


Core 24.4 Visualization

December 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	7
References	8
Skipped Critical Patch Updates	8
Discovered By	8
Frequently Asked Questions	8
Visualization	9
Visualization Tools	9
About the Visualization Tools	9
Map View	23
Traffic Table	33
Mesh View	39
App Groups	46
Work with the App Group Map	53
Reports	54
Work with Reports in the PCE	62
Work with the Visualization Tools	64
Dashboards	72
Servers and Endpoints Dashboard	72
Ransomware Protection for Servers Dashboard	74
Vulnerability Map	87
About the Vulnerability Map	87
Work with Vulnerability Maps	89

Legal Notice

Copyright © 2025 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

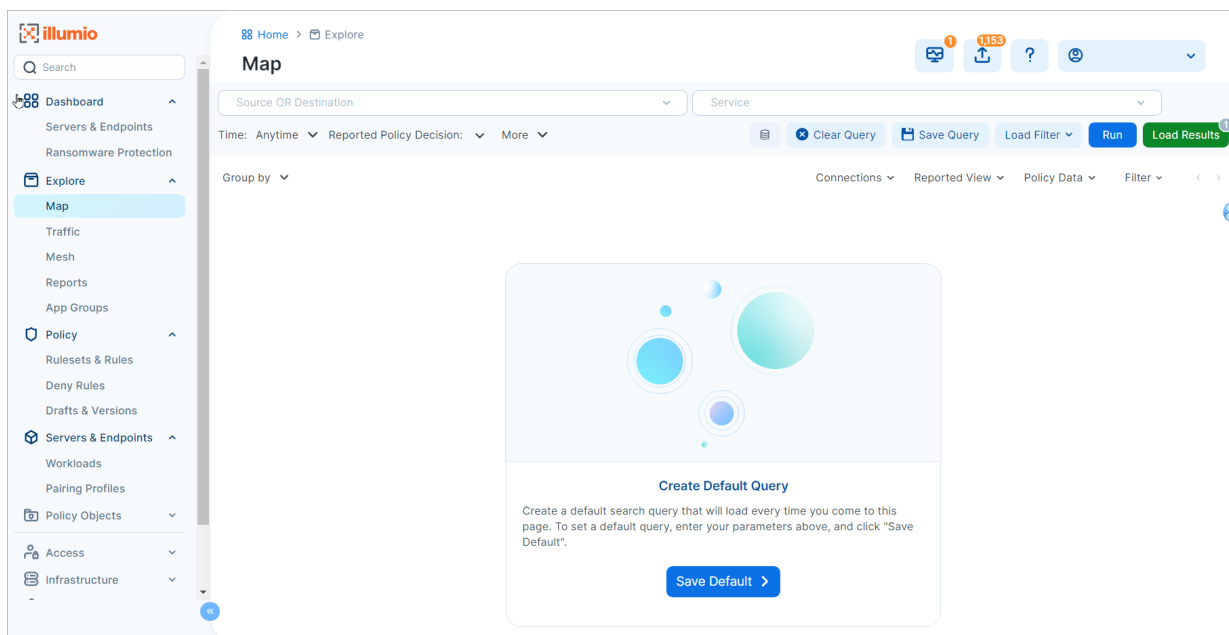
Visualization Tools

This section 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.

About the Visualization Tools

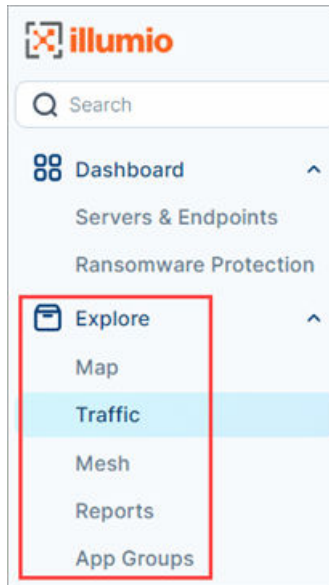
In the PCE UI, you can use the visualization tools to reveal the traffic flows in your network and to help you configure policies to secure your applications. These tools include the Map, Traffic table, Mesh, Reports, and App Groups.

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 from which you run your first query.



Types of Visualization Features

Select visualization tools from the Explore category in the left navigation.



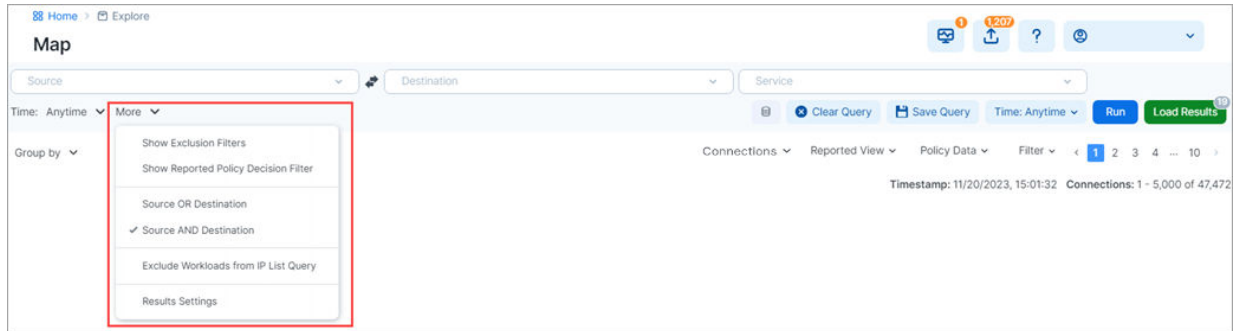
- **Map**
The Map depicts workloads that form logical groups (based on labels attached to workloads) and provides an understanding of the traffic flows between workloads. You select groups in the Map view to view details about that group and develop policy for the workloads in the group.
- **Traffic**
The Traffic table 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, the Mesh displays traffic flows as lists of destinations, sources, and the port. 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.
- **Reports**
The Reports feature allows you to generate four types of recurring reports:
 - Executive Summary reports
 - App Group Summary reports
 - Traffic Export
 - Rule Hit Count Report
 You can download reports in PDF and CSV formats and share them with people in your organization who don't have access to the PCE UI or PCE REST API.
- **App Groups**
App Groups allow application owners to see all workloads for an application instance, even when the workloads are not currently communicating with each other. This is helpful when building or validating security policies for traffic between workloads because it allows application owners to focus only on the workloads that belong to their applications, regardless of location.

Filters for the Visualization Tools

For each visualization tool (except Reports), traffic filters are available so you can show or hide different elements of your data and focus on what is most important to you.

More menu

To modify the filters, open the More menu.



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

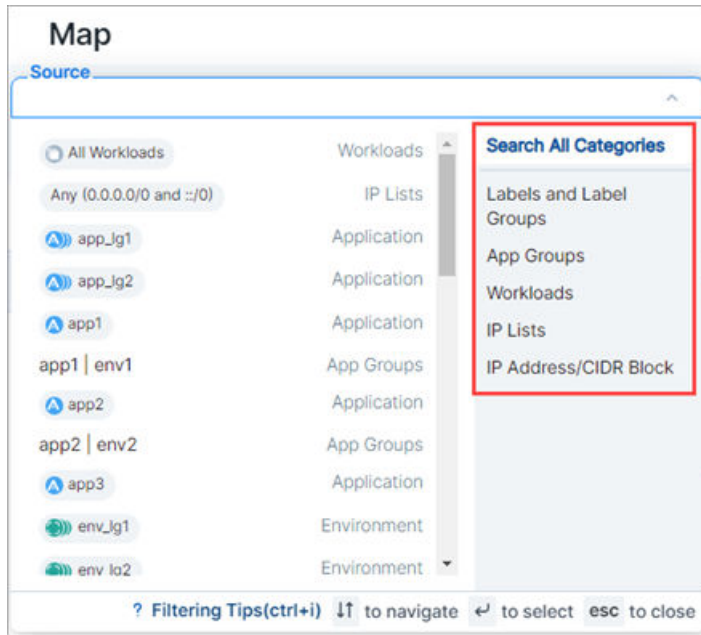
To search for traffic flows with a specific policy decision reported by the VENs, select the Show Reported Policy Decision Filter option. This option controls the type of policy decision (allowed, potentially blocked, blocked, or unknown) that the Traffic and Map views display.

Source and Destination filters

Depending on the visualization tool, the Source and Destination filters include some or all of the following query options:

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

Selecting the Search All Categories option avoids the need to first enter a category in the filters.



The Label and Label Groups category restricts the Map to show 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.

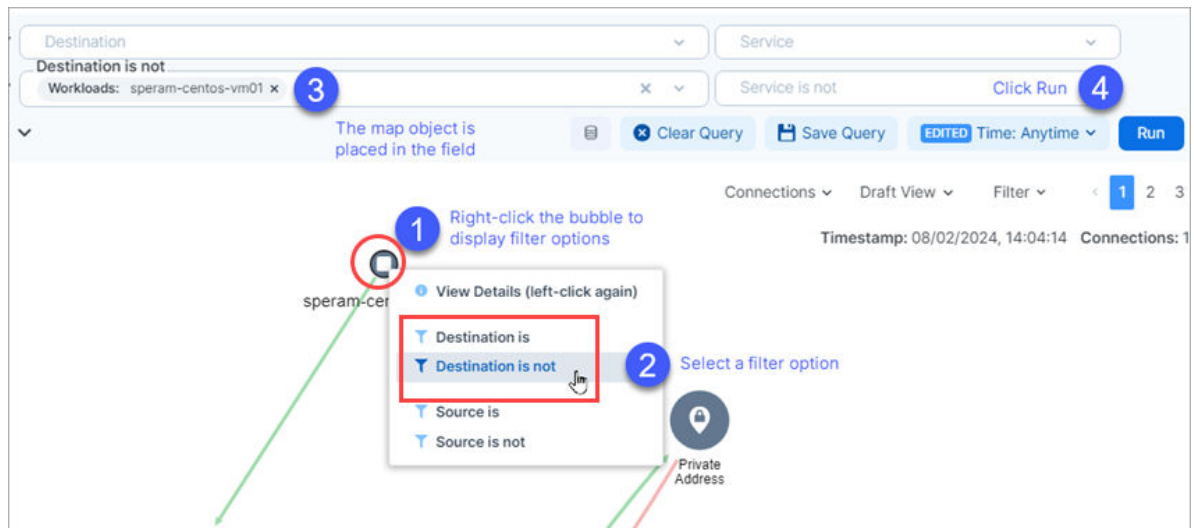
When you enter text in this filter, the PCE UI allows you to specify whether that text is a process name or a service. Once you make your selection, the UI reflects the option you chose.

is/is not filter options

You can easily modify filter results using the "is/is not" options available in Traffic and Map views. These are useful when you're looking at search results and you want to easily modify the query by adding or removing filtering criteria.

Map view

1. Right-click a bubble on the map for the list of options to appear.
2. Select whether to include or exclude the data from the source or destination search fields.
3. Notice your selection populating the search field.
4. Click Run to see modified query results.



Traffic view



TIP

This feature also applies to pills on the Traffic tab available in the Map's details panel.

1. Mouse over any data pill for the list of options to appear.
2. Select whether to include or exclude the data from the source or destination search fields.
3. Notice your selection populating the search field.
4. Click Run to see modified query results.

The screenshot shows the PCE UI search interface. At the top, there are filters for Destination, Service, and Workloads. A dropdown menu is open, showing options like 'Destination is not' and 'Source is not'. The main table displays traffic flow data with columns for Source, Destination, Destination Labels, Destination Port, and Flows/Bytes. A red box highlights the 'Destination is not' filter option, and a blue circle highlights the 'Run' button.

Example Search Using Filters

Before you write policy rules to either allow or block traffic, you should 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.

For example, using the visualization tools, you can run the following query:


Find 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."

1. In the Explore category in the left navigation, go to Map, Traffic, Mesh, or App Groups.
2. Click More and select Show Exclusion Filters. Exclusion filters allow you to exclude criteria from a search.
3. From the Destination drop-down list, enter or select the Environment label named "Development."
4. From the Destination is not drop-down list, enter or select the Role label named "Domain Controller."
5. From the Source drop-down list, enter or select the Environment label named "Production."
6. From the Source is not drop-down list, enter or select the Role label named "Domain Controller."
7. Under Service, leave the Service field blank (which means "any").
8. Under Service is not enter "80."
9. Under Time, select Anytime.
10. 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 left navigation in the Explore category, click any visualization tool (except Reports).
2. Choose More > Results Settings.
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*.

4. Click OK.

Load Results in Visualization Tools

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.

Results ×				
Name	Connections	Run At	Status	Actions
Time: Anytime	122,095* (5,000 displayed)	03/21/2024, 17:21	Completed	Export Delete
× Close				

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.

V-E Scores Comparison Tool

The Show Vulnerability Exposure (V-E) Score tool lets you see how the security of your app groups would change if you were to change their current enforcement mode. Columns in the App Group list and details pages provide a side-by-side comparison of the effect different enforcement modes would have on Vulnerability and Exposure (V-E) scores. A toggle allows you to simulate the switch between Full Enforcement and Visibility Only enforcement modes.

For details, see [Compare App Group V-E Scores by Enforcement Type \[51\]](#).

[Home](#) > [Explore](#)

App Groups

[Edit App Group Definition](#)
[Segment App Group](#)

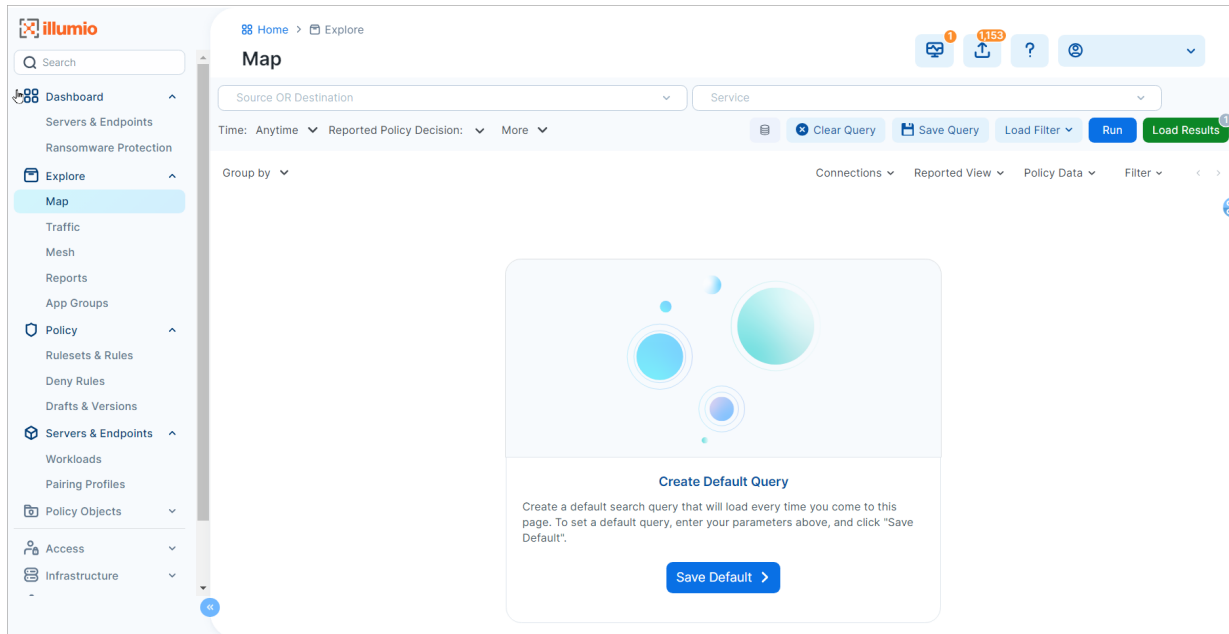
Show Vulnerability Exposure Score (V-E) Score in:
 Full Enforcement
Visibility Only
?

Visibility Only V-E Score	Current V-E Score	Name
34	6.1	app1 env1

About the Default Graph

In Core 22.5.x and earlier, the PCE cached the Illumination Plus queries (for the Map and Traffic table views) that you ran and saved them for 24 hours. 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.

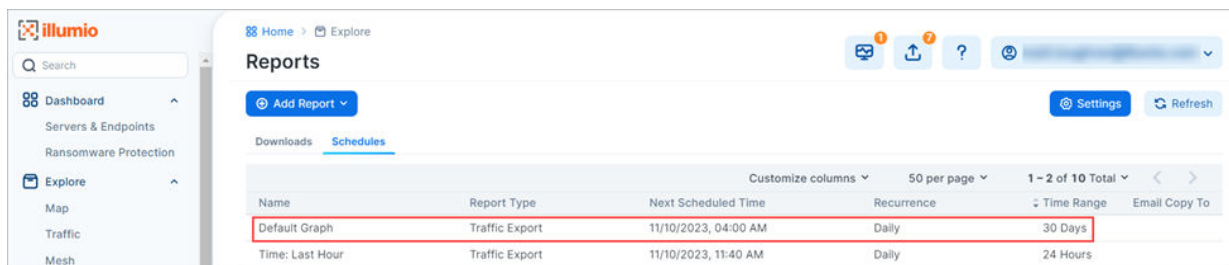
Beginning in 23.2.0, if you don't have a default graph in the PCE, the page below is your start page for the Map and Traffic pages.



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).

Later, 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.

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 to 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. You can export the query search results to either a comma-separated-value (.CSV) file or display them 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 left navigation, go to Map or Traffic from the Explore category.
2. Enter your query criteria in the fields. If you want to exclude criteria, select 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 Table 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 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.

