

This guide describes the Illumio Security Policy, including its policy objects. It provides guidance on designing a label schema and outlines recommended approaches for designing Illumio's security policy, including creating rules and rulesets.

"Show Me How" Videos and Animations

| | | |
|---|---|---|
| Policies and Enforcement | | |
| Apply a Staged Policy | Enforce a Workload Policy State | Add a Static Policy |
| Perform a Policy Check | Segment Multiple App Groups with the Policy Generator | Provision Changes |
| Services | | |
| Create a Service | Add a Virtual Service | Bind a Virtual Service to a Workload |
| Create an IP List | Generate an Export Report | |
| VENs and Workloads | | |
| Adding Unmanaged Workloads | VEN Suspension | Loopback Interfaces |
| Create an Unmanaged Workload from Blocked Traffic | Reject Connections | Set Workload Interfaces to Ignored |
| Schedule Label Assignments | Find and Review Matching Workloads | Export a Workload-Label Review List |
| Policy Rules | | |
| Search for Rules | Create Intra-Scope Rules with the Policy Generator | Create Rules Using IP Lists with the Policy Generator |
| Labeling Rules | | |
| Reorder Labeling Rules | Enable/Disable Labeling Rules | Remove Labeling Rules |

Table of Contents

| | |
|--|----|
| Overview of Security Policy | 8 |
| About the Illumio Policy Model | 8 |
| The Illumio Policy Model | 8 |
| Security Policy Guidelines | 8 |
| Enforcement States | 9 |
| Overview of Policy Objects | 9 |
| Policy Objects | 9 |
| About Policy | 10 |
| Understanding Policy | 10 |
| Staged Policies | 12 |
| Updating Policy Objects | 12 |
| Checking for Staged Policy Status on Workloads | 13 |
| Apply a Staged Policy | 13 |
| Static Policies | 14 |
| Static Policy Details | 15 |
| Use Cases for Static Policies | 15 |
| Scope-based and Scopeless Policies | 17 |
| Scopeless Policies | 17 |
| Scope-based Policies | 17 |
| Single Scope Policies | 17 |
| Advanced Scope Policies | 18 |
| Security Policy Objects | 19 |
| About Labels and Label Groups | 19 |
| Create a Label | 19 |
| Label Types | 20 |
| Label Groups | 21 |
| Use a Label Group in a Rule | 22 |
| Use a Label Group in a Scope | 23 |
| Labeling Workloads | 24 |
| Filtering Labels and Label Groups | 24 |
| System Default “All” Labels | 25 |
| Edit Labels for Multiple Workloads | 25 |
| Create a Service | 26 |
| View or Edit a Service | 27 |
| ICMP Services | 27 |
| IGMP Services | 29 |
| Filter the Services List | 29 |
| Virtual Services | 30 |
| Virtual Services Scenarios | 30 |
| How Virtual Services Work | 31 |
| Virtual Services in Rule Writing | 32 |
| Advanced Virtual Services Configuration | 34 |
| Filter the Virtual Services List | 35 |
| Add and Bind a Virtual Service | 35 |

| | |
|---|----|
| IP Lists | 37 |
| Overview of IP Lists | 37 |
| Examples of IP List Usage | 37 |
| Create an IP List | 38 |
| IP List Exclusions | 39 |
| Virtual Servers | 40 |
| Virtual Server Members and Labels | 40 |
| Policy for Virtual Server | 41 |
| Configure Virtual Servers | 42 |
| Virtual Server Load Balancers | 43 |
| Export Reports | 44 |
| Generate an Export Report | 44 |
| Workloads | 45 |
| Workload Setup Using PCE Web Console | 45 |
| Creating Managed Workloads by Installing VENS | 45 |
| Unmanaged Workloads | 45 |
| Workload Summary | 46 |
| Workload Enforcement States | 48 |
| Visibility Level | 49 |
| Workload Processes | 50 |
| Workload Rules | 50 |
| Workloads Blocked Traffic | 51 |
| Filter a View | 51 |
| Enforce a Workload Policy State | 52 |
| Set Workload Interfaces to Ignored | 52 |
| Compare Workload V-E Scores by Enforcement Type | 53 |
| How it works | 54 |
| Workload List pages | 54 |
| Update Workload Labels in Bulk | 56 |
| About the Export File | 56 |
| Procedure | 59 |
| Blocked Traffic | 62 |
| Overview of Blocked Traffic | 62 |
| Create an Unmanaged Workload from Blocked Traffic | 64 |
| Loopback Interfaces | 65 |
| VEN Administration on Workloads | 67 |
| VEN Suspension | 67 |
| Workloads and VENS | 67 |
| Manage Workloads and VENS | 67 |
| Enhanced Data Collection | 69 |
| Container Workloads | 72 |
| Pairing Profiles | 72 |
| Create Security Policy | 73 |
| Policies | 73 |
| Basic versus Scoped Policies | 73 |
| Policy Scope | 74 |

