




Illumio Core 21.5 Visualization User Guide

Published: 2024



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

Table of Contents

Legal Notice	4
Security Advisories	5
September 2024 Security Advisories	5
Ruby SAML gem component authentication bypass vulnerability	5
Severity	5
Affected Products and Patch Information	5
Resolution	5
References	6
Skipped Critical Patch Updates	6
Discovered By	6
Frequently Asked Questions	6
Modification History	7
September 2023 Security Advisories	7
Authenticated RCE due to unsafe JSON deserialization	7
Severity	7
Affected Products and Patch Information	7
Resolution	8
References	8
Skipped Critical Patch Updates	8
Discovered By	8
Frequently Asked Questions	8
Visualization	10
Illumination	10
About Illumination	10
Groups in Illumination	22
Virtual Servers in Illumination	27
Containers in the Illumio Core Maps	27
Work with Illumination	28
App Group Map	34
About the App Group Map	34
Work with the App Group Map	38
Explorer	40
About Explorer	40
Work with Explorer	45
Vulnerability Map	49
About Vulnerability Map	49
Work with Vulnerability Maps	51
Reports	56
About Reports	56
Work with Reports in the PCE	59

Legal Notice

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

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

Resources

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

Contact Information

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

Security Advisories

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

September 2024 Security Advisories

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

Ruby SAML gem component authentication bypass vulnerability

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

Severity

Critical: CVSS score is 9.9

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

Affected Products and Patch Information

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

Table 1. Products Affected by the Security Vulnerability

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

Resolution

Upgrade to the latest release for a given major version.

References

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

Skipped Critical Patch Updates

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

Discovered By

External Security Firm

Frequently Asked Questions

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

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

Modification History

- September, 2024: Initial Publication of CVE

September 2023 Security Advisories

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

Authenticated RCE due to unsafe JSON deserialization

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

Severity

Critical: CVSS score is 9.9

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

Affected Products and Patch Information

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

Table 2. Products Affected by the Security Vulnerability

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

Resolution

Upgrade to the latest release for a given major version.

References

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

Skipped Critical Patch Updates

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

Discovered By

External Security Firm

Frequently Asked Questions

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

- How long will the upgrade take?

The fix will be provided in a normal code release so this will take the same amount of time as any PCE upgrade.

- Were any Illumio customers impacted by this vulnerability?

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

Visualization

Illumination

Illumination provides rich dashboard capabilities that monitor interactions from unauthorized hosts, alerting you when suspicious behaviors occurs. It provides a visual “log” of every attempted and successful communication between managed workloads. It visualizes the communication between workloads and applications so that you can see any network traffic that is not covered by your rules.

About Illumination

Illumination provides a unique new way to reveal the traffic flows in your network and to help you configure policies to secure your applications.

How the Illumination Map Works

Illumination maps the outbound connections from workloads to unknown IP addresses to fully qualified domain names (FQDNs) or DNS-based names. For example, Illumination 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 address as seen by the workload. The FQDNs are also used in Policy Generator to propose DNS-based rules and are displayed in Explorer.

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











NOTE





You can take a static snapshot of existing, long-lived connections (for example, connections that must be active for six or twelve hours at a time) before pairing the workloads with the VEN. Any in-progress connections are captured to track the connection state, providing a static snapshot of established connections, including port and protocol information. In the Illumination map, these connections display in red until new matching traffic is observed by the VEN.

How to read the illumination map







Legend

ata ▾ Reported View ▾ Filter ▾ Legend ▾

WORKLOADS		GROUPS	
	Visibility Only		Visibility Only
	Selective Enforcement		Selective Enforcement
	Full		Full
	Unmanaged		
	Idle		

CONTAINER WORKLOADS	VIRTUAL SERVICES
	Visibility Only
	Selective Enforcement
	Full
	 Node

TRAFFIC LINKS

-  Loading rule data
-  Rules not calculated
-  Blocked
-  Allowed
-  Potentially Blocked
-  Cannot write Rules for Discovered Groups

Pay attention at the following:

- Workloads and groups inside full dark lines depict the **Full** enforcement mode.
- Workloads and groups inside dotted lines depict the **Selective** enforcement mode.
- **Visibility only** enforcement mode is depicted with a *light* full line and with no circle around it.

Traffic links are presented with arrows in different colors:

- **Green:** Traffic is allowed
- **Yellow:** Traffic is potentially blocked
- **Red:** Traffic is blocked
- **Grey:** Rules are not calculated
- **Dotted line:** Rules cannot be written for the discovered groups
- **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.

Workflow for Using Illumination Map

Illumination enables 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. See the *VEN Installation and Upgrade Guide* for Core 21.2.3 VEN for information about installing (also called pairing) VENs on workloads.
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). See "IP Lists" and "Rule Writing" in the *Security Policy Guide* for more information.
4. **Rule Testing:** Illumination gives you the power to test and evaluate your rules against existing traffic flows *without* enforcing the rules. Rules can be tested to ensure that legitimate traffic flows required by an application are permitted and malicious traffic is blocked. Exporting traffic summaries or using blocked traffic lets you know which traffic connections would be dropped if the rules were enforced. .
5. **Policy Enforcement:** When you are ready to implement the rules for a group, you can put the group into the enforced state. Leveraging Illumio's 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.

Illumination View Levels

Depending on the number of workloads you have paired, the Full map is displayed for 60 or fewer workloads and the Location view is displayed for more than 60 workloads.

The Illumination map provides four levels of detail, depending on how many workloads are in your environment:

- **Full map:** Displays up to 60 workloads
- **Location view:** Displays groups by location
- **Group view:** Displays all groups within a selected location group
- **Detail view:** Displays details about a selected group

If your organization has fewer than 300 workloads, the PCE does not cache data and all maps and views are current. If your organization has more than 300 workloads, the group and location views are cached for 5-60 minutes, depending on complexity (as determined by the number of workloads, flows, and IP lists), but the Detail view is always current.

**NOTE**

DHCP (UDP ports 67 and 68) and DNS (TCP/UDP port 53) traffic is not displayed in the Illumination map. DHCP and DNS traffic is allowed implicitly, so rules to allow this type of traffic are not required. DHCP or DNS traffic is included in the streaming flow summaries.

Full Map

When the PCE manages less than 60 workloads, you see them all in the Illumination map.

**NOTE**

The Role label filter is only available in the Full map.

If you have more than 60 Workloads, the following views are also available:

Location View (Groups by Location)

When you first view the map, you see the Location view, which shows all groups organized by their Location label. When you click a specific location, the Illumination map zooms in to display more detailed information. A cached view of the data is used, so when the Illumination map contains fewer than 60 workloads, the data displayed in the map is current. If the map has more than 60 workloads, the time required to refresh the display increases incrementally with each workload, so the wait time for the display might be longer.

In the Global view, circles represent groups that share the same Location label. Summaries in the circle indicate the number of groups in the location and the total number of workloads in all the groups. Up to 75 groups are displayed and up to 2,000 groups can be displayed per location.

**NOTE**

The Location view doesn't display workloads that are "discovered," meaning, they don't have any labels or they have labels except for the Location label. These types of entities are only displayed in the Group Detail view when there is traffic to workloads in the expanded Group.

Click one of the Location groups to go to the Groups view.

Group View (Expanded Groups)

The Group view displays all of the groups inside a selected Location group. The number at the center of each group represents the number of workloads inside that group. When any of the groups are communicating with each other, you see the traffic connection between them.

**NOTE**

In the Group view, rule coverage is only calculated when you have written rules that allow all traffic between the groups. When you are using Role-to-Role rules for all group roles or all group traffic, rule coverage might not be displayed accurately.

In the Group view, a cached view of the data is displayed to minimize rendering time. As a result, the displayed data might not reflect recent changes when you have more than 60 workloads (for example, there might be a short delay for new workloads or traffic or to reflect changes to IP lists and labels). If you have more than 60 workloads, the amount of time required to update the view increases incrementally based on the number of workloads in your organization. In the Group Detail view, the data is not cached so changes are reflected almost immediately.

**NOTE**

In the Group view and the Group Detail view, traffic lines are displayed in green only if there is a rule written using group labels to allow **all traffic** between these two groups. In all other cases, traffic lines are gray to indicate that rule coverage can only be determined once both groups are expanded. After expanding the groups, each traffic line is displayed in green or red depending on whether or not that traffic is explicitly allowed by rules.

Click a group to go to the Detail view.

Detail View (One or Two Groups in Focus)

The Detail view focuses on the selected group and shows you the group's constituent workloads grouped together by Role label. If any of the workloads inside the group are communicating with workloads in another group, you can expand the connected group to view details of traffic between roles. The view will focus on both groups and changes are reflected almost immediately.

When you click the Expand Roles icon, the workloads inside expand so you can see the traffic links for each connection. After the workloads are visible, you can start writing rules to allow the traffic between selected roles within or across groups by clicking on the traffic links and selecting the **Add Rule** link to allow the traffic between selected roles.

In the Detail view, inter-group traffic lines to and from discovered groups are displayed in gray while the group is collapsed, since rule coverage cannot be calculated until the discovered group is expanded.

From the Detail view, you can view the role and hostname if available. You can view all a group's constituent workloads by expanding the Role icon. To do this, click a Role icon and select Expand Roles from the command panel. After you do this, you can view details of traffic between workloads with same or different Role labels. You can view the hostname of

the workloads along with their Role labels by zooming in on the map. When the workload does not have an assigned Role ILabel, the hostname of the workload displays.

Discovered workloads (workloads without labels) or No Location workloads (workloads without a Location label) are not displayed in the Location or Group maps, only in the Detail view when a connected group is selected from the Groups view.

The Illumination map also provides a timestamp for traffic flows to identify when a particular traffic flow was last detected. This timestamp is included in the exported flow records and is displayed in the Detail view in the “Last Detected” field.

**NOTE**

When a workload first reports its traffic flow summary in the Illumination map, the traffic line displays in green as long as the outbound traffic from that providing workload is allowed. However, if the consuming workload does not allow the connection, the traffic line will turn red when the consuming workload reports its flow summaries.

**NOTE**

The maximum number of workloads, after which the Workload Summary Page is displayed in Illumination is 100,000.

Reported and Draft Views

The Illumination map provides two views into your organization: the Reported and Draft views.