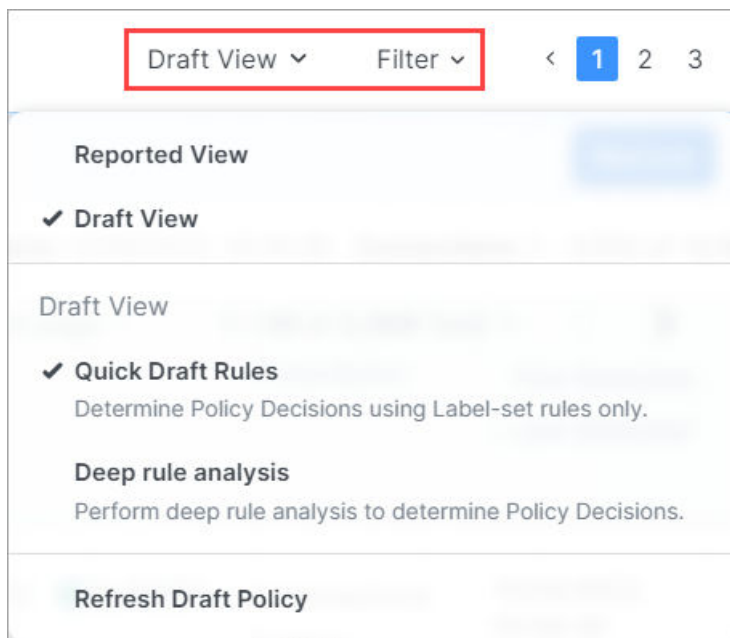
View Menu



IMPORTANT

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

When used with the options in the adjacent Filter drop-down, the View menu allows you to configure how the PCE UI displays your traffic data so you can see the connections between your groups with greater flexibility. The options on this menu are unaffected by how you've grouped traffic in your Map, Traffic, or Mesh pages.



From the View menu, select the following options:

- **Reported View**
For a description, see [Reported View \[21\]](#).
- **Draft View**
For a description, see [Draft View Options \[22\]](#).
- **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 UI.

Reported View

The Reported view presents your policy coverage as reported by your workloads so that you can examine the current state of your provisioned policy. This view provides visibility of 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.

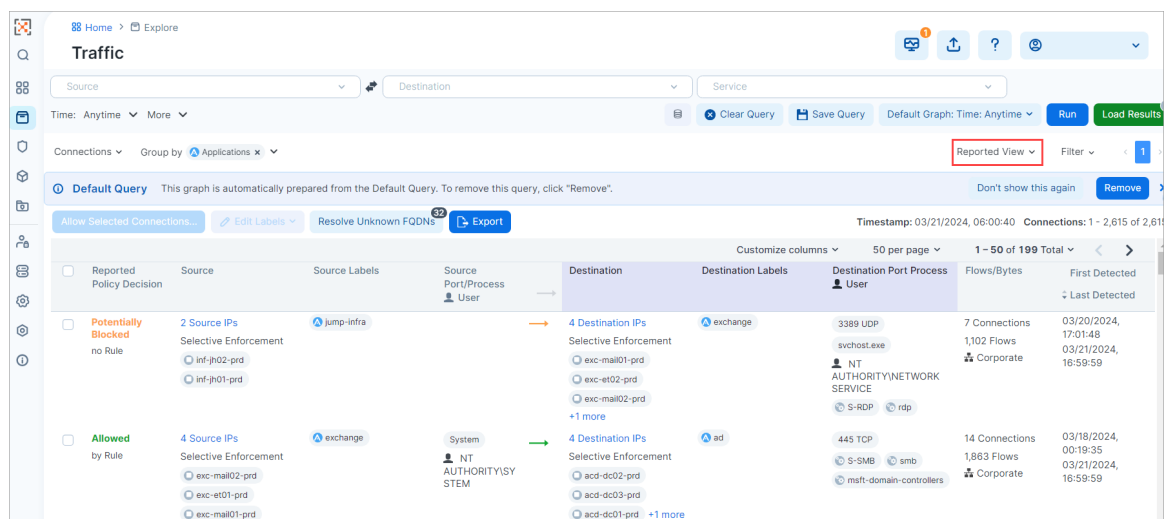
The Reported view is a read-only view. You can view all the rulesets that apply to the workloads from the Reported view but you must change to the Draft view to add rules. The Reported view does not immediately reflect the latest changes to the policy. It is updated only after you provision a change 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.

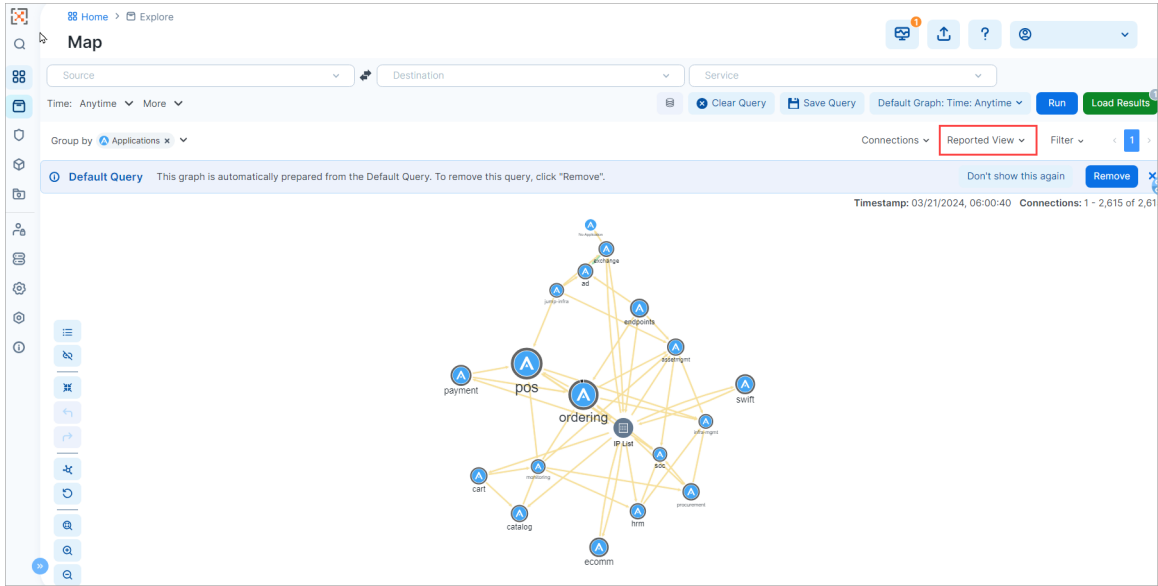
For each flow with a unique port/protocol, if there is a policy service created for that port/protocol, the name of that policy service displays in addition to the names of the actual services that reported the flows. The Reported view shows reported rule coverage for the latest reported flow with that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



Reported View (Map)



Draft View Options

The Draft view immediately presents 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. Because the PCE has to compute the expected coverage for each traffic flow, the Draft view can take longer to load than the Reported view, especially when you have a large number of workloads and traffic flows.

Draft View allows you to view either the rule that would permit traffic or to 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. Modify the view further using the options available in the adjacent Filters drop-down menu.

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.

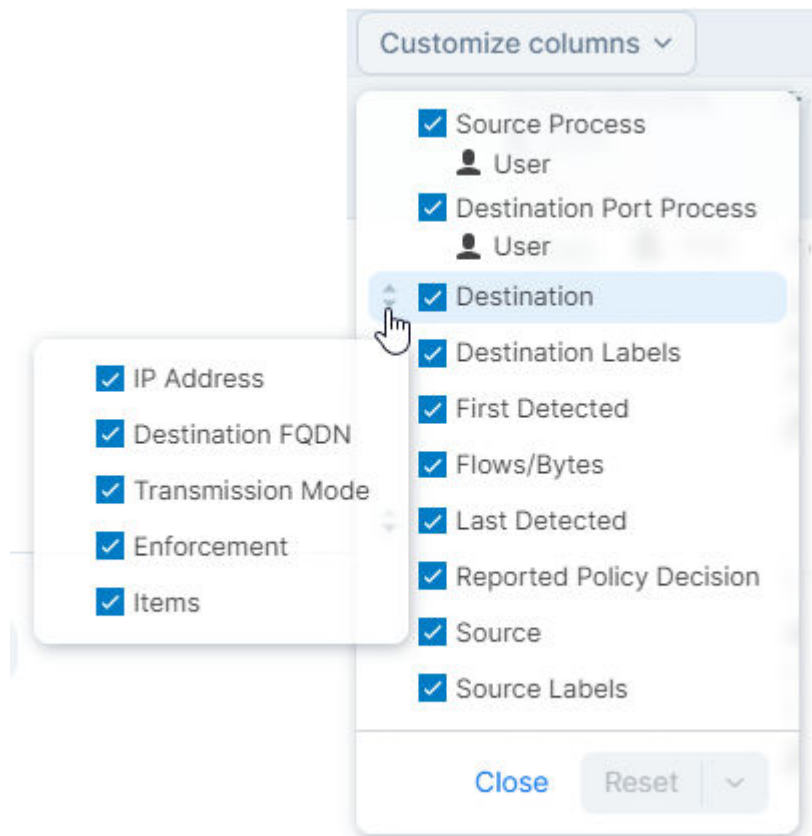
Customize Columns

You can customize columns in the following areas using the Customize columns menu:

- Explore > Traffic

- Servers & Endpoints > Workloads

You can further customize most columns by selecting the data that you want to appear within them. Mouseover the up and down arrows to the left of a column checkbox and select or deselect data within that column:



Customizing columns in this way doesn't 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, the Map 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.

Map View

Use the Map to visualize workloads that form logical groups (based on labels attached to workloads) and to better understand the traffic flows between workloads.

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. For information about installing (also called pairing) VENs on workloads, see the VEN Installation and Upgrade Guide.

The Map displays three different types of groups: a group based on a single label, an app group, or a common 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 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, app groups, etc. It also lets you split the view when in Map view mode by selecting items on the Map.

Configurable Grouping

The **Group by** menu allows you to specify different levels of grouping, such as grouping by types of labels and their order. You might want to group by OS and then by environment. If you do not specify a particular grouping, Illumio groups workflows that have the same set of labels. You can change your default grouping through the **Group by** 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

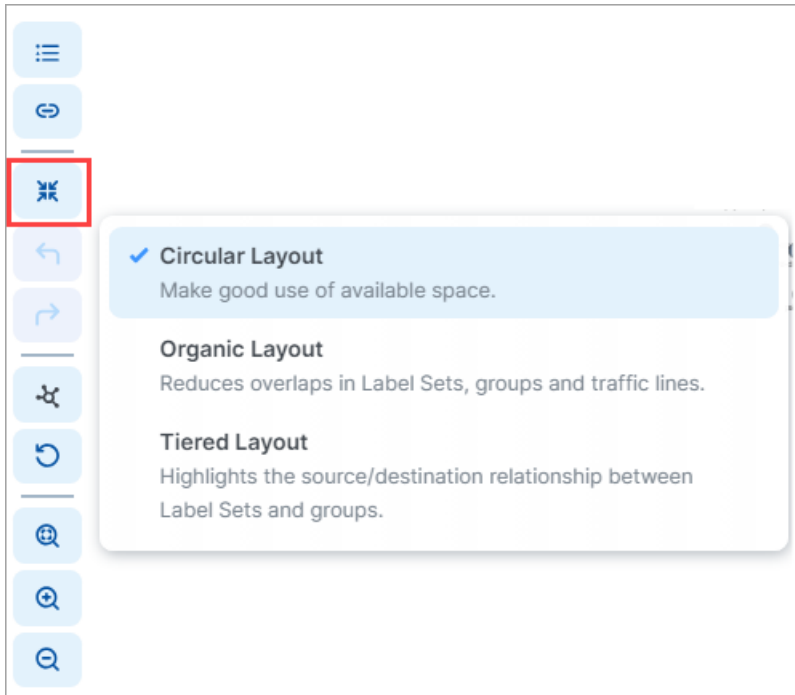
- Each group is a label set. Every workload which has the same set of labels is grouped into one of those label-sets.
- Mousing over a group in the Map displays a pop-up dialog box with the list of labels and the number of workloads using the labels.



- In the **Group by** drop-down list, you can drag and drop labels in the list to re-order how the Map displays groups. Labels at the top of the list control the prominence of those groups in the Map.
- The PCE UI 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 the PCE UI displays the Map:




Not every layout choice is good for your Map data. See the descriptions of each layout in the Layout 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 from top to bottom. The Tiered Layout option is better for smaller data sets than larger ones.


Show All Endpoints

In earlier Core releases, running a query in the Map revealed only endpoints that have traffic flows. Beginning in Core release 23.5, you can redraw the map to reveal all endpoints, including those with no traffic. Click the "link" button in the bottom left corner of the map.




How to Read the Map Symbols

There are two legends for the side panel, one for Policy Data mode and another for Vulnerability Data mode. You can use the drop-down selector above the panel to switch between these modes.

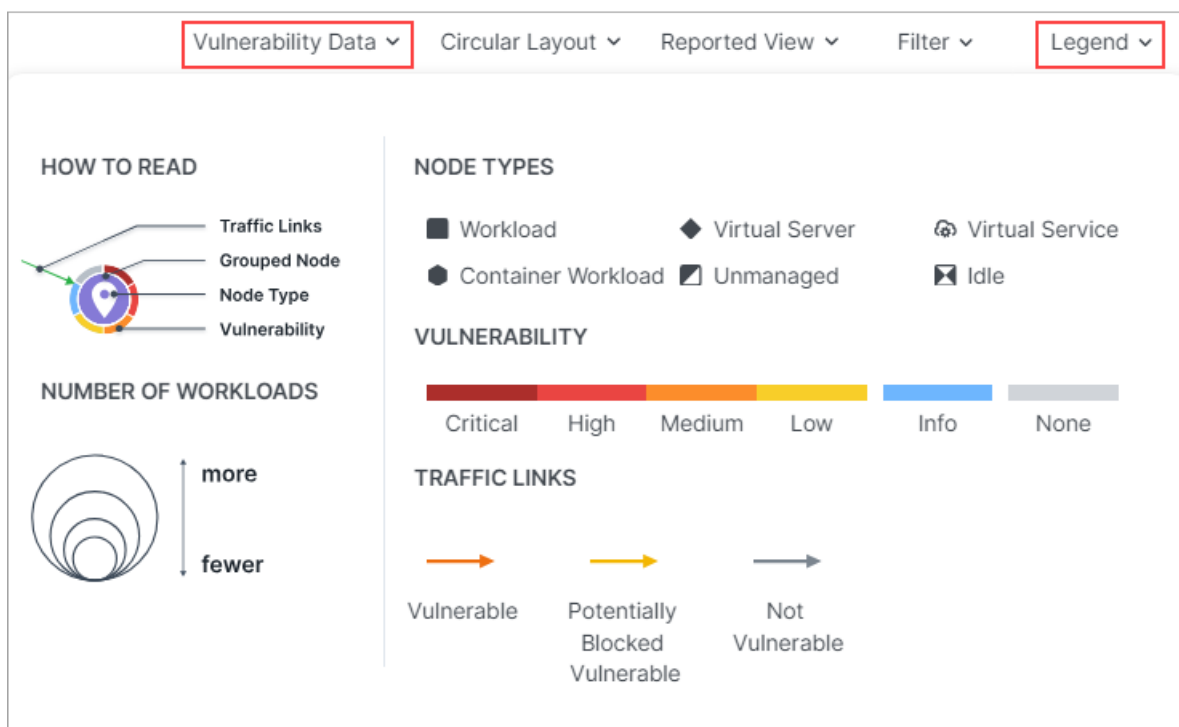


Legend - Policy Data



Legend - Vulnerability Data

For more about the Vulnerability Tab, see [Vulnerabilities Tab \[32\]](#).



Map Symbols Explained

Number of Workloads (Policy Data and Vulnerability Data modes)

The relative size of each node indicates the number of workloads in the node.

Enforcement (Policy Data mode)

Pay attention to how the Map groups designate the enforcement mode for groups:

- Workloads and groups inside fully dark lines are in FullEnforcement mode.
- Workloads and groups inside semi-dark lines are in SelectiveEnforcement mode.
- Workloads and groups inside light gray lines are in Visibility only mode.
- Workloads and groups not surrounded by any of the above-described lines are in Idle mode.
- The completeness of 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 (Policy Data mode)

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

- **Red:** Traffic is blocked
- **Yellow:** Traffic is potentially blocked
- **Green:** Traffic is allowed
- **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.
- **Grey:** Rules are not calculated

Traffic Links (Vulnerability Data mode)

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

- **Red:** Traffic is vulnerable
- **Yellow:** Traffic is potentially blocked and vulnerable
- **Grey:** Traffic isn't vulnerable

Vulnerability

When in Vulnerability Data mode, the color of each node's outer ring indicates the criticality of it's current vulnerability level.

- **Dark red:** Critical
- **Red:** High
- **Orange:** Medium
- **Yellow:** Low
- **Blue:** Info
- **Light gray:** None

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

Map Reported View

The PCE UI displays the traffic on the Map 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

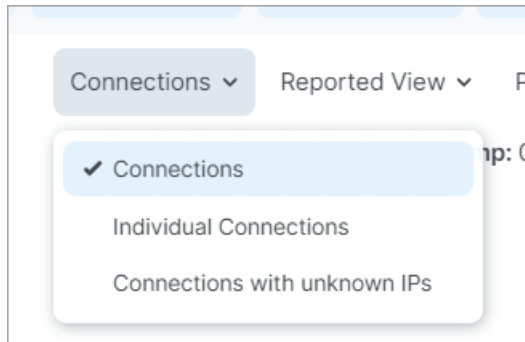
This view also displays the traffic using red, green, and orange lines to indicate whether the PCE has a rule to allow the connection that was reported by the VEN. This way, you can add rules and see their anticipated effect in real-time before the rules are implemented. In the Draft view of the Map,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

Filtering the Map

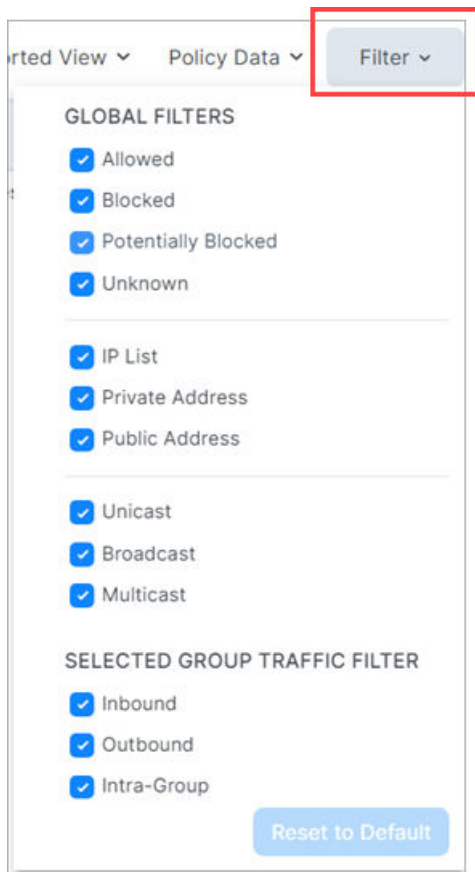
Connections Menu

When viewing the Traffic tab in on the Connections Menu allow you to view aggregated or individual connections.