| | |
|---|-----|
| Labels in scopes and rules | 78 |
| Manage Policies | 82 |
| Core Services Detector | 86 |
| Enabling Core Services Detection | 86 |
| Managing Core Services | 87 |
| Scanner Detection | 89 |
| About Rules | 91 |
| Types of Rules | 91 |
| Intra-scope Rules | 92 |
| Extra-scope Rules | 92 |
| Custom iptables Rules | 94 |
| Rule Writing | 96 |
| Permitted Rule Writing Combinations | 96 |
| Rules for Application Policies | 97 |
| Application Policy Rule Types | 97 |
| Allow Rules | 99 |
| Deny Rules | 99 |
| Implementing Deny Rules During the Transition to Allow Rules | 99 |
| Conflicted Rules panel | 100 |
| Policy Check and Rule Search | 100 |
| Perform a Policy Check | 101 |
| Rule Search | 101 |
| Stateful vs. Stateless Rules | 103 |
| Stateless Rules | 103 |
| FQDN-Based Rules | 104 |
| Benefits of FQDN-Based Rules | 104 |
| Features of FQDN-Based Rules | 105 |
| FQDN-Based Rule Requirements and Limitations | 106 |
| FQDN Visibility | 107 |
| Windows Process-Based Rules | 110 |
| Creating Services with System-Initiated Processes | 110 |
| Windows Environmental Variables | 111 |
| Rule-Based Labeling | 112 |
| Before you begin Rule-Based Labeling | 112 |
| Typical Labeling Rule Workflow | 112 |
| Step 1: Add a Labeling Rule | 113 |
| Step 2: Find and review matching workloads | 113 |
| Step 3: Assign labels to matching workloads | 113 |
| Work with Labeling Rules | 113 |
| Add a Labeling Rule | 113 |
| Find and Review Matching Workloads | 115 |
| Assign labels to matching workloads immediately | 116 |
| Schedule Label Assignments | 117 |
| Edit a Labeling Rule | 118 |
| Enable/Disable Labeling Rules | 120 |

| | |
|---|-----|
| Reorder Labeling Rules | 121 |
| Remove Labeling Rules | 122 |
| Export a Workload-Label-Review List | 123 |
| How Label Matching Works | 123 |
| Labeling Rule Terminology | 124 |
| Labeling Rule Matching Logic | 124 |
| Labeling Rule Examples | 125 |
| Example 1. Hostname Rule to match workloads that contain part of a specified host name | 125 |
| Example 2. OS Rule to match workloads running a specific operating system | 126 |
| Example 3. IP Address Rule to match workloads within a spe- cific IP address range: | 127 |
| Example 4. CIDR Block Rule to match workloads within a specific CIDR block: | 127 |
| Example 5. Rule with multiple attributes, each with a single value: | 127 |
| Illumio Policy Enforcement Model | 129 |
| Why Use Selective Enforcement? | 129 |
| Applying Selective Enforcement | 129 |
| How Selective Enforcement Works | 130 |
| Enforcement Progression Model | 130 |
| Use Cases and Limitations | 131 |
| Selective Enforcement Mode Limitations | 131 |
| Workload Enforcement States | 131 |
| Policy Exclusions | 132 |
| Policy Exclusions Overview | 132 |
| Policy Exclusions Support | 133 |
| Requirements and Restrictions | 135 |
| Adaptive User Segmentation | 137 |
| Configure a Microsoft directory application for use with Illumio AUS | 137 |
| Microsoft Entra ID | 137 |
| Microsoft Active Directory | 138 |
| Configure Microsoft Entra ID (Azure AD) for use with Illumio AUS | 138 |
| Key features for Microsoft Entra ID and Illumio AUS integra- tion | 139 |
| Prerequisites | 139 |
| Limitations for Entra ID and User Groups | 139 |
| Configure Entra ID for use with Illumio AUS | 140 |
| Set Up the VEN for Entra ID for AUS | 142 |
| Configure Microsoft Active Directory for use with Illumio AUS | 142 |
| Key features for using Illumio AUS integration with Microsoft Active Directory | 143 |
| Add Active Directory Groups to the PCE | 143 |

