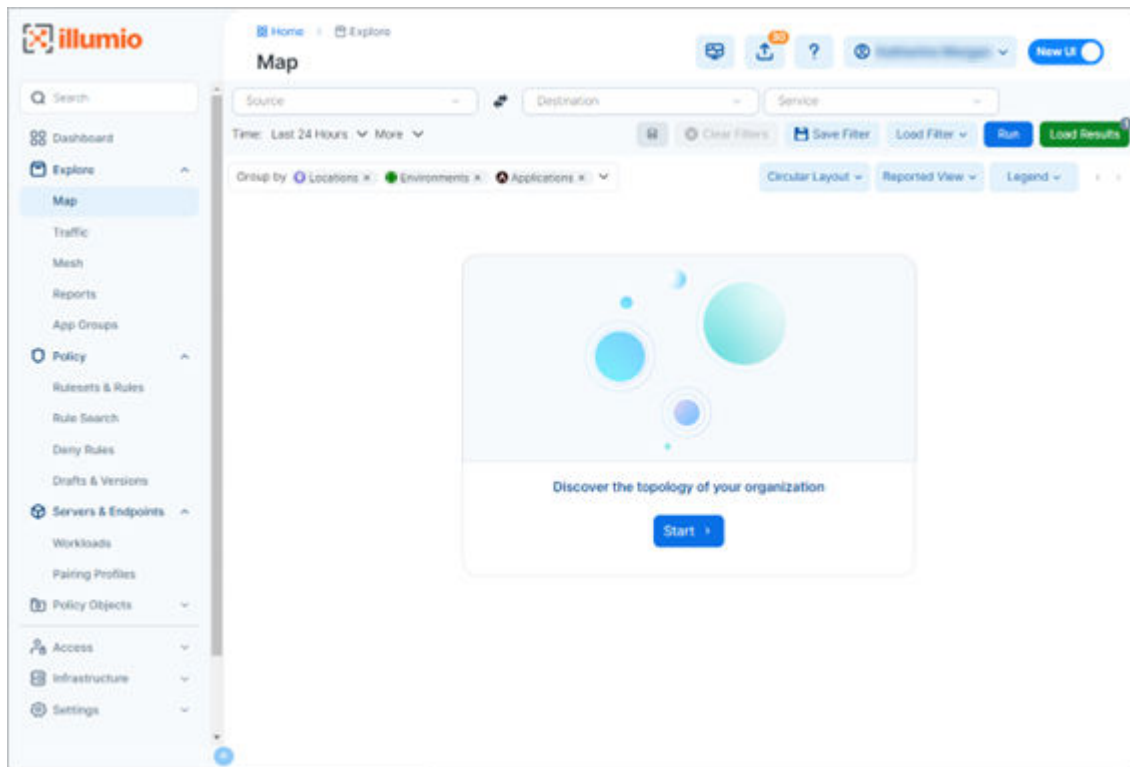
The load results process runs in the background to increase the speed that view pages display. Using this feature is optional, though recommended.

Switching between the Map and Traffic table doesn't reload your data. Instead, the PCE UI switches immediately to that view.

About the Default Graph

In Core 22.5.x, the PCE cached the Illumination Plus queries (for the Map and Table views) that you ran and were saved for a 24-hour period. Caching your query results allowed the PCE to display Illumination Plus pages quickly. To view and access your cached queries, you clicked Load Results at the top-right corner of the Map page. The Results page appeared.

In 23.2.0, if you don't have a default graph in the PCE, the page below is your start page for the Map and Traffic pages.



When you click Start, the PCE creates a map or traffic table based on the values you have in the filters at the top of the page. The PCE saves this query with those filters as the default graph. The graph expires in 24 hours; however, the PCE saves the default graph as a scheduled report that runs every 24 hours (between 12:00 midnight and 8:00 AM).

Then, when you return to the Map or Traffic page, the PCE loads that saved default graph, unless you already have another graph (different filters) displayed. You won't see this Start page again, unless you delete the default graph.

This page now appears when you click Load Results in the Map page to display the entry for the Default Graph:

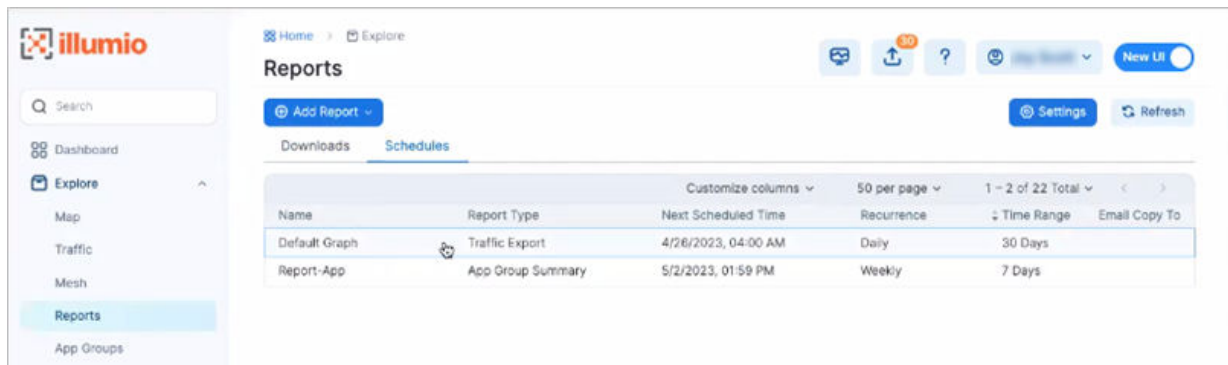
Results

Name	Connections	Run At	Status	Actions
Default Graph: Time: Last Month	23,960* (10,000 displayed)	05/17/2023, 08:08	Completed	<div>Export</div> <div>Delete</div>

* Number of matching connections exceeds the configured maximum. [\(Edit\)](#)

Close

When you open the Reports feature from the left navigation and select the Schedules tab, you see the scheduled report for the Default Graph.



IMPORTANT

Not all Illumio users can access the Default Graph scheduled report. You must have the correct Access permissions. See the PCE Administration Guide for information.

Tips for Using the Default Graph

- To change the query that the PCE runs for the Map and Traffic page:
- Go to the Reports page and select a different saved query.
- Delete the default graph by clicking Load Results in the Map or Traffic page and clicking Delete in the Load Results dialog box. Then, navigate to the Map or Traffic page so that the Start page appears. Click Start to create a default graph.
- Click the Schedule Time field and select a new time to change when the default graph report runs each 24 hours. However, you must have the correct permission to edit the Default Graph (RBAC roles and permissions).

Asynchronous Queries

You can run asynchronous queries for your filters. You first set up your filters and then run an asynchronous query.

Asynchronous queries allow you initiate multiple queries in parallel and view the results of the queries later. Going offline during a query does not result in lost query results. Whether you remain online or offline, the results of asynchronous queries will be preserved for a period of 24 hours. In addition, while a query is in progress, you can work in other areas of the product. The query search results can be exported to either a comma-separated-value (.CSV) file or displayed in the PCE UI. Depending on the size of the query, the results might take time to display.

In the visualization tools, you can run multiple queries and change or retain the default file name for exported results.

- **Multiple Queries:** You can run multiple queries, including running some in the background.
 - If there is only one query, the results of that query will display when the query completes.
 - If there are multiple queries, you can select the result that you want to view by clicking the number beside the **Load Results** button.

- If identical queries are run within a minute of each other, only one query will be processed. The results of the oldest query will be displayed.
- **Default File Name:** The system assigns a default file name based on your query field names (Source, Service, or Destination) in the filter. The exported file will have the same name.
- Giving filters a unique name will help you identify your filters when you want to rerun a query. This name will also appear as your report name.
- You can also specify or change a filter name as needed.

**NOTE**

Handling Duplication Flows in Queries

A database query that spans multiple days can contain duplicate flows if the flow is repeated.

Run Asynchronous Queries

Asynchronous job queries are easy to initiate and can be run in parallel, which means that before the first query completes, a second query can be initiated. In the following example, two queries are initiated: the first, with Production-only entries, and the second, with Production and Staging entries.

To run an asynchronous query:

1. From the PCE UI left navigation, choose **Map** or **Traffic** from the **Explore** category.
2. Enter your query criteria in the fields. If you want to exclude criteria, browse to **More > Show Exclusion Filters**.
You can enter a Source, Destination, or Service, or merely indicate Production in the Destination column.
3. Click **Run** to begin the query process.
4. In the confirmation dialog box, click **Hide**.
5. Enter the next search criteria based on a new Destination; for example, Production and Staging.

Given support for asynchronous queries, you will see a number appear next to the **Load Results** button, indicating the number of simultaneous queries being processed

**NOTE**

Depending on the size of the queries, your second query could complete before your first query.

You will see the results of your two queries, one with Production-only entries and a second with Production and Staging entries.

6. . At any time, can click the **Load Results** button to view what queries were run.
Viewing results from past queries will not re-initiate a query. It displays cached query results. When you select a result, notice that the filter changes automatically, and displays new results.

Global Queries for Superclusters



IMPORTANT

Configuration for a Supercluster deployment does not apply to Illumio Core Cloud customers; you must be an Illumio Core On-Premises customer to configure your Illumio deployment as a Supercluster.

Global queries leverage the capabilities of asynchronous job queries for every region in a Supercluster. When you have a Supercluster and you initiate a query from the Supercluster leader, the Illumination Plus Table view displays results from all its PCE members. Queries run from a Supercluster member only show flows reported by VENs paired to that member.



NOTE

In a Supercluster, a query run on the leader PCE can return 200,000 results for each PCE in the Supercluster, including the leader. For example, in a Supercluster with four regions, the maximum is 800,000, and in a stand-alone PCE, it is 200,000.

When logged in to a member PCE on a Supercluster, the limits are the same as for any SNC or MNC. In every case, the maximum number of results that can be shown in the PCE web console is 100,000 results. If more than 100,000 results are retrieved, the full results are available as a downloaded CSV file, and the first 100,000 are available in the PCE UI.

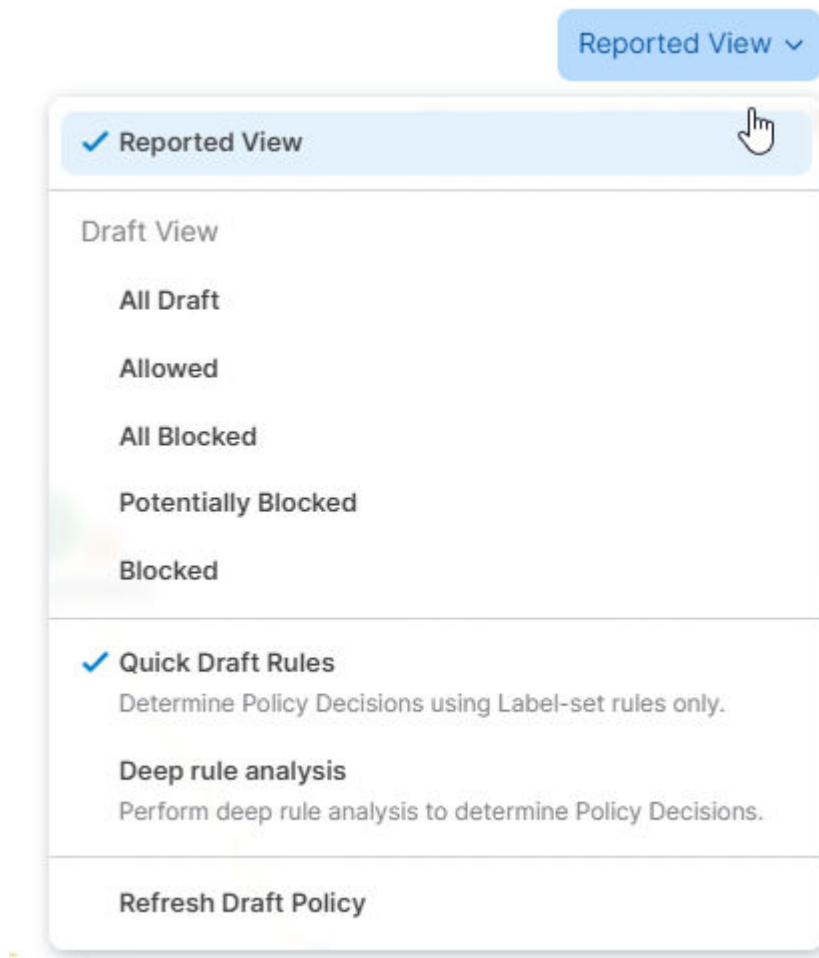
View Menu in the Map and Traffic Pages



IMPORTANT

The View menu only appears when you are in the Map and Traffic pages. The Mesh always displays traffic flows based on the Reported view. You cannot switch to the Draft view for the Mesh.

Using the View menu, you configure how the PCE UI displays your traffic data. The options on this menu are unaffected by how you've grouped traffic in your Map or Traffic pages. This menu provides flexibility in how you see the connections between your groups.



From the View menu, select the following options:

- **Reported View**
For a description, see [Reported View \[22\]](#).
- **Draft View Options** – All, Allowed, All Blocked, Potentially Blocked, Blocked
In the Draft view, you can choose all connections, or filter by the policy state (allowed, potentially blocked, or blocked). For a description, see [Draft View Options \[23\]](#).
- **Quick Draft Rules**
Provides a fast way to analyze your environment and display results in your views because it determines policy decisions based on label-set rules only.
- **Deep Rule Analysis**
Returns additional rulesets that the Quick Draft Rules option won't detect. However, displays results more slowly than using Quick Draft Rules due to the deeper analysis of rulesets. This option will find any rules written directly for workloads versus created by using labels. It can combine two rules that use IP lists; for example, workload "A" has connections to IP addresses in an IP list ("IP list B"). IP list B connects to another workload C. Deep analysis shows when rules have been optimized so that workload A can connect to workload C.
- **Refresh Draft Policy**
if you've written rules after the draft policy was last run, you can force it to refresh in the PCE web console.

Reported View

The Reported view visualizes your policy coverage as reported by your workloads, so that you can examine the current state of your provisioned policy. This view provides visibility for the actual traffic handling (rather than the expected traffic handling provided by the Draft view) and loads more quickly, especially when you have a large number of workloads and traffic flows. The Reported view helps you to understand your traffic patterns.

The Reported view is a read-only view. You can view all the rulesets that apply to the workloads from the Reported view, but you must change to the Draft view to add rules. The Reported view does not immediately reflect the latest changes to the policy. It is updated only after you provision a change to the policy and when new traffic flows that use the updated policy are reported from the VEN.