Filter drop-down

Options in the Filter drop-down allow you to control which traffic information is displayed on the Map. This is useful for controlling the overall complexity of the visual information, making it easier to focus on the types of traffic you're interested in at any given time.



The Filter dropdown presents two types of filters:

Global Filters

These filters allow you to control the display of traffic for everything on the Map, whether selected or not.

Selected Group Filters

These filters allow you to control the display of traffic only for the selected group on the Map.

Panels in the Map



TIP

Use the drop-down selector above the panel to switch between the **Policy Data** and **Vulnerability Data** modes.


When you click an object in the Map, a side panel opens on the right that contains a number of tabs.

Summary Tab

The Summary tab displays information about the selected object. To view the Summary tab, click an item on the Map. The information displayed depends on the type of object you clicked and how deeply you've drilled into the object. For example, when you click a group in the Map, the Summary tab displays the labels in use, the number of workloads and virtual services, and the enforcement level. In general, the deeper you drill into an object, the more detailed information that is displayed in the side panel.

Increase VEN Traffic Update Rate

By default, VENs update traffic on the Illumination map every 10 minutes. You can temporarily increase the update frequency to once per minute. After 10 minutes, the default update rate of once every 10 minutes resumes. To use this feature, click a workload in the map to display the panel. In the **Summary** tab, click the option **Increase Traffic Update Rate**.



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 Traffic tab displays the following columns.

- Policy Decisions (reported and draft)
- Source Labels
- Destination Labels
- Destination Port Processes

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

- Source Port/Process User
- First Detected
- Flows/Bytes
- Last detected

See [Customize Columns \[22\]](#) 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:

- Connectivity
- V-E (vulnerability) score
- Enforcement
- Visibility
- Name
- Policy Sync status
- Ransomware Exposure
- Protection Coverage Score
- Labels
- When the policy was last applied

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.

Virtual Services Tab

The Virtual Services tab displays a list of all Virtual Services in the selected group. A drop-down selector allows you to filter the list by **Virtual Services with Traffic** or **All Group Virtual Services**. The list provides following information for each virtual service:

- Name
- Provision Status
- Service/Ports
- Addresses
- Labels
- Workloads / Container Workloads
- Description

You can add or remove columns by using the Customize columns drop-down list. See [Customize Columns \[22\]](#) for more information.

Vulnerabilities Tab



TIP

Use the drop-down selector above the panel to switch between the **Policy Data** and **Vulnerability Data** modes.

The Vulnerabilities Tab appears in the right panel when you're in Vulnerability Data mode (see note above). It details risk due to vulnerabilities. The workload with the most vulnerabilities appears at the top of the list. You can sort the V-E score column by vulnerability score. You can then define your patch priority based on the most critical score.

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
- Vulnerability score of the most severe network-accessible vulnerability on the workload
- East-West exposure
- Northern Exposure (Internet exposure)
- Number of workloads exposed to this vulnerability
- Associated port and protocol
- CVE-IDs (a unique identifier for the vulnerability)
- Name of the vulnerability

The East-West Exposure Score is recalculated whenever the rules associated with the workload are changed.

For more details about the Vulnerability Map and how to work with it, see [About Vulnerability Map](#). [87]

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 the Traffic Table](#) [67]
- [Create Unmanaged Workloads from IP Addresses](#) [66]
- Traffic exploration
- View the details about policy affecting each connection
- View the ransomware protection details

About the Traffic Table

You can use the Traffic table to query the PCE's traffic database to analyze traffic flows for auditing, reporting, and troubleshooting. You can query for traffic flows between workloads or hosts, labeled workloads, or IP addresses, and you can restrict the query 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 query 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 queries.

Queries

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

- **Source:** 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.
- **Destination:** 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.



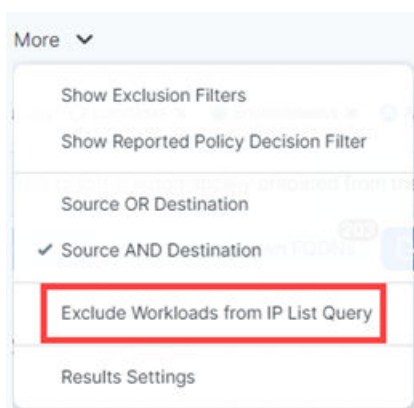
NOTE

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

- **Service:** 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 the search are used as a search criteria and the ones you add to exclude data are not used in the query. If you do not specify a value, all ports, protocols, port ranges, processes, and services are included in the query.
- **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:** (To enable this drop-down, click **More** and then select **Show Reported Policy Decision Filter**). To query for flows with a specific policy decision reported by the VEN, select the type of policy decision.

For more information, see [Deny Rules and the Traffic Table \[38\]](#) in this topic.

- **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 Source or Destination 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

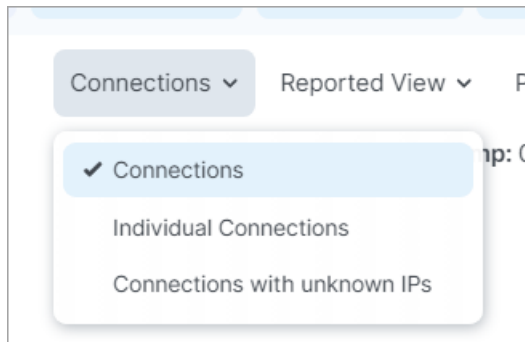
To gather the results of the current query in a .CSV file, click Export.

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.

Ways to View Query Results



Options in the Connections menu allow you to view traffic connections in the Traffic table as:

- Aggregations of multiple connections, for a more concise view
- Individual connections, for a more granular view
- Connections with unknown IP addresses

View Traffic as Aggregated Connections

Select the Connections option for an aggregated, more concise view of the traffic flows.



TIP

Contrast this method with the [View Traffic as Individual Connections \[36\]](#) below.

Note the following in the image below:

- **Connections** is selected in the Connections menu.
- **Common Set of Labels** is selected in the **Group by** menu. (This **Group by** selection is optional and shown here merely as an example grouping method. You can select this and/or any other combination of grouping options in the **Group by** menu, or no option at all.)
- The Source column shows an aggregation of 8, commonly-labeled IP addresses.
- The Flow/Bytes column shows 8 connections.
- The total row count for the page is 11.

Traffic

Source: redisjob x and Staging x and London x and CentOS x

Destination: Service: 8080 TCP x

Time: Last 24 Hours | Reported Policy Decision: | More

Connections | Group by: Common Set of Labels | Edit Labels | Resolve Unknown FQDNs | Export

Timestamp: 03/22/2024, 12:31:28 | Connections: 1 - 84 of 84

1 - 11 of 11 Total

Source	Source Labels	Destination	Destination Labels	Destination Port Process User	Flows/Bytes	First Detected	Last Detected
Potentially Blocked no Rule	10.29.82.254 10.29.82.72 10.29.82.28 10.28.0.153 10.28.81.176 10.28.36.60 10.28.81.160	redisjob London CentOS	Staging	1 Destination IP Full Enforcement	app_4070 role_4 8080 TCP http-alt root Jenkins CI Service - 8080 TCP WildFly	8 Connections 12 Flows 569.8 KB → 740.1 KB ←	03/21/2024, 01:28:26 03/21/2024, 19:49:09

View Traffic as Individual Connections

Select the Individual Connections option for a more detailed, granular view of the traffic flows.



TIP

Contrast this method with [View Traffic as Aggregated Connections \[35\]](#) above.

Note the following in the image below:

- **Individual Connections** is selected in the Connections menu.
- The **Group by** menu does not appear when **Individual Connections** is selected.
- The Source column shows several rows of individual, commonly-labeled IP addresses.
- The Flow/Bytes column shows that each row represents a single connection
- The total row count for the page is 50

Traffic

Source: redisjob x and Staging x and London x and CentOS x

Destination: Service: 8080 TCP x

Time: Last 24 Hours | Reported Policy Decision: | More

Individual Connections | Allow Selected Connections... | Edit Labels | Resolve Unknown FQDNs | Export

Timestamp: 03/22/2024, 12:31:28 | Connections: 1 - 84 of 84

1 - 50 of 84 Total

Reported Policy Decision	Source	Source Labels	Source Port/Process User	Destination	Destination Labels	Destination Port Process User	Flows/Bytes	First Detected	Last Detected
Potentially Blocked no Rule by Destination	solr-s14 10.29.82.254 Selective Enforcement	redisjob Staging London CentOS solr-s		discovery-s1 10.28.147.138 Full Enforcement	app_4070 role_4	8080 TCP http-alt root Jenkins CI Service - 8080 TCP WildFly	2 Flows 160.1 KB → 182.5 KB ← Corporate	03/21/2024, 05:25:11 03/21/2024, 16:58:15	
Potentially Blocked no Rule by Source	solr-s45 10.29.82.72 Selective Enforcement	redisjob Staging London CentOS solr-s	http-alt root	10.36.66.162 10.36.66.162	app1	8080 TCP http-alt root Jenkins CI Service - 8080 TCP WildFly	1 Flow 8.9 KB → 61.6 KB ← Corporate	03/21/2024, 23:47:21 03/21/2024, 23:47:21	
Potentially Blocked no Rule by Source	solr-s41-new 10.28.0.153 Selective Enforcement	redisjob Staging London CentOS solr-s	http-alt root	AWS - US East (Ohio) 10.8.251.138	Application12345 Azure AWS - US East (Ohio)	8080 TCP http-alt root Jenkins CI Service - 8080 TCP	1 Flow 25.6 KB → 84 KB ← Corporate	03/21/2024, 16:58:37 03/21/2024, 16:58:37	

View Connections with Unknown IP Addresses

From the Connections menu, select the **Connections with Unknown IPs** option to see a list of connections found by the query that your organization has not turned into unmanaged workloads. You can

easily create unmanaged workloads from these connections using the **Create Unmanaged Workloads** option above the list. For more information, see [About Unmanaged IP Addresses \[65\]](#).

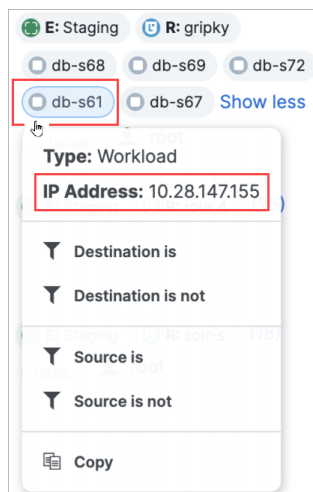
The screenshot shows the 'Traffic' section of a network visualization tool. A dropdown menu 'Connections with unknown IPs' is open, showing the 'Create Unmanaged Workloads' option. Below this is a table of connections. The table has columns for IP Address, FQDN, Destination Port/Process, Transmission, Direction, Workloads, and Flows. Three rows are selected, and a 'Create Unmanaged Workloads' button is visible above the table.

IP Address	FQDN	Destination Port/Process	Transmission	Direction	Workloads	Flows
10.28.81.162		8080 TCP	Unicast	Inbound	1	1
10.40.113.124		22 TCP	Unicast	Inbound	28	28
10.60.222.125		8983 TCP	Unicast	Inbound	23	28
		2181 TCP				
		3306 TCP				
		+4 more				
10.40.113.106		8983 TCP	Unicast	Inbound	7	7
		6380 TCP				
10.29.82.128		9042 TCP	Unicast	Inbound	5	6
		8080 TCP				

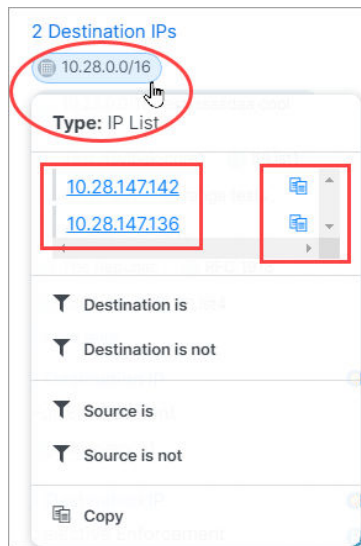
Mouse-over Policy Objects

You can mouse-over workloads, IP addresses, and IP lists in the table to access information and functionality.

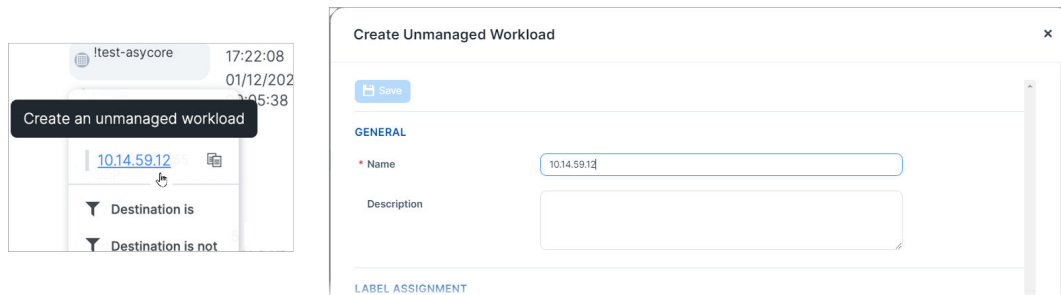
- View IP address associated with a workload.



- (Not available in all versions of the PCE) View and copy IP addresses in an IP list.



- (Not available in all versions of the PCE) Right-click an IP address to create an unmanaged workload.



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 Destination 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

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 to understand their impact before provisioning them.

**TIP**

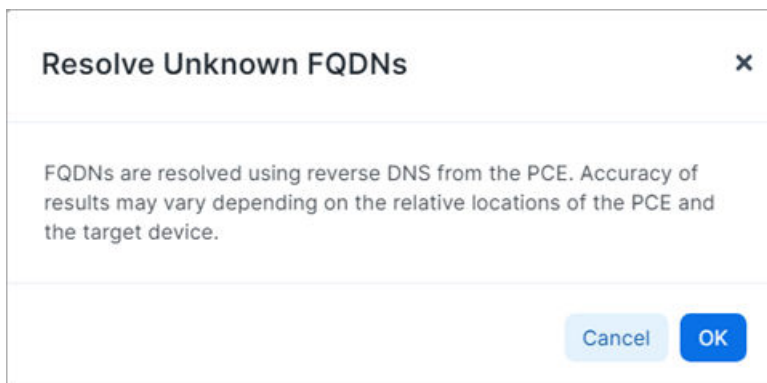
To view the details about a Deny Rule, click the linked text for traffic allowed across the rule (“Allowed”) or blocked by a Deny Rule (“Blocked”) while in a **Draft** view of the Traffic table. The **View Policy** dialog box opens. Then, click the **Deny Rules** tab.

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 then click OK.



2. Click Export adjacent to the Resolve Unknown FQDNs button.

**NOTE**

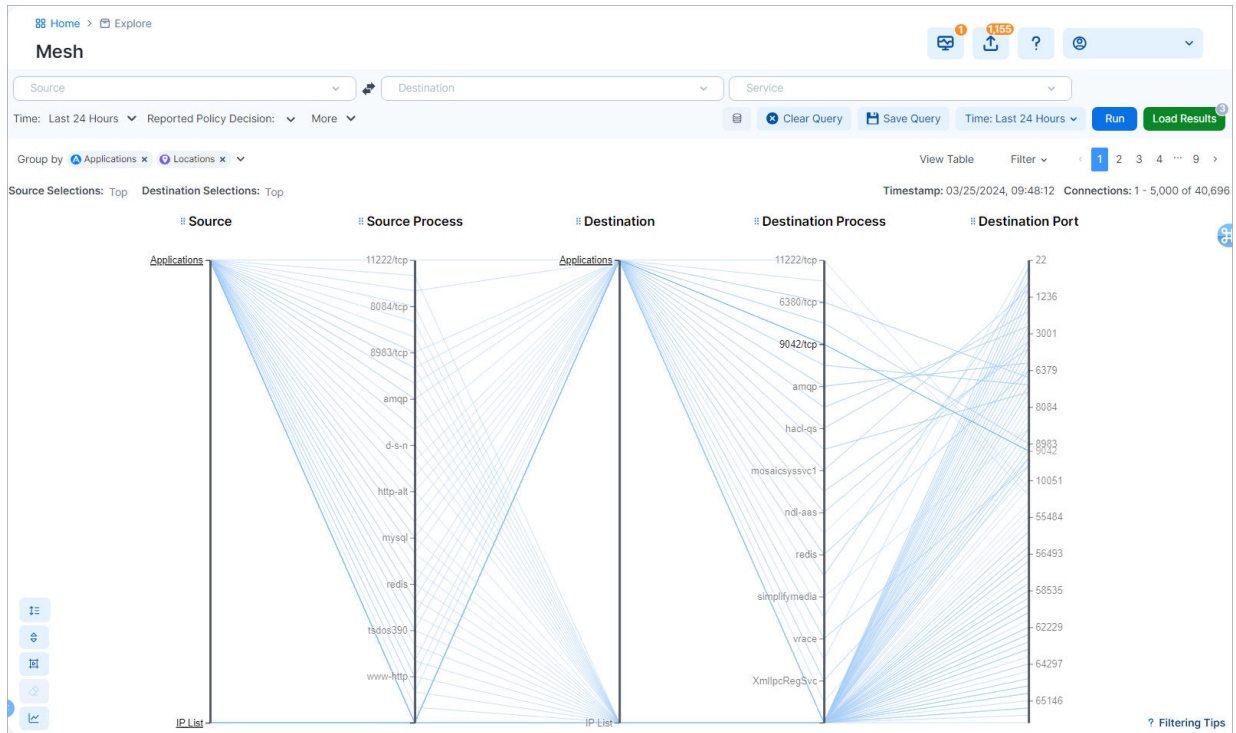
If you do not find relevant information, clear cached FQDN values and reload the results.