| | |
|---|-----|
| Add User Group-Based Rules for AUS | 144 |
| About the Policy Generator | 146 |
| Overview of Policy Generator | 146 |
| Policy Generator Prerequisites and Limitations | 147 |
| Create Intra-scope Rules with the Policy Generator | 147 |
| Create Extra-scope Rules with the Policy Generator | 149 |
| Create Rules Using IP Lists with the Policy Generator | 150 |
| Segment Multiple App Groups with the Policy Generator | 152 |
| Policy Generator Wizard | 154 |
| About Provisioning | 155 |
| How Provisioning Works Internally | 155 |
| Full Provisioning | 156 |
| Selective Provisioning (Quick Provision) | 156 |
| Provisioning in Static Policy Mode | 156 |
| Versioning, Restore, and Revert | 156 |
| Versioning Features | 157 |
| Restore vs Revert Action | 157 |
| Policy Versions | 157 |
| Restore Policy | 158 |
| Provision Changes | 158 |
| Revert Provisioned Changes | 159 |
| Provision Note Setting | 159 |
| Segmentation Templates | 161 |
| Catalog Retrieved from Support Portal | 161 |
| Features of Segmentation Templates | 161 |
| Segmentation Template Prerequisites and Limitations | 163 |
| Editing Segmentation Templates | 164 |
| Editing Policy Object Names or IDs | 164 |
| Deleting or Editing Policy Objects | 164 |
| Segmentation Templates Installation and Upload | 165 |
| Install a Segmentation Template | 165 |
| Upload a Segmentation Template | 165 |
| Update a Segmentation Template | 166 |
| Uninstall a Segmentation Template | 167 |
| Secure Workload Connections | 168 |
| SecureConnect | 168 |
| SecureConnect Overview | 168 |
| SecureConnect Use Cases | 169 |
| SecureConnect Features and Enforcement | 169 |
| SecureConnect Rules and Visibility-Only State | 169 |
| AdminConnect | 172 |
| Features of AdminConnect | 172 |
| AdminConnect Prerequisites and Limitations | 173 |
| Legal Notice | 174 |

Overview of Security Policy

This section describes the security policies, configurable rules that protect network assets from threats and disruptions.

Illumio Segmentation for Data Centers relies on security policy to secure communications between workloads.

About the Illumio Policy Model

Illumio offers a distinct approach to managing security policies for workloads from traditional network security policies. Traditional policies rely on network-specific details like VLANs, zones, and IP addresses, tying security directly to network infrastructure.

In contrast, Illumio uses a multidimensional labeling system to classify and define workload functions. Each workload receives labels based on four dimensions: role, application, environment, and location. These labels enable users to set clear, functional security policies, removing ambiguity from policy definitions.

Users define rules and rulesets using these labels to specify how workloads within their organization interact. The Policy Compute Engine (PCE) then translates these functional, label-based security policies into specific firewall rules applied at the workload's operating system level.

The Illumio Policy Model

Illumio allows you to manage your security policies using adaptive or static policies. The Illumio policy model allows you to choose how to implement security policies.

Security Policy Guidelines

The following guidelines are recommendations on how to create your security policy in PCE

Creating a security policy is an iterative process; following these recommendations will provide a broad initial policy, which can then be incrementally improved until a sufficiently robust policy is established.

When creating your security policy:

1. Refine your initial policy to strengthen it by narrowing overly broad access.
2. Use the Visibility Only enforcement to verify and enact your policy.

Enforcement States

After creating a policy, you can preview its potential effects using Illumination's Draft View. This view displays the changes that will occur once the policy is enforced.

- **Visibility only:** Policies are refined initially until most traffic lines appear green in Illumination. In this state, no traffic is blocked, allowing verification of policy accuracy. Any new, unaddressed traffic appears as a red line.
- **Selective enforcement** . This state enables partial enforcement of policies, targeting specific applications or processes. It helps address vulnerabilities rapidly by temporarily enforcing security rules, while the remaining services and ports remain unaffected.
- **Full enforcement:** Gradually implementing full enforcement can minimize disruption by starting with less critical workloads, stabilizing them, and progressively including more sensitive systems. This phased approach reduces potential issues to a manageable number of workloads.

Overview of Policy Objects

The Illumio Policy Compute Engine (PCE) includes several objects for defining security policies:

Policy Objects

The Illumio Policy Compute Engine (PCE) includes several objects for defining security policies:

- **Labels and Label Groups:** Group similar labels together and use the label groups in rule writing.

- **Services:** This allows you to define or discover existing services on your workloads. When a workload is paired with the PCE (with a VEN installed), it is scanned for any running processes displayed in the Services list.
- **Virtual Services:** This allows you to label processes or services on workloads. Virtual services can be used directly in rules, or the labels applied to virtual services can be used to write rules.
- **IP Lists:** Create IP lists (allowlists) to define IP addresses, IP ranges, and CIDR blocks that will be allowed access to your applications.
- **Virtual Servers and Load Balancers: Add F5 Load Balancer configurations to the PCE so you can create a policy for workloads to manage traffic through load balancers.**
- **Pairing Profiles** are explained in Configurations in VEN Installation and Upgrade Guide . They allow you to apply specific properties to workloads as the key pair with the PCE, such as applying labels and setting workload enforcement.
- **User Groups:** You can import Active Directory User Groups to write user-based rules for adaptive segmentation.