**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.

Reported View

The Reported view visualizes your policy coverage as reported by your workloads, so you can examine the current state of your provisioned policy. This view displays the traffic using red 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

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

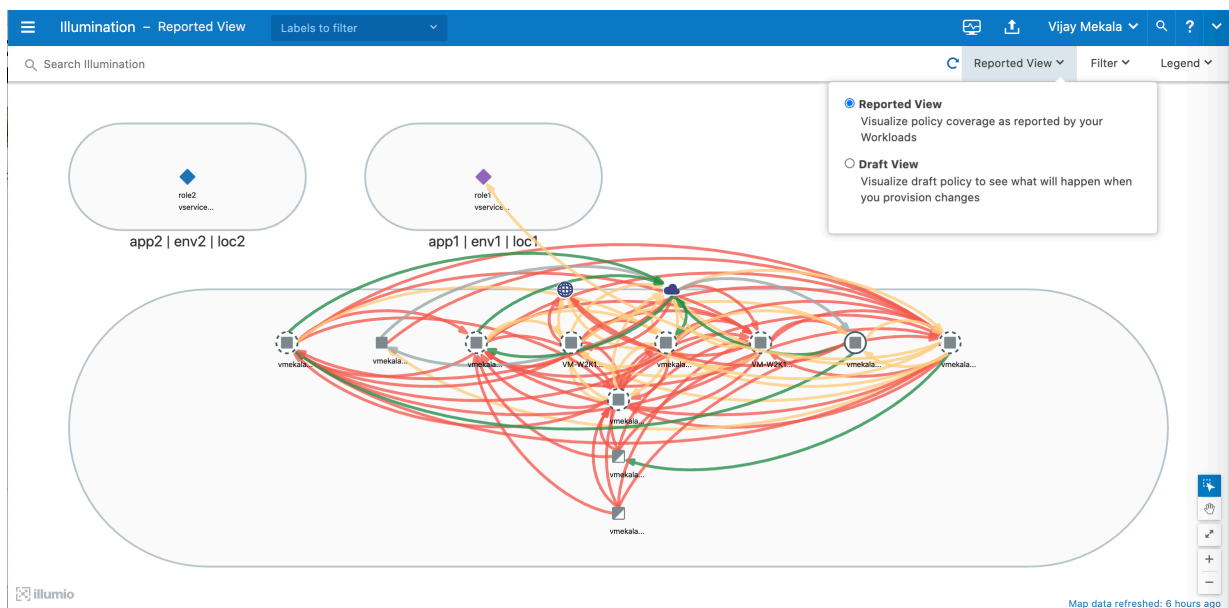
This view provides visibility for the actual traffic handling (rather than the expected traffic handling provided by the Draft view) and loads more quickly, especially when you have a large number of workloads and traffic flows.

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.

The Reported view is a view-only map. 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.

In Reported view, rule coverage (the number of connections that have been included in rules) is not supported for traffic between unmanaged workloads. The Draft view always provides accurate rule coverage for traffic between unmanaged workloads.

Reported View



Draft View

The Draft view visualizes the potential impact of your draft policy so that you can examine what will happen when you provision your changes. This view displays the traffic using red or green lines to indicate whether the PCE has a rule to allow the connection that was reported by the VEN. Specifically:

- 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.

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. This map can take longer to load than the Reported view, especially if you have a large number of workloads and traffic flows, since the PCE has to compute the expected coverage for each traffic flow.


In Draft view, you can either view the rule that would permit traffic (turning the color of the line from red to green) or add a rule to allow a specific flow. In this view, you can immediately see the impact of the latest changes to the active or draft policy as they are reflected in the color of the traffic lines.



NOTE

In Draft view, rule coverage is now supported for Windows process-based services. You can analyze the effect of the policy change and edit the rule before you provision it. Rules written for specific user groups are not included in rule coverage. The Reported view always provides accurate rule coverage for process- or user group-based rules.

Draft View



The default view in the Illumination map is the Reported view. The current view is identified in the upper left corner, next to the Select Labels to filter views field (either “Illumination - Reported View” or “Illumination - Draft View”). The different colors of the backgrounds help you quickly identify the current view.

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

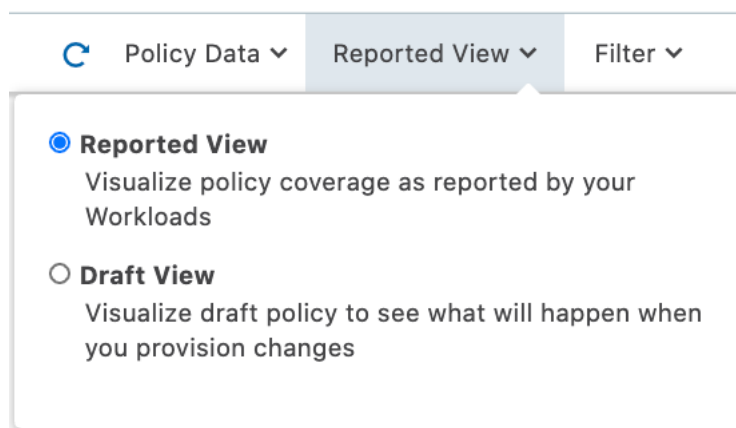
configured policy. The modeling might not work correctly for the following types of rules configured on the PCE:

- **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.

Changing Views

You can switch between the two views by selecting the view from the top right corner of the PCE web console.



Command Panel

The command panel in the Illumination map displays information about a selected traffic flow. To view the command panel, select a traffic line on the Illumination map.



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 Illumination map. The process or service name that was part of the most recently reported connection is displayed.

The command panel displays the following information:

- Services between the provider and the consumer

- Rule coverage for each service, based on the draft rules in Draft view and as reported by the VEN in the Reported view
- Providers
- Consumers
- When this traffic was last detected, based on the last detected timestamp for the selected service

The command panel displays up to 64 ports per IP address per traffic link, up to 5 overlapping IP lists per IP address, and up to 500 IP addresses per traffic link.

If there are no rules associated with this traffic flow, a link to add a rule is displayed in the Draft view.

The Reported view helps you to understand your traffic patterns. You can view all rulesets with rules that apply to the selected traffic link, but you cannot add a rule in the Reported view. If you click the View Rulesets link, the Rulesets page displays.

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, instead of 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 command panel. If there is no policy service that matches that port/protocol, either a separate entry displays for each service name reported in the flows or “Unknown” displays when the service name could not be captured (for example, an outbound flow).

While creating a service, if the traffic is to a Windows server, Illumination automatically picks a Windows-based service to create and display the process name. You have the option to add the path later. The traffic link turns green when you have a rule that uses a service that is process-based.

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. To hide entries that have potentially outdated rule coverage, use the Traffic filter “Hide traffic detected before last policy provision time.” In both cases, the Draft view shows the calculated rule coverage for port/protocol, without considering flow’s service names or rules with process name/service name based policy services.

**NOTE**

If there are multiple rules allowing traffic, only one rule is displayed in the Traffic Link panel.

Illumination Filters

In the Illumination map, you can select one of several traffic filters so you can show or hide different elements of the map and focus on what is most important to you. By default, all options are selected.

a ▾
Reported View ▾
Filter ▾
Leg

TRAFFIC LINKS

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

☒ Allowed
☒ Blocked

☒ Potentially Blocked

☒ Unknown
☒ Intra-Group

☒ FQDN
☒ Internet

☒ IP List

☐ Broadcast
☐ Multicast

☒ ICMP

Filter by Time

Anytime
Now

☐ Since Last Provision

WORKLOADS

☒ Visibility Only
☒ Idle

☒ Full Enforcement
☒ Unmanaged

☒ Selective Enforcement

[Close](#)
[Reset to Default](#)

Label Filters

This filter restricts the Illumination map to only those entities that have the labels you enter in the filter at top of the map. For example, you might want to filter the Illumination map to show only those workloads with the Web Role label.



NOTE

The Label filter does not filter the selected group. Only the connected groups are filtered. The Role Label filter is only available in the Full map.

Workload Enforcement Filter

Using the Workloads section of the filter select to see only those workloads that are in a particular enforcement state. For example, you might want to see only those workloads that are in Full Enforcement, so you can write rules for them.

Traffic Links Filter

You can also filter the Illumination map to show or hide network traffic. For example, you can hide or show all traffic inside a group, traffic among groups, or even hide *all* the traffic that is currently being allowed by your security policy.

You can also select to show or hide traffic generated by specific services. For example, you might not want to see traffic for common running services that would clutter the Illumination map with traffic not relevant to your policy.

You can also choose to display only recent traffic (as defined by the timestamp). For example, you can hide all traffic that has not been observed since the last time a policy was provisioned by checking the Hide traffic detected before last policy provision time checkbox.



NOTE

This filter is only available in the Reported view. For more information, see [Reported and Draft Views \[15\]](#).

The Traffic Volume slider at the bottom allows you to filter by the number of traffic connections. This filter is only available in the Full map. When other filters are selected, the number of traffic connections displayed by the slider does not change.

Illumination Impact on Workloads

Illumination provides rich visibility into your workloads (traffic flows and running services) without blocking traffic, so you can build and test policies before you enforce them. Illumination implements industry standard security settings on those workloads in Illumination mode. The following table describes how Illumination impacts your workloads.

Category	Impact
What gets installed	<p>Windows and Linux</p> <p>Illumio daemons are started and running. An additional daemon is started when SecureConnect is enabled. An Illumio user is created to run some of the daemons.</p> <p>Linux</p> <p>iptables kernel module is installed (if not already on the target system) and rules are installed in the kernel iptables. Any existing iptables rules are removed.</p> <p>Windows</p> <p>For Windows workloads that are paired with the PCE, the VEN takes control from Windows Firewall and install a Windows Filtering Platform (WFP) callout driver. Illumio rules are added as WFP filters and the WFP callout driver captures information about traffic that matches the WFP filters.</p>
Changes to target Workload's global security settings	<p>The following basic security settings are implemented on workloads in Visibility Only:</p> <ul style="list-style-type: none"> • syncookie for TCP DOS/Scan Protection is enabled • Sysctl IP forwarding is turned off • SMURF sysctl attack protection is enabled • A persistent (ultra low bandwidth) TCP connection to the PCE for lightning bolts is established

Category	Impact
Traffic that is dropped	<p>The following traffic is dropped when you pair a workload in Visibility Only:</p> <ul style="list-style-type: none"> • Late RST, RST-ACK, FIN-ACK packets tagged as INVALID • Illegal TCP packets with illegal flag combination (SYN/FIN, SYN/RST, FIN/PSH/URG, FIN w/o ACK, NULL, X-MAS (all flags) are dropped
Traffic that is Allowed	<p>The following traffic is allowed in Visibility Only:</p> <ul style="list-style-type: none"> • Inbound/outbound multicast traffic is allowed. • Inbound/outbound non-IP protocols are allowed.
IPv6	<p>All IPv6 traffic is allowed by default but can be blocked. For more information, see “Allow or Block IPv6” in the PCE Administration Guide.</p>
SecureConnect	<p>When enabled, SecureConnect creates IPsec connections between workloads in a ruleset’s rules.</p> <p>When you have SecureConnect enabled on a workload, the following applies:</p> <ul style="list-style-type: none"> • On Linux, SecureConnect fails when the VEN detects a conflicting IPsec process. • On Windows, the VEN takes over all Connection Security Rules on the host.