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

Mesh View

The Mesh view displays traffic flows along several vertical axes representing Destinations, Sources, ports, and processes. Mesh helps illustrate the overall topology of the network, offering an intuitive layout of network interaction.

You can click data points on the axes to focus on specific flows, sort and filter query results in several ways, drill down into data points for more granular detail, isolate ranges of data points with the Brush tool, by and to filter, and then go to the Table view to write rules.



About the Mesh



NOTE

The Mesh only supports Reported view. Draft view is not supported.

The Mesh shows multiple data dimensions simultaneously, providing a clear picture of how each data point interacts with others in their environment. Mesh offers the following:

- **High-dimensional data representation:** Mesh portrays high-dimensional data. Every dimension is represented as a vertical axis, and data points are depicted as lines that intersect these axes. Network traffic data is presented as source, destination, port, and processes.
- **Cluster identification:** In Mesh, patterns (like clusters or outliers) can emerge when multiple lines follow a similar path across the axes. This can help you quickly identify certain patterns or anomalies in the network flow data.
- **Simple representation:** Mesh provides a structured, organized view of traffic flows, especially when the dimensions (or axes) are logically ordered.
- **Compactness:** Mesh can represent a vast amount of data in a compact space by adjusting the spacing and ordering of its axes.
- **Enhanced data structuring:** Mesh methodically places data points along the axes in sequential order, not randomly. This neat ordering produces to a crisper, more comprehensible visualization.

Customize the Mesh Display

The **Group by** aggregation tool available in the Map and Traffic is also available in the Mesh view. You can select and apply multiple **Group by** parameters. You can also specify the priority of parameters by dragging them up or down in the list. You will see your top group based on your selection and you can

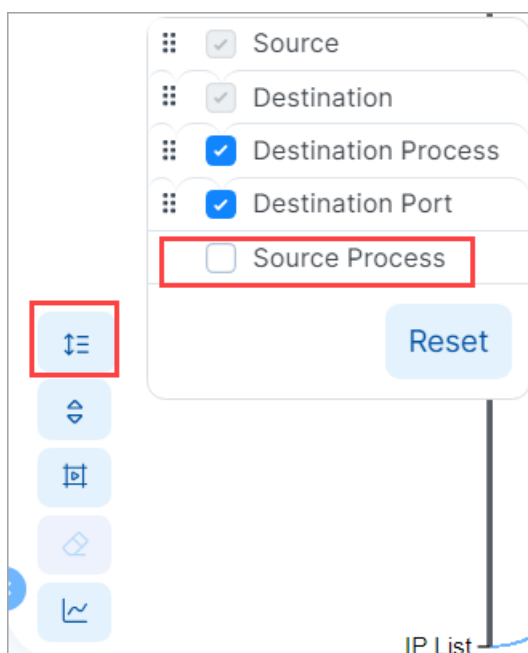
drill down through the groupings. The hierarchy of the parallel coordinates in the mesh is based on your selected grouping.

Other customization options include:

Add or remove axes

You can remove or add axes, keeping only those dimensions that are helpful for your current exploration. For example, except for the mandatory source and destination axes which you cannot remove, you might want to add or remove the other axes to narrow down the visualization and emphasize or de-emphasize certain data points.

In the image below, deselecting Source Process removes that axes from the Mesh.



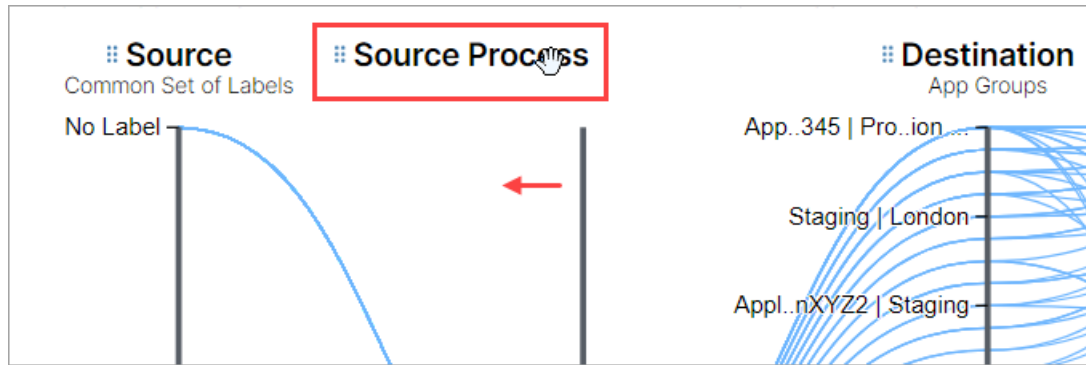
Reorder axes

You can reorder the Mesh axis columns to change the Mesh display. For example, you may want to move the axis you are most interested in to the center of the mesh and move less important data to the sides of the mesh.

There are two ways to reorder the axes:

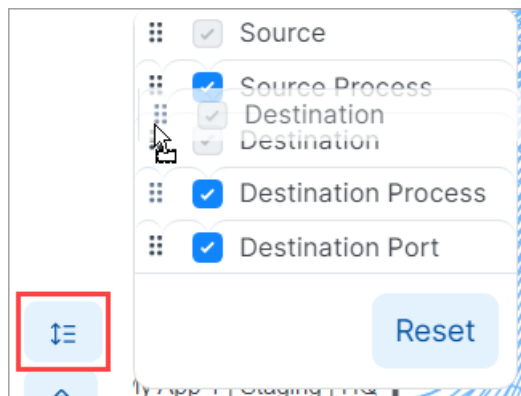
- **Dragging**

Click an axis heading and drag the axis left or right.



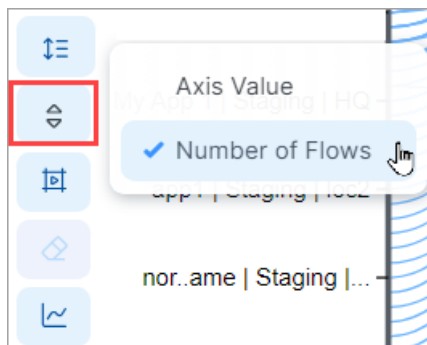
- **Specify the order in a list**

Click the **Customize Axes** icon in the Tool Cluster to display axes in a list, and then use the handles to select and drag axes up or down within the list.



Change the sorting order of Mesh data

You can sort the Mesh data by the **axis value** or the **number of flows**.



- **Axis Value**


In ascending order, orders the ports numerically and Source and Destination columns alphabetically.

- **Number of Flows**

Each axis is sorted according to its amount of traffic. The data point with the most flows appears at the top of the axis.

Change the Line Style

You can click the Line Style icon in the tool cluster to change mesh lines from curved (default) to straight.

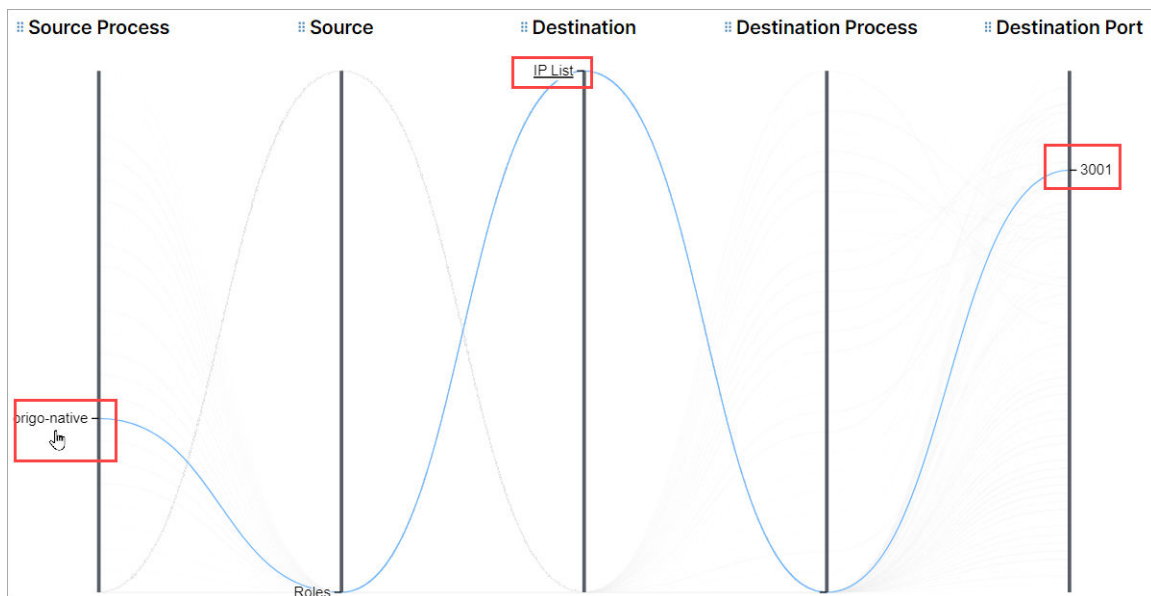


The Mesh provides several ways to navigate your data.

Explore the Mesh View Data

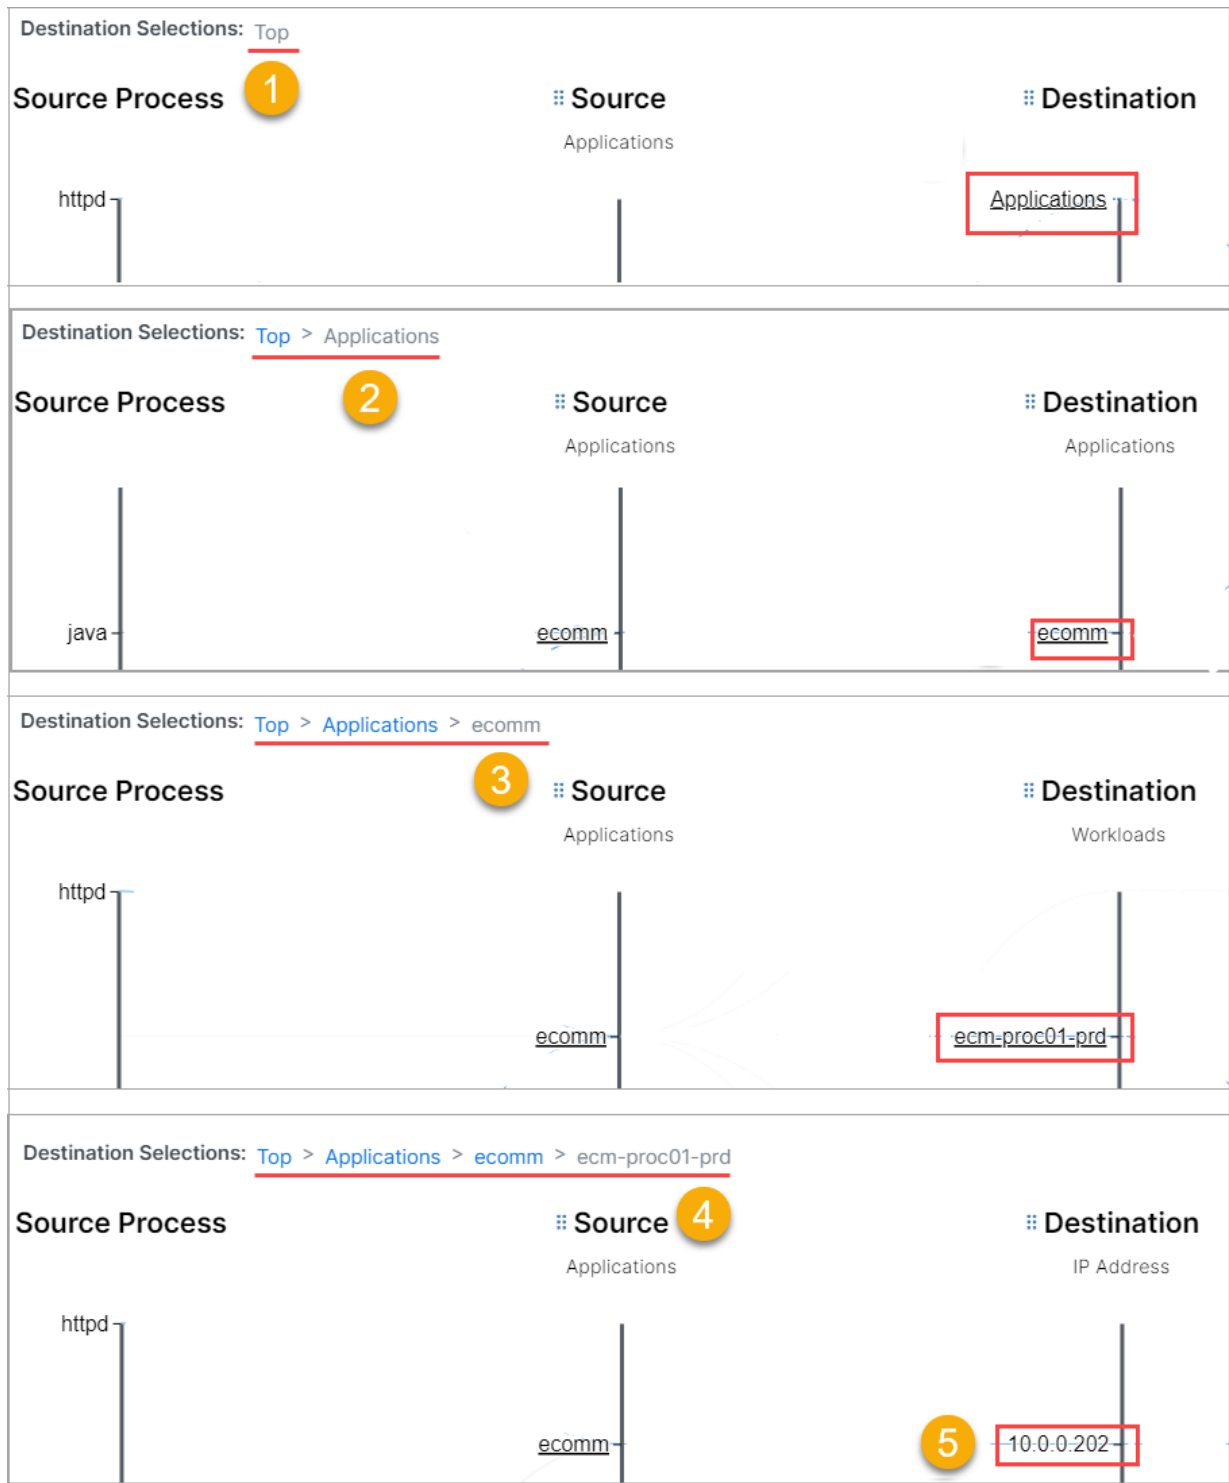
Basic mouseover

You can mouseover each vertical axis along data points to view specific connections between that point and points in other axes.



Drill in for more granular data





Click underlined data points to dive into deeper layers, traversing from grouped labels down to the granularity of individual IP addresses. Breadcrumbs above the mesh help you understand and traverse back through the layers.



Brush to define and isolate data points

The Brush tool allows you to isolate one or more data points on one or more axes. This simplifies your view of the data so that you can focus on the connections you're interested in.


You change Brush tool modes by clicking the Brush icon on the Tool Cluster.

-  Brush is off, ready to be activated.
-  **Activated.** When the Brush is activated, the icon turns dark blue and the cursor changes to crosshairs. By placing the crosshairs on the axis and dragging (brushing) up or down, you can select one or more data points on an axis. The selection is represented by the gray area on the axis. The crosshairs remain in effect while the Brush is in Activation mode.
-  **Deactivated.** When you have finished selecting areas on the axes using the crosshairs, you can click the icon to deactivate the crosshairs and restore the regular cursor. Now you can interact with the Mesh while the brushed portion of the axes, defined in gray, remain. You can still slide the gray bars up or down while in Deactivated mode.
-  **Clear Brush.** Click to clear the brush and exit Brush Mode.

Once the Brush is activated, you can:


- **Mouseover an axis line**

Source Process




- Click and drag up or down an axis line to **select and isolate one or more data points**. The selected portion is represented by a gray area on the axis.


Source Process



- To **highlight connections within specific flows**, apply the brush feature to multiple axes.



- To **turn off the crosshairs** and restore the regular cursor, click the brush icon again. You can **interact with the Mesh** while the brushed portion of the axes, defined in gray, remain.



When done isolating a range of data points on axes, click **Clear Brush** in the Tool Cluster to reset the Mesh.

App Groups

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 and how to configure them.

App Group features allow application owners to see all workloads for an application instance in a single App Group, even when the workloads are not currently communicating with each other. This is helpful when building or validating security policies for traffic between workloads because it allows application owners to focus only on the workloads that belong to their applications, regardless of location.

Ways to View App Groups

App Groups List Page

To view a list of all the App Groups in your PCE, click App Groups in the left navigation. The filterable list page presents high-level information about the listed App Groups as well as hyperlinks to each App Group's rules and map.

Two options are provided at the top of the App Groups list:

- **Edit App Group Definition:** Allows you to configure whether App Groups comprise Application and Environment labels or Application, Environment, and Location labels. This is a global setting for all App Groups.
- **Segment Multiple App Groups:** Allows you to apply nano-segmentation (also known as ringfencing) on multiple App Groups using the Policy Generator. Nano-segmenting App Groups allows all workloads to communicate across all services within each App Group. For related information, see .
- **Protection Coverage Score** scores are shown for each app group on the List Page.