The Reported and Draft views handle unmanaged workloads differently. In Draft view, rule coverage (the connections that have been included in draft rules) has limited support for traffic between unmanaged workloads. The Reported view always provides accurate rule coverage for traffic between unmanaged workloads.

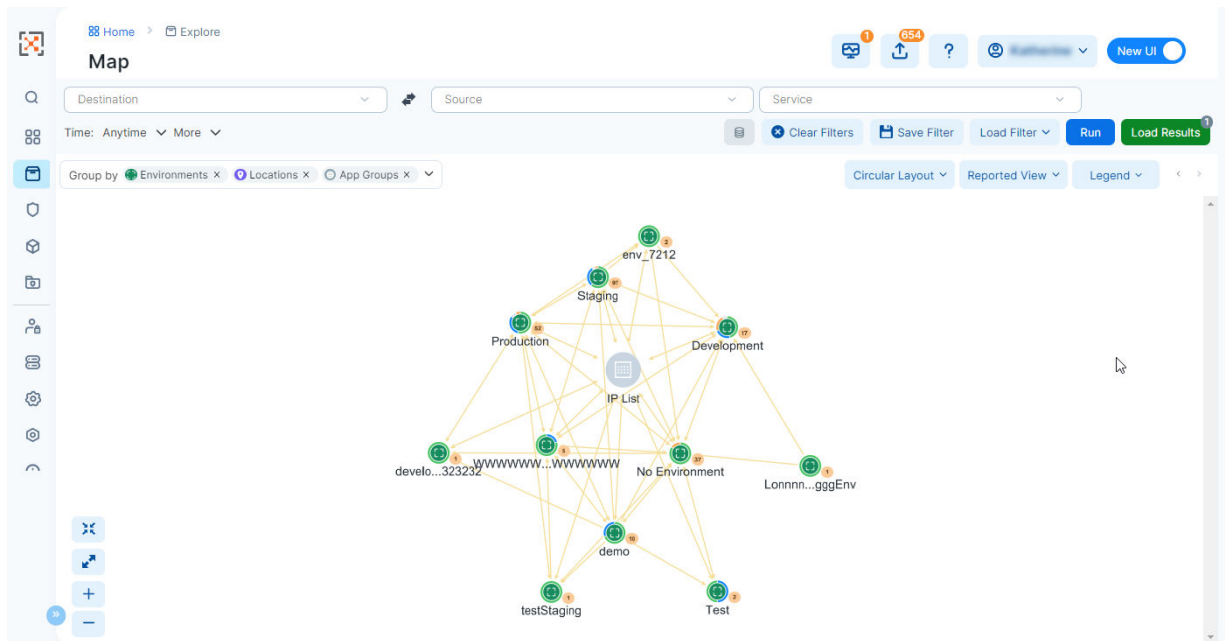
For each flow with a unique port/protocol, if there is a policy service created for that port/protocol, the name of that policy service displays, in addition to the names of the actual services that reported the flows. The Reported view shows reported rule coverage for the latest reported flow with that port/protocol in the right side panel.

Different services can be running on the same port at different times or on different interfaces. The Reported view shows reported rule coverage of each flow separately, as well as its timestamp. In both cases, the Draft view shows the calculated rule coverage for traffic. For Windows, it looks at the port, protocol, the process name (but not the process path), and the Windows service name. For Linux, it looks at only the port and protocol.

Reported View (Traffic)

The screenshot displays the 'Traffic' section of the Illumio interface, specifically the 'Reported View'. The top navigation bar includes 'Home' and 'Explore' links. The main header shows 'Traffic' with search filters for 'Destination', 'Source', and 'Service'. Below the header, there are buttons for 'Allow Selected Connections...', 'Resolve Unknown FQDNs', and 'Export'. The table below lists traffic flows with columns for 'Reported Policy Decision', 'Destination Labels', 'Destination Port Process', 'Source Labels', 'Flows/Bytes', and 'First Detected'. The first row shows a flow from 'solaris' to 'db-d' on port 3306 TCP, with a policy decision of 'Potentially Blocked' and 'no Rule'. The source is 'root' and the process is 'mysql'. The flow is associated with 'Full Enforcement' and 'db-d41'. The source labels include '7 Source IPs', 'Full Enforcement', 'autojob-s22', 'autojob-s44', and 'autojob-s46'. The flow statistics show 7 Connections, 223 Flows, 6.8 MB sent, and 6.5 MB received. The first detected time is 04/30/2023, 19:41:14.

Reported View (Map)



Draft View Options

The Draft view immediately visualizes the potential impact of your draft policy. This view helps provide an understanding of the expected traffic handling (rather than the actual traffic handling provided by the Reported view) and considers both recently provisioned policy and draft policy. The Draft view can take longer to load than the Reported view, especially when you have a large number of workloads and traffic flows, since the PCE has to compute the expected coverage for each traffic flow.

In Draft view, you can either view the rule that would permit traffic or add a rule to allow a specific flow. In this view, you can immediately see the impact of the latest changes to the active or draft policy.

Limitations of Draft View

The Draft view is the result of a “what-if” analysis conducted by the PCE. It is a modeling tool that depicts whether flows known to the PCE will be allowed or blocked, based on the configured policy. The modeling might not work entirely correctly for the following types of rules configured on the PCE:

- **Process-based rules:** Process-based rules are written using the process name or service name that sends or receives the traffic on the workload.
- **User-based rules:** User-based rules allow administrators to leverage the Microsoft Active Directory User Groups to control access to computing resources.
- **Custom iptables rules:** Custom iptables rules are configured on each workload and can include processes that are not known to the PCE.
- **System rules:** The VEN has implicit rules to permit necessary traffic (for example, rules permitting DHCP and DNS outbound traffic on the workload).

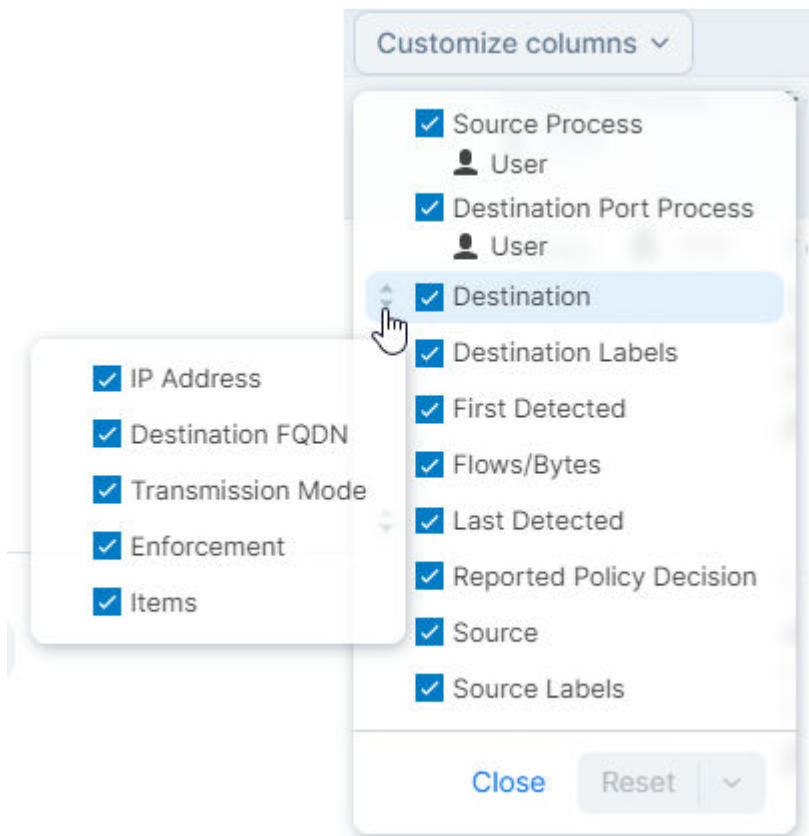
In most cases, the Reported view provides an accurate representation of what will be allowed or blocked by the VEN, so the Reported view should be used to verify your changes.

Customize Columns

In the following visualization tools, you can customize the columns from the default display:

- Traffic
- Map > Traffic tab
- Map > Workloads tab

Columns in these areas are customizable from the **Customize columns** menu. Most columns can be further customized by setting what data will appear within that column. Hover over the up and down arrows to the left of a column checkbox and select or deselect data within that column:



Customizing the columns that appear in the Tables does not impact how you create your rules or the data that they contain.

How the Map Works with FQDNs

The visualization tools map the outbound connections from workloads to unknown IP addresses to fully qualified domain names (FQDNs) or DNS-based names. For example, Illumination Plus could display that the outbound connections from a workload are going to `maps.google.com` instead of 100s of different IP addresses. The FQDNs used are reported by the VEN to the PCE in the flow summaries. The VEN learns about the FQDNs by snooping the DNS responses on the workloads, which is the FQDN for the IP addresses as seen by the workloads.

The Map visualizes the workloads that form logical groups (based on labels attached to workloads) and provides an understanding of the traffic flows between workloads.

About the Map

Use the Map feature to visualize workloads that form logical groups (based on labels attached to workloads) and provides an understanding of the traffic flows between workloads.



IMPORTANT

The map feature is available in the PCE Classic UI and the PCE New UI.

To access this feature in each of the PCE UIs:

- In the Classic UI, choose **Illumination Plus** from the left navigation; select **Map** from the left drop-down list on the page toolbar.
- In the New UI, choose **Map** under the **Explore** category of the left navigation.

Other than the differences in the navigation, the functionality of the Map is comparable across both PCE UIs.

Grouping in the Map

Groups in the Map represent a collection of workloads or services that communicate with each other and for which you can write rules. Groups are displayed in the Map after you pair workloads. See the VEN Installation and Upgrade Guide for information about installing (also called pairing) VENs on workloads.

The Map displays three different types of groups: a group based on a single label, an app group, or a label set. A label set is a group of entities that have the same set of labels.

Once you pair VENs to create workloads, the PCE analyzes the workload data reported by the VENs. Based on the traffic flows among your workloads, the Map organizes them into groups. A group could represent an instance of an application running in your data center, such as an HRM application running in the Test environment in your North America data center; or a group could represent a Web store in Production with its web workloads hosted in AWS and its databases hosted in your private data center.

The Map lets you group by labels, locations, etc. It also lets you split the view when in Map view mode by selecting items on the Map.

Configurable Grouping

In the Map, grouping is implemented flexibly as you run your queries. Using the **Group by** menu, you can add different levels of grouping, such as grouping by types of labels and their order. You might want grouping by OS and then by environment. If you do not specify a particular grouping, Illumio will group workflows by the default, which is by workloads with the same set of labels. You can change your organization's default grouping using the same drop-down menu.

**NOTE**

For optimal scale and performance, if there are two connections with the same source workload, destination workload, destination port, and protocol but the process or service names are different, the two connections are combined in the Map. The process or service name that was part of the most recently reported connection is displayed.

Tips for Grouping in Your Map

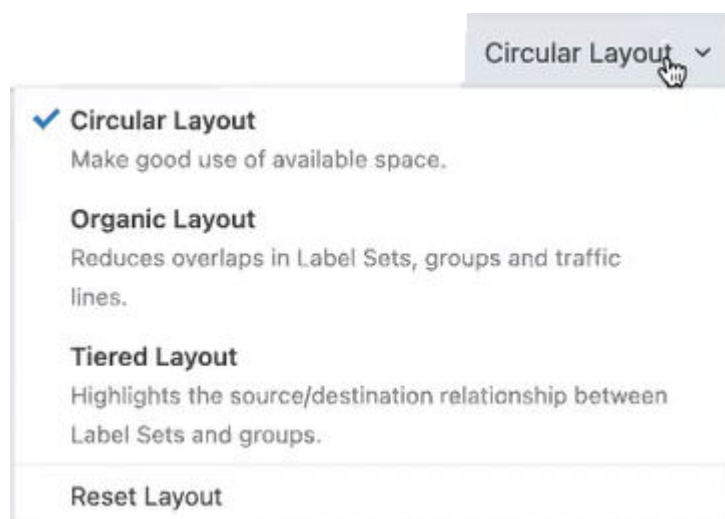
- Each group is a label set. Every workload that has the same set of labels will be grouped into one of those label-sets.
- Hovering over a group in the Map displays a pop-up dialog box with the list of labels and number of workloads using the labels.



- In the Group by drop-down list, you can drop and drag labels in the list to order how your Map displays the groups. Labels at the top of the list control the prominence of those groups in the Map.
- The PCE web console displays the groups in your Map using the colors you've selected for your labels. Use these colors to help orient yourself on the Map.

Map Layout Options

You can choose how you the PCE web console to display your Map:



Not every layout choice is good for your Map data. See the descriptions of each layout in the drop-down menu.

For example, the Organic Layout option attempts to organize groups so that the workloads that are connected are grouped together and displays less cross traffic. Workloads that are communicating are grouped together on one side of the Map and the traffic links aren't crossing as much.

The Tiered Layout option provides a sense of traffic flow top to bottom. The Tiered Layout option is better for smaller data sets than larger ones.

How to Read the Map Symbols

Legend - New UI

The legend is displayed in a panel with tabs for 'Circular Layout', 'Reported View', and 'Legend'. The 'Legend' tab is active. It contains four columns of symbols and their corresponding meanings:

NODE EXAMPLES	ENFORCEMENT	TRAFFIC LINKS	POLICY DECISIONS BY DENY RULES
Collapsed Group	Visibility Only	Blocked	Blocked
Workload	Selective Enforcement	Potentially Blocked	Potentially Blocked
Virtual Server	Full Enforcement	Allowed	Allowed
Virtual Service	Mixed Enforcement	Loading rule data	
		Rules not calculated	

Legend - Classic UI



Map Symbols Explained

Pay attention to the way that the Map groups designate the enforcement mode for groups:

- Workloads and groups inside full dark lines depict the FullEnforcement mode.
- Workloads and groups inside light blue lines depict the SelectiveEnforcement mode.
- Workloads and groups inside light orange lines depict the Visibility only mode.
- The ring around a group denotes the proportions of different enforcement states

As you navigate into the groups, you notice that the workloads also have borders indicating their enforcement modes.

Traffic links are presented with lines and arrows in different colors:

- **Green:** Traffic is allowed
- **Yellow:** Traffic is simulated blocked
- **Red:** Traffic is blocked
- **Grey:** Rules are not calculated
- **Gradient arrows:** The light color is next to the source and dark next to the destination. Gradient arrows are used while the rule data is still loading from the traffic.

When you click a group in the Map, the PCE web console highlights the links to and from that group using the colors defined above.