Groups in Illumination

Groups in the Illumination 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 Illumination map after you pair workloads. For information about installing (also called pairing) VENs on workloads, see the “VEN Installation and Upgrade Guide” for Core 21.2.3 VEN.

Illumination Group Detail Levels

You can choose one of three levels of detail in Illumination for enforced workloads in a group.

These levels allow you to control how much data the VEN collects from a workload when enforced, so you can control resource demands on workloads:

- **High detail:** The VEN collects connection details (source IP, destination IP, protocol and source port and destination port). This option applies to both allowed and blocked connections. This option provides rich Illumination detail but requires some system resources from a workload.
- **Low detail:** The VEN only collects the blocked connection details (source IP, destination IP, protocol and source port and destination port), including all packets that were dropped. This option provides less Illumination detail but also demands fewer system resources from a workload than high detail.
- **No detail:** The VEN does not collect any information about traffic connections. This option is only available for workloads that are in the enforced state. This option provides no Illumination detail and demands the least amount of system resources from a workload.

Types of Groups in Illumination

Once you pair workloads, the PCE analyzes the workload data reported by the VENs. Based on the traffic flows among your workloads, Illumination organizes them into groups. A group could represent an instance of an application running in your datacenter, such as an HRM

application running in the test environment in your North America datacenter; or a group could represent a web store in production with its web workloads hosted in AWS and its databases hosted in your private datacenter.

In cases where more than 100 workloads are paired, groups are displayed in different levels of detail in the Illumination map. For information about the different view levels of groups in Illumination, see [Illumination View Levels \[12\]](#).

Group enforcement modes

Groups on the Illumination map are in the following enforcement modes:

- **Full:** Rules are enforced for all inbound and outbound services. Traffic not allowed by the segmentation rule is blocked. This was previously known as the "enforced" mode.
- **Selective Enforcement:** The new enforcement mode where rules are enforced only for the selected inbound services when workload is within the scope of the Selective Enforcement rule.
- **Visibility Only:** no traffic is blocked by policy. This was previously the so-called "illuminated" mode.

Groups in Full Enforced Mode

When you are ready to enforce the rules you have written, place the group into the **Full** state. When you put a group into the Full state, all traffic flows permitted by rules are allowed and all other traffic is blocked.

The line around the group is a thick full line.

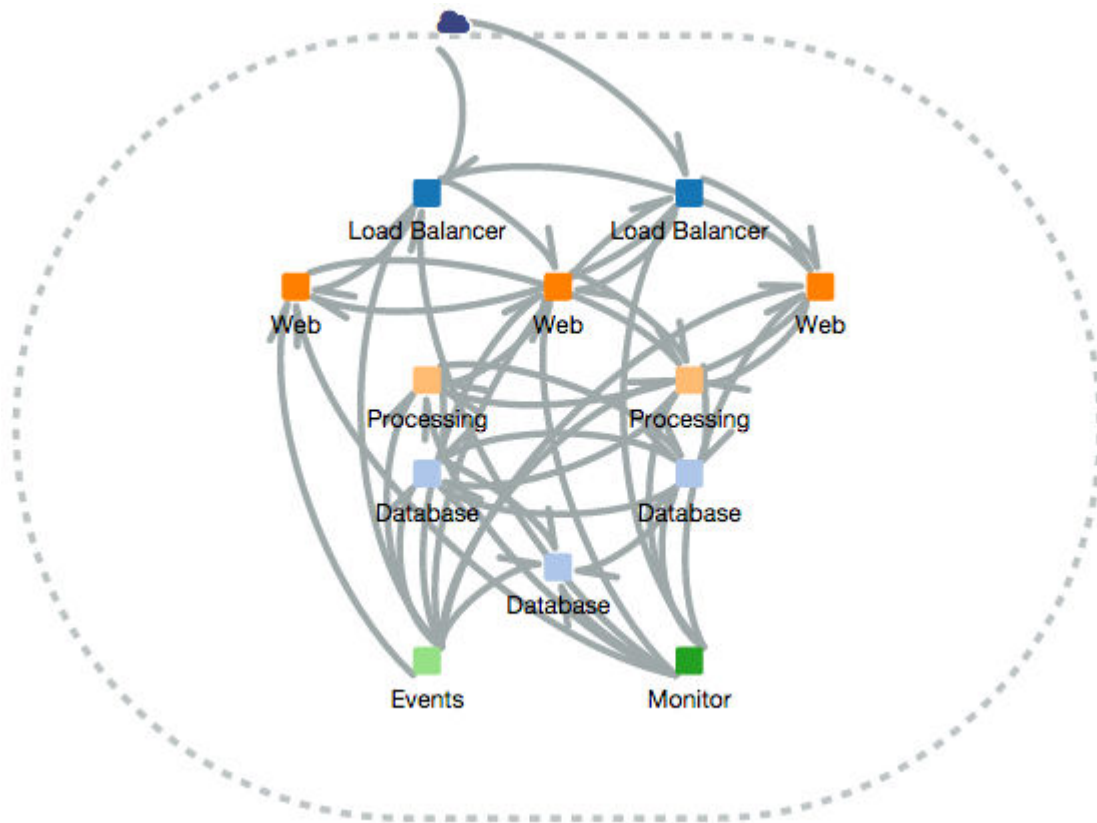
Groups in Visibility Only Mode

When you have written rules for the traffic flows in the group, you can place the group into the Visibility Only state to view all traffic that will be blocked when the group is put into the Full enforced mode. To change the enforcement for a group, select the group and select Selective Enforcement from the command panel.

In the Visibility Only mode, all traffic is still allowed, even traffic flows *not* permitted by your rules. You can view all traffic that will be blocked by going to the Blocked Traffic page and selecting the Potentially Blocked Traffic filter.

Discovered Group without rules

When a Group is first "discovered" by the PCE, its boundary is indicated by a gray dashed line and all traffic lines are gray because rules cannot be written for the group yet.



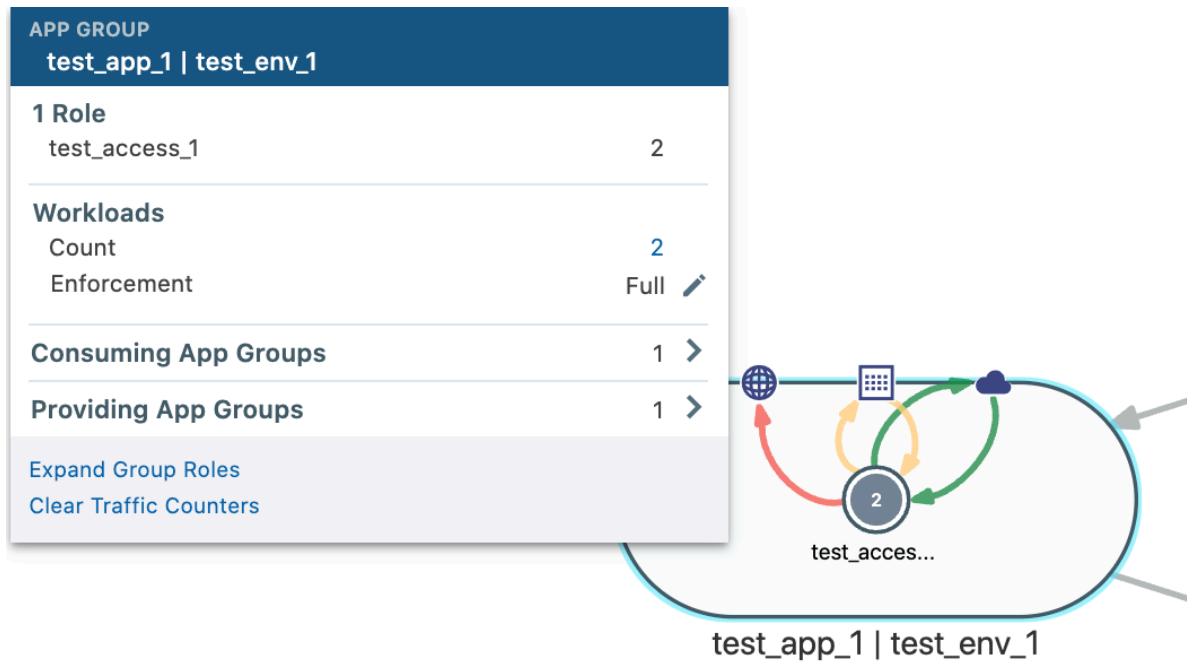
View Group Details

After you pair workloads and apply at least one label to them, the Illumination map puts all workloads that communicate together into a group.

You can view a group's details to view or change the labels assigned to the workloads, change the enforcement of the workloads, or unpair or pair new workloads.

To prepare a group for rules:

1. From the PCE web console menu, choose **Illumination Map**.
The Illumination map appears.
2. Select a group by clicking inside the group (but not on any workloads).
The command panel for the group appears.



3. The Group details page appears. It shows all the workloads that share the same scope.
4. Make changes to workloads or rules and click **Save**.

Expand or Collapse Group Roles

When you drill down into a group detail in the Illumination map, multiple workloads that share the same Role label are collapsed together to save space.