App Groups Details Page

To view an individual App Group, click an App Group in the list to view its **Details Page**. The details page includes several tabs:

- **Explore:** Offers three visualization options ([detailed below \[47\]](#)) for exploring App Group traffic.
- **Members:** Lists the App Group member applications; presents high-level information about each member.
- **Rules:** Lists the inbound and outbound Essential Service Rules that the App Group requires.
- **Policy Generator:** Simplifies the Illumio policy creation process by recommending the optimal security policy for your App Groups. You can use it to accelerate security workflows and reduce the risk of human error while creating security policy. For details, see .
- **Vulnerabilities:** Shows the different vulnerability exposure scores for the selected App Group based on the ports, protocols, and workloads to which it is exposed.
- **Ransomware Protection (Preview version):** Presents a protection coverage score, workload exposure score, and recommended actions widget. A table of Risky Services is also provided.

App Group Visualization Options

App Group Traffic Table




The App Groups Traffic table displays details about the App Group in a traditional table format, including the traffic, group members, rule coverage, and ransomware protection statistics associated with that App Group.

You can use the Traffic table to query and analyze the PCE's traffic database for auditing, reporting, and troubleshooting. You can query for traffic flows between workloads or hosts, labeled workloads, or IP addresses, and you can restrict the query 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.

For more information, see [Traffic Table \[33\]](#).

App Groups Mesh View




The Mesh view displays traffic flows along several vertical axes representing Destinations, Sources, ports, and processes.

You can hover over data points on the axes to focus on specific flows. You can also sort and filter query results to view them in a number of ways. You can drill down into deeper layers of data, traversing from grouped labels down to the granularity of individual IP addresses. You can use the Brush tool to isolate and investigate separate data points, and then go to the Table view to write rules.

For tips 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.

For more information about the Mesh, see [Mesh View \[39\]](#).

App Groups Map View



The App Groups Map view displays the workloads and traffic in your data center. The Map 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.

For more information, see [Map View \[23\]](#).

You can search for specific App Groups and see the associated workloads, traffic, and rule coverage between members in the group and other Source and Destination App Groups that provide or consume its services, as well as rule coverage for the traffic between App Groups.

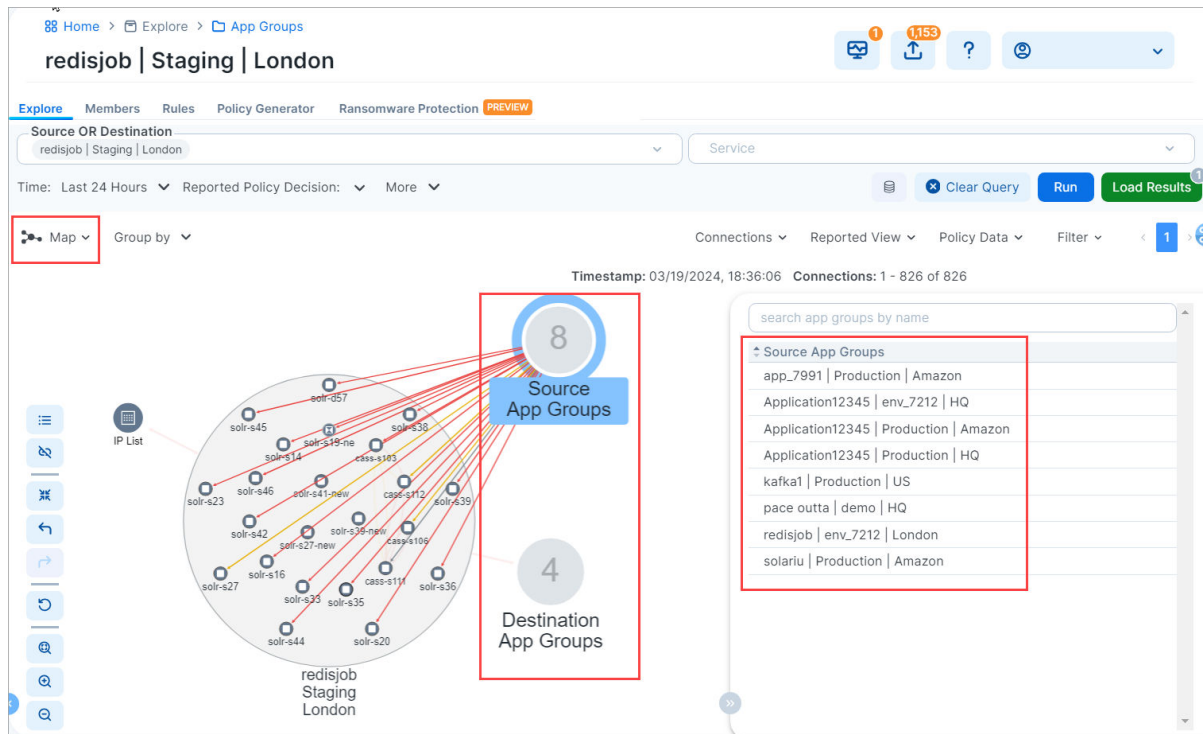
- **Source App Groups:** Use services provided by the current application
- **Destination App Groups:** Provide services used by the current application



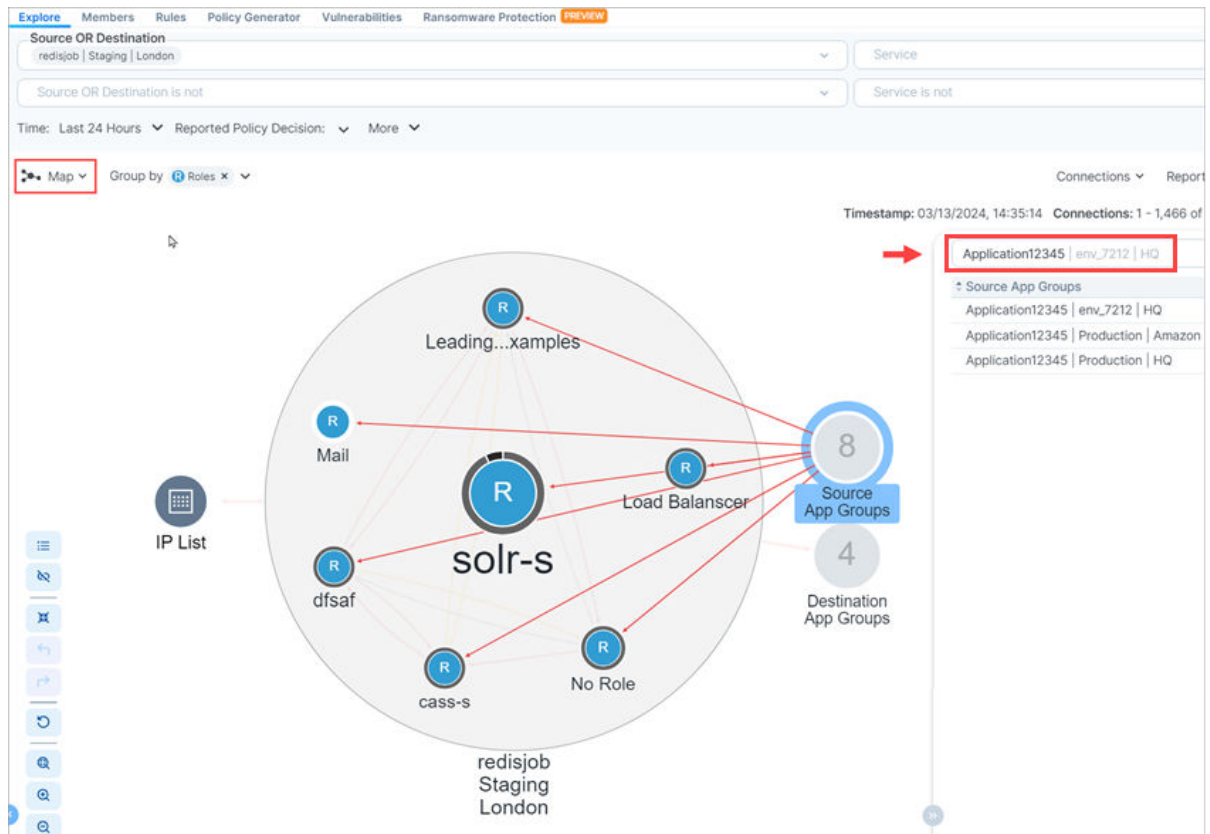
NOTE

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

- When you click an App Group in the Map, the workloads and their associated traffic in that App Group displays, as well as a pop-up list of other App Groups communicating with that App Group either as the source or destination of services.
- Connected to the App Group by arrows are the Source App Groups that initiate connections to this application instance and the Destination App Groups that provide services for this application instance. To view a list of the source or destination App Groups, single-click its circular representation on the map. A panel opens which displays the name of each App Group along with its Environment and Location label.



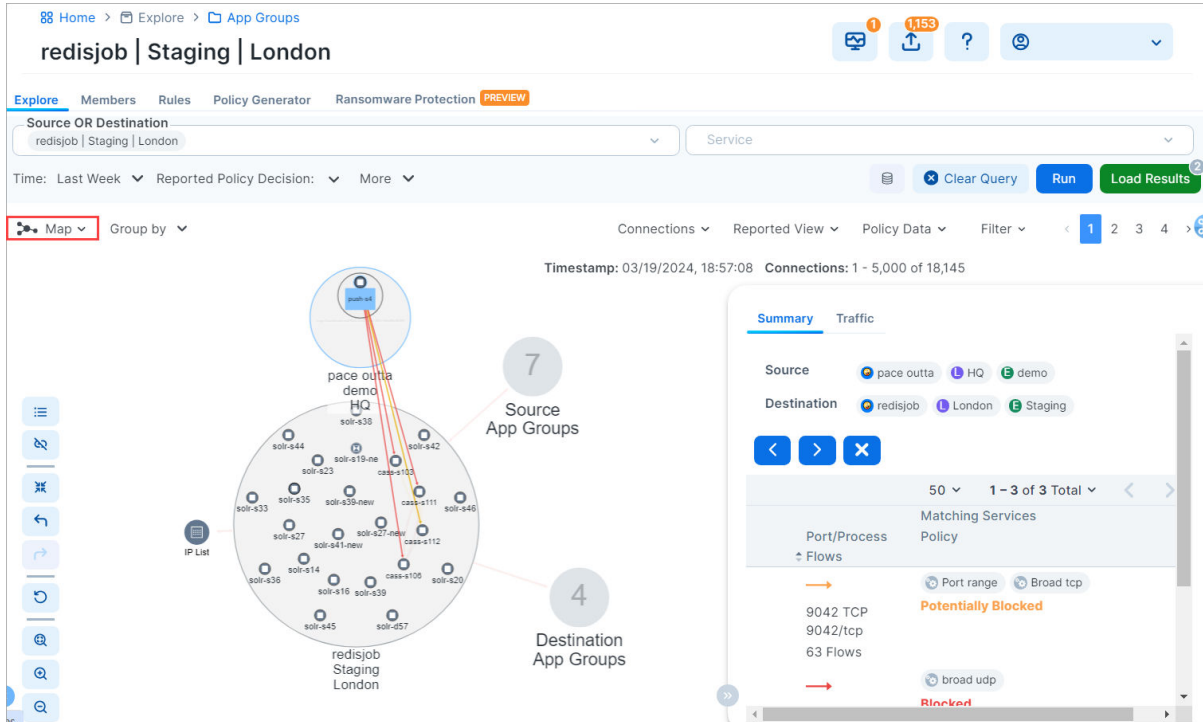
You can also search for App Groups connected to a given App Group. This is convenient for filtering a long list of connected App Groups.



NOTE

If the App Group does not have any connections, the Destination and Source App Groups do not display.

- When you click a Source or Destination App Group in the panel, lines representing the traffic links between the App Groups display in either red for blocked traffic or green for allowed traffic. Source App Groups display above the original App Group and Destination App Groups display below the original App Group.
- If an expanded Source or Destination App Group is currently displayed in the App Group Map, you can view the next or previous connected App Groups by clicking the **Left** or **Right** arrows in the panel.
- 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.



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 Destination App Group to allow other Source App Groups to access your application workloads.

Compare App Group V-E Scores by Enforcement Type

The Show Vulnerability Exposure (V-E) Score tool lets you see how the security of your app groups would change if you were to change their current enforcement mode. Columns in the App Group list and details pages provide a side-by-side comparison of the effect different enforcement modes would have on Vulnerability and Exposure (V-E) scores. A toggle allows you to simulate the switch between Full Enforcement and Visibility Only enforcement modes.



NOTE

This option allows you to simulate the switch between Full Enforcement and Visibility Only modes. It doesn't change the actual enforcement mode of your app groups.

How it works

- The PCE displays V-E scores in the UI based on ransomware and vulnerability statistics it previously calculated and stored in a database.
- If the stored data is stale (4 hours or older), the PCE recalculates the statistics and updates the V-E scores in the UI.
- Toggling the Full Enforcement/Visibility Only options provides a side-by-side comparison of the effect of the different enforcement modes.
- Because the PCE calculates and re-checks for new data periodically, the information in the UI may not immediately reflect the current V-E score.
- API responses include the complete vulnerability data set for the different enforcement modes. V-E data for all modes is pre-processed and stored in a database to eliminate the performance impact that would result from frequent recalculation.

- A V-E score is the calculated value based on the Vulnerability Score and Exposure Score = $\sum f(VS, ES)$. It can be shown for an individual vulnerability on a port for a single app group or as a summation of all the V-E Scores for an App Group, role, or workload.

App Group List pages

On App Group list pages, two adjacent columns show the following:

- Full Enforcement / Visibility Only V-E Score: Depending on the item's current enforcement mode, this column matches the Current V-E Score column or changes to show a different V-E score obtainable if the actual enforcement mode were changed.
- Current V-E Score: The most recently calculated V-E score of the .

88 Home > Explore

App Groups

Segment App Group

Select properties to filter view

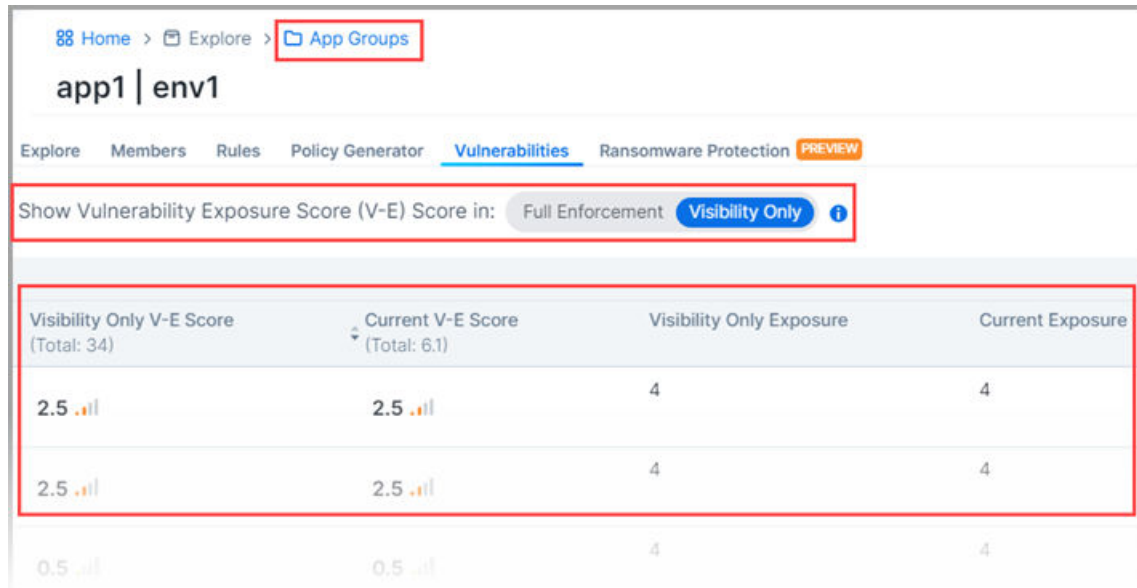
Show Vulnerability Exposure Score (V-E) Score in: Full Enforcement **Visibility Only**

Visibility Only V-E Score	Current V-E Score	Name
34 .lll	6.1 .lll	app1 env1

App Group Details pages

On the Vulnerabilities tab of App Group details pages, four adjacent columns show the following:

- Full Enforcement / Visibility Only V-E Score: Depending on the item's current enforcement mode, this column matches the Current V-E Score column or changes to show a different V-E score obtainable if the actual enforcement mode were changed.
- Current V-E Score: The most recently calculated V-E score of the app group.
- Full Enforcement Exposure: Depending on the item's current enforcement mode, this column either matches the Current Exposure column or changes to show a different exposure score obtainable if the actual enforcement mode was changed.
- Current Exposure: The current exposure score of the app group.



Work with the App Group Map

There are two types of App Groups: Destination App Groups (formerly Providing) and Source App Groups (formerly Consuming). Destination App Groups provide service to an application instance and Source 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 (Destination) or consume (Source) its services, and segmentation rule coverage for the traffic between App Groups.

How App Groups are Created and Associated

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 no App Group with those labels exists, an App Group 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 Number of Matching App Group Labels

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

- Application and Environment labels only (the Location label is ignored). This is the default.
- 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 with two or three labels as described above:

1. In the left navigation, in the Explore section, click App Groups.
2. Click Edit App Group Type.
3. Select either **Application | Environment** (default) or **Application | Environment | Location**.
4. Click Save.

Caveats

- When the App Group Configuration setting is changed, the list of “Most Recently Viewed App Groups” is cleared.
- If you have a large number of workloads in your organization, it may take up to five minutes to regenerate the Map.


Reports

The PCE allows you to generate, download, and manage several types of recurring reports:

- Executive Summary
- App Group Summary
- Traffic Export
- Rule Hit Count

Reporting in the PCE

The PCE UI menu includes a *Reports* option. Generated reports appear on the **Downloads** tab.