About Policy

Rules form the core of Illumio's security policy. A policy is a set of rules that define permitted network traffic. Create the rules using labels that identify your workloads.

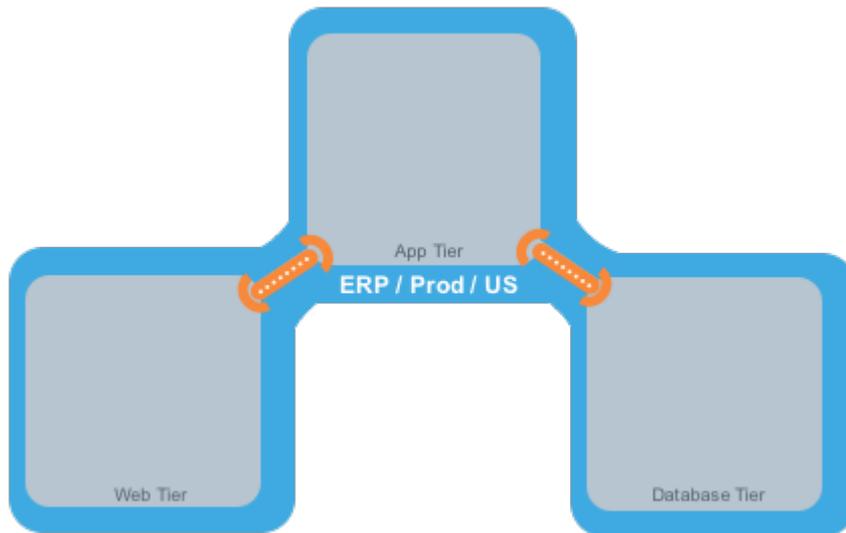
Understanding Policy

allow list model for security policy uses rules to define the allowed communication between two or more workloads. For example, if you have two workloads that comprise a simple application—a web server and a database server—you must write a rule that describes this relationship to allow these two workloads to communicate.



NOTE

The order in which the rules are written or any possible overlap between rules does not affect the allowlist model since each rule permits some traffic between workloads.



The relationships between the tiers (or workloads, as they are known in Illumio Segmentation for Data Centers):

- The Web workload can initiate communications with the App workload (Web → App).
- The App workload can initiate communications with the Database workload (App → Database).

The relationship in the diagram above is expressed as two separate rules:

- The Web workload can initiate communications with the App workload.
- The App workload can initiate communications with the Database workload.

To build your network security policy, create a policy for each workload. Use labels to identify your workloads and scopes, allowing you to apply policy to multiple workloads simultaneously.

**NOTE**

Illumio recommends creating no more than 500 rules per policy; otherwise, the PCE web console will not be able to display all the rules.

Suppose you want to create a policy with more than 500 rules. In that case, Illumio recommends splitting the rules across multiple policies or using the REST API, which allows you to create unlimited rules per policy.

Staged Policies

"Staged" status denotes that policy updates are delivered to the VEN but have not yet been enforced for workloads matching a Static Policy filter. To activate this status, create a matching policy. Implementing a static policy sends new OS-level firewall rules but requires manual activation by clicking "Apply Policy." Even with a static policy, manual provisioning is needed to update the PCE's security policy, with changes applied by the VENs after activation.

**TIP**

The orange badge on the Provision button indicates the number of changes you need to provision.

Updating Policy Objects

Provisioning is essential for changes in Illumio policy objects like services, IP lists, and label groups. Illumio enables reusable objects in rules, simplifying security policy maintenance. Updating a policy object automatically updates all associated rules, removing the need to modify each rule individually.

Provisioning changes to policies and objects results in the PCE saving your security policy as a new version. It recalculates OS-level firewall rules for impacted workloads and prompts VENs to download the updated rules.

Checking for Staged Policy Status on Workloads

To identify workloads with staged OS-level firewall rules, visit the Workloads page and assess each VEN's Policy Sync column status. Filtering workloads by this column reveals which ones have a staged rule.

Workloads Page Overview:

- **Active (Syncing):** The PCE dispatches a new policy to the VEN, usually completed within seconds. Workloads using adaptive or static policies may display as active (syncing) during policy updates.
 - **Staged:** The VEN has received the latest OS-level firewall rules but has not implemented them.
 - **Active:** The VEN has received, applied to, and acknowledged all policies sent from the PCE. A green dot icon denotes active workloads.
- For more information about the VEN Policy Sync states, see “VEN Policy Sync” in VEN Installation and Upgrade Guide.

Workload Details page

The Workload details page shows how and when workloads received staged policies.

- The General section specifies if the workload uses static policy (Policy Update Mode field) and shows the policy staging timestamp (Last Policy Staged field).
- The VEN section lists the Policy Sync state as active (syncing), staged, active, error, warning, or suspended.
- For workloads using adaptive policy, the General and VEN sections do not display these fields.