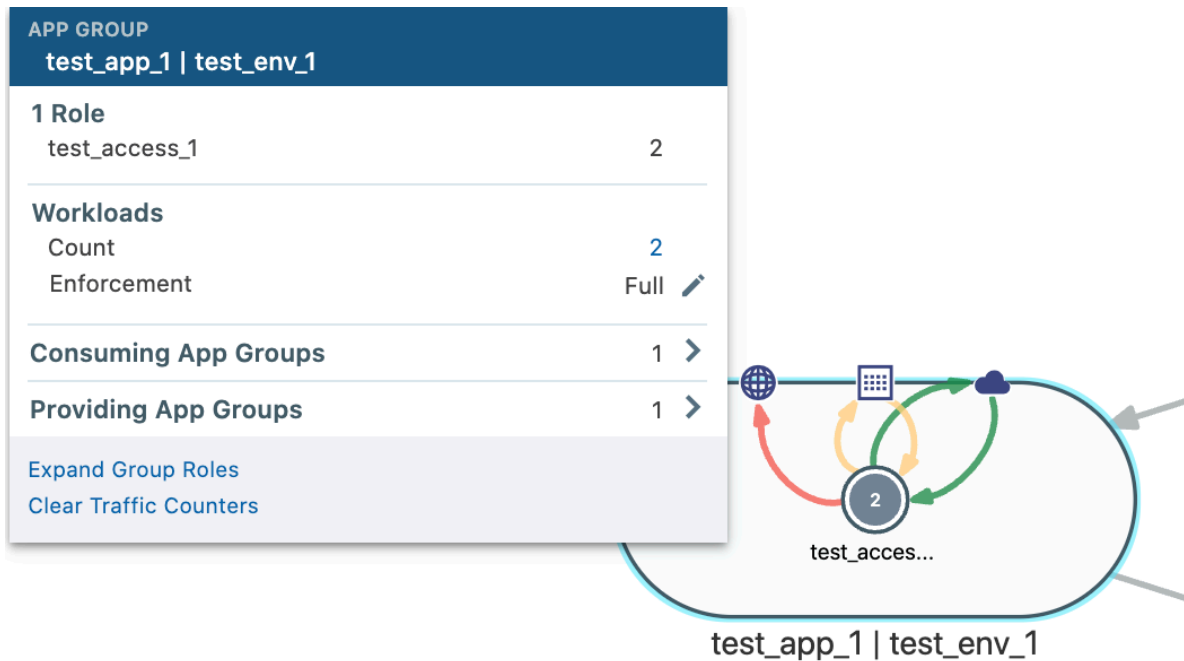
You can easily expand the workloads by selecting them and clicking **Expand Role** in the command panel.



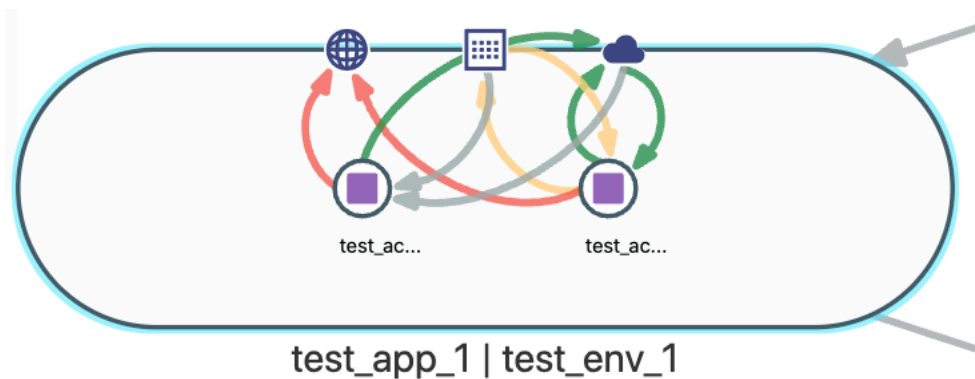
NOTE

You can expand up to 200 workloads per collapsed role and up to two roles.

1. Open the group command panel, same as above.



- Click **Expand Group Roles** to see the roles inside the group.



- Workloads that share the Role label are expanded:

Add or Remove Workload to or from a Group

In the Illumination map, you might see workloads that don't belong to a group, such as management or monitoring services that run in your network but are not relevant to the policy you want to build.

You can remove workloads from a group by simply dragging them out. Conversely, if you notice that a workload that should be included in a group but is not, you can simply drag it into the group.



NOTE

You can only add or remove a workload to or from groups that have been prepared for rule writing.

When you add a workload to a group, the workload inherits the Application, Environment, and Location labels associated with the group.

When you remove a workload from a group, the workload's Application, Environment, and Location labels are removed from the workload.

Virtual Servers in Illumination

Illumio Core supports enforcement activation on F5 Load Balancers using Local Traffic Manager (LTM) or Advanced Firewall Manager (AFM) modules on F5 BIG-IP systems. Each virtual server on a BIG-IP system is represented as a policy endpoint in the PCE, which computes policy for each virtual server and displays the virtual servers in the Illumination map, Location view, and the App Group Map.

How Virtual Servers Appear in Illumination

A virtual server is identified by a set of labels. The consumers and providers for a virtual server can be assigned different labels, which could place them in the same group or a different group in Illumination. Providers are allowed to have an incomplete label set (for example, only a Location label), so the providers can be in all groups within the specified location. As a result, a single virtual server can have providers in any group or in any number of groups in Illumination.

See "Load Balancers and Virtual Servers" in the *Security Policy Guide* for more information.

Based on their labels, the virtual servers are shown in the Full map. If the workload is in an enforced state, the traffic lines are displayed in green when the traffic is allowed by rules or in red when the traffic is blocked. To view more details, select a virtual server from the Illumination map or from the App Group Map command panel.

From the Illumination map, you can also add a segmentation rule for the incoming and outgoing traffic links from the virtual server by selecting the traffic line and clicking **Add Segmentation Rule**.

Prerequisites for Virtual Server Display

To display virtual servers in Illumination properly, the following prerequisites must be met:

- A virtual server from the server load balancer has been discovered and is under PCE management (meaning, it exists as an object in the PCE).
- There are traffic flows from a consumer to the VIP of the virtual server or there are traffic flows from the VIP to a pool member (backend server).
- Labels have been assigned to the virtual server.

Containers in the Illumio Core Maps

Illumio Core provides visibility and segmentation for containerized applications. You can use the Illumio Core maps (Illumination, the App Group, the Vulnerability Map, and Explorer) to gain visibility into your containers so that you can build a segmentation policy for your containerized applications that use OpenShift and Kubernetes.

Most of the containerized applications connect to non-containerized applications, such as core services and databases running on bare metal. The Illumination map provides visibility across different types of workloads in order to build your segmentation policy.

For more information about deploying containers, see Illumio Core for Kubernetes and Open-Shift.

View Container Workload Traffic

When you deploy containers in your organization, you can view container workloads and traffic associated with them in Illumination, the App-Group Map, the Vulnerability Map, and Explorer. Container workloads and regular workloads with same labels can be collapsed into single role and provide more granular view.

- Illumination map provides visibility into network connectivity and policy coverage for applications irrespective of the type of workloads the applications run on, containers, virtual machines, or bare metal servers.
- Containers are displayed in Illumination, App-Group Map, Vulnerability Map, and Explorer only if they are currently running.
- Containers are displayed in the maps in the same way that collapsed roles for regular workloads are displayed.
- Containers are supported in Illumination Full map as well as scaled versions of Illumination map.
- Containers are supported on the App Group Maps in the same way.
- The command panel displays container details in the same way as it does for regular workloads.
- Explorer displays the detailed flow summaries to and from containers.
- HREF for containers is included in the events and fluentd summaries streamed out of the PCE.

Work with Illumination

Illumination provides rich visibility into your workloads (traffic flows, running services) without blocking traffic, so you can build and test policies before you enforce them.



NOTE



In previous releases, this feature was referred to as “Segmentation Rulesets.” In Illumio Core 21.5.0 and later releases, this feature is now referred to as “Rulesets.” Not all image in this guide are updated to reflect this name change.

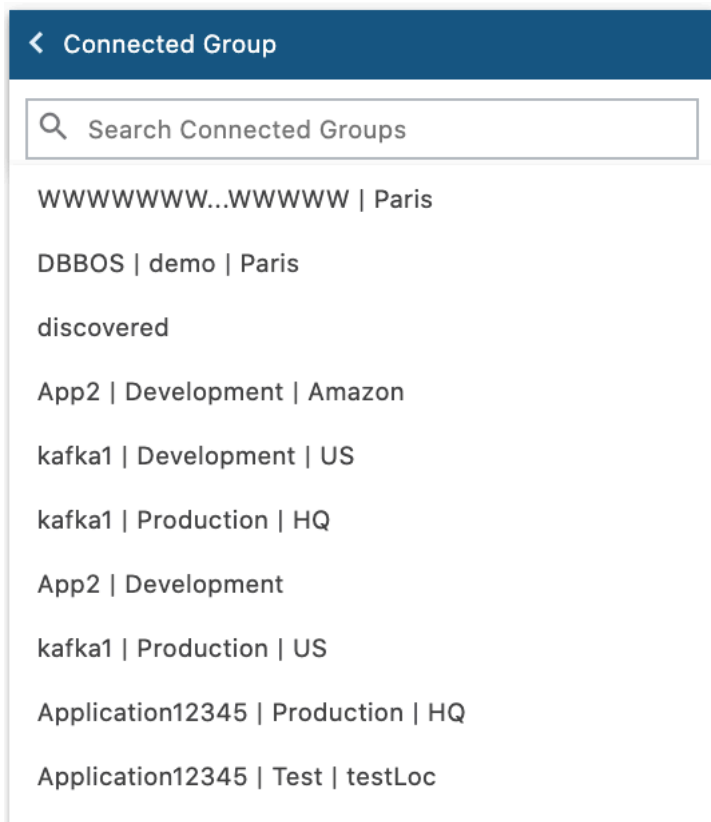
Search in Illumination

You have the ability to search for specific locations and groups in the Illumination map at a global level. This feature is especially helpful when managing dozens or even hundreds of locations or groups in the Illumination map.

1. Click on the magnifying glass in the upper left corner next to **Search Illumination**.



2. Select the location where you want to perform search.
3. Click inside the location to open the command panel.

4. Click the group for which you want to see the details.

5. The group's command panel shows the following:
 - Workloads associated with the group
 - Workload's enforcement (with an option to edit)
 - A link to locate groups connected to the selected group in the Illumination (Connected Groups)
 - Links that allow you to view associated rulesets, to start Policy Generator, clear traffic counters or expand group roles
6. Click the arrow next to Connected Groups.