Name	Report Type	Generated By	Status	Action
Weekly Executive Summary	Executive Summary	Generated At 11/10/2023, 12:41 PM	Completed	Download
App Group Summary	App Group Summary	11/10/2023, 12:40 PM	Completed	Download
Daily Traffic Report	Executive Summary	11/10/2023, 12:39 PM	Completed	Download
Time: Last Hour	Traffic Export	11/10/2023, 11:46 AM	Completed	Download
Default Graph	Traffic Export	11/10/2023, 03:06 AM	Completed	Download
Rule Hit Count	Rule Hit Count	11/10/2023, 03:06 AM	Completed	Download

Reports are created in either PDF or CSV format, depending on the report type. You can download reports and share them with people in your organization who don't have access to the PCE UI or PCE REST API.

While the data in the reports is not customizable, you can configure the time range of the data that the reports are generated from and the frequency at which they are run.

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
- **Quarterly:** At the start of the next quarter after the report was added, then every three months thereafter. For Example, if the report was added in November, the first report is generated December 31st including data from October 1st - December 31st. The next report is generated March 31st including data from January 1st - March 31st.

The PCE only retains reports for a maximum of 7 days, but there is no limit to the number of reports you can create. Generated reports include data for provisioned security policy, managed and unmanaged workloads, and provisioned policy objects. Reports 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](#). [87]

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 Source 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. For more information, see [App Groups \[46\]](#) in this guide and "Role-based Access for Application Owners" in the PCE Administration Guide, respectively.

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. For more information, see [Configure App Groups \[54\]](#).

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.

Traffic Export Report

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

**NOTE**

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

Rule Hit Count Report

The Rule Hit Count Report provides the following:

- **Policy Compliance:** Generate a Rule Hit Count Report to provide evidence that security controls are in place and working effectively, demonstrating compliance to auditors.
- **Redundancy Removal:** Identify unused or less-used rules so you can remove or modify them to reduce redundancy and clutter in your implementation.
- **Troubleshooting:** When network issues arise, identify the rules that were in effect during the relevant traffic flow, allowing you to resolve problems faster and more efficiently.

For more information, see the topics [Work with Reports in the PCE \[62\]](#) in this guide and "Rulesets and Rules" in the REST API Developer Guide.

Requirements

- PCE Version:
 - SaaS: 24.2.0 or later
 - On-prem 23.5.10 or later
- VEN Version: 23.2.30 or later

Specifications

- Support for up to 25k VENs.
- Support for up to 75k total rules.
- The VEN can report a maximum of 100 rule IDs for each reported flow entry. If there are more than 100 rule ID matches for a flow, the rule IDs are truncated.
- No support for Superclusters.
- Only active rules are counted.
- Essential rules (rules necessary for the Illumio platform to function) are not counted.
- The report includes each rule's hypertext reference attribute (HREF). The HREF maps directly to a rule in the PCE UI, but clicking the HREF does not redirect you to the specific rule. It merely loads the JSON object of the rule.
- VENs report to the PCE the hit count of all the overlapping rules for a flow.
- VEN enablement for this feature makes use of label scopes similar to firewall co-existence and SecureConnect.
- Rule count data is retained for 90 days, after which the oldest data is dropped.
- Last Hit timestamps are retained for the life of the PCE.
- The report includes the active rule IDs within the rule sets you specified when you configured the report, plus all the deny rules.
- Hit Count values reflect the total number of hits recorded during the configured time range.
- Due to PCE policy optimization, some rules that weren't written to overlap may end up overlapping. For example:
 - Given two flows:
 - A → B on TCP/443
 - A → C on TCP/443

- Although the flow from A → B on TCP/443 never overlaps with the flow from A → C, due to policy optimization, the rule counter for both rules may increment.

Procedure

STEP 1: Enable Rule Hit Count

The PCE and VENS require enablement through the Illumio REST API. For details, see "Enable Rule Hit Count" in the REST API Developer Guide.

STEP 2: Create Rule Names in the PCE

To ensure that there are names in the Rule Name column, you need to add a Note through the PCE UI for each rule that will be captured in the report. The Rule Hit Count Report populates the Rule Name column by pulling the Note contents from the specified rules.

A	
1	Date Of Export: 2024-02-01
2	Exported By:
3	Start Date: 2024-01-02T00:00:00Z
4	End Date: 2024-01-31T23:59:00Z
5	Rule Name
6	workloads to workloads on 18000 TCP/UDP
7	workloads to workloads on 17200 TCP/UDP
8	workloads to workloads on 17000-17500 TCP/UDP
9	workloads to workloads on 16200 TCP/UDP RS2
10	workloads to workloads on 16000-16500 TCP/UDP RS
11	workloads to workloads on 16200 TCP/UDP RS1_2
12	workloads to workloads on 16000-16500 TCP/UDP RS
13	workloads to workloads on essential service UDP
14	extra scope rule role1 talks to role2
15	extra scope rule role1 talks to role2
16	any iplist ams to label
17	overlapping rule with all scope rules in app2-env2 rules
18	extra scope rule role2 talks to role1
19	any iplist ams to label

- Go to **Policy > Rulesets & Rules**.
- Click a ruleset to open its details page.

Rulesets and Rules

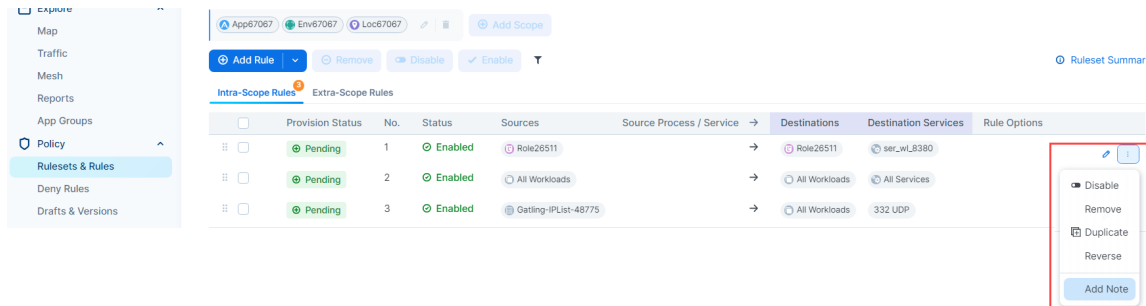
[Rulesets](#)
[Rule Search](#)

⊕ Add
Start Policy Generator
⬆ Provision
⬅ Revert
⊖ Remove
⏻ Disable
✓ Enable


<input type="checkbox"/>	Provision Status	Status	↕ Name
<input type="checkbox"/>	⊕ Pending	⊙ Enabled	App67067 Env67067 Loc67067
<input type="checkbox"/>	⊕ Pending	⊙ Enabled	App70053 Env70053
<input type="checkbox"/>	⊕ Pending	⊙ Enabled	Global


- To provide a name that the Rule Hit Count Report can use, add a note to a rule.

- a. Click the three vertical dots adjacent to the rule, and then click **Add Note**.



- b. Enter a name for the rule. The name you enter appears in the Rule Name column when you generate the Rule Hit Count Report.



- c. Click the **Save** icon in the upper right corner. A word bubble icon  appears next to the pencil icon.

STEP 3: Generate the Rule Hit Count Report

See [Add a Report \[62\]](#).

How Rule Hit Counts are calculated

The following example scenarios help explain how rule hit counts are calculated and reported.

Scenario 1

Flow: Workload A → Workload B on TCP/443 (reported by both sides)

Enforcement Mode: n/a

Rules	Count	Comments
Workload A → Workload B on TCP/443	2	Both workloads reported the flow and this rule is executed by both of them.
Workload A → Any IP List	1	Only workload A executes this rule.
Some IP List Covering A → B	1	Only workload B executes this rule.

Scenario 2

Flow: Workload A → Workload B on TCP/443 through a network enforcement point that blocks A → B (so only reported by A)

Enforcement Mode: n/a

Rules	Count	Comments
Workload A → Workload B on TCP/443	1	Because A has a VEN on it and it allowed the flow and B hasn't reported it.
Workload A → Any IP List	1	Because A has a VEN on it and it allowed the flow.
Some IP List Covering A → B	0	Because A has a VEN on it and it allowed the flow.

Scenario 3

Flow: Workload A → Workload B on TCP/445

Case 1 Enforcement:

- Workload A Enforcement Mode - Visibility and TCP/445 is not allowed outbound
- Workload B Enforcement Mode - Full

Rules	Count
Allow Any (0.0.0.0/0) → Workload B on all services	1

Case 2 Enforcement:

- Workload A Enforcement Mode - Full and TCP/445 is not allowed outbound
- Workload B Enforcement Mode -Full

Rules	Count
Allow Any (0.0.0.0/0) → Workload B on all services	0

Case 3 Enforcement:

- Workload A Enforcement Mode - Selective
- Workload B Enforcement Mode -Full

Rules	Count
TCP/445 is blocked outbound on A via boundary	1
Allow Any → Workload B on all services	0

Scenario 4

Flow: Workload (Endpoint) C → Workload (Server) B on TCP/443

Endpoint A - Label:Loc1 (IP address: 10.3.2.4/24 → subnet = 10.3.2.0/24 == 10.3.2.0 → 10.3.2.255)

Server B - Label:App1

Endpoint C - Label:Loc2 (IP address: 10.3.2.7/24 → subnet = 10.3.2.0/24 == 10.3.2.0 → 10.3.2.255)

Behavior:

- Endpoint C will drop the flow if it's in Enforcement Mode (because there's no rule allowing outbound)
- Server B will accept a flow from either Endpoint A or Endpoint C if the flow makes it to server B

Case 1 Enforcement:

Endpoint C Enforcement Mode - Full

Rules	Count	Comments
Loc1 Endpoints (Use WL subnets) → App1	0	Endpoint C will drop the flow because there is no outbound rule.

Case 2 Enforcement:

Endpoint C Enforcement Mode - Selective

Rules	Count	Comments
Loc1 Endpoints (Use WL subnets) → App1	1	Endpoint C will allow the flow because there is no boundary. Server B will allow the flow because Endpoint C is in the same subnet as Endpoint A. The report indicates that the Loc1 rule was hit, but the flow is coming from a Loc2 Endpoint.

Scenario 5 (PCE rule optimization)

Flow: Workload A → Workload B on TCP/443

If the address of workload B and workload C overlap, then PCE rule optimization could merge the following rules resulting in the second rule also being incremented.

Rules	Count	Comments
Workload A → Workload B on TCP/443	2	Both workloads report the flow. 2
Workload A → Workload C on TCP/443	2	The reported flow could potentially contain this rule ID as well because of PCE rule optimization.

Work with Reports in the PCE

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

Add a Report

1. From the left navigation under the **Explore** category, click **Reports**.
2. Click **Add Report** and select the type of report you want to add.
3. Configure the report settings:

Executive Summary

- **Add Report:** To generate the report immediately after you click **Save**, select **Now**. To generate the report later, select **Later** and then specify the time and date in the **Scheduled Time** field.
- **Recurrence:** Select how frequently you want the PCE to run the report.
- **Time Range:** Select the time range for the report.
- **Name:** Specify a name that describes the purpose of the report. Report names can be from 2-255 characters and contain special characters.

App Group Summary

- **App Group:** Select the application for which you want to generate the report.
- **Add Report:** To generate the report immediately after you click **Save**, select **Now**. To generate the report later, select **Later** and then specify the time and date in the **Scheduled Time** field.
- **Recurrence:** Select how frequently you want the PCE to run the report.
- **Time Range:** Select the time range for the report.
- **Name:** Specify a name that describes the purpose of the report. Report names can be from 2-255 characters and contain special characters.

Traffic Export

- **Saved Filter:** You can run a previously saved Traffic filter and export the results to a CSV file on a recurring schedule. If you edit the filter, subsequent recurrences of the exported file continue to use the original version of the filter.
- **Run Query:** To generate the report immediately after you click **Save**, select **Now**. To generate the report later, select **Later** and then specify the time and date in the **Scheduled Time** field.
- **Recurrence:** Select how frequently you want the PCE to run the report.
- **Name:** Specify a name that describes the purpose of the report. Report names can be from 2-255 characters and contain special characters.

Rule Hit Count



NOTE

Requires 23.2.30-VEN or later

- **Rule Sets:** Select the ruleset(s) whose details you want to capture in the report.
 - **Add Report:** To generate the report immediately after you click **Save**, select **Now**. To generate the report later, select **Later** and then specify the time and date in the **Scheduled Time** field.
 - **Recurrence:** Select how frequently you want the PCE to run the report.
 - **Time Range:** Select the time range for the report.
 - **Name:** Specify a name that describes the purpose of the report. Report names can be from 2-255 characters and contain special characters.
4. Choose whether to have a copy of the report emailed to you.
 5. Click **Save**.

Manage Reports

To download a report:

1. From the left navigation under the **Explore** category, click **Reports**.
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 (maximum 7 days).

1. From the left navigation under the **Explore** category, click **Reports**.
2. Click **Settings** in the top right corner of the page.
3. In the **Retention** field, enter the number of days you want the PCE to retain PDF files.
4. Click **Save**.

To edit the settings for a report:

Only reports configured to recur appear in the **Schedules** tab, and only these reports can be edited.

1. From the left navigation under the **Explore** category, click **Reports**.
2. Click the **Downloads** tab.
3. Click the row for the report you want to modify.
4. Change settings.
5. Click **Save**.

To remove a report:

Removing a report from the **Schedules** tab prevents 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 left navigation under the **Explore** category, click **Reports**.
2. Click the **Schedules** tab.
3. Click the row for the report you want to remove.
4. Click **Remove** in the dialog box and again in the confirmation message.

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 UI, 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.

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:

1. **Group discovery:** When you pair workloads, the VEN introspects those Workloads and determines their open ports, running services, and traffic flows.
2. **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.
3. **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). The Map also proposes suitable rules for you to use or modify if you do not want to manually create rules from scratch. See in the *Security Policy Guide* for more information.
4. **Rule Testing:** The Map allows you to test and evaluate your rules against existing traffic flows *without* enforcing the rules. You can test Rules 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 UI. When you export the file, it lists IP addresses.



NOTE

The DNS names are not displayed on the Map 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

You can select IPv4 or IPv6 addresses displayed in the Map from the internet, IP lists, or traffic links. You can change the default interface and hostname if needed and you can add labels 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, click one of the following groups: IP List, Private Addresses, or **Public Addresses**. The right-side panel for the object opens. If you click an IP List group, the Summary tab displays the IP addresses in each IP list. Similarly, if you click a Private or Public IP Address group, the Summary tab displays the list of IP addresses.
2. If you're viewing an IP list in the panel, click the name of an IP list to expand it in the panel. The panel 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 adjacent to the IP address for which you want to create an unmanaged workload.

The screenshot shows the Palo Alto Networks Map interface. The main area displays a network diagram with various nodes and connections. On the right, a panel is open showing the 'Summary' tab for a selected IP list. The panel includes a 'Create Unmanaged Workloads' button, which is highlighted with a red box. Below this button, there is a table with columns for 'Unmanaged Source IP' and 'Destination Port/Process'. The first row in the table shows the IP address '5.128.4.80' with a checked checkbox in the 'Unmanaged Source IP' column. The 'Destination Port/Process' column lists several services: '1299 UDP hp-sci', '22 UDP', '9396 TCP', and '+21 more'.

4. Click **Create Unmanaged Workloads**.
5. In the **Assign Labels** dialog box, click the drop-down list and select the labels you want to assign to the unmanaged workload.
6. Click **Confirm**.
7. [Optional] Recalculate your map with the newly created unmanaged workload by clicking **Recalculate** in the confirmation dialog box.
8. In the left navigation, click **Workloads**.
9. In the Workloads list, locate the new unmanaged workload you created. Identify the unmanaged workload by its name, which is its IP address.
No enforcement information is available for the new unmanaged workload because it doesn't yet have a VEN installed on it.
10. Click IP address in the Workload list.
11. Click Edit and enter the workload information.

12. Click Save.

Add Rules for Traffic Using the Traffic Table

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

In the Traffic table, 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 Traffic table.)

To add rules for traffic flows:

1. In the left navigation, click **Traffic**.
2. Select **Connections** in the Connections filter.
3. From the **Group by** drop-down, select **Common Set of Labels** and then click **Apply**.
4. From the **Draft View** drop-down, select a Draft View.
5. Using the checkboxes, select traffic flows that you want to write rules for.
The **Allow Selected Connections** button becomes available and includes the number of allowable connections for which the PCE can write rules.
6. Click **Allow Selected Connections**.



NOTE

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


7. 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.

8. As needed, edit the scopes for your ruleset and then click 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.

9. (Optional) To control how the PCE uses services in the rules, click **Settings**.



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

Automatically merge similar rules together when

☒ The Source and Destination are common

☐ The Destination and Services are common

☐ The Source and Services are common

☐ Never

OK

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

When you select **Allow the use of 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.

- 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.

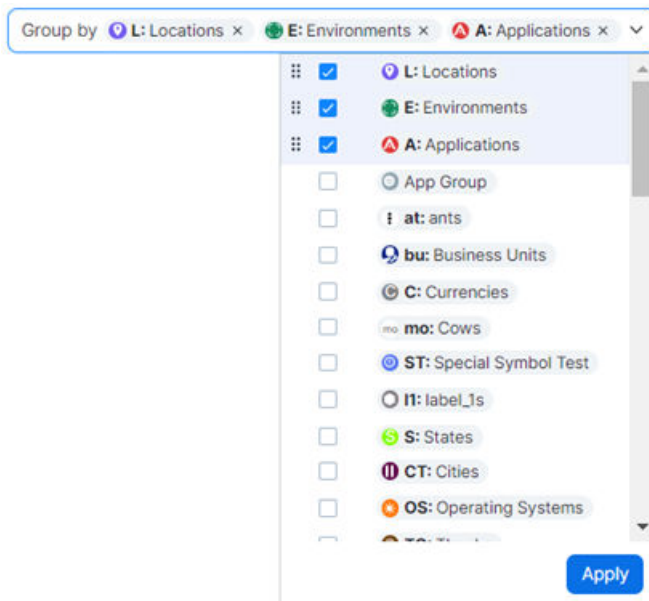
- One you're satisfied with the ruleset selected and the rules within the ruleset, click **Save** or **Save and Provision**, depending on whether you want to immediately provision the ruleset.