Map Reported View

In the Map, the PCE UI displays the traffic using red, orange, or green lines to indicate whether the VEN had a rule that allows the traffic when the connection was attempted.

- A green line indicates that the VEN had an explicit rule to allow the traffic when the connection was attempted
- A red line indicates that the VEN did not have an explicit rule to allow the traffic when the connection was attempted
- An orange line indicates that no explicit rule exists, but because of the enforcement state of the workloads, the traffic is not blocked when provisioned.



NOTE

When a policy change occurs, only flows that are created after the policy change are displayed in red or green based on the new policy. Flows created before the policy change might continue to be displayed in red or green using the old policy.

If multiple rules allow traffic between entities, only one green line is displayed.

Rules created for existing or live traffic don't change the color of the traffic lines in the Reported view, even when they are provisioned, until new traffic is detected.

Map Draft View

This view also displays the traffic using red, green, and orange lines to indicate whether the PCE has a rule to allow the connection that was reported by the VEN. This way, you can add rules and see their anticipated effect in real-time before the rules are implemented. In the Draft view of the Map, line colors have the following meanings:

- A green line indicates that the PCE had an explicit rule (in either a draft or an active policy) to allow traffic when the connection was attempted.
- A red line indicates that the PCE did not have an explicit rule (in either a draft or an active policy) to allow traffic when the connection was attempted.
- An orange line indicates that no explicit rule exists, but because of the enforcement state of the workloads, the traffic will not be blocked when the rules are provisioned.

Panels in the Map

When you click an object in the Map, the PCE UI displays a side panel on the right that contains three tabs (dependent on the object you clicked): Summary, Connections, and Workloads.

Summary Tab

The Summary tab for the Map displays information about a selected object. To view the Summary tab, select an item, such as a traffic line, on the Illumination Plus Map. The Map has a few types of Summary tabs:

- Traffic detail
- Group detail
- Workload/Virtual Service/Container Workload/Virtual Server

SummaryTrafficWorkloads

Add Rule

LABELS

Labels

Production

GENERAL

Workloads51

Virtual Services1

ENFORCEMENT

Visibility Only3 Workloads

Selective Enforcement16 Workloads

Full Enforcement28 Workloads

For example, when you click a group in the Map, the Summary tab displays what the labels are, how many workloads there are, how many virtual services, the enforcement level, and you get detailed links depending on how far you drilled into the group.

Traffic Tab

The Traffic tab is a summary version of the main Traffic table and filtered by what you've selected in the Map.

The Traffic tab appears regardless of what you select in the Map: group types, workloads, IP lists, private addresses, public addresses, or links. By default, the Connections tab displays the following columns.

- Policy Decisions (reported and draft)
- Consumer Labels
- Provider Labels
- Provider Port Processes

You can add additional columns by selecting options from the Customize columns drop-down list:

- Consumer Processes
- Flows/Bytes
- First Detected
- Last detected

See [Customize Columns \[23\]](#) for more information.

Workloads Tab

The Workloads tab displays a list of all workloads in the selected group and the following information for each workload:

- Workload name
- The V-E (vulnerability) score
- Enforcement mode
- Labels

As you drill in and out of the groups in the Map, the Workloads tab adjusts to show the workloads in the super set group.

Traffic Table

The Traffic table in the visualization tools displays search results in a traditional table format. You can use the Traffic table in the following ways:

- To write rules for specific connections; see [Add Rules for Traffic Using Illumination Plus \[41\]](#)
- Create unmanaged workloads from IP addresses; see [Create Unmanaged Workloads from IP Addresses \[39\]](#)
- Traffic exploration
- View the details about policy affecting each connection
- View the ransomware protection details.

About the Traffic Table

Using the Traffic table, you can query the PCE's traffic database to analyze traffic flows for auditing, reporting, and troubleshooting. You can search for traffic flows between workloads or hosts, labeled workloads, or IP addresses, and you can restrict the search by specific port numbers and protocols.

The VEN decorates the flow summary logs with DNS names when it sends them to the PCE. In the Traffic table, the PCE appends the DNS names to the flow logs so that auditors and SOC analysts can look at these DNS names instead of performing reverse look-ups on random IP addresses.

When you want to search for traffic flows on a regular basis, you can save that filter and it appears under your *Saved* filters in the **Load Filter** drop-down list. You can save up to 100 filters. You can make changes to an existing Saved filter and save the modified query. The Traffic table also displays your ten most recent searches.

Searches

When you search data in the Traffic table, you are searching traffic flows between sources and destinations over a specific time period over a specific port and protocol. A search consists of the following elements:

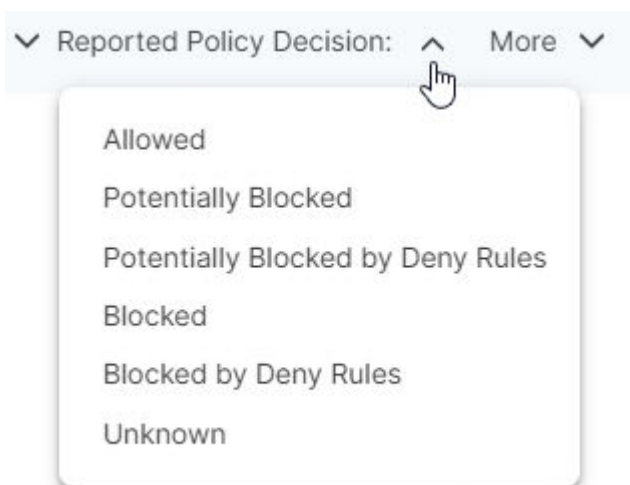
- **Destination:** Enter workloads, IP addresses, or labels that are consuming the service provided in the traffic flow. The entries you add in the filter that includes the data are used as a search criteria and the ones you add in the a field that excludes data are not used in the search.
- **Source:** Enter workloads, IP addresses, or labels that are providing the service in the traffic flow. The entries that you add to include the data are used as a search criteria and the ones you add to exclude the data not used in the search.



NOTE

You can choose to search either “Destination *And* Source” or “Destination *Or* Source” by selecting the option from the **More** menu.

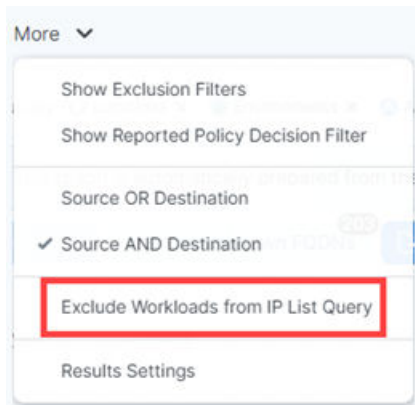
- **Services:** Enter port and protocol, port ranges, process, Windows services, or policy services. Enter port numbers and protocol types to search for traffic flows whose destination port values and protocols match the search criteria. The entries you add to include in the search are used as a search criteria and the ones you add to exclude data are not used in the search. If you do not specify a value, all ports, protocols, port ranges, processes, and services are included in the search.
- **Time:** Select how far in the past (last hour, day, week, month, or anytime) or specify a custom time range. The custom time filter displays all the flows between the selected from-to date-time stamp.
- **Reported Policy Decision:** Select the type of policy decision to search for flows with a specific policy decision reported by the VEN.



See [Deny Rules and the Traffic Table \[35\]](#) in this topic for more information.

- **Exclude Workloads from IP List Query:** (Available in the **More** drop-down menu.) This setting applies to queries that contain an IP list in the Consumer or Provider fields. It specifies whether known managed and unmanaged workloads are excluded from the query results.

When selected (the default setting), managed and unmanaged workloads are excluded from query results when their IP addresses are within the range of one of the IP lists in the query. When this option is not selected, workloads are not excluded from the query results.



Export Query Results

In the Table view, click **Export** to gather your data in a CSV file for the results from the current query.

To export results from previous queries, click **Load Results** to display queries from the past 24 hours. Click the **Export** button in the **Action** column for the results you want to save as a CSV file.

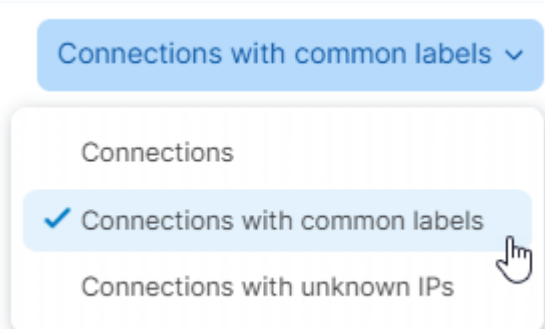
The exported CSV file uses a separate column for each label type, and the column data is alphabetized.

If you are an Illumio Core customer who has upgraded to 22.5.0 and are using Illumination Plus, be aware that the format of exported CSV files has changed from previous releases of Illumination Classic. You should update any scripts that you used for processing these CSV files.

View by Connections with Common Labels

In the Traffic view, you can view aggregated results of the Destination and Source labels for traffic flows or view all traffic flows for a query.

To choose the type of view you want, select the option from the **Connections with common labels** drop-down menu (New UI):



Label-Set Connections drop-down menu (Classic UI):



IMPORTANT

The Classic UI uses the terminology "Label-Set Connections" for this feature.

Using this feature, you can see a more concise view of your traffic flows.



IMPORTANT

This setting is important because to write rules from the Traffic table, you must be viewing the Traffic table using the **Connections with common labels** option. The **Allow Selected Connections** button in the Traffic table is disabled until you choose this setting.

The view for **Connections with common labels** displays the Draft rules based on the label queries; whereas the view for **Connections** displays the workload-to-workload rules, which can take longer to display the list but can be more accurate. Toggling back to the **Connections with common labels** option after displaying the individual connections does not reload the page so that the page displays quickly.

View Policy Details from the Traffic Table

The Traffic table includes a Policy Decision column (either Reported or Draft depending on the view selected), which indicates whether traffic flows are allowed, blocked, or potentially blocked based on your policy.

When you see traffic flows that are potentially blocked, it could mean that you haven't created rules for those flows or you have rules written for the flows, but the provider workload enforcement is set to Visibility Only for those flows.

Clicking a link for Allowed traffic opens the **View Policy** dialog box. When applicable, the dialog box displays in separate tabs all your policy, including Deny Rules, rules, and Essential Service rules that apply to the selected traffic flow

Deny Rules and the Traffic Table



NOTE

In the Classic UI, Deny Rules are still referred to as Enforcement Boundaries.

Deny Rules are displayed in Draft and Reported views of the Traffic table. When you view your traffic flows in the table, you see whether traffic is blocked by a Deny Rule or allowed through a Deny Rule. Viewing this information is useful to determine where Deny Rules are in place and understand their impact before provisioning them.



TIP

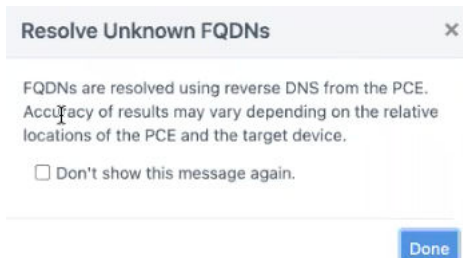
To view the details about a Deny Rule, click the linked text for traffic allowed across the rule (“Allowed”) or blocked by a Deny Rule (“Blocked”) while in a **Draft** view of the Traffic table. The **View Policy** dialog box opens. Then, click the **Deny Rules** tab.

You can obtain the following information:

- A Deny Rule is blocking a traffic flow.
- Traffic is potentially blocked by a Deny Rule.
A Deny Rule is in place, but the workload is still in visibility-only mode. The traffic won't be blocked by the rule until you move it into selective enforcement mode.
- A Deny Rule is in place, but an allow rule is allowing traffic through the Deny Rule.

Resolve Unknown FQDNs

1. Click **Resolve Unknown FQDNs** to export FQDN information for unknown IP Addresses and **Done** from the confirmation dialog box.



2. Click **Export**. This button appears next to Resolve Unknown FQDNs.



NOTE

Clear cached FQDN values and reload the results if you do not find relevant information.

Depending on the number of draft rules, the data might be slow to load. Once it loads, columns called Draft Policy Decision and Reported Policy Decision will be populated with data and will appear in the exported zip file.

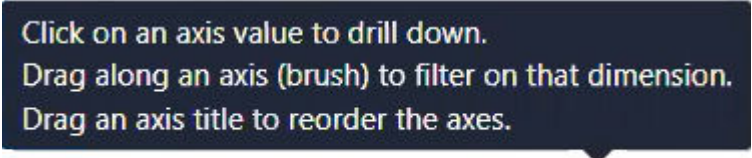
Mesh View

The Mesh view displays traffic flows as a vertical list of Destinations, Sources, and the port being used in the flows.

About the Mesh

You can click any item in the results to focus on specific flows. You can also sort the results to view results based on port number or number of traffic flows. From the Mesh view, you can drill down to filter, brush to filter, and then go to the Table view to write rules.

If you need prompting on how to filter the data in your Mesh view, click the **Filtering Tips** link in the bottom-right corner of the page for a pop-up tooltip.



Click on an axis value to drill down.
Drag along an axis (brush) to filter on that dimension.
Drag an axis title to reorder the axes.

 **Filtering Tips**

The Mesh view always displays traffic flows based on the Reported view. You cannot switch to the Draft view for the Mesh. The **View** menu only appears when you are in the Illumination Plus Table and Map views.

Mesh View Limitation

In this release, you cannot filter in Mesh view to view only allowed or blocked traffic. The Mesh view does not allow you to toggle your data between the Reported view and the Draft view. The Mesh view only displays reported view data.

Customize the Mesh View Display

The *Group by* field that you use with the Map view is also used in the Mesh view. You will see your top group based on your selection and you can drill down through the groupings. The hierarchy of the parallel coordinates in the mesh is based on your selected grouping.

You can reorder the axis columns in the Mesh view by clicking an axis heading and dragging it left or right. You might want to reorder the Mesh axis columns to change the Mesh display; for example, you might drag the axis you are most interested in to the center of the mesh and less important data to the sides of the mesh.

You can sort the Mesh data by the axis value which displays the ports numerically low (at the top) to high (at the bottom); for example, port 22 at the top and 10051 at the bottom of the axis. Along the axes, the app groups are in alphabetical order.

You can also sort by number of flows. Each axis is sorted based on which group has the most traffic. The page displays the “tick” with the most flows at the top of the axis.




You can use the brush feature to select along the top of an axis to select and filter by the highest ports.

Navigate the Mesh View Data

In the Mesh view, you can step down each vertical axis along the axis ticks to view specific connections between that tick and the data in other axes. Stepping through the axes ticks is reflected in the breadcrumbs above the Mesh to pinpoint your location as you view the data.