Apply a Staged Policy

See Static Policy Prerequisites, Limitations, and Caveats before you complete this task.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Apply+a+Staged+Policy.mp4>

1. From the PCE web console menu, select **Workloads**.
The Workloads page appears.

2. (Optional) Filter by Policy Mode:

Use the Workload property filter: **Policy Update Mode > Static Workloads**

3. To apply a staged policy to specific workloads:**NOTE**

- Select the workloads with staged changes.
- Click **Apply Policy** to enforce the staged rules.

4. To apply a policy to all workloads with a staged policy, select **Apply Policy > Update All Workloads**.**NOTE**

If you filter workloads by label and select Update All Workloads, the PCE applies the staged updates to all workloads matching that label scope, not just the workloads displayed on the PCE web console page.

The Apply Policy dialog box displays the number of workloads to which the staged policy will be applied.

5. Click **OK**.

The VEN applies the staged policy and displays the status of the update.

Static Policies

A static policy lets admins stage updates for workloads based on labeled scopes. These workloads will receive new firewall rules but won't apply them until manual approval is granted (via UI or API's Apply Policy action).

While adaptive security is ideal for safeguarding most workloads from threats, it's crucial to customize the application of new or modified OS-level firewall rules to specific scenarios. By configuring workloads with static policies, you dictate when VENs implement new rules received from the PCE.

Consider static policies a security control setting rather than a policy type. They allow you to manage the precise application of new firewall rules to workloads, creating an audit trail of rule application by Illumio users.

Static Policy Details

Before configuring your workloads to use a static policy, review the prerequisites, limitations, and Illumio recommendations.

Prerequisites

Access requires a Global Organization Owner or Global Administrator role.

VENs on impacted workloads must run version 17.2 or later to support static policies.

Static Policy Limitations

- Label groups must be provisioned before incorporation into the static policy.
- Immediate security updates occur in specific scenarios: new workload pairing, VEN tampering detection, or offline VENs returning online.
- Offline-then-online VENs may lead to out-of-sync rules compared to continuously online VENs.

Recommendation:

Hold security policy updates in the draft state until they are final to prevent immediate application by VENs. For optimal performance, PCE sends up to 5,000 firewall updates until completion.

Regenerate response.

Static Policy Recommendations

Implement static policies for specific cases under experienced user supervision. While the system typically updates security policies dynamically, configuring workloads with static policies alters this behavior, potentially causing inconsistencies.

Limit static policies to maintain operational efficiency as per business needs.

Use Cases for Static Policies

While the PCE typically updates security policies dynamically, there are instances where you may need to regulate when OS-level firewall rules updates apply to workloads. Here are some examples:

- **Business-Critical Application Policies:** Some organizations require explicit control over security updates for critical applications, setting specific dates and times for these updates to ensure oversight and compliance.
- **Maintenance Window Policies:** IT teams often establish policies for applying security updates during maintenance windows to minimize application downtime and mitigate risks. This approach may involve staggered upgrades to maintain application availability.
- **Environment-Specific Security Policies:** Central security teams may choose to use static policies for certain environments and adaptive policies for others. For instance, development environments might follow adaptive policies based on labels, while production environments necessitate static policies for stricter control.

See **Caveats** for guidance on choosing when to configure workloads with a static policy.

Static Policy Workflow Example

- Retail app security team configures static policy for production database tier.
- Automated scaling adds web servers during a demand spike.
- PCE updates the security policy for web servers connecting to the database tier.
- During the maintenance window, the team applies staged policy changes.
- VENS receive and apply the latest OS-level firewall rules.

Applying a Static Policy

Default Setting: Adaptive security applied across all roles, applications, environments, and locations.

Customization: Add a static policy to control OS-level firewall rule updates for workloads.

Configuration: Designate workloads by setting the Policy Update Mode in Security Settings. Define roles, applications, environments, and locations for static policy application. Multiple scopes can be added without overlap. Label groups are not supported; use separate scopes for multiple labels of the same type.

See [Static Policy Prerequisites, Limitations, and Recommendations \[15\]](#) before you complete this task.

Add a Static Policy

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Add+a+Static+Policy.mp4>

1. From the PCE web console menu, choose **Settings** > **Security** > **Static Policy**
2. To define the scope, click **Add**.
A dialog box appears, where you set the scope of the static policy.
3. Select labels to select workloads for a static policy (Role, Application, Environment, Location).
4. Click **OK**.
The static policy is listed.
5. Click **Provision** from the PCE web console toolbar.

Scope-based and Scopeless Policies

This section explains the differences between scopeless and scope policies.

Scopeless Policies

Scopeless policies are used broadly across diverse workloads. Require caution to prevent unintended communications. For example, a Default policy may open specific ports for all workloads.