After saving your ruleset and rules, the PCE UI reloads your data so that the Traffic table 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.

- In the left navigation, click **Map**.
- Verify the criteria by which the group was established.
In the **Group by** filter, select the grouping criteria.

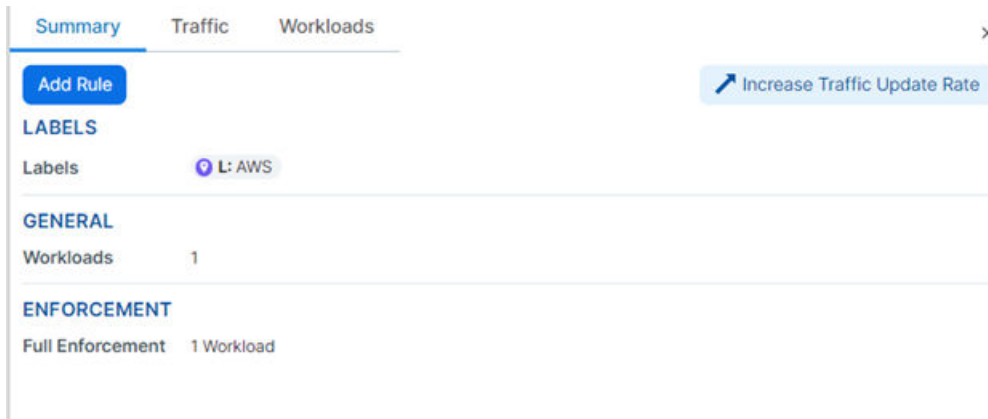


3. Keep the current selection (Locations, Environments, Applications), or add or remove the grouping criteria.
4. Once you have the desired selection, click **Apply**.
The group is now established according to your needs.
5. Now put the cursor over the group that you want to change (here it is AWS).

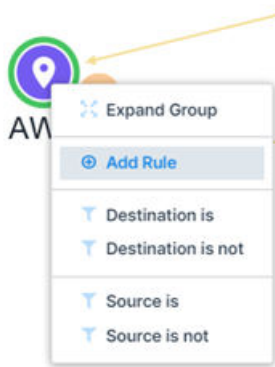


The pop-up dialog on the left shows the selected group's stats.

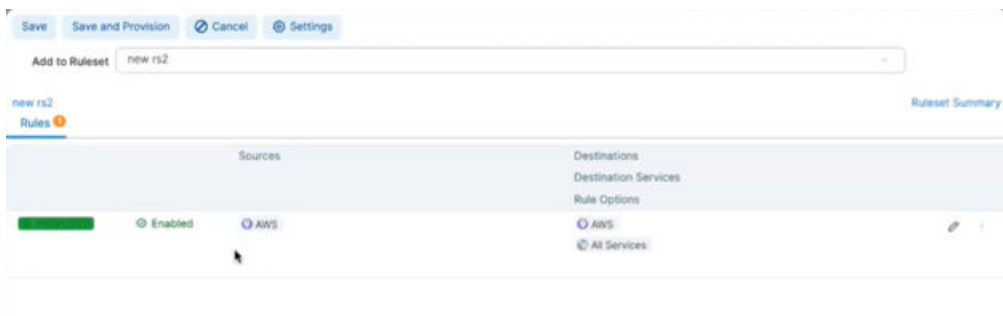
6. You can also click on the group to see its stats that show in the right panel.



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



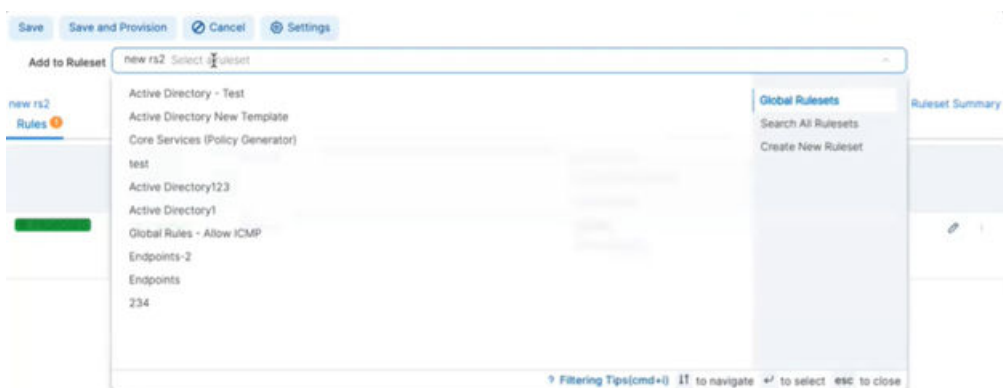
8. Choose the ruleset to which you are adding the new rule. For example, the ruleset named **new rs2**.




9. Select **Rule Options**.


For example, you can select **All Services**.

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



11. After you have added the rule, click **View Rule** to view it.





Everything inside that Rule communicates with 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.

Dashboards

This section describes the Server & Endpoints and Ransomware Protection dashboards. Together, the dashboards provide a set of widgets for Server and Endpoint VEN statistics as well as broad visual information about ransomware protection readiness, risk exposure, and protection coverage.

Servers and Endpoints Dashboard

The Servers and Endpoints Dashboard displays summary statistics and status information for Server VENs and Endpoint VENs.

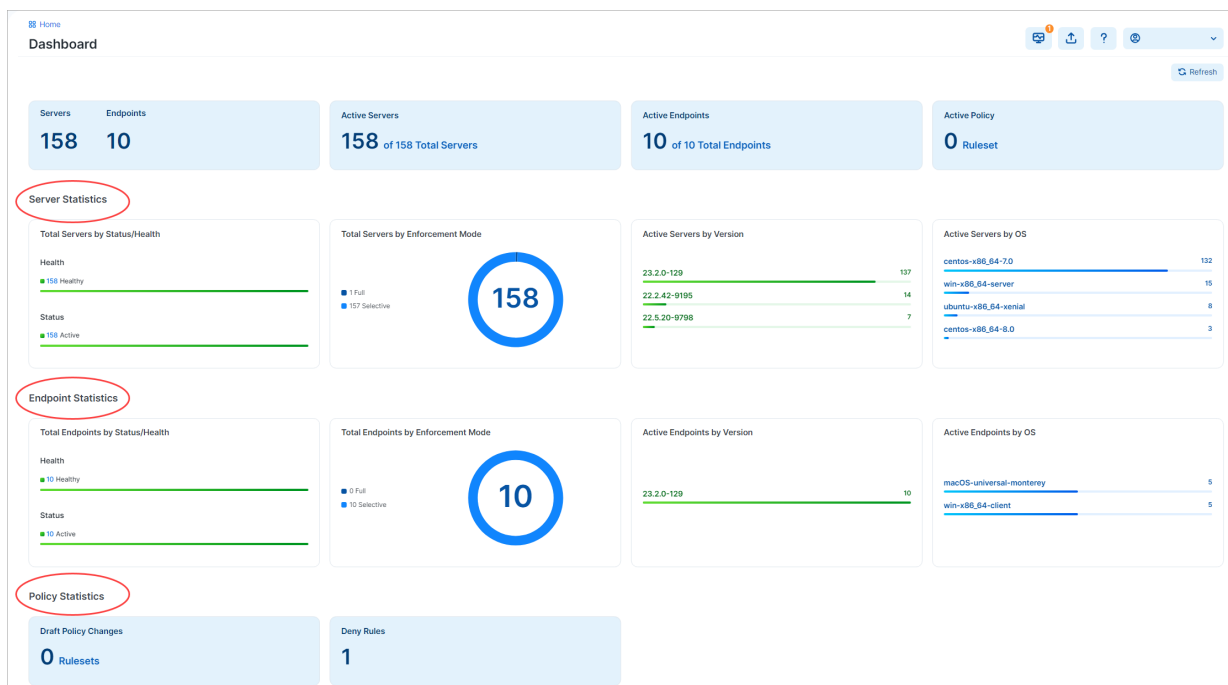
**NOTE**

One of the following global user roles are required to use the Ransomware Protection Dashboard:

- Global Org Owner
- Global Administrator

Working with the Servers and Endpoints Dashboard

To access the Servers and Endpoints Dashboard, click Dashboard > Servers and Endpoints Dashboard in the left navigation.



Servers and Endpoints Statistics

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

The Server Statistics and Endpoint Statistics sections of the Dashboard present several widgets summarizing statistics and status. You can click hyperlinks on widgets to redirect to relevant areas of the product, often with the appropriate filters applied to the target resource.

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

Policy Statistics

The Policy Statistics section includes widgets detailing the following:

- The number of Draft Policy Changes. Clicking the widget redirects to the Draft Changes page, filtered by rulesets.
- The number of Deny Rules. Clicking the widget redirects to the Deny Rules page.

Ransomware Protection for Servers Dashboard

The Ransomware Protection dashboard provides broad visual information about ransomware protection readiness, risk exposure, and protection coverage.



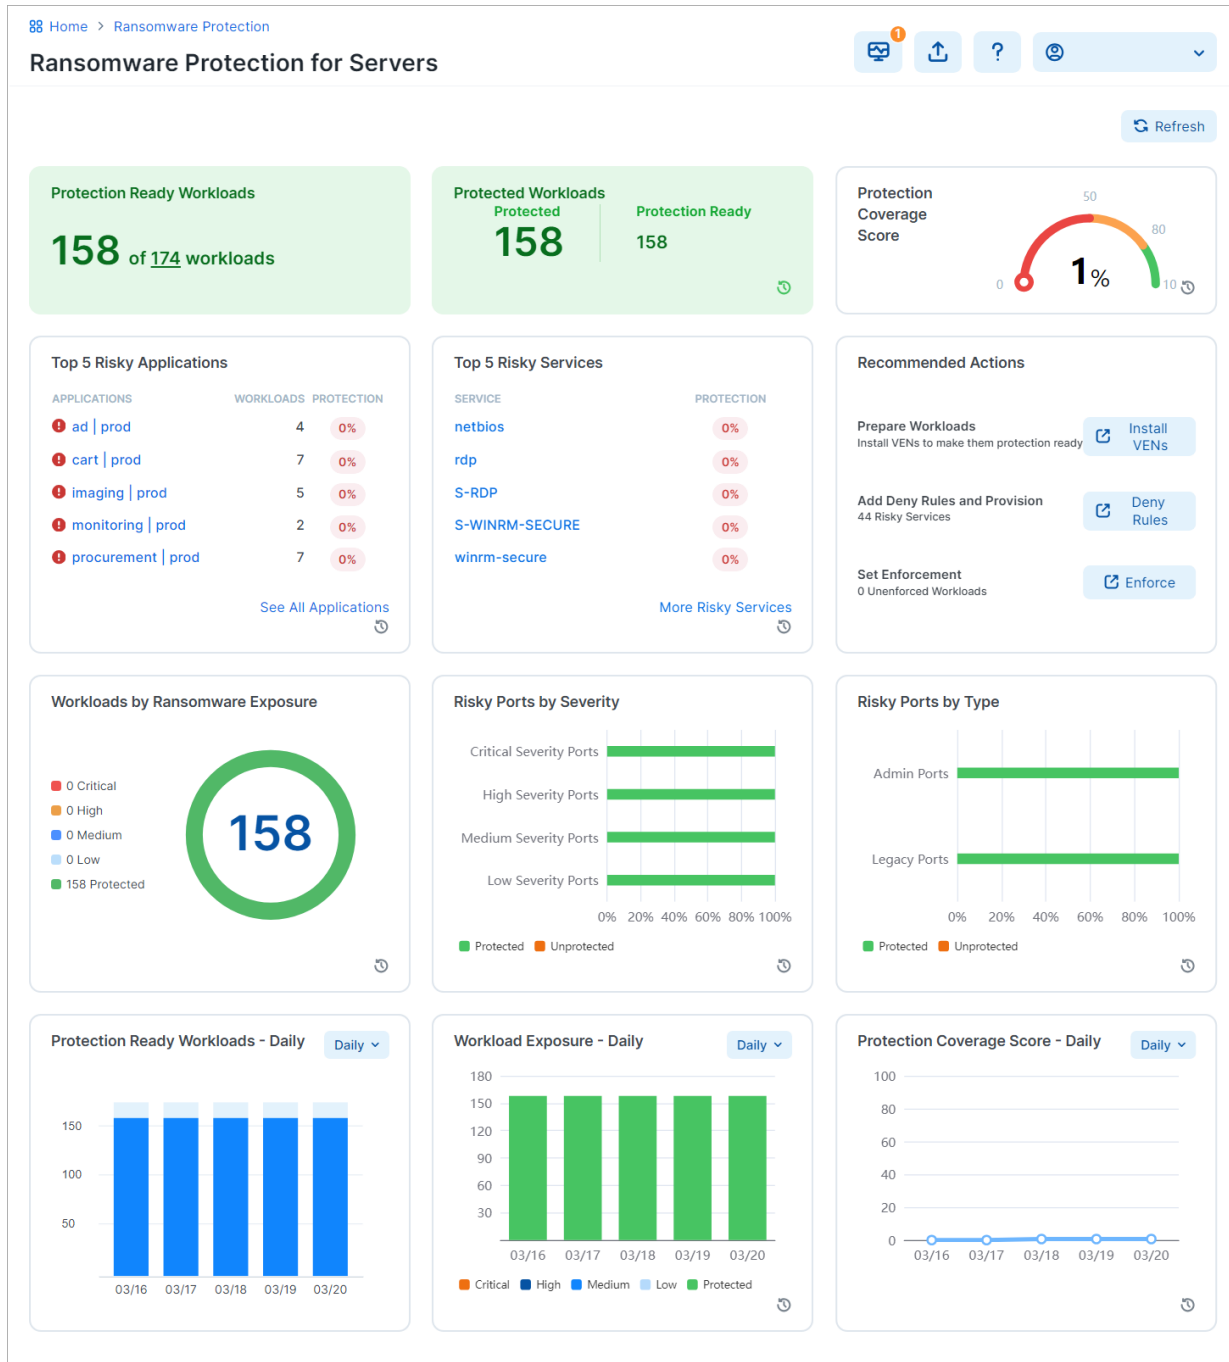
NOTE

One of the following global user roles are required to use the Ransomware Protection Dashboard:

- Global Org Owner
- Global Administrator
- Global Viewer

About the Dashboard

To access the Ransomware Protection Dashboard, click Dashboard-> Ransomware Protection in the left navigation.




Dashboard Layout

The Dashboard includes multiple columns and widgets.

Refreshing the widget information

The widgets that include small clock icons are auto-refreshed every four hours. To learn about the auto-refresh schedule, click on the clock icon. The widgets with no clock icons are refreshed when users click Refresh.



Widget color changes

Widgets change colors to show the percentage of the achieved coverage:

- Red: indicates coverage between 0 and 50%
- Yellow: indicates coverage between 50% and 80%
- Green: indicates coverage between 80% and 100%

Getting more information from the Dashboard

Click **Info** (?) to learn about the Dashboard functions.

Widget Types


The Ransomware Protection Dashboard presents several types of widgets.

Protection Readiness Widgets


These widgets show workload protection readiness.

Protection Ready Workloads

A workload is protection-ready when there is a VEN installed on it and it is configurable to enforce Illumio security policies.



You can change the number of workloads requiring protection by clicking the underlined number on the widget and entering a new target number of workloads. This non-underlined number indicates the number of such workloads that are protection-ready.



Protection Ready Workloads over time

This widget provides a chart showing the number of Protection Ready workloads over 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%).

You can view protection readiness over time: Daily, Weekly, Monthly, and Quarterly.



Workload Protection Exposure Widgets

These widgets display information about the workloads protection exposure.

Protected Workloads

A workload is protected when it has policies on all the ransomware-risky services / ports and the policies are enforced.




To be counted as a protected workload, the workload must be in Selective Enforcement or Full Enforcement mode.

In this example, out of 423 that are protection-ready, 182 workloads are protected. Because the percentage of protected workloads is 43%, the widget color is light red.

Workloads by Ransomware Exposure

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

A workload is assessed according to its exposure to the services commonly exploited by ransomware. For a list of these services, see the [Ransomware-Risky Services Table \[83\]](#).




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

- The service is blocked by enforcement boundary in Selective Enforcement, or
- The workload is in Full Enforcement, regardless of whether there is or is not a rule for that service.

Workloads Exposure Over Time

The Workload Exposure widget shows, over the selected time period, the percentage of existing workloads that are or are not protected from the ransomware. Unprotected workloads are further grouped in their exposure categories as Critical, High, Medium, and Low.

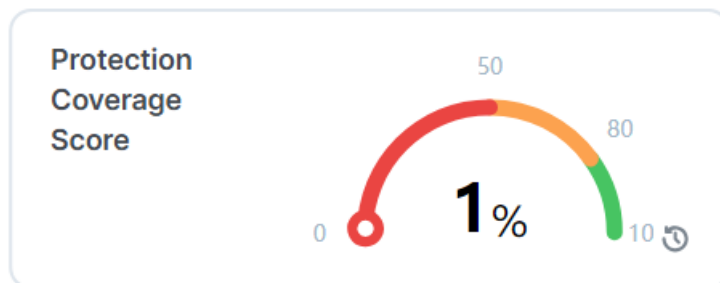


Protection Coverage Widgets

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.



Protection Coverage Score over Time

This widget displays the percent of the ransomware protection coverage over a time period: Daily, Weekly, Monthly, and Quarterly. In each case, it displays the last data point of the period.

To help illustrate the protection coverage trends, five percentage data points are used: 20%, 40%, 60%, 80%, and 100%.

When you mouseover the widget, the pop-up shows ransomware protection during the target period.