To brush the axes lines and filter Mesh data:

1. Hover over an axis line, so that the black line changes color to orange.

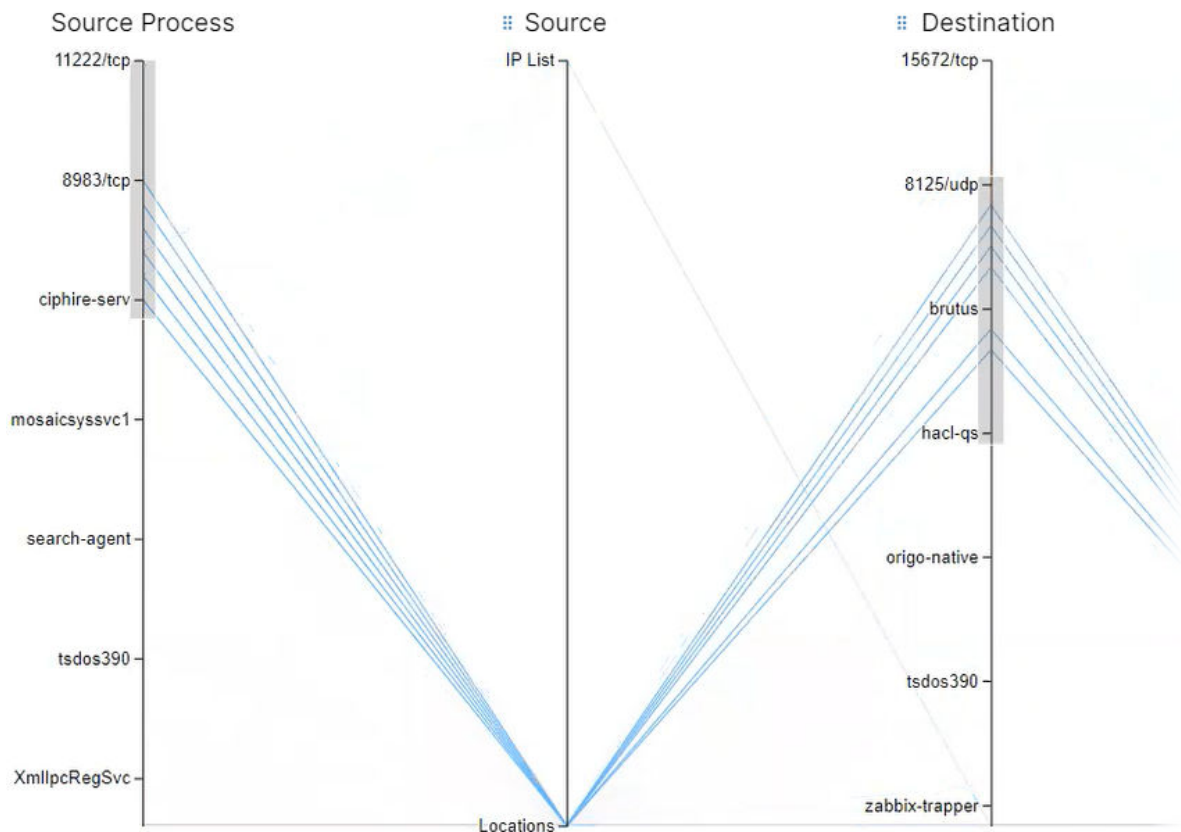


2. Click and drag down an axis line to apply a "brush" and select portions of the data.

Destination Process



- Using the brush feature on multiple axes, you can select areas that have connections to highlight certain flows.



- When done filtering Mesh data, click **Clear Brush** in the top-right menu bar to reset the Mesh.

Work with the Visualization Tools

You can use the visualization tools to perform the following tasks.

Workflow for Using the Visualization Tools

The visualization tools enable you to build security policies for your workloads by following this workflow:

- Group discovery:** When you pair workloads, the VEN introspects those Workloads and determines their open ports, running services, and traffic flows. See the VEN Installation and Upgrade Guide for information about installing (also called pairing) VENs on workloads.
- Prepare group for rules:** Prepare a group for rules by applying labels to each workload in the group so you can write policies for them.
- Rule writing:** After you have prepared the group for rule writing, you can begin to write rules for the workloads in the group. This requires writing rules to allow communication between workloads across groups, between workloads in the same group, or between workloads and other entities outside the group (for example, the Internet or an IP List). Illumination Plus will also propose suitable rules for you to use or modify if you do not want to manually create rules from scratch. See "IP Lists" and "Rule Writing" in the *Security Policy Guide* for more information.

- 4. Rule Testing:** Illumination gives you the power to test and evaluate your rules against existing traffic flows *without* enforcing the rules. Rules can be tested to ensure that legitimate traffic flows required by an application are permitted and malicious traffic is blocked. Exporting traffic summaries or using blocked traffic lets you know which traffic connections would be dropped if the rules were enforced. .
- 5. Policy Enforcement:** When you are ready to implement the rules for a group, you can put the group into the enforced state. Leveraging the Illumio Core allowlist policy model, any traffic flows that are not explicitly allowed by a rule are dropped. If a legitimate application flow is broken or an intrusion occurs, you can configure notifications to alert you.

About Unmanaged IP Addresses

From the Map, you can quickly create unmanaged workloads from IP addresses. A reverse DNS lookup is done on the IP addresses to obtain and display the server name for the unmanaged workload. The server names are only displayed in the PCE web console. When you export the file, it lists IP addresses.



NOTE

The DNS names are not displayed in Illumination Plus for Illumio Core Cloud customers.

When you select an IP address in the Map that is not currently associated with another policy object, it automatically populates the IP address into an unmanaged workload with the following values:

- A default interface of eth0
- The hostname, which is the IP address by default

IPv4 or IPv6 addresses displayed in the Map can be selected from the internet, IP lists, or traffic links. The default interface and hostname can be changed if needed and labels can be added to the unmanaged workload.

Until new traffic for the unmanaged workload is observed, the traffic lines are not displayed for the unmanaged workload. The traffic lines in the Map are updated after new flows are reported by the PCE.

If you try to create an unmanaged workload from an IP address where an unmanaged workload already exists, an error message is displayed.

After you convert an unmanaged IP address to an unmanaged workload, you can use it in your policy; for example, you want to allow one of your hosts to communicate with a managed workload. A reverse DNS lookup is done on the IP addresses listed under the Destination column and you see the name of the server instead of the IP address.

Create an Unmanaged Workload from an IP Address

The Map includes groups for unmanaged IP addresses. First, the PCE maps IP addresses to an IP list; then, if the IP address is in the RFC set of IP addresses, those IP addresses

appear in the private IP address group. Lastly, the Map contains a public IP address group that encompasses all the rest of the IP addresses that are part of the Internet. You can create unmanaged workloads for each type of IP address.

1. In the Map view, click one of the following groups: **IP List**, **Private Addresses**, or **Public Addresses**.

The right-side panel for the object opens. For example, locate and click the IP List group. When you click the IP List group, the Summary tab in right-side panel displays the IP addresses broken down into each of the IP lists that they match.

For private and public IP addresses, the tab displays the list of IP addresses.

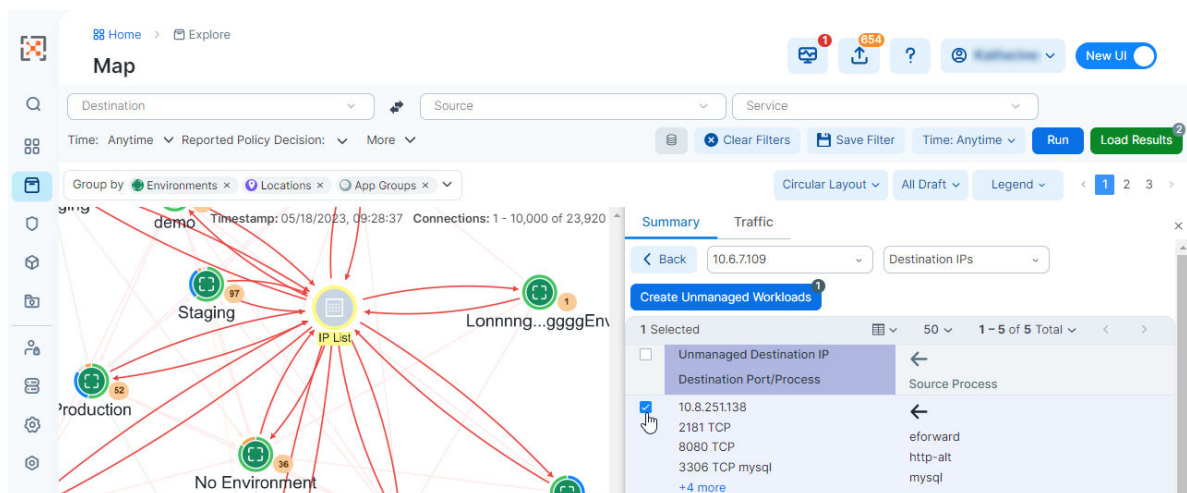
2. For an IP address in an IP list, click a row to expand the IP addresses in that IP list. The panel refreshes and displays any unmanaged IP addresses that are communicating with your managed workloads. (For public and private IP addresses, you can skip this step.)



NOTE

If you have a reverse DNS lookup, the server name is used instead of the IP address.

3. Select the checkbox that you want to create an unmanaged workload for. The **Create Unmanaged Workloads** button becomes enabled.



4. Click **Create Unmanaged Workloads**. The Assign Labels dialog box appears.
5. From the drop-down list, select the labels to assign to the unmanaged workload and click **Confirm**.

A dialog box appears indicating that the unmanaged workload was created.

6. [Optional] Recalculate your map with the newly created unmanaged workload by clicking **Recalculate** in the confirmation dialog box.

To complete the configuration of the unmanaged workload, perform the rest of the steps.

7. From the PCE UI navigation, choose **Servers & Endpoints > Workloads** (New UI).

The Workloads page appears.

In the Workloads list, locate the new unmanaged workload you created. Identify the unmanaged workload by its name, which is its IP address.

The new unmanaged workload does not list any information for its enforcement because it does not have a VEN installed on it.

8. To complete the configuration for the unmanaged workload, click its IP address in the Workload list.

The Unmanaged Workload page appears.

9. Click **Edit** and complete the workload information. For information about the fields for unmanaged workloads, see "Add an Unmanaged Workload" in the Security Policy Guide.

10. Click **Save**.

.

Add Rules for Traffic Using the Traffic Table

You can use Illumination Plus to add rules for traffic flows by selecting traffic flows and then allowing the selected connections.

In the Table view, you can only write rules for one page of traffic flows at a time. You must click through each page. (This limitation matches the way other tasks are performed in the Table view.)

To add rules for traffic flows:

1. From the PCE web console main menu, choose **Explore > Traffic** (New UI).
2. From the Traffic table, select **Connections with common labels** and a Draft view: **All Draft**, **All Blocked**, **Potentially Blocked**, or **Blocked**.
3. Using the checkboxes, select traffic flows that you want to write rules for. The **Allow Selected Connections** button changes color from pale to bright blue and includes the number of allowable connections for which the PCE can write rules.
4. Click **Allow Selected Connections**.



NOTE

Under certain conditions the button won't be enabled; for example, you've only selected traffic flows that are already allowed. When this occurs, either select other traffic flows or click the **Edit Labels** button to modify the traffic flows.

The page refreshes and displays proposed rulesets or rules depending on whether you have enabled basic or advanced modes for rule writing. See "Basic and Advanced Modes for Rules" in the Security Policy Guide for a distinction between these modes.

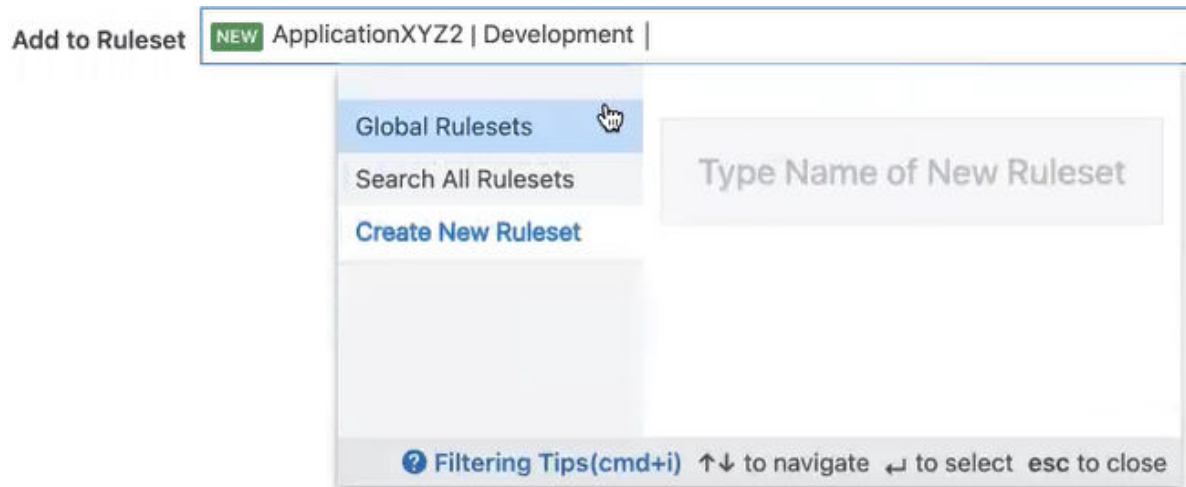
When you are using the basic mode for rule writing, the page contains only a list of proposed rules, and you aren't able to add scoped rulesets. You can only select global rulesets.

When advanced mode for rule writing is enabled (so that you can create scoped rules), the page contains tabs for relevant intra-scope and extra-scope rules for the ruleset. The PCE chooses the proposed ruleset based on the scope of the traffic flows you selected.

For example, you have selected two traffic flows that have the same set of labels so that they fall within the same scope. When you have a ruleset that already has that scope, the PCE defaults to that ruleset. Therefore, the PCE displays a list of options that match that

scope. Alternatively, you select a third traffic flow that has different labels from the first two traffic flows, the PCE will display the global rulesets as an option to add the rules to.

5. Either accept the default ruleset or select a different ruleset to add the rules to.




When in advanced rule writing mode, the **Add to Ruleset** drop-down menu contains these categories: rulesets appropriate for the scope, global rulesets, and the ability to search all rulesets, and create a new ruleset.

When you elect to create a new ruleset, the **Add Ruleset** dialog box appears. Select the **Add Scopes** checkbox to see all the scopes that are common to the selected traffic flows you are adding rules for.