Scope-based Policies

Scope-based policies can be broad or specific and are the preferred method for writing policy rules.

Scope-based policies restrict the broad application of rules, thereby limiting the impact of mistakes. However, the restrictive scope also limits the scope of how broadly rules can be written.

Single Scope Policies

Single-scope policies are preferable for rule-writing, balancing restrictions with flexibility. Enhance precision and decrease the risk of broad impact from errors.

Advanced Scope Policies

Advanced scope policies can be grouped into multi-scope policies and single-scope policies.

- **Multi-Scope Policies**
Apply rules to multiple workload groups in stages, ensuring policy application to one scope before moving to the next.
- **Single-scope Policies**
Refine rule application to specific workloads, promoting communication within a chosen scope.

Security Policy Objects

This chapter provides information about labels and label groups in Illumio Segmentation for Data Centers.

It covers creation, editing, and filtering of labels and label groups, label workloads, and the use of label groups in rules and scopes.

It also explains how to work with services, virtual services, virtual servers, IP Lists, and how to read export reports.

About Labels and Label Groups

This chapter provides information about labels and label groups in Illumio Segmentation for Data Centers:

It covers creation, editing, and filtering of labels and label groups, label workloads, and the use of label groups in rules and scopes.

Create a Label

1. From the PCE web console menu, choose **Policy Objects > Labels**.
2. On the Labels page, click **Add**.
3. Enter a label name (such as "Web") and select a label type (such as "Role").
4. Click **Save**.

Rules for Naming Labels

If you create an additional label type with a space in the 'key' (such as **ven type**), you can not group by that label type.

The label will initially display, but it will be unchecked when you add a check mark and apply the new label.

To make sure any new label types are properly added, keep the following rules in mind:

Table 1. Label Names

| Style | Name | Key | Correct Yes/No |
|------------|------------|-----------|----------------|
| Ven Type | Ven Type | ven type | No |
| new_label | new_label | newLabel | Yes |
| t e s t | t e s t | t e s t | No |
| name space | name space | namespace | Yes |
| keyspace | keyspace | key space | No |

Make sure that the Key does not contain spaces.

Label Types

| Label | Description |
|------------------------|--|
| Role | Describes a workload's function, like Web or Database. |
| Application | Describes the application a workload supports, allowing service relationships to be defined. |
| Environment | Describes the workload's stage in the product lifecycle, such as QA or production. |
| Location | Describes workload location, like Germany or AWS data centers. |
| Flexible labels | Custom label types reflect unique workload characteristics and accommodate business needs. Limited to 20 labels. |

Additional Note

Each workload can have only one label per type, promoting clear communication boundaries.

For workloads spanning roles or services across boundaries, like a database server serving multiple applications, create distinct role labels for clarity and effective communication.

Create a Label Type

Illumio Segmentation for Data Centers provides the default label types Role, Application, Environment, and Location. You can define custom label types to reflect additional characteristics of the workloads in your installation. Create

any label type that meets your organization's business needs. For example, you might label workloads according to their operating systems. The maximum number of labels is 20.

To create a new label type:

1. From the PCE web console menu, choose **Settings > Label Settings**.
2. On the Label Settings page, click **Add**.
3. Enter a unique Key. The PCE will use this key to identify the label internally, such as OS.
4. Enter singular and plural versions of the Display Name (Operating System and Operating Systems).
5. Enter a label type initial, a one- or two-character unique initial to be displayed with the icon (for example, OS).
6. Choose an icon.
7. Choose foreground and background colors to be used when the label is displayed.
8. Click **Save**.

The new label type will appear in the web console UI wherever the default label types appear, such as in the Type dropdown selector when creating a new label.

Label Groups

Label groups streamline policy creation by grouping common labels for efficient rule application. Each Label Groups list page can hold up to 10,000 groups, with individual pages supporting 10,000 members. Utilize filters to locate labels or groups.

Example: Workloads across Dallas, New York, and Washington can be collectively managed by creating a "US" Location label group. This method avoids separate rules for each location.

Displayed label group details include provision status, name, type (e.g., Role, Application), current policy use, last modified timestamp, and modifying user.

Policy Calculation Using Label Groups

Label groups can be nested, so it is essential to understand how they can impact policy.

**NOTE**

You cannot assign a label group to a workload - only individual labels can be applied to workloads. Label groups can only be used in policies.

Create a Label Group

Create label groups when you want to combine several labels that share common characteristics into a single label category. After the labels are added to a Label Group, you can use the label group in a rule.

1. From the PCE web console menu, choose **Policy Objects > Label Groups**.
2. On the Label Groups page, click **Add**.
3. In the Add Label Group page, choose the label type and enter a name for the label. You cannot create a label group name that already exists, regardless of its alphabetic case. For example, you cannot create a new label group named "WINDOWS" if the label group name "Windows" already exists.
4. Click **Save**.
5. In the Members tab, click **Add**. Use the dropdown list to find existing labels. You can also enter a label name to create a new label, click **Save**, then add the new label to the group. You can add as many labels (or label groups) of the same type to the group as desired.
6. Click **OK**.