The VEN uploads the traffic flow data to the PCE every 10 minutes. The “Increase VEN Update Rate” option increases the rate at which the information is uploaded, which helps you see the flow more frequently. When you click “Increase VEN Update Rate” from the App Group window, the data gets uploaded every 30 seconds for the next 10 minutes. After 10 minutes, it resets to the default value of uploading the data every 10 minutes.

Clear Traffic Counters

To draw the traffic patterns on the Illumination map, the PCE stores all traffic data that flows between workloads.

When you want see only current traffic data and purge traffic connections that might have stopped, you can clear all traffic counters between two workloads, for all the workload communication in the group, or for your entire organization:

- To clear all traffic between two workloads, select a traffic link and click **Clear traffic counters** in the command panel.
- To clear traffic for an entire group, select the group and click **Clear traffic counters** in the command panel. In the dialog, select **For this Group only** and click **OK**.
- To clear traffic for an entire group, select the group and click **Clear traffic counters** in the command panel. In the dialog, select **For entire Organization** and click **OK**.

Write Rules In Illumination

You can write rules for traffic inside of groups by selecting traffic links and specifying the traffic flows you want to allow in a rule. This means that only the traffic that you permit between workloads is allowed and all other undefined traffic is blocked.

This method is for writing rules on an individual basis. To write larger sets of rules for your workloads, Illumio recommend using rulesets. See "Create a Ruleset" in the *Security Policy Guide* for more information.



NOTE

Inter-group traffic links to and from discovered groups from the selected group on the Details view are displayed in gray.


To write rules for traffic from the Group Detail view or Full map:

When you write a rule from the Group Detail view, you can create one of two kinds of rules, depending on the ruleset:

- When you click a line inside a group, the generated rule is an Intra-Scope Rule.
- When you click a line between two groups, the generated rule is an intra-scope rule if the scope is All | All | All for the selected ruleset.
- For the same line, if the ruleset scope matches the Provider labels, the generated rule is an Extra-Scope Rule.

To write rules for workloads:

1. In the Illumination Draft view, select one of the traffic links between the workloads. In the command panel, workloads and enforcement are displayed.



2. To allow this traffic, click **Add rule** in the command panel. An Add rule panel appears.

Add Segmentation Rule

Add to Ruleset
rule_set_1

Scopes
test_ap...
test_en...
test_pla...

Consumers
All Workloads

SecureConnect
On
Off

Service
All Services

Providers
All Workloads

Cancel
Save

3. You can edit the rule and choose to add it to the recommended ruleset.
4. Click **Save**.
The traffic link turns green to indicate a permitted flow for the workloads.

Write a Group-Level Rule

In the Illumination map, you can write a rule that allows all workloads in the group to communicate with each other. You can write this type of rule from any Illumination view level: Global, Groups, and Detail view levels. This feature is useful if your goal is to ringfence a particular application instance or high value asset, when you want only the group's members to communicate with each other, but you want to separate it from everything else in your organization.

To write a group-level rule:

1. From the PCE web console menu, choose **Illumination**.
The Illumination map appears.
2. Select the group you want to write a group-level rule for.
The command panel appears.
3. In the command panel, click **Add Rule**.
The command panel opens an **Add Rule** dialog box. In the dialog box, you can write a rule that allows all workloads and other entities to communicate with each other for all services.
4. In the **Add Rule** dialog box, select a ruleset to add the rule to, or leave the default when one is already configured for the group, and complete the rest of the rule fields.
(Optional) Select SecureConnect when you want to encrypt traffic between workloads.

Set Group Enforcement

After you have written rules for a group, you can change the group's enforcement, which determines how a group's rules affect the communication among the group's workloads.

You can choose from the following enforcement states for the workloads inside of a group:

- **Visibility only**. In this state the PCE displays the flow of traffic to and from the workload, providing insight into the datacenter and the applications running in it. No traffic is blocked in this state.
- **Full Enforcement**. A state of a workload in which all ruleset rules are enforced and all traffic flows that are not allowed by the rules are blocked.
- **Selective Enforcement**. Selective enforcement applies only to managed workloads; it does not apply to NEN-controlled or other unmanaged workloads. It controls which ports or services are enforced on workloads.

To set group enforcement:

1. From the PCE web console menu, choose **Illumination**.
The Illumination map appears.
2. Find the group for which you want to change enforcement and click inside the group.
The command panel appears.

Illumination – Reported View Labels to filter ▼

Search Illumination


GROUP
test_app_1 | test_env_1 | test_place_1

1 Role
test_access_1 2

Workloads
Count 2
Enforcement Full ✎

View Segmentation Rulesets
Start Policy Generator
More ▼

App Group Map 🔗



test_app_1 | test_env_1 | test_place_1

- Click the edit tool next to Enforcement.
- The pop-up dialog explains that editing affects only the workloads in this group. Click **Continue**.
- The group's command panel allows you to select the enforcement from the dropdown list.

GROUP
test_app_1 | test_env_1 | test_place_1

1 Role
test_access_1 2

Workloads
Count 2
Enforcement Full ▼

Visibility Only
Selective
Full

Cancel Save

View Segmentation Rulesets
Start Policy Generator
More ▼

App Group Map 🔗

- Select the new enforcement type and click **Save**

Create Unmanaged Workloads from IP Addresses

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



NOTE

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

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

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

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

Until new traffic for the unmanaged workload is observed, the traffic lines are not displayed for the unmanaged workload. The traffic lines in Illumination 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.

App Group Map

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

About the App Group Map



NOTE

In previous releases, this feature was referred to as “Segmentation Rulesets.” In Illumio Core 21.5.0 and later releases, this feature is now referred to as “Rulesets.” Some images might still display the previous feature name.

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

The Illumination map visualizes the workloads and traffic in your datacenter, which takes time to render with large-scale deployments. However, some users such as application owners prefer to think about their datacenter in terms of traffic between workloads that belong to different application instances, rather than between physical locations.

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

The App Group Map visualizes the network traffic by organizing it based on App Groups. App Groups can either be a set of Application and Environment labels or a set of Application, Environment, and Location labels.