6. As needed, edit the scopes for your ruleset and click the Save icon:



After clicking the Edit icon, the scope field become editable:



If you remove all the scopes from the ruleset, the labels for the scope appear in the rules.

7. [Optional] To control how the PCE uses services in the rules, click the **Settings** button.

Settings

☐ Allow traffic on all Services

☒ Allow the use of Services that include more than one port/protocol

If no existing Service matches a port/protocol

☒ Use the port/protocol in a rule

☐ Create a new Service

You can choose to use all services or services that include more than one port/protocol. .

When selecting to use services that include more than one port/protocol, the PCE doesn't require an exact match on the service. For example, you want to use service TCP 3306, but the PCE contains TCP/UDP 3306. Selecting this option enables the rule to use TCP/UDP 3306 as a matching service. When the PCE doesn't have a matching service, you can choose to use the port/protocol in the rule or create a new service. By default, the PCE creates the rules by using the port/protocol.

8. As needed, edit the proposed rules, and save your changes by clicking the **Save** icons at the end of the rows.



NOTE

When you edit rules and if any overlap exists between rules due to your changes, the PCE will optimize the rules so that duplicates are eliminated. For each duplicate rule that isn't provisioned, the PCE displays a label in left column "Proposed Delete" and will delete that rule.

9. Once you're satisfied with the ruleset selected and the rules within the ruleset, click **Save** or **Save and Provision**, depending on whether you want to immediately provision to ruleset.

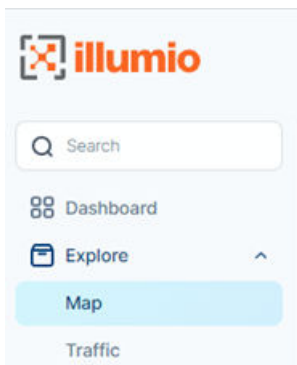
See "Provisioning" in *Security Policy Guide* for information.

After saving your ruleset and rules, the PCE web console reloads your data so that the Table view and Map view reflect the changes.


Write a Ringfencing Rule

Using the Map view, you can quickly create a ringfencing rule by adding that rule to a new ruleset within the scope of the selected group.

1. In the left-hand menu select the Map view (New UI).



2. Verify by which criteria the group has been established.
Look at *Group By* selection and apply any changes that you might need.



3. Keep the current selection (Locations, Environments, Applications), or add or remove the grouping criteria.

Once you have the desired selection, click **Apply**.

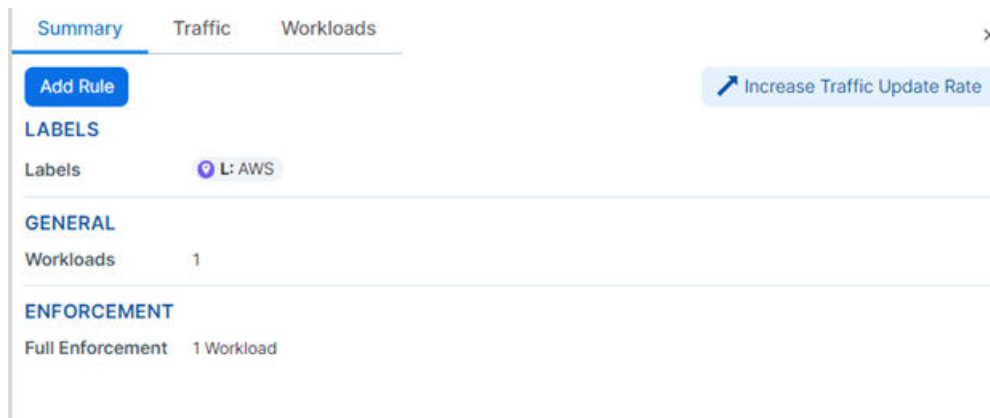
The group is now established according to your needs.

4. Now put the cursor over the group that you want to change (here it is AWS).




The pop-up dialog on the left shows the selected group's stats.


5. You can also click on the group to see its stats that show in the right-hand panel.



6. Now click on the group where you are adding the rule and then on **Add Rule**.



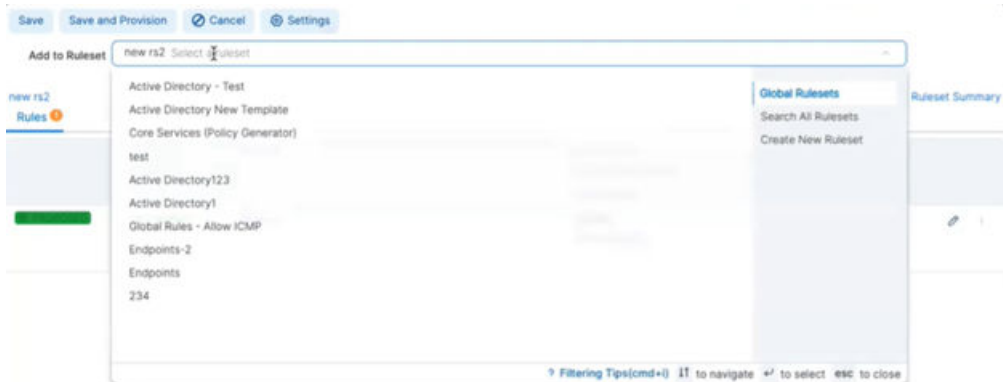
7. Choose to which ruleset you are adding the new rule, for example the ruleset named **new rs2**.




8. Select **Rule Options**.


You can, for example, select **All Services**.

9. Add a rule that is *All Services* to *All Services*.



- 10 After you have added the rule, click on **View Rule** to view it.





Everything inside that Rule talks to each other.

Monitor Traffic Database Size and Receive Alerts

You can monitor your database traffic usage and be alerted when you are close to capacity.




NOTE

The storage information is based on your customer organization limit and not the overall capacity of the PCE for your environment.

To monitor traffic database size:

1. In the PCE UI left navigation, choose **Explore > Traffic**.
The Traffic page appears.
2. From the top status bar, hover over the database icon:



A pop-up window appears, which displays the how many more days of data you can store your traffic data in the Illumio Core cloud. You receive an alert when your disk space is within 15% of your available space.

App Group Map

An App Group is a logical grouping of workloads associated with an application instance, which is defined by the labels assigned to the workloads in it. This section describes the types of App Groups, the App Group Map, and how to configure App Groups.

About the App Group Map

The App Group map visualizes all the App Groups in your PCE to help you quickly access specific workloads based on the App Group to which they belong. You can also view the traffic with rule coverage considering Windows process-based services.

The Illumination and Illumination Plus Maps visualize the workloads and traffic in your data center, which takes time to render with large-scale deployments. However, some users such

as application owners prefer to think about their data center in terms of traffic between workloads that belong to different application instances, rather than between physical locations.

The App Group Map is designed to provide visualization for application owners by showing all workloads for an application instance in a single App Group, even when they are not currently communicating with each other. This feature allows application owners to focus on the workloads that only belong to their applications, regardless of location, when building or validating security policies for traffic between workloads.

The App Group Map visualizes the network traffic by organizing it based on App Groups.

The App Group Map displays all the App Groups in your PCE to help you quickly access specific workloads their traffic based on the App Group to which they belong. For each chosen App Group, you can view:

- **Consuming App Groups:** Use services provided by the current application
- **Providing App Groups:** Provide services used by the current application

You can search for specific App Groups and see the associated workloads, traffic, and rule coverage between the workloads in that App Group, other App Groups that provide or consume its services, and rule coverage for the traffic between App Groups.

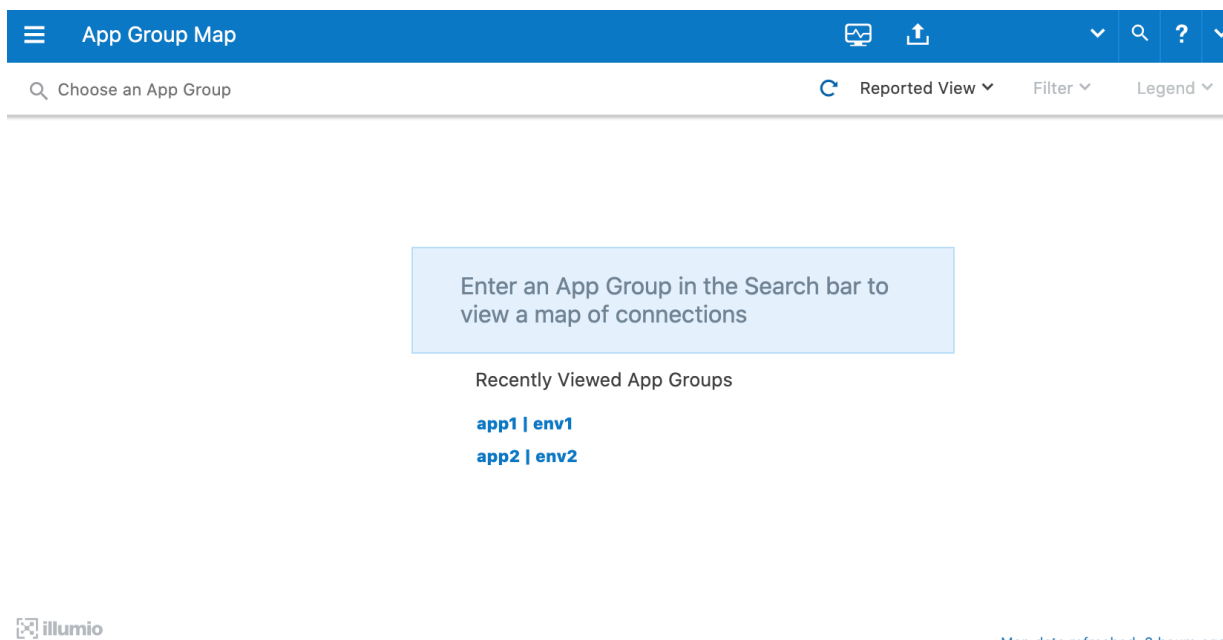
App Group Views

The App Group Map initially displays a search bar that allows you to search for a specific App Group. When you have previously used the App Group page, a list of recently viewed App Groups is also displayed.



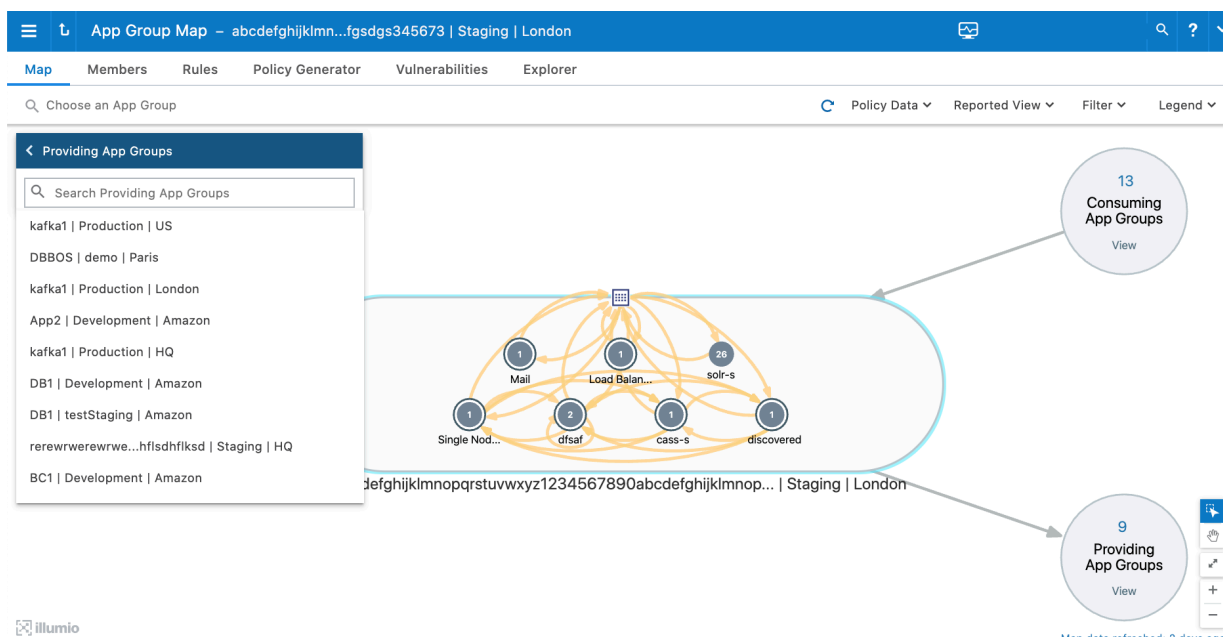
NOTE

If you click an App Group that contains more than 1,000 workloads, you see an alert message and the workloads are not displayed.



When you select an App Group (either from the list of recently viewed App Groups if it exists or from the drop-down list in the App Group search bar), the workloads and traffic for the workloads in that App Group displays, as well as a list of other App Groups communicating with that App Group either as providers or consumers of services.

Above the App Group, you see a link to the App Groups that initiates connections to this application instance. Below the App Group, you see a link to the App Groups that provide services for this application instance.



To view the consuming or providing App Groups, click **View**. A pop-up window displays the name of each App Group, its Location label, and the number of workloads it contains.

From this pop-up window, you can click **Close** to close it or select an App Group to display it in the App Group Map.

**NOTE**

If the App Group does not have any connections, the Providing and Consuming App Groups do not display.

When you select a Consuming or Providing App Group, an oval representing the expanded App Group displays in the App Group Mmap. Lines representing the traffic links between the App Groups are displayed in either red for blocked traffic or green for allowed traffic. Consuming App Groups display above the original App Group and Providing App Groups display below the original App Group.