You cannot create a label group name that already exists, regardless of its case. For example, you cannot create a new label group named "WINDOWS" if the label group name "Windows" already exists.

Use a Label Group in a Rule

Using a label group in a rule expands into multiple rules, and cross-communication is allowed.

For example, the Non-Prod label group is used again here, but in the rule, not the scope, which allows for cross-communication.

Scope:

- App: HRM
- Env: All
- Loc: US

Rule:

- Sources: Non-prod DB
- Services: MySQL
- Destinations: Non-prod DB

This means “allow MySQL from Non-Prod DB to Non-Prod DB for the HRM application in all environments located in the US,” and would allow the following communication:

- HRM | Dev | US | DB ← HRM | Dev | US | DB
- HRM | Dev | US | DB ← HRM | QA | US | DB
- HRM | Dev | US | DB ← HRM | Stage | US | DB
- HRM | QA | US | DB ← HRM | Dev | US | DB
- HRM | QA | US | DB ← HRM | Stage | US | DB

Use a Label Group in a Scope

When you use a label group in a scope, it is expanded into multiple scopes. Cross-communication is not allowed.

For example, to create a scope that applies to all environments except production, first create a Non-Prod label group consisting of labels for the Dev, QA, and Stage environments.

Scope:

- App: HRM
- Env: Non-prod
- Loc: US

Rule:

- Sources: DB
- Services: MySQL
- Destinations: DB

This means “workloads in all Non-Prod environments (Dev, QA, and Stage) can communicate within their environments with the DB using MySQL,” and would allow the following communication:

- HRM | Dev | US | DB ← HRM | Dev | US | DB

The following communication would not be allowed, since the Environment labels are different and cross-communication is not allowed:

- HRM | Dev | US | DB ← HRM | QA | US | DB
and
- HRM | Dev | US | DB ← HRM | Stage | US | DB

Labeling Workloads

You assign labels to workloads to specify their role, application, environment, location, and any custom categories you’ve defined (flexible labels like OS). Once labeled, these tags enable you to create rules based on the applied labels.

To label a workload:

- Automatically assign labels during pairing by adding labels in the pairing profile.
- Manually add labels on the workload summary page.

To apply labels:

1. Go to **Servers & Endpoints > Workloads** in the PCE console.
2. Click the workload you want to label, click **Edit**, scroll to **Label Assignment**, and then select the labels you want to apply to the workload.
3. Click **Save** to confirm the label assignment.

Filtering Labels and Label Groups

You can use the property filter at the top of the Policy Objects > Labels or Label Groups pages to find the label or label groups you want.

You can filter by label type and exact label name on the Labels page. Similarly, you can filter by label name, description, provision status, and type on the Label Groups page.

For example, select **Type: Location** in the Label property filter to see only location labels.

System Default “All” Labels

Upon initially logging into the PCE as the organization owner, the following default labels are available:

| Label | Description |
|--------------|--|
| Role | Web, Database, API, Mail, Single Node App, Load Balancer |
| Environment | Production, Stage, Dev, Test |
| Applications | None |
| Location | None |

Default Environment, Application, and Location labels are set to “All,” simplifying the creation of comprehensive policies covering All Applications, Environments, and Locations.

To prevent confusion for policy creators, Illumio advises against naming labels as “All Applications,” “All Environments,” or “All Locations” (verbatim). If labels with these names are attempted, like “All Applications,” an “HTTP 406 Not Acceptable” error will occur.

These default labels can be edited or removed at any time as needed.

Edit Labels for Multiple Workloads

You can add, modify, or remove labels on multiple workloads. This approach saves time when you want to apply or remove the same label or set of labels to more than one workload at a time.

**NOTE**

Remember that label changes do not require provisioning, so mass label changes can potentially have a major impact on your rulesets, rules, and overall security policy.

1. From the PCE web console menu, choose **Workloads and VENs > Workloads**.
2. Select the workloads you want to change labels from the left side of the Workloads list.
3. From the top of the Workloads list, click **Edit Labels**.
A dialog box appears asking if you are sure you want to edit labels for multiple workloads.
4. Click **OK**.
5. You can add or remove labels assigned to the selected workloads in the Edit Labels dialog box. The top of the dialog indicates how many workloads will be affected by the label change. Depending on the assigned labels, you have three general options:
 - When the selected workloads share the same label of a specific type (for example, Role), you can change the current label by clicking the little **X** on the label to remove it. Then, you can type or select a new label assignment.
 - When the selected workloads have different labels of the same type, faded text in the Label field indicates that they contain multiple labels of that type. You can click in the Label field to add a new label.
 - When you remove a label assignment, that label is removed from all selected workloads.
6. When you are finished, click **OK**.