Table for 10 total address spaces:

Enforcement Mode	Policy	blocked_peer_set_count	Coverage %
Selective Enforcement	No deny or allow	0	0%
	allow (no deny)	0	0%
	Deny	10	100%
	Deny and allow	5	50%
Full Enforcement	No allow rules	10	100%
	Allow	5	50%

Weight assigned for protection coverage score:

Protection	Weight assigned
Critical	40
High	30
Medium	20
Low	10

Coverage score example

Protection coverage score calculation for four ports

Ports						Policy	Idle	Visi- bili- ty	Selec- tive En- force- ment	Full En- force- ment
SMB	S-SMB	TCP	445	Criti- cal	40	No rules	Un- pro- tected	Un- pro- tected	0	100%
VNC	S-VNC	TCP	5900	High	30	Deny rules	Un- pro- tected	Un- pro- tected	100%	100%
POP3	S- POPV3	TCP	110	Low	10	Al- lowed rules	Un- pro- tected	Un- pro- tected	0	50%
FTP Data	S-FTP- DATA	TCP	20	Medi- um	20	Deny rules and al- low rules	Un- pro- tected	Un- pro- tected	50%	50%
Protection Coverage Score							0%	0%	40%	85%

According to the table above, here is how the protection coverage was calculated:

- Selective Enforcement = $(40 * 0 + 30 * 100\% + 10 * 0 + 20 * 50\%) / (40+30+10+20) = 40\%$
- Full Enforcement = $(40 * 100\% + 30 * 100\% + 10 * 50\% + 20 * 50\%) / (40+30+10+20) = 85\%$

Top 5 Risky Applications and Services

This section provides a summary of risky applications and services.

Risky Applications

This widget displays the top 5 riskiest applications in your environment. Application risk is based on the Protection Coverage Score that appears in the App Group List. For an app-level risk assessment and remediation recommendations, click any application in the list to redirect to its details page.

For an app-level risk assessment and remediation recommendations, click any application in the list to redirect to its details page.

Top 5 Risky Applications		
APPLICATIONS	WORKLOADS	PROTECTION
! app_525 env_525 loc_525	3	0%
! app_6066 env_6066 loc_...	4	0%
! app_6343 env_6343 loc_...	3	0%
! app_7991 Production Am...	1	0%
! App2 Staging London	3	0%
See All Applications ↻		

Risky Services

The PCE automatically assigns default ransomware protection settings on certain services deemed to be at risk of ransomware penetration and lateral movement. These services and their default risk assessment are listed in the [Ransomware-risky services table \[83\]](#). Based on this default risk assignment, the top 5 riskiest services in your environment are displayed in a dedicated widget on the Ransomware Dashboard.

Click any service in the list to redirect to its details page. From there you can edit or remove the service, or navigate up one level to add new services.

Top 5 Risky Services	
SERVICE	PROTECTION
IST Common POPv3	0%
IST Common SSH	0%
IST Common Telnet	0%
IST Microsoft LLMNR (port based)	0%
IST Microsoft SMB	0%
More Risky Services ↻	

To address the unique conditions in your environment, you can change the default ransomware risk assessment on a per-service basis by going to **Policy Objects > Services** and changing the Severity as shown in the following image.

Home > Policy Objects > Services

S-MDNS

MODE: EDIT

Save Cancel

GENERAL

* Name: S-MDNS

Description: Multicast DNS

RANSOMWARE PROTECTION

Severity: Medium

* OS Exposure: Linux x Windows x

* Port Type: Legacy

Ransomware-Risky Services Table

The list of services at risk of ransomware penetration and lateral movement is provided in this table to help you assess ransomware exposure on your Enterprise Services. All new organizations created after the release Illumio Core 23.2 have services created and tagged with the metadata as system default. Organizations created before the release Illumio Core 23.2 with services that have exact match of protocol and port numbers will be tagged with the ransomware risk metadata.

Customers should work with Illumio Support to review and revise their services objects to match the list below for accurate assessment.

Service	Service Name	Protocol	Port Number	Severity	Category	OS
HTTP	S-HTTP	TCP	80	Medium	Legacy	Linux, Windows
LLMNR	S-LLMNR	UDP	5355	Medium	Legacy	Linux, Windows
NFS	S-NFS	TCP/UDP	2049	Medium	Admin	Linux
RDP	S-RDP	TCP/UDP	3389	Critical	Admin	Windows
MSFT RPC	S-RPC	TCP	135	Critical	Admin	Linux, Windows
SMB	S-SMB	TCP/UDP	445	Critical	Admin	Linux, Windows
SSH	S-SSH	TCP/UDP	22	Medium	Admin	Linux
WinRM	S-WINRM	TCP	5985	Critical	Admin	Windows
WinRM Secure	S-WINRM-SECURE	TCP	5986	Critical	Admin	Windows
FTP Data	S-FTP-DATA	TCP	20	Medium	Legacy	Linux, Windows
FTP Control	S-FTP-CONTROL	TCP	21	Medium	Legacy	Linux, Windows
METASPLOIT	S-METASPLOIT	TCP/UDP	4444	Low	Legacy	Linux, Windows
Multicast DNS	S-MDNS	UDP	5353	Medium	Legacy	Windows
NetBIOS	S-NETBIOS	UDP	137, 138	High	Legacy	Windows
		TCP	137, 139			
POP3	S-POP3	TCP	110	Low	Legacy	Linux, Windows
PPTP	S-PPTP	TCP/UDP	1723	Low	Legacy	Linux, Windows
SSDP	S-SSDP	UDP	1900	Medium	Legacy	Windows
SunRPC	S-SUNRPC	TCP/UDP	111	Low	Legacy	Linux
TeamViewer	S-TEAMVIEWER	TCP/UDP	5938	High	Admin	Linux, Windows
Telnet	S-TELNET	TCP/UDP	23	Medium	Admin	Linux, Windows
VNC	S-VNC	TCP/UDP	5900	High	Admin	Linux, Windows


Service	Service Name	Protocol	Port Number	Severity	Category	OS
WSD	S-WSD	TCP/UDP	3702	Medium	Legacy	Windows

Risky Ports Widgets

These widgets illustrate risky ports in your environment.

This widget shows the percentage of ransomware-risky ports in your environment according to their level of severity (Critical, High, Medium, and Low). Each category of risky ports has a different total on each workload, and therefore, across the environment.

To illustrate the protection coverage by severity, five percentage data points are used: 20%, 40%, 60%, 80%, and 100%. Colored bars depict the percentage of protected (green) and unprotected (orange) ports.




Risky Ports by Type

This widget shows the percentage of ransomware-risky ports in your environment by type, administrative or legacy.

To help illustrate the protection coverage by port type, five percentage data points are used: 20%, 40%, 60%, 80%, and 100%.

Colored bars depict the percentage of protected (green) and unprotected (orange) ports.



Recommended Actions Widget

This widget presents links for securing your workloads so that you can more easily address the risks revealed in the other widgets.

Recommended Actions

Prepare Workloads
Install VENS to make them protection ready

[Install VENS](#)

Add Deny Rules and Provision
44 Risky Services

[Deny Rules](#)

Set Enforcement
0 Unenforced Workloads

[Enforce](#)

Workload Ransomware Protection for Server 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 \[74\]](#) 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.							
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. 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.

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 two main APIs are: `time_series` and `risk_summary`. To learn about APIs used to power the Ransomware Protection Dashboard, see "Ransomware Protection Dashboard APIs" in the REST API Developer Guide.

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 the 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.

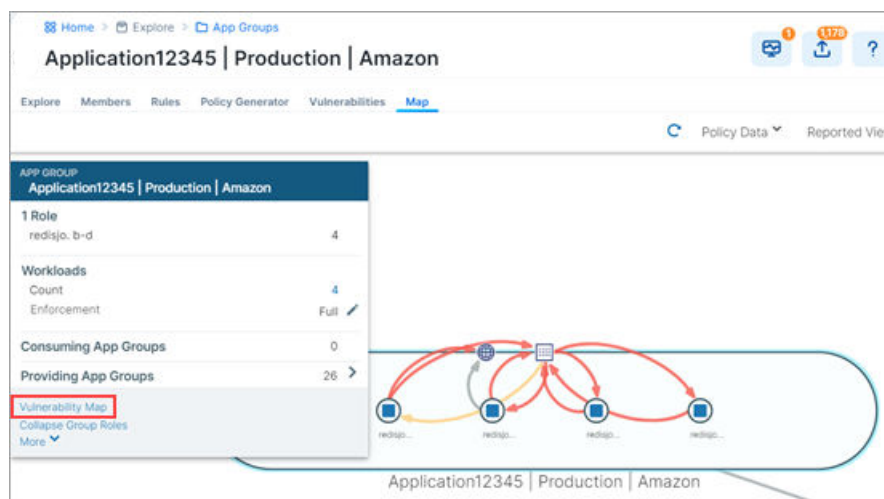
Work with Vulnerability Maps

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 **App Groups > Map** as described in [View and Mitigate Vulnerabilities \[90\]](#).



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.

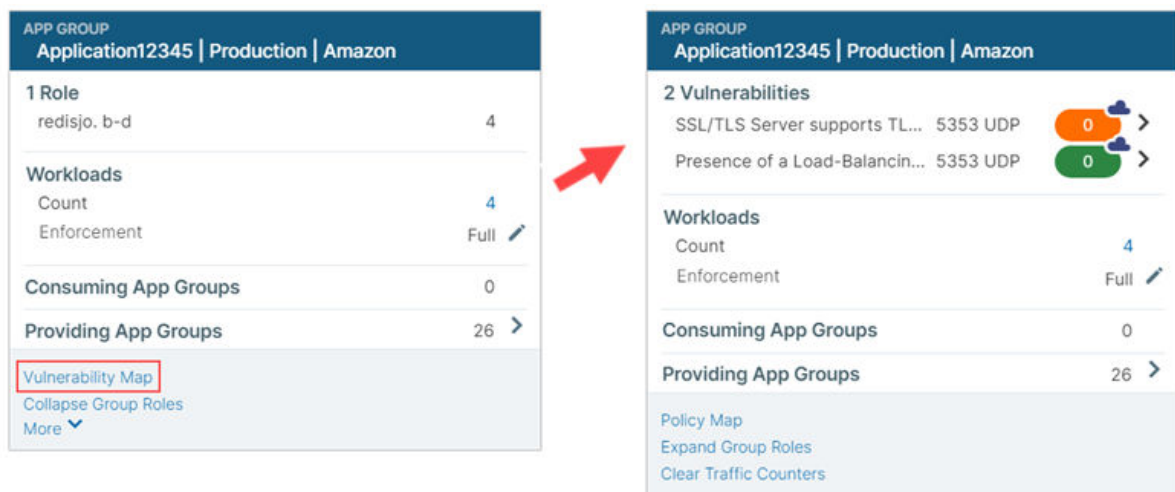
To view and mitigate vulnerabilities:

1. In the left navigation, under **Explore**, click App Groups.
2. Click an **App Group** in the list.
3. Click the **Map** tab.



4. On the Map, click on the App Group.
5. From the command panel, click **Vulnerability Map**.

The command panel shows the different vulnerability exposure scores for the selected App Group based on the ports, protocols, and workloads to which it is exposed. It is overlaid with the App Group Map. You see the Destination and Source App Groups and the vulnerable applications that are being accessed.



NOTE

The Cloud icon denotes Northern Exposure.

6. 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 includes settings for viewing Vulnerability Exposure Score and Traffic. Use the slider to change the criticality of the vulnerabilities you want to view.

TRAFFIC LINKS

- ☒ Standard Services
- ☐ Custom Services ([Edit](#))

☒ Intra-Group ☒ FQDN

☒ Internet ☒ IP List

☐ Broadcast ☐ Multicast

☒ ICMP

VULNERABILITY EXPOSURE SCORE

- ☒ All Vulnerabilities
- ☐ Exposed Vulnerabilities

VULNERABILITY TRAFFIC


Info Low Medium High Critical

☒ Blocked Vulnerable Traffic

WORKLOADS

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

7. After identifying the vulnerabilities, you can mitigate them by writing a security policies to reduce the risk to your datacenter .
- a. Click the Policy Generator tab above the Map.



- b. In Policy Generator, click either **Replace Intra-Scope Rules** or **Start with Intra-Scope**. 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.

Auto Level allows you to 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:

- **Reduced:** Exposure to the port is minimized to a reduced set of workloads, which still keep your applications up and running.
- **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.

- Complete the fields in the Policy Generator wizard.

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

- **Before includes:** Current provisioned policy
- **After includes:** All draft policy



NOTE

- Click Save after reviewing your policy.

Vulnerabilities Tabs

A Vulnerabilities Tab is provided in the **Workload Details** page, the **App Groups** page, and in the **Map**. In each location the tab displays risks due to vulnerabilities. The workload or App Group with the most vulnerabilities appears at the top of the list. You can sort the V-E score column by vulnerability score. You can then define your patch priority based on the most critical score.

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). All vulnerabilities for the workload are sorted in order of severity. The following information is provided:

- 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. This score is recalculated whenever the rules associated with the workload are changed.
- Internet exposure
- Type of traffic on that port
- Name of the vulnerability

In Workload Details

- Click **Workloads** in the left navigation.
- Click a workload in the Workloads List page.
- Click the **Vulnerabilities Tab** on the workload's details page.

feqe-data-exp-train4-vm25							
Summary Processes Rules Deny Rules Blocked Traffic Vulnerabilities Ransomware Protection							
						Total V-E Score 52	
V-E Score	Vulnerability Score	E/W Exposure	Northern Exposure	Provided Traffic (Reported)	Port/Protocol	CVE-IDs	Name
3.2	6.9	1	☁	Blocked	22 TCP	CVE-2013-2566 CVE-2015-2808	SSL/TLS use of weak RC4 cipher
3.2	6.9	1	☁	Blocked	22 TCP	CVE-2016-2183	Birthday attacks against TLS cipher...
3.2	6.9	1	☁	Blocked	22 TCP		SSL/TLS Server supports TLSv1.0
3.2	6.9	1	None	None	68 UDP	CVE-2016-2183	Birthday attacks against TLS cipher...
3.2	6.9	1	None	None	68 UDP	CVE-2013-2566 CVE-2015-2808	SSL/TLS use of weak RC4 cipher
3.2	6.9	1	None	None	68 UDP		SSL/TLS Server supports TLSv1.0
3.2	6.9	1	None	None	21 TCP		SSL/TLS Server supports TLSv1.0

On the Workload Details page, the Processes tab shows the V-E score of each process that is communicating over the network port.

feqe-data-exp-train4-vm25		
Summary	Processes	Rules Deny Rules Blocked Traffic Vulnerabilities Ransomware Protection
<i>i</i> These are processes that are listening for incoming connections - Last updated at 10/26/2023, 12:38:50		
V-E Score	Process Name	Process path
11	sshd	/usr/sbin/sshd
11	systemd-networkd	/lib/systemd/systemd-networkd

In App Groups

1. Click **App Groups** in the left navigation.
2. Click an App Group in the App Groups List page.
3. Click the **Vulnerabilities Tab** on the App Group's details page.

kafka1 | Development | Amazon

ExploreMembersRulesPolicy GeneratorVulnerabilitiesRansomware ProtectionPREVIEW

Total V-E Score 56

Customize columns							50 per page	1 - 13 of 13 Total
V-E Score	Vulnerability Score	E/W Exposure	Northern Exposure	Workloads	Port/Protocol	CVE-IDs	Name	
21	7.8	19		1	52828 TCP		Web Server HTTP Protocol Versions	
14	6.9	19		1	52828 TCP		SSL/TLS Server supports TLSv1.0	
14	6.9	19		1	31337 TCP		SSL/TLS Server supports TLSv1.0	
2.2	3.7	19		1	31337 TCP		Presence of a Load-Balancing Device Detected	
2.2	3.7	19		1	8383 TCP		Presence of a Load-Balancing Device Detected	
2.2	3.7	19		1	123 UDP		Presence of a Load-Balancing Device Detected	
N/A	7.8		None	1	9092 TCP		Web Server HTTP Protocol Versions	
N/A	6.9		None	1	26188 TCP		SSL/TLS Server supports TLSv1.0	
N/A	6.9		None	1	25 TCP		SSL/TLS Server supports TLSv1.0	
N/A	3.7		None	1	9092 TCP		Presence of a Load-Balancing Device Detected	
N/A	3.7		None	2	25 TCP		Presence of a Load-Balancing Device Detected	
N/A	3.7		None	1	46755 TCP		Presence of a Load-Balancing Device Detected	
N/A	3.7		None	2	8081 TCP		Presence of a Load-Balancing Device Detected	

In the Map

1. Click **Map** in the left navigation.
2. Use the fields and/or the Group by feature to select the objects you want to visualize.
3. Click a node on the map to open the right panel.
4. In the drop-down selector above the panel, select **Vulnerability Data** mode.

Vulnerability Data ▾

Circular Layout ▾

Reported View ▾

Policy Data

Show traffic based on Rules

✓ Vulnerability Data

Show severity and exposure of workload vulnerabilities and when traffic is inbound to a vulnerable port.

5. Click the **Vulnerabilities Tab** in the panel.

Summary	Traffic	Workloads	Vulnerabilities	Customize columns ▾				50 per page ▾	1 – 42 of 42 Total ▾	<	>
V-E Score	Vulnerability Score	E/W Exposure	Northern Exposure	Workloads	Port/Protocol	CVE-IDs	Name				
3.2	6.9	1		2	22 TCP	CVE-2013-2566	Name Does Not Match Server FQDN				
3.2	6.9	1		2	22 TCP	CVE-2015-2808	SSL/TLS use of weak RC4 cipher				
3.2	6.9	1				CVE-2016-2183	Birthday attacks against TLS ciphers with 64bit block size vulnerability (Sweet32)				
3.2	6.9	1									
3.2	6.9	1		2	22 TCP		SSL/TLS Server supports TLSv1.0				
3.2	6.9	1		2	22 TCP		SSL Certificate -				