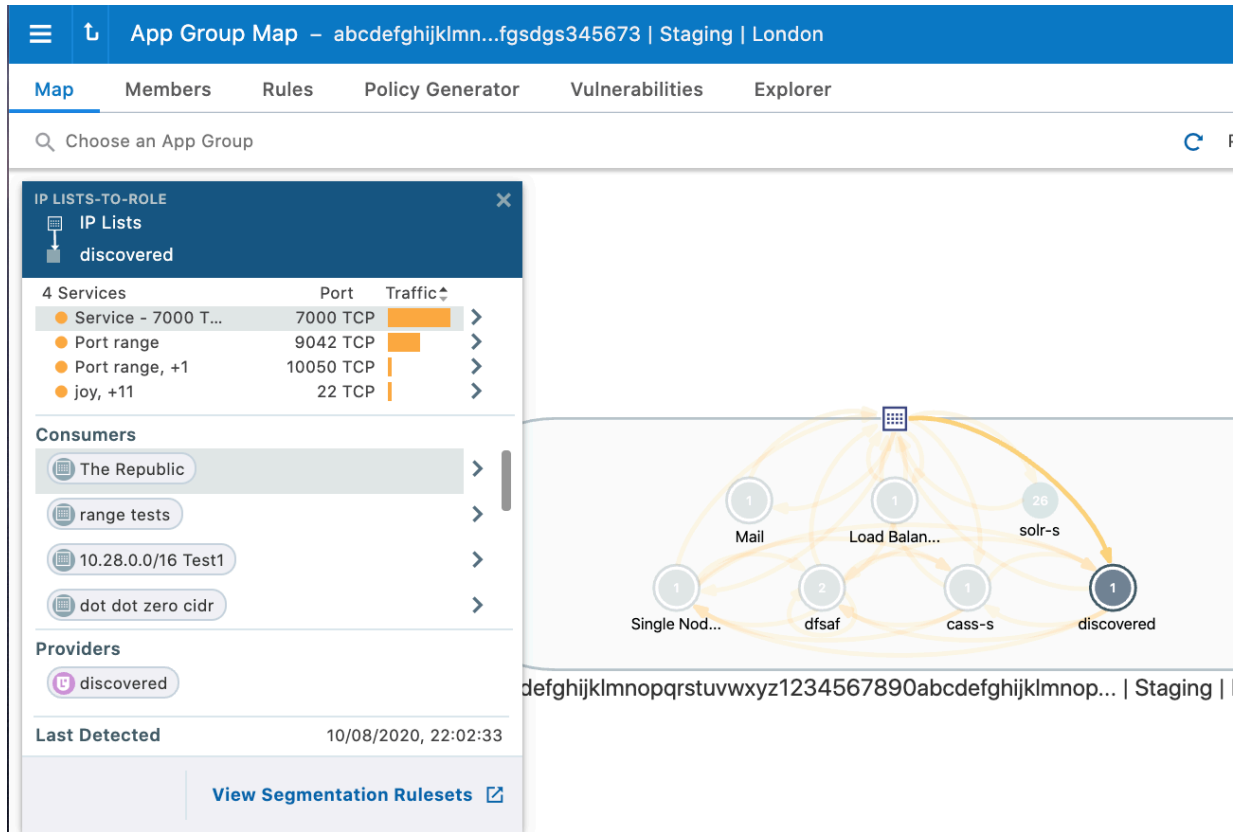
If an expanded Consuming or Providing App Group is currently displayed in the App Group Map, the link in the App Group's circle changes from **View** to **Next**. Click **Next** to view the next connected Consuming or Providing App Group.

When you select an App Group, the list of all observed services between any workloads in that App Group displays. When you click a specific line between two workloads, all services between the selected workloads display.

When you have virtual servers, you can view their details in the App Group Map command panel in both Reported and Draft views.

When you select a traffic line between two App Groups and click **Create Ruleset**, the auto-populated name is a combination of the labels for the selected App Group.

When a ruleset already exists for this traffic, click **View Ruleset** to display it.



NOTE

In previous releases, this feature was referred to as “Segmentation Rulesets.” In Illumio Core 21.5.0 and later releases, this feature is now referred to as “Rulesets.” This image still displays the previous feature name.

Application owners can write both intra- and extra-scope rules to allow others to use the application instance. However, as an application owner, you can only write rules when you are the owner of the Providing App Group to allow other Consuming App Groups to access your application workloads.

Work with the App Group Map

There are two types of App Groups: Providing App Groups and Consuming App Groups. Providing App Groups provide service to an application instance and Consuming App Groups rely on those services to run the application instances.

You can search for specific App Groups and see the associated workloads, traffic and segmentation rule coverage between the workloads in that App Group, other App Groups that provide or consume its services, and segmentation rule coverage for the traffic between App Groups.

App Group Creation and Association

An App Group is created when:

- A new workload is added or discovered and there are no existing App Groups using the workload's labels
- A label is changed on an existing workload and there are no existing App Groups using that label combination

A workload is associated with an App Group when:

- A workload is paired or unpaired with the PCE
- A label is changed on a workload

When a new workload is added, it is associated with any existing App Group that uses the workload's labels. When an App Group with those labels does not exist, it is created and associated with the workload.

When the App Group uses a different Location label but has the same Application and Environment labels as an existing App Group, a new App Group using the Application, Environment, and Location labels is created and associated with the workload.

Configure App Groups

App Groups are created automatically based on workload labels and the App Group Type setting. App Groups can be configured to require two or three matching labels.

There are two ways to configure App Groups:

- App Groups formed by the Application and Environment labels
- App Groups formed by the Application, Environment, and Location labels



NOTE

If the Application | Environment option is selected, the workloads displayed in the Illumination map and the App Group map are not the same and there is no link to return to the Illumination map.

To specify whether App Groups should be created based on the Application and Environment labels or the Application, Environment, and Location labels:

1. In the PCE web console menu, choose **Settings > App Group Configuration**.
The App Group list page appears. The page displays the type of group (two labels or three labels) and the number of workloads per App Group.
2. Click **Set App Group Type**.
3. Select the appropriate radio button (either Application and Environment or Application, Environment, and Location). The default option is Application and Environment.
4. Click **Save**

Caveats

- When the App Group Configuration setting is changed, the list of “Most Recently Viewed App Groups” is cleared.
- When you have a large number of workloads in your organization, it can take up to five minutes to regenerate the Illumination map. To confirm the request has not timed out or is still pending, check the Network tab.

Vulnerability Map

You can visualize vulnerabilities across datacenters and clouds through a real-time Vulnerability Map. The vulnerability and threat data from the Qualys Cloud Platform is integrated with Illumio application dependency mapping to show potential attack paths in real time.

About Vulnerability Map

Vulnerability management and micro-segmentation are foundational security controls of a successful cybersecurity strategy. The Illumio Vulnerability Map combines Illumio’s App Group Map (an application dependency map) with vulnerability data from [Qualys Cloud Platform](#) to provide insights into the exposure of vulnerabilities and attack paths across your applications running in datacenters and clouds. This enables application security teams, vulnerability management teams, and segmentation teams to understand not only the vulnerability of a workload but more importantly the paths that bad actors can leverage to exploit vulnerabilities.

The Vulnerability Map integrates application dependencies and network flows with the vulnerabilities on the host that are exposed on communicating ports.

Vulnerability Terminology

- **Vulnerability:** A generic vulnerability that can exist on any workload (or port and protocol), for example, Apache heart bleed.
- **Detected Vulnerability:** The instance of a vulnerability that exists on a workload, for example, Apache heart bleed existing on workload X on port 80.
- **Vulnerability Report:** A report containing the detected vulnerabilities.
- **Vulnerability Score:** The summation of severities of the vulnerabilities for an App Group, role, or workload where the individual vulnerability scores range between 0 and 10.
- **Exposure Score:** The E/W Exposure Score combined with the Internet Exposure. It is a score of how many workloads can use the vulnerable port on a workload based on the provisioned rules.
- **Vulnerability Exposure Score (V-E Score):** A calculated value based on the Vulnerability Score and the Exposure Score = $\sum f(VS, ES)$. It can be shown for an individual vulnerability on a port for a single workload or as a summation of all the V-E Scores for an App Group, role, or workload.
- **East-West (E/W) Exposure Score:** A count of workloads that can use a vulnerable port with the currently provisioned rules, and whether the vulnerability is exposed to the internet.
- **Internet Exposure:** Indicates whether a vulnerable port is exposed to traffic from the internet. Internet Exposure is enabled by the rules allowing inbound traffic on that port.
- **Severity:** Represents a range of Vulnerability Score values.

- 0 = Info
- 0.1 to 4.0 = Low
- 4.1 to 7.0 = Medium
- 7.1 to 9.0 = High
- 9.1 to 10 = Critical

You can select the severity level you want to consider when showing which traffic is going to the vulnerable ports.

Benefits of the Vulnerability Map

The Vulnerability Map has the following benefits:

- Visibility into the potential attack paths that could be exploited by a bad actor.
- The East-West exposure score calculates how many workloads can potentially exploit vulnerabilities.
- You can apply vulnerability-based micro-segmentation as a compensating control to reduce East-West exposure.

The East-West Exposure Score shows you how vulnerable a workload is to exploitation from other workloads in your datacenter. It is displayed per workload and is a calculation of how many workloads can potentially exploit individual vulnerabilities on any given workload that has a VEN. The lower the score, the smaller the chance that a bad actor can exploit vulnerabilities. This insight can be used to prioritize and generate precise micro-segmentation policies as a compensating control and help prioritize patching efforts.



NOTE

Vulnerabilities exposed over network ports can be exploited by remote bad actors. You can write security policies in the Illumio Core to eliminate or constrain exposure to such vulnerabilities. However, the Vulnerability Map does not include the local vulnerabilities (those not exposed over network ports) in its calculation, because there is no network exposure due to them.

Vulnerability Map Usage

In most organizations, vulnerability management is performed through scanners that scan infrastructure to identify vulnerabilities and provide reports. In some cases, there is no patch for zero-day vulnerabilities. Illumio Core vulnerability-based micro-segmentation gives security teams the ability to focus on where they are most vulnerable—inside their datacenter and cloud, leveraging micro-segmentation as a compensating control.

For example, consider the increased East-West traffic (server-to-server traffic within your datacenter) that the cloud brings with it. This creates many new attack surfaces. Combining vulnerability and threat data from the Qualys Cloud Platform and Illumio's application dependency mapping yields a vulnerability map that displays connections to vulnerabilities between and within applications. Using the Vulnerability Map you can see which of your workloads are highly vulnerable to attacks and can reduce the vulnerability score to make those workloads more secure.

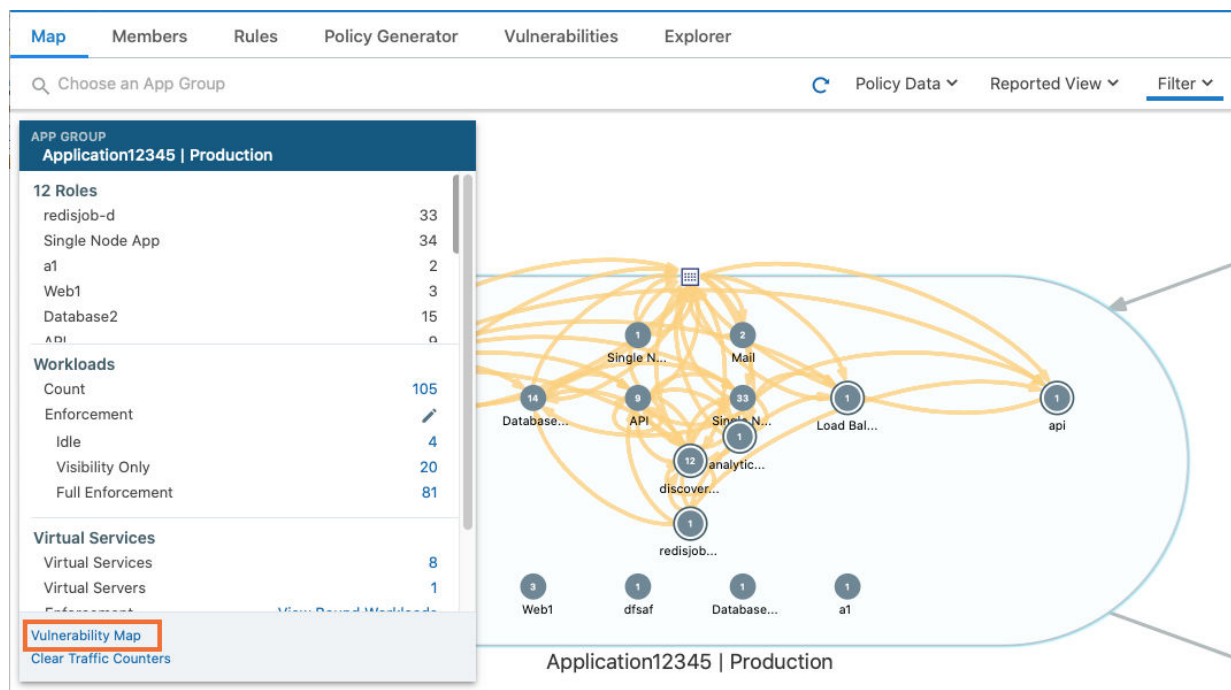
Work with Vulnerability Maps

The Vulnerability Map is a separately licensed feature of Illumio Core. The licensing is based on the number of workloads. The license is required to import Qualys report data into the Illumio PCE.

For information about obtaining the Illumio Core Vulnerability Map license, contact Illumio Customer Support.

Enable the Vulnerability Map

When you obtain the license, you will receive information about how to apply the license on the PCE and enable the feature.



After the Vulnerability Map is enabled, access it from the App Group Map by clicking **Vulnerability Map**.



NOTE

The Vulnerability Map is supported for VEN versions 16.9 and later.

Caveats

- A maximum of 100,000 vulnerabilities can be detected per organization.
- A maximum of 100 vulnerabilities can be detected per workload.
- The Vulnerability Map is not supported in Supercluster implementations.
- The exposure score is calculated on the first firewall sync for a given workload. When a PCE is restarted:

- Vulnerability Score and Exposure Score are not available until the firewall sync occurs.
- The scores are not available when a workload is offline.
- Vulnerabilities can only be imported using the PCE CLI Tool.