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

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

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

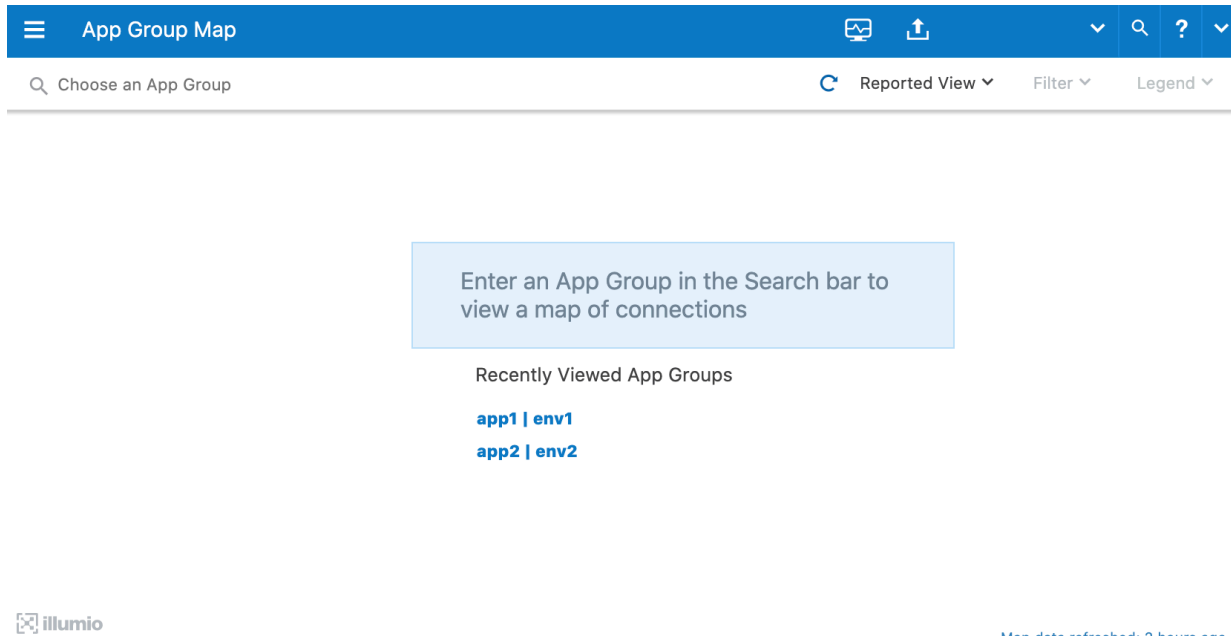
App Group Views

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




NOTE

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



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

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



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

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

**NOTE**

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

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

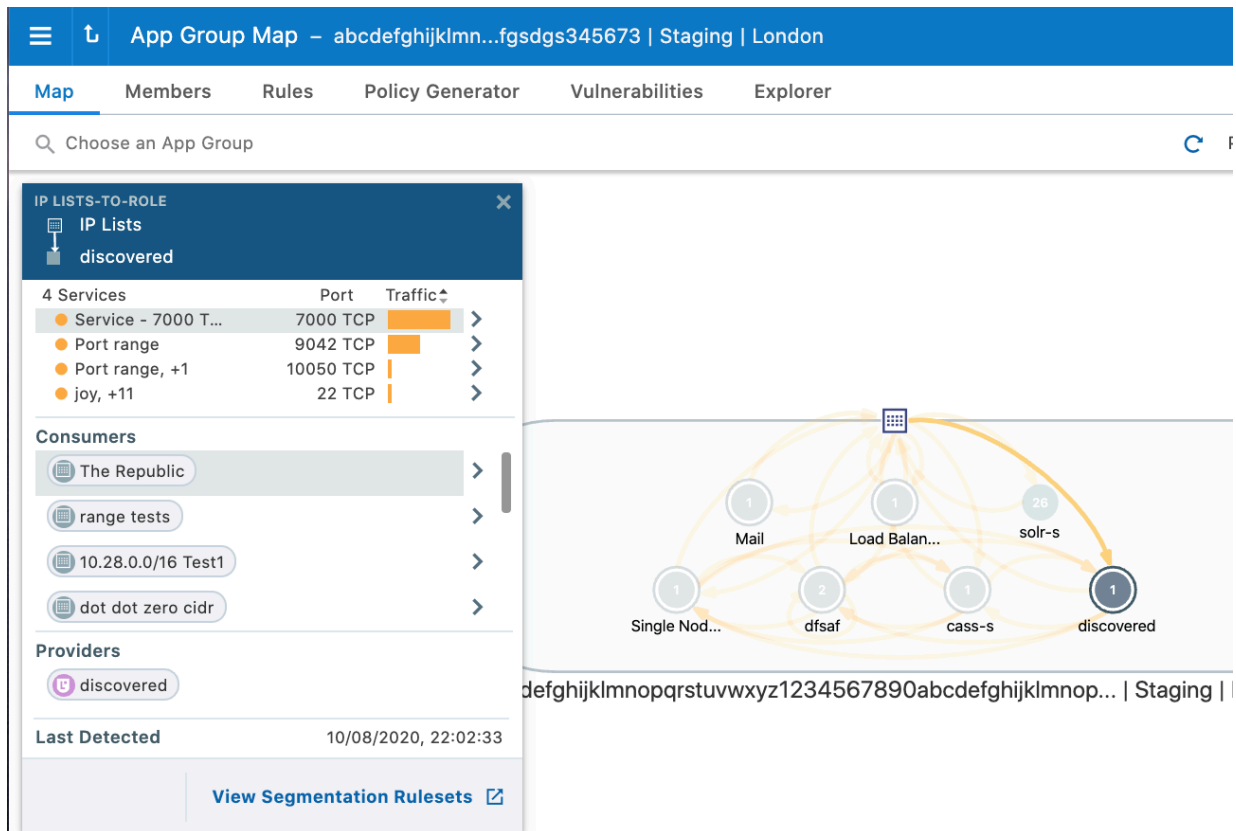
If an expanded Consuming or Providing App Group is currently displayed in the App Group Map, the link in the App Group's circle changes from **View** to **Next**. Click **Next** to view the next connected Consuming or Providing App Group.

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

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

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

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



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

Work with the App Group Map

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

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

App Group Creation and Association

An App Group is created when:

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

A workload is associated with an App Group when:

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

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

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

Configure App Groups

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

There are two ways to configure App Groups:

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



NOTE

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

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

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

Caveats

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

Explorer

Explorer allows you to analyze traffic flows for auditing, reporting, and troubleshooting purposes. It provides the ability to view traffic by time, port, consumers, providers, and services. It is not an interactive tool, so you cannot write rules using it.

About Explorer

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

The VEN decorates the flow summary logs with DNS names when it sends them to the PCE. Explorer 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.

Explorer relies on traffic flow data stored in the PCE traffic database. When a single disk is used for all PCE storage, the default traffic database disk usage limit is in effect. When the amount of traffic flow data exceeds the limit, older data is pruned, and is no longer available in Explorer. To increase the amount of storage available for traffic flow data, you can set up a separate disk for the traffic database and set runtime parameters as required for a two-storage-device configuration. For more information, see "PCE Storage Device Partitions" in the *PCE Installation and Upgrade Guide*.

Explorer Searches

When you search data using Explorer, you are searching traffic flows between providers and consumers over a specific time period over a specific port and protocol.

When you want to search for particular traffic flows on a regular basis, you can save that filter and it appears under *Favorites*. You can make changes to an existing favorite and save the modified query. Explorer also displays your ten most recent searches. In Explorer, you can also see the effect of currently configured and unprovisioned policies on the traffic that was observed in the past.

An Explorer search consist of the following elements:

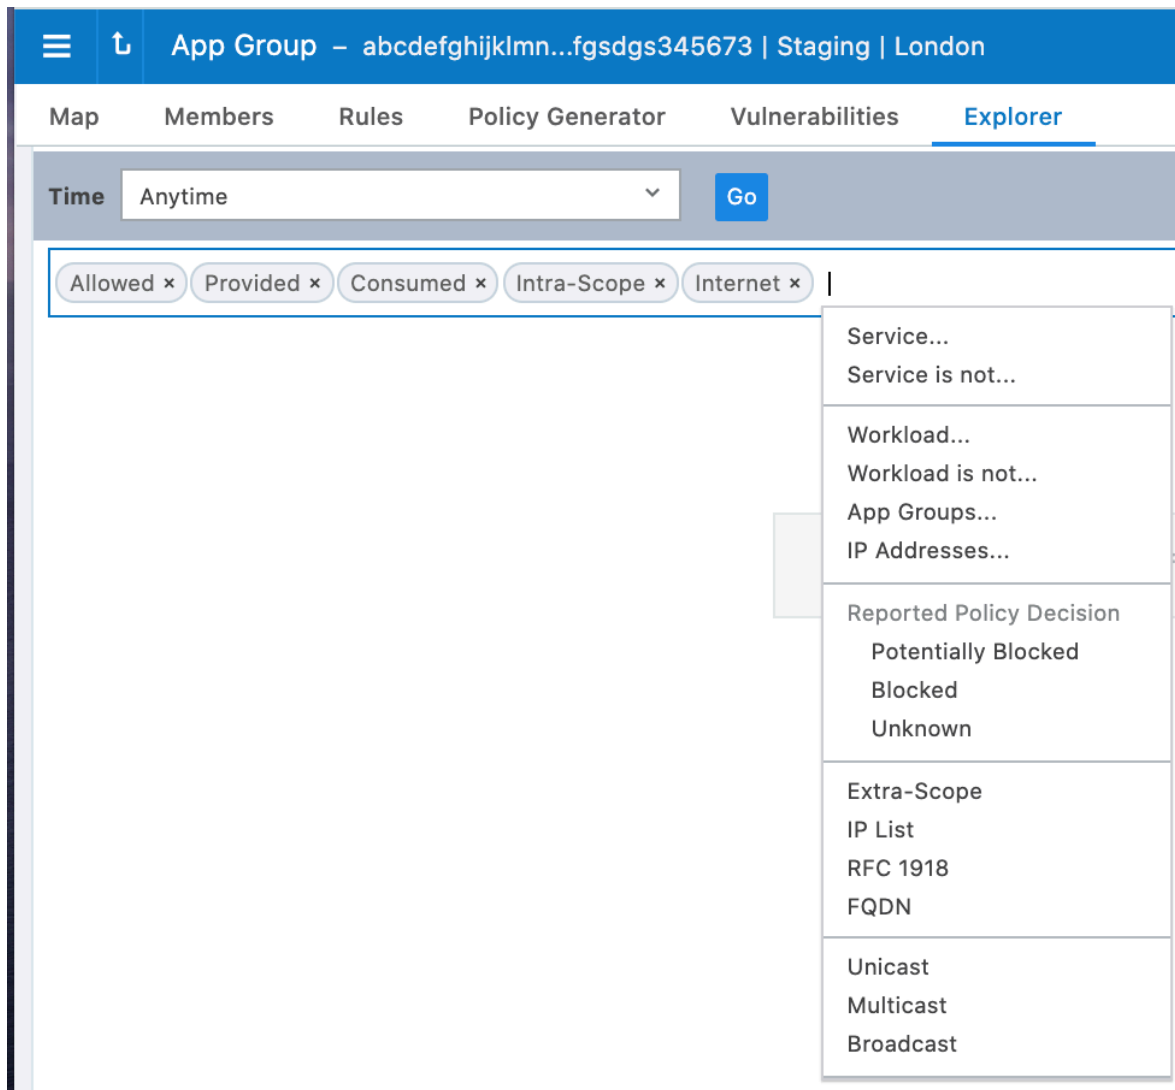
- **Consumers:** Enter workloads, IP addresses, or labels that are consuming the service provided in the traffic flow. The entries you add in the Include field are used as a search criteria and the ones you add in the Exclude field are not used in the search.
- **Providers:** Enter workloads, IP addresses, or labels that are providing the service in the traffic flow. The entries you add in the Include field are used as a search criteria and the ones you add in the Exclude field are not used in the search.



NOTE

You can choose to search either "Consumers And Providers" or "Consumers Or Providers" by clicking the settings icon.


- **Services:** Enter port and protocol, port ranges, process, Windows services, or policy services. Enter port numbers and protocol types to search for traffic flows whose destination port values and protocols match the search criteria. The entries you add in the Include field are used as a search criteria and the ones you add in the Exclude field are not used in the search. If you do not specify a value, all ports, protocols, port ranges, processes, and services are included in the search.
- **Time:** Select how far in the past (last hour, day, week, or month, or anytime) or specify a custom time range. The custom time filter displays all the flows between the selected from-to date-time stamp.
- **Reported Policy Decision:** Select the type of policy decision (allowed, potentially blocked, blocked, or unknown) to search for flows with a specific policy decision reported by the VEN.
- **Connection State:** The following traffic flow states are displayed under the “Connection State” column only in the exported table.
 - **Active:** The flow is in progress.
 - **Closed:** The flow in each direction is recorded and the connection (TCP only) is closed.
 - **Static:** The flow came from a static illumination “snapshot” of the current connection, from an idle VEN, or from data imported from another source.
 - **Timed Out:** No packets in either direction were received for a long time for this session and it is being timed out.
- Additionally, the following two “Blocked” traffic states are also listed:
 - **New:** Dropped TCP packet that contains a SYN and is associated with a new connection
 - **Invalid:** Dropped TCP packet that does not contain a SYN and is dropped by firewall as invalid



Explorer Search Results

The Explorer search results can be displayed in the following three formats:

- **Parallel Coordinates:** Displays traffic flows as a vertical list of Consumers, Providers, and the port being used in the flows. You can click any item in the results to focus on specific flows. You can also sort the results to view based on port number or number of traffic flows.
- **Table:** Displays search results in a traditional table format. This view includes a column named Policy which indicates if the flow was allowed, blocked, or potentially blocked based on your policy. Flows that are potentially blocked could mean that there is no segmentation rule written for the flow or there is a segmentation rule written for the flow, but the provider workload's enforcement is set to Visibility Only.
- **Unmanaged IP Addresses:** Displays all connections to or from hosts that are unmanaged IP addresses. This view is useful for discovering the IP addresses of hosts that you want to managed with the PCE, either as managed workloads or unmanaged workloads. You can create unmanaged workloads in this page by selecting one or more of the IP addresses and clicking **Create Unmanaged Workloads**.
- **Unmanaged FQDNs:** Displays all connections from the workload that are unmanaged FQDNs.