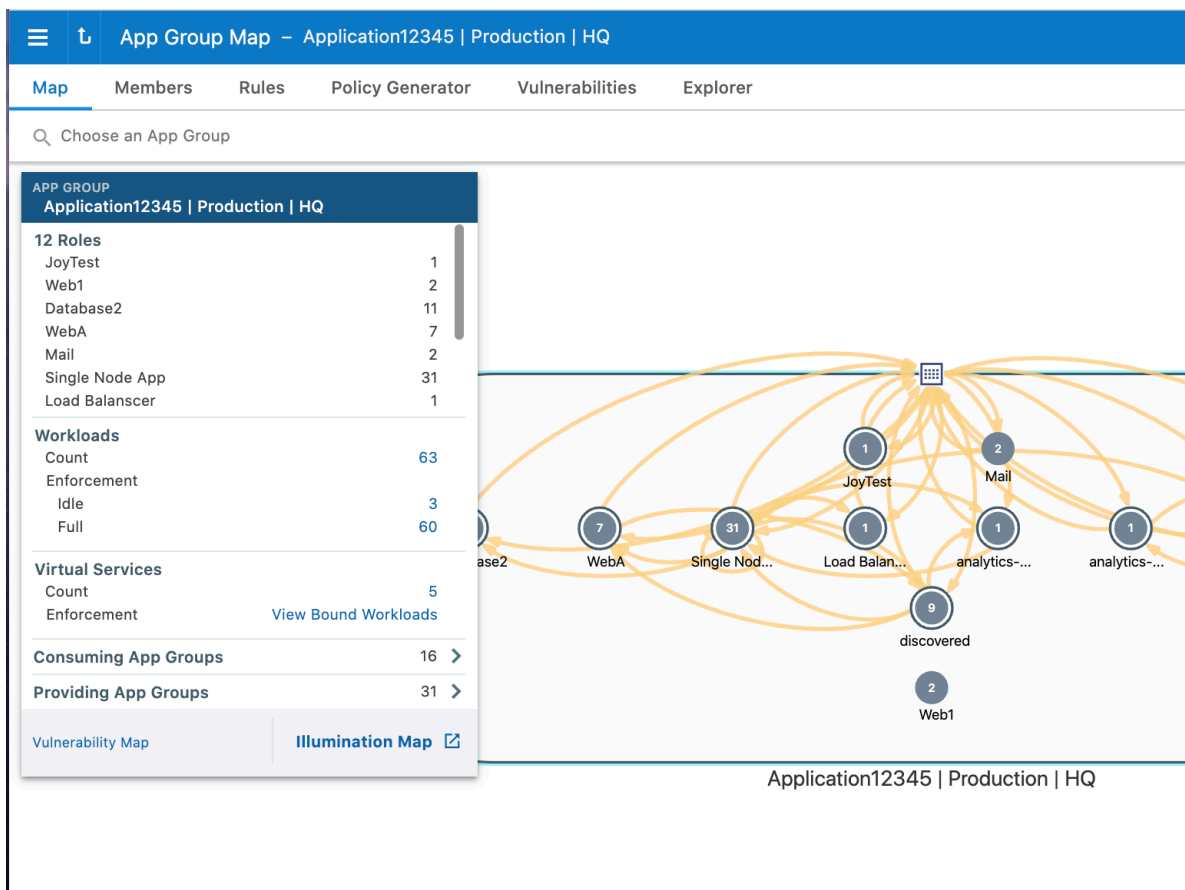
View and Mitigate Vulnerabilities

The Vulnerability Map in your PCE is disabled by default. Vulnerability information is available for traffic flows, workloads, roles, and App Groups.

To view and mitigate vulnerabilities:

1. In the PCE web console menu, in the upper left corner click on **Choose an App Group**.
2. From the pop-up list, select the App Group you want to work with

The command panel shows the different vulnerability exposure scores for the selected App Group, because of the port and to which workloads it is exposed. It is overlaid with the App Group Map. You see the Providing and Consuming App Groups and the vulnerable applications that are being accessed.



3.



NOTE

The Cloud icon denotes Northern Exposure.

4. To refine how you view the vulnerabilities for the selected App Group, select the **Filter** in the top-right corner of the map.

The Filter contains settings to view Vulnerability Exposure Score and Traffic. Based on your preference, you can set the slider to view only critical or high vulnerabilities or all of them.

The screenshot shows the 'Filter' menu in the Illumio Core 23.2 Visualization User Guide. The menu is open, displaying options for Traffic Links, Vulnerability Exposure Score, Vulnerability Traffic, and Workloads. The background shows a network diagram with a node labeled 'Load Balan...'.

TRAFFIC LINKS

- ☒ All Services ☐ Selected Services ([Edit](#))
- ☒ Intra-Group ☒ FQDN
- ☒ Internet ☒ IP List
- ☒ Broadcast ☒ Multicast
- ☒ ICMP

VULNERABILITY EXPOSURE SCORE

- ☒ All Vulnerabilities ☐ Exposed Vulnerabilities

VULNERABILITY TRAFFIC

(Slider from Info to Critical)

- ☒ Blocked Vulnerable Traffic

WORKLOADS

- ☒ Visibility Only ☒ Idle
- ☒ Full Enforcement ☒ Unmanaged
- ☒ Selective Enforcement

5. After identifying the vulnerabilities, you can constrain them to reduce the risk to your datacenter by writing a security policy.
 - a. Click **Policy Generator** in the menu in the left to open the Policy Generator.
 - b. In Policy Generator, select **Auto level** to automatically generate policy and set the Severity (slider) to the level of vulnerabilities that you want to constrain to.



NOTE

To see the Auto Level option, you must first import the vulnerability license and vulnerabilities.

Using **Auto Level**, you can write broad rules while minimizing the vulnerability exposure:

- Roles with no vulnerabilities: Role < All Services < All Workloads
- Roles with traffic to vulnerabilities: Role < All Services < Role
- Roles without traffic to vulnerabilities: Role < Specified Services < Role

You can also see the number of vulnerabilities for each workload:

- **Eliminated:** The port is not exposed to any other workload

- You can pick and choose the flows for which you want to include the policy.

The Preview page shows the before and after Vulnerability Exposure Scores, where:

- **After Includes:** All draft policy

6. Click **Save** after reviewing your policy.

Vulnerabilities Tab for Workload Details

The Workloads list page is enhanced to display risk due to vulnerabilities. The workload with the most vulnerabilities is listed at the top.

[illegible]

The Workload detail page includes a Vulnerabilities tab. You can click the V-E score column to sort the vulnerabilities based on the vulnerability score. You can then define your patch priority based on the most critical score.

Workload - solr-s41							
Summary	Processes	Rules	Blocked Traffic	Vulnerabilities			
~V-E Score	Vulnerability Score	E/W Exposure	Northern Exposure	Provided Traffic (Reported)	Port/Protocol	CVE-IDs	Name
33	7.8	48	None	None	123 UDP		Web Server HTTP Protocol Versions
33	7.8	48	None	Potentially Blocked	10050 TCP		Web Server HTTP Protocol Versions
33	7.8	48	None	None	8081 TCP		Web Server HTTP Protocol Versions
23	6.9	48	None	None	34571 TCP		SSL/TLS Server supports TLSv1.0
3.5	3.7	48	None	Potentially Blocked	22 TCP		Presence of a Load-Balancing Device Detected
3.5	3.7	48	None	None	32000 TCP		Presence of a Load-Balancing Device Detected
3.5	3.7	48	None	None	25 TCP		Presence of a Load-Balancing Device Detected

You can see the highest severity type for the workload and the total number of vulnerabilities associated with the workload. The port and protocol is mapped to a vulnerability (if it exists). Under the Vulnerabilities tab, all the vulnerabilities for the workload are sorted in order of severity. You can see the following information for each vulnerability:

- Total V-E score of the workload
- V-E score of the highest accessible network port of the workload
- Vulnerability score of the most severe network accessible vulnerability on the workload
- East-West exposure
- Internet exposure
- Type of traffic on that port
- Name of the vulnerability

Under the Processes tab, you can see V-E score of each process that is communicating over the network port. The East-West Exposure Score is recalculated whenever the rules associated with the workload are changed.

Reports and Statistics

This feature provides two types of recurring reports:

- Executive Summary reports
- App Group Summary reports
- Dashboard for VEN statistics

Ransomware Protection Dashboard for Servers

In Illumio Core 23.2.0, a new dashboard gives you broad, visualized information about ransomware protection readiness, risk exposure, and protection coverage statistics.

Working with the Ransomware Protection for Servers Dashboard

You can access the Dashboard by clicking Dashboard in the left menu.



In this release, only the following global user roles are allowed to use the Ransomware Protection Dashboard:

- Global Org Owner
- Global Administrator
- Global Viewer

Only managed server workloads are included in the Dashboard statistics. Endpoints and container workloads are not included.

In the upper-right corner, you can see the *Refresh* button, which brings new fresh data from the system.



Four of the dashboard widgets are auto-refreshed at a regular interval and are not refreshed by the Refresh button. The last refresh time is indicated by the tool tip over the Clock icon.

Left Section of the Dashboard

The two widgets show the following stats:

Protection Ready Workloads

A workload is protection-ready when there is a VEN installed on the workload and can be configured to enforce Illumio security policies.

Users can optionally enter the target number of workloads requiring protection, which can be edited at any time. This widget indicates the number of such workloads compared to all available target workloads.



In the example above, 51 workloads are protection-ready.

Protection Ready Workloads (daily, weekly, monthly, quarterly)

This widget shows the number of Protection Ready workloads for a selected period of time.

In each of the selected views, the number of Protection-Ready Workloads is represented as a percentage of the available target workloads (100%).

The resolution might be Day, Week, Month, and Quarter.

Middle Section of the Dashboard

The middle sections shows the following stats:

Protected Workloads

A workload is protected when it has policies on all the ransomware risky services / ports and the policies are enforced: the workload has to be in Selective Enforcement or Full Enforcement mode.

Workloads by Ransomware Exposure

A workload is assessed with its exposure to the common services exploited by ransomware. The risk of each service is classified into four severities: Critical, High, Medium, and Low.

For more details, see "Services" in the Security Policy guide.

A workload is protected for the service in these two cases:

- The service is blocked by enforcement boundary in Selective Enforcement or
- The workload is in Full Enforcement, whether there is rule or no rule for that service.

This widget shows the number of workloads by their ransomware exposure (Critical, High, Medium, Low, and Protected) across the organization.

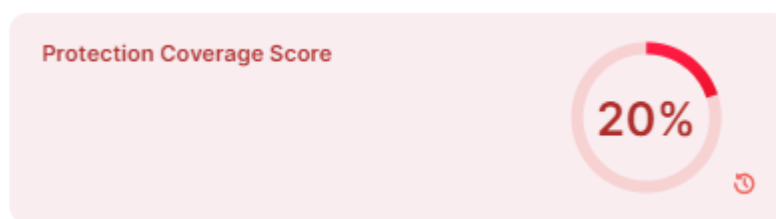
Right Section of the Dashboard

The right section shows the following stats:

Protection Coverage Score

The Protection Coverage Score is a metric used to measure the effectiveness of security policies in protecting workloads. It indicates the percentage of the entire possible attack surfaces that are actively protected by security policies. For example, a policy that allows all workloads as source will have a lower coverage score compared to a policy that only allows a small number of source workloads.

Protection coverage score takes all the protection ready workloads into consideration across the organization.



The color of the widget changes from red to yellow and then to green as the protection coverage score increases.

Risky Ports by Severity

This widget shows how many ransomware-risky ports, categorized by their severity (Critical, High, Medium, and Low) are in the system. Each category of ransomware-risky ports has a different total on each workload and hence across the system.

API Support for the Ransomware Protection for Servers Dashboard

The Dashboard uses several APIs to aggregate various data from the system and helps you focus on the data you are interested in.

The Dashboard is powered by the two main APIs: `time_series` and `risk_summary`.

For the complete list of REST APIs that are used to power the Ransomware Protection Dashboard, see "Ransomware Protection Dashboard APIs" in the *REST API Developer Guide*.

View Workload Ransomware Protection for Servers Details

The Ransomware Protection tab provides detailed protection information for the workloads regarding each of the ransomware-risky services.

Information about the ransomware risk is then aggregated into the Ransomware Protection Dashboard for the system-side ransomware risk analysis.

Summary Processes Rules Deny Rules Blocked Traffic <u>Ransomware Protection</u>								
🔔 Enforcement: Full Workload enforcement is set to Full. Maximum protection is in place.								
● Customize columns 50 per page 1 - 30 of 30 Total								
Service	Port/Protocol	OS	Severity	Port Status	Protection	Active Policy	Draft Policy	
S-RDP	3389 TCP	Windows	Critical	Listening	Protected (Blocked)	Blocked	Blocked	
S-SMB	445 TCP	Linux, Windows	Critical	Listening	Protected (Blocked)	Blocked	Blocked	
S-WINRM	5985 TCP	Windows	Critical	Listening	Protected (Blocked)	Blocked	Blocked	
S-SSDP	1900 UDP	Windows	Medium	Inactive	Protected (Blocked)	Blocked	Blocked	
S-SMB	445 UDP	Linux, Windows	Critical	Inactive	Protected (Blocked)	Blocked	Blocked	

The Severity and Port Type are designated per each ransomware-risky service.

For more details, see "Services" in the Security Policy guide.

Here is the explanation for the data provided in the Ransomware Protection table:

- **Severity:** Severity of the ransomware risk, which can be Critical, High, Medium or Low.
- **Port Status:** Port status can be Active or Inactive.
 - **Listening:** Listening means there is a running process on that port.
 - **Inactive:** Inactive means there is no process running on the port. The same information is also provided on the Processes tab.
- **Port Type:** The port type can be Admin or Legacy.
 - **Admin:** Admin refers to the service and ports are used for common administrative tasks.
 - **Legacy:** Legacy means that ports are used for legacy protocols.
- **Protection:** Protection types are:
 - **Protected (Blocked):** When port is blocked by deny rules in Selective Enforcement or blocked with no allow rules in Full Enforcement. No ransomware can propagate through that port.
 - **Unprotected:** The port is exposed to ransomware exploits.
 - **Protected (Allowed by Policy):** When there are allow rules intentionally policing the traffic. Only the trusted sources are allowed to access the port and hence the risk of

lateral movement for ransomware is reduced. The workload has to be either in Selective Enforcement or Full Enforcement for the policy to be enforced.

- The Port status does not affect the protection state.
- **Active Policy** and **Draft Policy**: Indicates whether there is an Active or Draft policy to protect that particular port and the corresponding action.

Dashboard Access for Supercluster Users

Starting from the release 23.2.10, access to the Ransomware Protection Dashboard was extended to supercluster users in the following way:


- The Supercluster Leader is able to see the Ransomware Dashboard icon in the side menu and to visit dashboard page after clicking on it;
- Supercluster members are **not** able to see the Ransomware Dashboard icon in the side menu and have no access to the dashboard page.
- Both the Supercluster Leader and members can view the workload information in the workload page and the service page.

VEN Dashboard

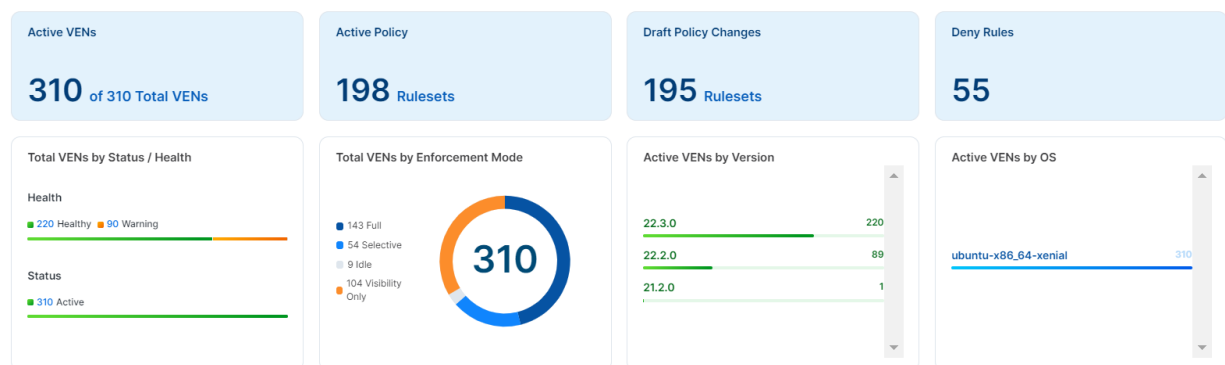
Illumio provides a set of dashboard widgets to give you broad, visualized information about VEN statistics.

Working with the VEN Dashboard

You can access the Dashboard by clicking on the Dashboard button in the left menu.



The VEN Dashboard is located above the new Ransomware Protection Dashboard.



The Dashboard uses an API to aggregate various data from the system and helps you focus on the data you are interested in.

For more information about the API support, see "VEN Dashboard APIs" in the *REST API Developer Guide*.

In this release, only two user roles are allowed to use the VEN Dashboard:

- Global Org Owners
- Global Administrators

The VEN Statistics section of the Dashboard contains several widgets to display summary statistics or status. To get new fresh data, click the Refresh button at the top of the page. To see more details, click the widget, and the list page is displayed with the appropriate filter to see the resources.

The upper four sections show stats about:

- Active VENs (how many VENs are active out of the total number of VENs)
- Active Policy (number of rulesets)
- Draft Policy Changes
- Deny Rules

In the lower sections, the VEN Statistics part of the Dashboard includes the following widgets:

This documentation assumes you are using the latest Illumio UI. If your screen does not match the descriptions below, click the New UI toggle.

Total VENs by Status/Health

- VEN counts by status (stopped, suspended, uninstalling, and active statuses)
- VEN counts by health (error, warning, and healthy)

Total VENs by Enforcement Mode

- VEN stats by enforcement mode: full, visibility only, idle, and selective enforcement.

Active VENs by Version

- VEN stats by versions (number of VENs per version)

Active VENs by OS

- VEN stats by currently employed versions (number of VENs per OS)

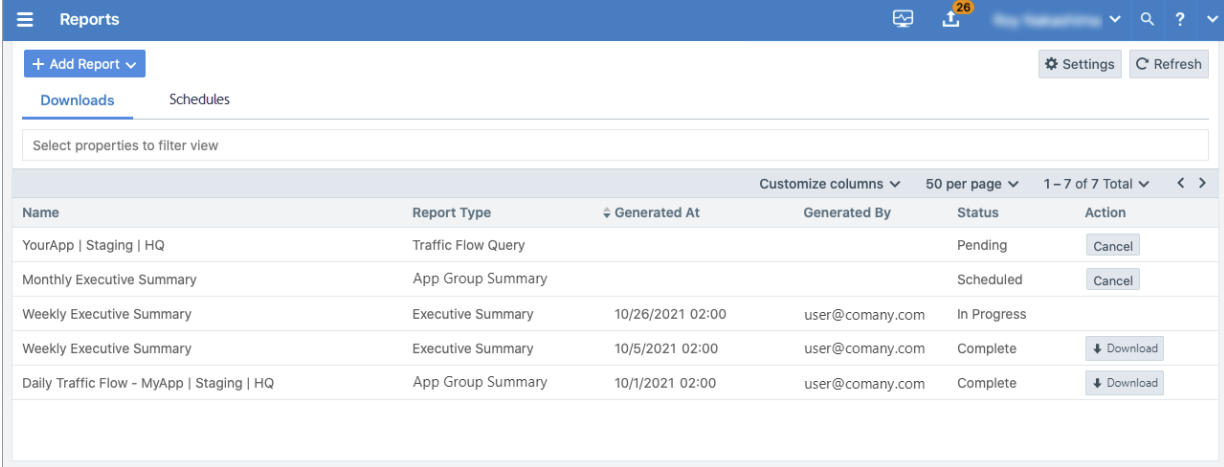
About Reports

The PCE includes the ability to generate, download, and manage two types of recurring reports: Executive Summary reports and App Group Summary reports.

Reporting in the PCE

The PCE web console menu includes a *Reports* option. When you choose the Reports option, the Reports page appears. This page includes two tabs: **Downloads** and **Schedules**. Gener-

ated reports appear on the **Downloads** tab. By default, the list is sorted in descending order by the **Generated At** time.



The screenshot shows the 'Reports' section of the Illumio interface. The 'Downloads' tab is active, displaying a table of reports. The table has columns for Name, Report Type, Generated At, Generated By, Status, and Action. There are 7 reports listed, with the first two being 'YourApp | Staging | HQ' and 'Monthly Executive Summary'. The last two reports, 'Weekly Executive Summary' and 'Daily Traffic Flow - MyApp | Staging | HQ', have 'Download' buttons next to them.

Name	Report Type	Generated At	Generated By	Status	Action
YourApp Staging HQ	Traffic Flow Query			Pending	Cancel
Monthly Executive Summary	App Group Summary			Scheduled	Cancel
Weekly Executive Summary	Executive Summary	10/26/2021 02:00	user@comany.com	In Progress	
Weekly Executive Summary	Executive Summary	10/5/2021 02:00	user@comany.com	Complete	Download
Daily Traffic Flow - MyApp Staging HQ	App Group Summary	10/1/2021 02:00	user@comany.com	Complete	Download

Because Illumio provides the reports as downloadable PDF and CSV files with an email option, you can share them with people in your organization who don't have access to the PCE web console or PCE REST API.

The data in the reports is not customizable. However, you can configure the time range of the data that the reports are generated from and the frequency at which they are run. Both types of reports include when a specific report was generated, which Illumio user generated it, and the PCE version from which the data was obtained.

Recurring reports are run on the following schedule:

- **Daily:** Midnight each day
- **Weekly:** At midnight on the first Saturday after the report was added, then weekly at Saturday midnight
- **Monthly:** Midnight on the last day of month after the report was added, then monthly on the last day at midnight

The PCE does not cap the number of reports you can create, the only the length of time you can retain them. Generated reports include data for provisioned security policy, managed and unmanaged workloads, and provisioned policy objects. They do not include changes you have made to your environment but haven't provisioned.

Executive Summary Reports

Executive Summary reports are high-level by design. They provide information to decision makers, such as an organization's CIO or VP of IT, about the overall deployment of Illumio within the organization's computing environment. These reports are intended to provide more business-oriented information than tactical data.

Executive Summary reports give the decision makers a snapshot into how Illumio policy enforcement is progressing and can display the return on investment (ROI) for purchasing and deploying Illumio software.

Executive Summary reports answer the following questions for decision makers:

- How are we progressing in deploying security policy into our environment?
- How many of our workloads are being managed by Illumio (VENs are installed on the hosts but they aren't in enforcement mode)?
- How quickly is enforcement progressing over time (the number of workloads that have moved into the enforcement mode over the report's specified time range)?
- What potentially dangerous traffic is Illumio blocking that wouldn't have been blocked without Illumio Core, resulting in a security risk.
- What sort of vulnerabilities do our workloads have? Vulnerability information is provided as a V-E score that is the sum of all app groups.



IMPORTANT

To include app group and workload vulnerability data in the Executive Summary report, you must have purchased a license for the Vulnerability Map feature. The Vulnerability Map is a separately licensed feature of Illumio Core. The licensing is based on the number of workloads. The license is required to import Qualys report data into the Illumio PCE. For information about obtaining the Illumio Core Vulnerability Map license, contact Illumio Customer Support.


For more information about Vulnerability Maps, see [Vulnerability Map. \[52\]](#)

Tips for Reading Executive Summary Reports

Executive Summary reports provide high-level information for decision makers. They are meant to show trends and patterns in your roll out of Illumio Core into your data center environment.


For example, an executive who has approved deploying Illumio Core might want to know how many of their workloads are being managed (enforced) by Illumio policy. The Workloads by Enforcement Mode graph shows the trend for how quickly enforcement is progressing over time and the percentage of workloads in deployment versus enforcement.

Workloads by Enforcement Mode



The Provider Workloads by Policy Decision graph can help confirm when the rules you have created for your data center look viable and you can start enforcing policy on your workloads. This example graph shows a trend you want to see; and visually represents how you initially had workloads deployed but not in enforcement.

Provider Workloads by Policy Decision



App Group Summary Report

Illumio Core contains many features designed for application owners; such as the App Group Map and role-based access (RBAC) for applications owners. See “App Group Map” in the Visualization Guide and “Role-based Access for Application Owners” in the PCE Administration Guide, respectively, for information.

App Group Summary reports are designed for application owners (for example, members of your business applications group like your Oracle or ServiceNow app admins) or other people in your organization who need to understand the security of you applications, such as IT security auditors (for example, auditors of PCI or HIPA systems).

You create App Group Summary reports by application; meaning, each report provides data for only one application defined by a set of labels. Whether you choose 2 labels (application and environment) or 3 labels (application, environment, and location) for a report depends on how you have configured the PCE to define app groups. See [Configure App Groups \[51\]](#) for information.

Using the App Group Summary report, application owners or IT security auditors can accomplish the following goals:

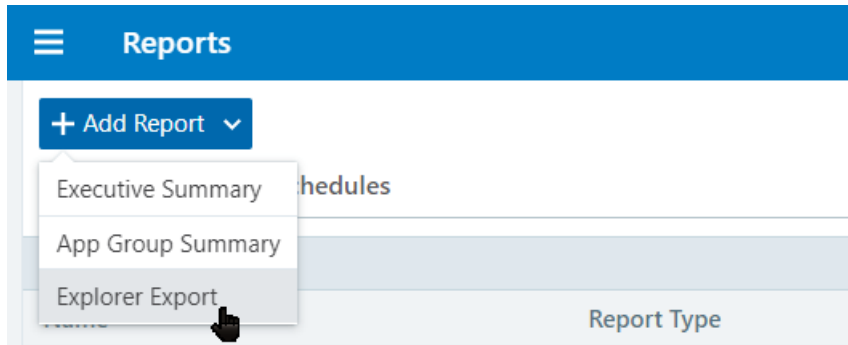
- Examine which inbound and outbound services interact with a specific application. Having a clear picture of all traffic into and out of an application is important for accessing the security posture of the application.
- Understand whether connections are normal for an application and monitor the application’s health and status over time. For example, you can create a weekly report to monitor the state of an application over time and detect any changes in inbound or outbound network services.
- Fulfill compliance auditing requirements. For example, you can run a report every 30 days and review the report to ensure the application connection status matches with the application’s baseline.
- Establish a connection baseline for an application and use that baseline to create security policy (rules or selective enforcement rules) for the application. See “Rules” and “Rule Writing” in the Security Policy Guide for information.
- After creating security policy (rules) for an application in the PCE, see the impact of the Illumio security policy on the application.

Explorer Report

You can run a previously saved Explorer filter and export the results to a CSV file on a recurring schedule.

**NOTE**

If you edit the filter, subsequent recurrences of the Explorer Export will continue to use the original version of the filter.



Work with Reports in the PCE

This topic describes how to manage your reports in the PCE web console.

REST API to Generate Reports

In Illumio Core 21.2.0, Illumio previewed the Reporting feature by providing the ability to generate an Executive Summary report for your managed environment. In addition to the PCE web console, you can use the Illumio REST API to generate and manage reports. In 21.2.0 and any on-prem PCE before Illumio Core Release before 22.2.0, you can generate and manage reports through the Illumio REST API by editing the `runtime_env.yml` file.

1. `# sudo vi /etc/illumio-pce/runtime_env.yml`
2. Add: `reporting_enabled: true`.
3. Restart the PCE.

For information about using the Illumio REST API to manage reports, see the REST API Developer Guide.

Add a Report

1. From the PCE web console menu, choose **Reports**. The Reports page appears.
2. Click **Add Report** and select the report type from the drop-down menu.
A dialog box appears so that you can configure the report.
3. Configure the following report settings and click **Save**:
 - **Name:** Specify a name that describes the purpose of the report. Report names must be from 2-255 characters and contain special characters.
 - **Recurrence:** From the drop-down list, select how frequently the PCE will run the report.
 - **Time Range:** From the drop-down list, select the time range for the report (the time range differs by report type).
 - **App Group:** (App Group Summary report only) Select the application that you want to generate the report for.

The new report appears on the **Recurrence** tab.

Manage Reports

Perform the following tasks to manage how you generate reports for your organization and computing environment.

To download a report:

1. From the PCE web console menu, choose **Reports**. The Reports page appears.
2. Click the **Downloads** tab.
3. In the row of a completed report, click the **Download** button.

To set the retention period for all reports:

You can configure globally how long the PCE retains the PDF files generated for each report you add. You can only retain PDF files up to 7 days in the PCE. By default, reports are configured to be retained 7 days.

1. From the PCE web console menu, choose **Reports**. The Reports page appears.
2. Click **Settings** in the top right corner of the page.
3. In the Retention field, specify the number of days to retain PDF files.
4. Click **Save**.

To edit the settings for a report:

1. From the PCE web console menu, choose Reports. The Reports page appears.
2. Click the Recurrence tab.
3. Click the row for the report you want to modify.
4. Change the recurrence rate, time range, or report name.
5. Click Save.

To end the recurrence of a report:

Removing a report from the Recurrent tab stops the report from running again. Existing PDF files generated for the report remain in the PCE until the global retention period expires and they are deleted by the PCE.

1. From the PCE web console menu, choose **Reports**. The Reports page appears.
2. Click the **Recurrence** tab.
3. Click the row for the report you want to stop being regenerated.
A dialog box appears prompting you to confirm that the report won't be generated again.
4. Click **Remove**.

Manage Reports by Using the Illumio REST API

Beginning in Illumio Core 21.2.0, Illumio previewed the Reporting feature by providing the ability to generate an Executive Summary report for your managed environment. In addition to the PCE web console, you can use the Illumio REST API to generate and manage reports. In 21.2.0 and any on-premises PCE before Illumio Core release 22.2.0, you can generate and manage reports through the Illumio REST API by editing the `runtime_env.yml` file.

1. `# sudo vi /etc/illumio-pce/runtime_env.yml`
2. Add: `reporting_enabled: true`.
3. Restart the PCE. See the PCE Administration Guide for further information.