Flow Collection

In a data center that contains mostly Windows servers, certain types of broadcast and multi-cast traffic comprise a large percentage of total traffic, which can degrade the functionality and usefulness of the PCE. To resolve this, you can use the PCE web console to set per-org filters or aggregation rules to observe the ongoing traffic in your organization and filter out or aggregate the traffic based on destination address, subnet, protocol, and disposition (unicast, broadcast, or multicast).

You can configure the PCE traffic collector to drop or aggregate certain types of broadcast and multicast traffic based on the following criteria:

- Destination address (IP address or CIDR block)
- Destination port/protocol
- Transmission Type (broadcast or multicast)



NOTE

Only users with Organization Owner roles can set the Flow Collection.

To set the flow collection:

1. From the main navigation menu of the PCE web console, click **Settings > Flow Collection**.
2. Click **Add**.
3. Select the *Action*.
 - **Aggregate**: Protocol, Destination IP Address, and Destination Port fields are not supported. Select the Transmission type and click Save.
4. Select the *Transmission* type. You can only enter a single port for each filter. Multiple ports and port ranges are not supported.

5. After reviewing your selections, click **Save**.
6. You can edit a filter by clicking on the row in the Flow Collection Filter page.

Enforcement Boundaries in Explorer Draft View

Prior to Illumio Core 21.4.0, Enforcement Boundaries weren't reflected in Illumination or Explorer. When you moved a workload into Selective Enforcement mode, the Illumination map displayed a dotted line around the workload but didn't indicate whether an Enforcement Boundary was in place. In Explorer reported view, you could detect that traffic was blocked; however, Explorer did not distinguish between traffic that was blocked because of full enforcement or because an Enforcement Boundary was in place.

In Illumio Core 21.4.0 and later releases, Enforcement Boundaries are displayed in Explorer draft mode. In Explorer, when you view your traffic flows, you see a visual indication whether traffic is blocked by an Enforcement Boundary or allowed through an Enforcement Boundary. Viewing this information in Explorer draft mode is useful to determine where Enforcement Boundaries are in place and understand the impact of the boundaries before provisioning them.

In particular, you can obtain the following information:


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



TIP

In the Draft Policy Decision column, click the text for traffic allowed across a boundary ("Allowed") or blocked by a boundary ("Blocked") to view the details about the boundary.

Explorer indicates these states with the following icons:



Work with Explorer

You can use Explorer to search for information about your organization and to create unmanaged workloads. You can specify IPList (CIDR and FQDN) in the traffic_queries API and get results of all entities that match that IPList. If there are multiple IPLists matching a source or destination IP, the top five are displayed. You can also optionally specify (only available through API) a flag to obtain workloads, whose IP addresses are part of the specified IPList and have flows to/from the IP address. Explorer returns 5 matching IP lists by default, which can be expanded to 50.

Explorer Search Example

One preliminary method of creating policy is to make sure that different environments of your datacenter are segmented from each other. For example, you can separate Development or Testing environments from your Production environments. Before you write policy rules to either allow or block this traffic, you want to determine if there are any traffic flows between them.



Using Explorer you can query, for example, the following:

"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'"

Example search using Explorer:

1. In the PCE web console menu in the upper left corner, choose **Explorer**. The Explorer page appears.
2. Under Consumers, enter or select the Environment label named "Development" from the Include drop-down list.
3. Under Consumers, enter or select the Role label named "Domain Controller" from the Exclude drop-down list.

4. Under Providers, enter or select the Environment label named “Production” from the Include drop-down list.
5. Under Providers, enter or select the Role label named “Domain Controller” from the Exclude drop-down list.
6. Under Port/Protocol, leave the Include field blank (which means “any”) and under Exclude enter “80.” One of the options is also ICMPv6.
7. Under Time, select Anytime.
8. Click **Go**.

The results appear when the search criteria is met.

Asynchronous Queries

Asynchronous queries allow you initiate multiple queries in parallel and view the results of the queries at a later time. Prior to Release 21.2.0, going offline during a query would result in lost query results. Starting with this release, whether you remain online or offline, the results of asynchronous queries will be preserved for a period of 24 hours. In addition, while a query is in progress, you can work in other areas of the product. The query search results can be exported to either a comma-separated-value (CSV) file or displayed in the Explorer Web Console. Depending on the size of the query, the results may take time to display.

In this release, Explorer enables you to run multiple queries and allows you to 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 may select the result that you wish to view by clicking the number beside the **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 (Consumer, Service, or Provider) in the filter. The exported file will have the same name.
 - Giving filters a unique name will help you identify your filters when you wish to rerun a query. This name will also appear as your report name.
 - You can also specify or change a filter name as desired.

Handling Duplication Flows in Queries

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

Prior to Release 20.3, these duplicate flows were merged together outside the database, and could have resulted in fewer results being returned to the user interface.

From Release 20.3 and later, duplicate flows spanning multiple days are merged in the database, allowing more unique flows to be returned.

Run Asynchronous Queries in Explorer

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 a query, proceed as follows:

1. From the main menu of the Web Console, navigate to **Explorer**.

2. Enter your query criteria in the **Include** field.
You can enter a Consumer, Provider, or Service, or merely indicate Production in the Provider column.
3. Click **Go** to begin the query process.
4. To process a parallel query, click **Go** again.
5. In the confirmation dialog box, click **Hide**.
6. Enter the next search criteria based on a new Provider. For example, Production and Staging.
7. Press **Go**.
8. Given support for asynchronous queries, you will see a number appear next to the **Results** button, indicating the number of simultaneous queries being processed.
Depending on the size of the queries, your second query may complete before your first query.
9. To view the results of your queries, click the **Results Available** pop-up that will appear in the bottom-right corner of the PCE Web Console.
You will see the results of your two queries, one with Production-only entries and a second with Production and Staging entries. At any time, you may click the **Results** button to view what queries were run.
Viewing results from past queries will not re-initiate a query. It will display cached query results. When you select a result, notice that the filter changes automatically, and displays new results.


View and Modify Query Result Settings

You can view results from the **Results** button, **Results Available** pop-up, or the **gear** drop-down in the upper-right corner of the Web Console.

In the **Results** window, when you see an asterisk in the Connections column, it indicates that there are more entries in the database than what you requested in your query.

To view additional results, you must increase the values in the **Results Settings** menu. You can access this menu from multiple locations; from the gear drop-down or your user profile drop-down.

1. From the upper-right corner of Explorer, click the **gear** drop-down to view the Results Settings window.



Results Settings [X]

For queries spanning a long time period, increasing the maximum number of connections provides more complete results at the expense of slower performance. If the number of connections returned from the database exceeds the maximum displayed in Explorer, all connections can be viewed in the exported CSV file.

Maximum Number of Connections

Displayed in Explorer: Valid range: 1 - 100,000

Returned from Database per Region: Valid range: 1 - 200,000

Cancel [OK]

- From Results Settings, change the value for maximum connections.

You can change the number of connections for what can be 'Displayed in Explorer' or 'Returned from the Database per Region'. Up to 100,000 can be 'Displayed in Explorer' and up to 200,000 can be 'Returned from the Database per Region.'

Export Query Results

There are multiple locations from where you can export results. Depending on the location from which you export results, the number of results that are shown may differ. The number of results from Explorer may differ from the number of results returned by the PCE web service.

In the Web Console, you can display additional data if you include draft rules and FQDN look ups.

Export from the Explorer Web Console


From the central pane of the Explorer window, you can export query results with enhanced information.

- From the Reported View drop-down, choose **Draft View**.

This will include Draft Policy Decision (All Draft, Blocked, or Allowed) entries in your exported file. If you do not make a selection, you will receive information on the Reported Policy Decision results.

Reported Policy Decision	Draft Policy Decision	Reported by
Potentially Blocked	Blocked	Provider
Potentially Blocked	Blocked	Provider

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



Resolve Unknown FQDNs [X]

FQDNs are resolved using reverse DNS from the PCE. Accuracy of results may vary depending on the relative locations of the PCE and the target device.

☐ Don't show this message again.

[Done]

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

**NOTE**

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

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

Export from the Results Button

Once your query completes and results are available, you can export these results to a CSV file.

1. Click **Results** to view the list of your queries.
2. Click the query result list item to view the results for a particular report.
3. Click **Export** to gather your data in a CSV file.

When you create a direct export from the Results list, you will receive Reported Policy Decision entries in the report. This report will not contain Draft Policy Decision entries or FQDN information.

Vulnerability Map

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

About Vulnerability Map

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

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

Vulnerability Terminology

- **Vulnerability:** A generic vulnerability that can exist on any workload (or port and protocol), for example, Apache heart bleed.
- **Detected Vulnerability:** The instance of a vulnerability that exists on a workload, for example, Apache heart bleed existing on workload X on port 80.
- **Vulnerability Report:** A report containing the detected vulnerabilities.
- **Vulnerability Score:** The summation of severities of the vulnerabilities for an App Group, role, or workload where the individual vulnerability scores range between 0 and 10.

- **Exposure Score:** The E/W Exposure Score combined with the Internet Exposure. It is a score of how many workloads can use the vulnerable port on a workload based on the provisioned rules.
- **Vulnerability Exposure Score (V-E Score):** A calculated value based on the Vulnerability Score and the Exposure Score = $\sum f(VS, ES)$. It can be shown for an individual vulnerability on a port for a single workload or as a summation of all the V-E Scores for an App Group, role, or workload.
- **East-West (E/W) Exposure Score:** A count of workloads that can use a vulnerable port with the currently provisioned rules, and whether the vulnerability is exposed to the internet.
- **Internet Exposure:** Indicates whether a vulnerable port is exposed to traffic from the internet. Internet Exposure is enabled by the rules allowing inbound traffic on that port.
- **Severity:** Represents a range of Vulnerability Score values.
 - 0 = Info
 - 0.1 to 4.0 = Low
 - 4.1 to 7.0 = Medium
 - 7.1 to 9.0 = High
 - 9.1 to 10 = Critical

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

Benefits of the Vulnerability Map

The Vulnerability Map has the following benefits:

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

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



NOTE

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

Vulnerability Map Usage

In most organizations, vulnerability management is performed through scanners that scan infrastructure to identify vulnerabilities and provide reports. In some cases, there is no patch for zero-day vulnerabilities. Illumio Core vulnerability-based micro-segmentation gives securi-

ty teams the ability to focus on where they are most vulnerable—inside their datacenter and cloud, leveraging micro-segmentation as a compensating control.

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

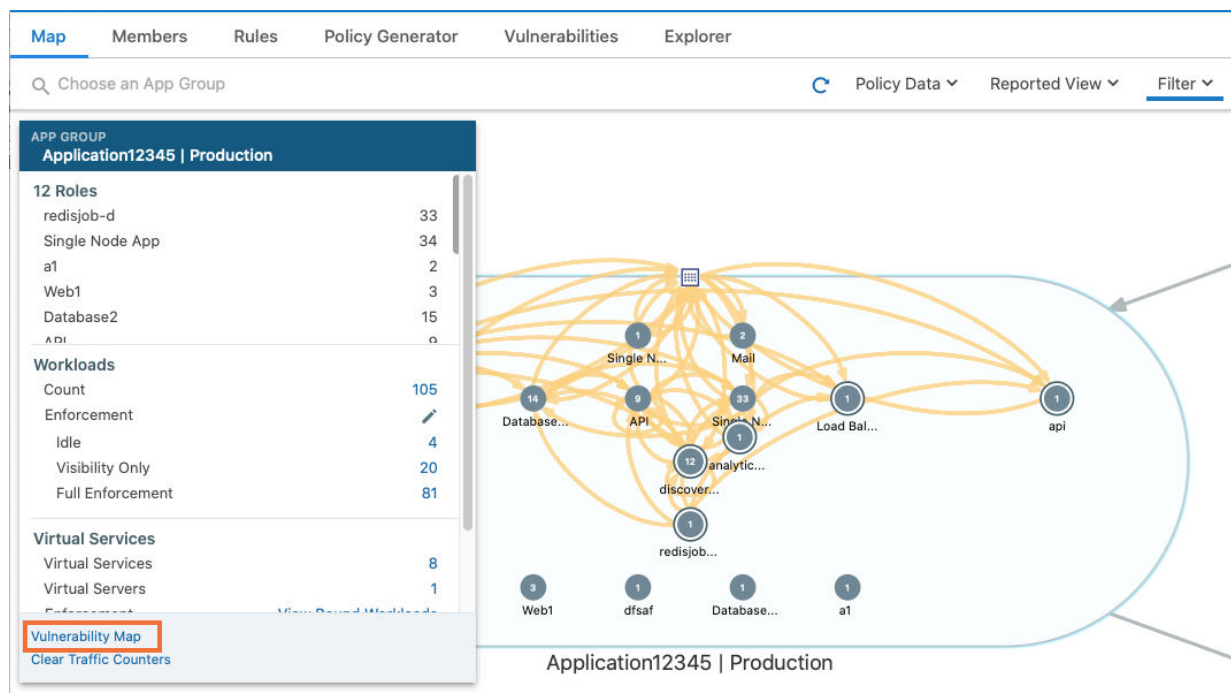
Work with Vulnerability Maps

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

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

Enable the Vulnerability Map

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



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

**NOTE**

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

Caveats

- A maximum of 100,000 vulnerabilities can be detected per organization.
- A maximum of 100 vulnerabilities can be detected per workload.
- The Vulnerability Map is not supported in Supercluster implementations.
- The exposure score is calculated on the first firewall sync for a given workload. When a PCE is restarted:
 - Vulnerability Score and Exposure Score are not available until the firewall sync occurs.
 - The scores are not available when a workload is offline.
- Vulnerabilities can only be imported using the PCE CLI Tool.

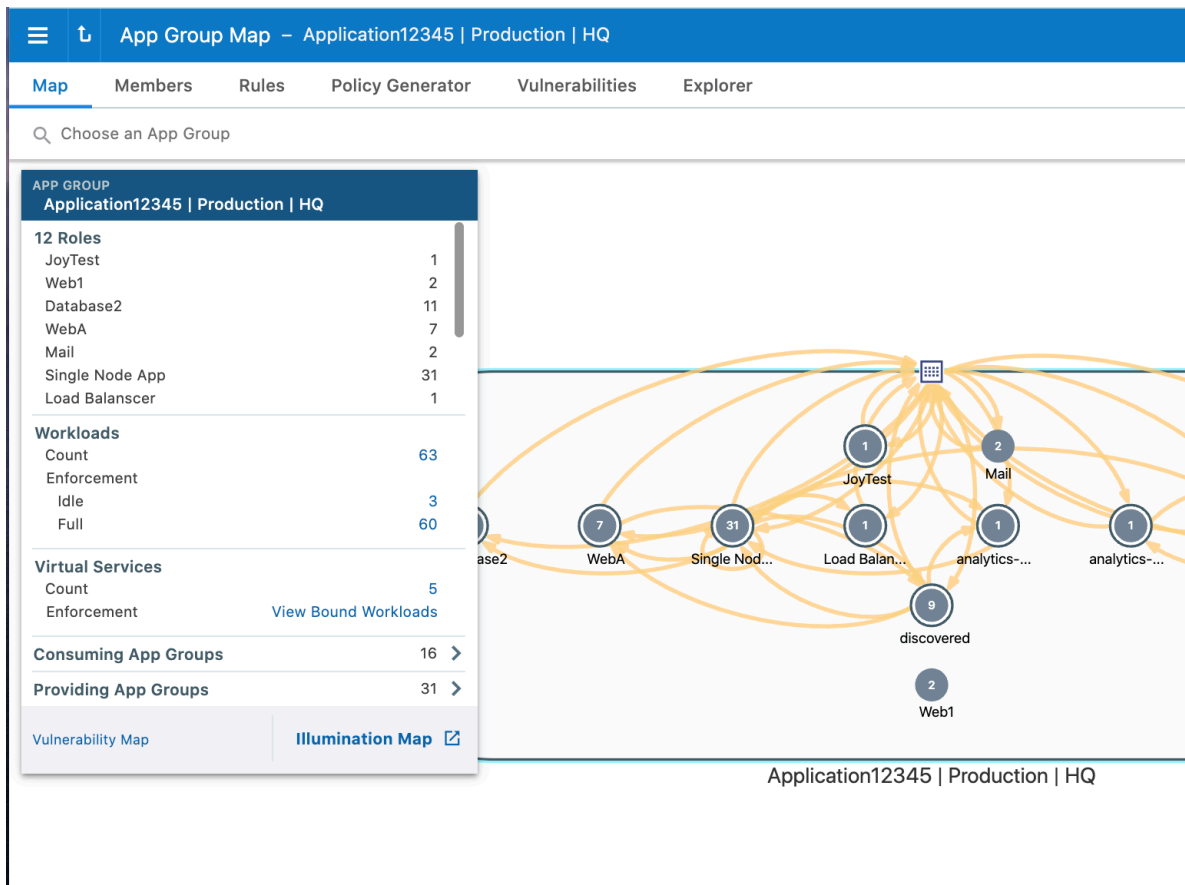
View and Mitigate Vulnerabilities

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

To view and mitigate vulnerabilities:

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

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



3.

**NOTE**

The Cloud icon denotes Northern Exposure.

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

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

The screenshot shows the Illumio Core 21.5 Visualization interface. At the top, there are tabs for 'Vulnerability Data', 'Reported View', 'Filter', and 'Legend'. The 'Filter' tab is selected. Below the tabs, a network diagram is visible with a node labeled 'Load Balan...'. A slider for 'Vulnerability Exposure Score' is set to 'Info'. The 'Filter' menu is open, showing the following options:

- TRAFFIC LINKS**
 - ☒ All Services ☐ Selected Services (Edit)
 - ☒ Intra-Group ☒ FQDN
 - ☒ Internet ☒ IP List
 - ☒ Broadcast ☒ Multicast
 - ☒ ICMP
- VULNERABILITY EXPOSURE SCORE**
 - ☒ All Vulnerabilities ☐ Exposed Vulnerabilities
- VULNERABILITY TRAFFIC**
 - ☒ Blocked Vulnerable Traffic
- WORKLOADS**
 - ☒ Visibility Only ☒ Idle
 - ☒ Full Enforcement ☒ Unmanaged
 - ☒ Selective Enforcement

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



NOTE

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

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

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

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

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

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

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

- c.** Complete the fields in the Policy Generator wizard.

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

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

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

Vulnerabilities Tab for Workload Details

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

[illegible]

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

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

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

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

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

Reports

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

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

To enable reporting on SaaS for releases prior to Illumio Core Release 22.2.0, you must submit a Illumio support ticket.

This feature provides two types of recurring reports:

- Executive Summary reports
- App Group Summary reports

About Reports

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

Reporting in the PCE

The PCE web console menu includes a “Reports” option. When you choose the Reports option, the Reports page appears. This page includes two tabs: Downloads and Recurrence. Generated reports appear on the “Downloads” tab. By default, the list is sorted in descending order by the “Generated At” time.

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

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

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

Recurring reports are run on the following schedule:

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

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

Executive Summary Reports

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

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

Executive Summary reports answer the following questions for decision makers:

- How are we progressing in deploying security policy into our environment?
- How many of our workloads are being managed by Illumio (VENs are installed on the hosts but they aren't in enforcement mode)?

- How quickly is enforcement progressing over time (the number of workloads that have moved into the enforcement mode over the report's specified time range)?
- What potentially dangerous traffic is Illumio blocking that wouldn't have been blocked without Illumio Core, resulting in a security risk.
- What sort of vulnerabilities do our workloads have? Vulnerability information is provided as a V-E score that is the sum of all app groups.



IMPORTANT

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


For more information about Vulnerability Maps, see [Vulnerability Map. \[49\]](#)

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 is enforcement progressing over time and percentage of workloads in deployment versus enforcement.

Workloads by Enforcement Mode



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

Provider Workloads by Policy Decision



App Group Summary Report

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

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

You create App Group Summary reports by application; meaning, each report provides data for only one application defined by a set of labels. Whether you choose 2 labels (application and environment) or 3 labels (application, environment, and location) for a report depends on how you have configured the PCE to define app groups. See [Configure App Groups \[39\]](#) for information.

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

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

Work with Reports in the PCE

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

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

Add a Report

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

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

Manage Reports

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

To download a report:

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

To set the retention period for all reports:

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

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

To edit the settings for a report:

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

To end the recurrence of a report:

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

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