Create a Service

When you create a rule, you can select a service to indicate the allowed communication between workloads and other entities.

When you create a service, that service becomes available to use in a rule.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Create+a+Service.mp4>

To create a service from the Services page:

1. From the PCE web console menu, choose **Policy Objects > Services**.
2. On the Services page, click **Add**.
3. Enter the service name and description (optional).
4. Under Attributes, choose whether you want to create a port-based or Windows service-based service.
5. In the Port and/or Protocol section, click **Add** and enter the ports, using a space to separate them from the protocol. To enter a range, separate the port numbers by a hyphen. You can also copy and paste lists of services from another source here.
6. When the service uses any UDP ports, enter them as well.
7. Click **Save**.

View or Edit a Service

To view or edit an existing service:

1. Click the name of the desired service. Various attributes can filter the list. See [Filter the Services List \[29\]](#) for details.
2. Go to **Policy Objects > Services>** to view information about the service, including its general data, attributes, and, if appropriate, the external data for the service and ransomware protection details.
3. Double-click on the Service to view the Service page and then **Edit** to enter edit mode.
 - GENERAL: Change the Name or Description of the service.
 - RANSOMWARE PROTECTION:
 - Select severity: None, Low, Medium, High, or Critical
 - OS Exposure: Select one or more OSes
 - Port Type: Admin or Legacy
 - ATTRIBUTES:
 - Operating Systems: All Operating Systems: port-based
 - Service Destinations: Add or Remove port and/or protocol

ICMP Services

ICMP can be added as a service for detailed inbound or outbound policy creation, commonly used for traceroute and path MTU discovery.

- Export ICMP traffic in JSON, CEF, or LEEF format.
- Blocked ICMP services won't appear in the Blocked Traffic list, resulting in a silent connection drop.

- Supported ICMP types/codes range from 0 to 255, allowing specific rule configurations.

The following table describes the correct format for each type of supported ICMP rule:

| Example | Format | Meaning in Rule |
|----------------------|--|--|
| ICMP (on a new line) | Protocol name only | Allow all ICMP traffic |
| 3 ICMP | Type = 3 Protocol name = ICMP | All ICMP traffic of type 3 (Destination Unreachable) is allowed regardless of the code used in the rule. |
| 3/6 ICMP | Type = 3 Code = 6 Protocol name = ICMP | Only type 3 and code 6 ICMP traffic is allowed. |
| 3 ICMP, 6 ICMP | Type 3 of ICMP, Type 6 of ICMP | Only type 3 and type 6 ICMP traffic is allowed regardless of the code used in the rule. |



TIP
Use this format to add as many types as you need.

ICMP traffic is displayed in Explorer, similar to TCP/UDP traffic.

You can see ICMP traffic flows in Illumination and the App Groups Map. You can choose to conceal them by using the filter in Illumination.

You can also create and update services that use the ICMP protocol using the REST API.

Caveats

- ICMP is not supported for virtual services.
- ICMP rules allow all types but lack granular control or specific multicast addresses.

- For IPv6 functionality on Windows VENS, specific ICMPv6 types (e.g., Router Solicitation, Router Advertisement) must be managed separately in security rules.

The ICMPv6 types that are required in those rules are as follows:

| ICMPv6 Message | ICMPv6 Type |
|--------------------------------|-------------|
| Router Solicitation Message | 133 |
| Router Advertisement Message | 134 |
| Neighbor Solicitation Message | 135 |
| Neighbor Advertisement Message | 136 |

IGMP Services

Internet Group Management Protocol (IGMP) can be added as a service for detailed inbound or outbound policy creation. It is mainly used for multicast and does not need a specific range.

- Export IGMP traffic in JSON, CEF, or LEEF format.
- Create or update services using the IGMP protocol via the REST API.



NOTE

When IGMP is used in a rule, all IGMP types are permitted, and granular control for specific multicast addresses is lacking.

IGMP isn't supported in the Illumination map.

Filter the Services List

The property filter at the top lets you filter the Services list by entering a service name, description, port, protocol, and provision status (draft or active).

| Services | | | | | |
|---|--|--------------------------------------|--|----------------------|------------------|
| <input type="checkbox"/> Provision Status <input type="checkbox"/> Name <input type="checkbox"/> Port/Protocol <input type="checkbox"/> Last Modified On <input type="checkbox"/> Description | | | | | |
| <input type="checkbox"/> | All Services | ALL | | 12/01/2020, 11:09:12 | Unknown |
| <input type="checkbox"/> | ICMP | ICMP, ICMPv6 | | 12/01/2020, 11:09:12 | Unknown |
| <input type="checkbox"/> | ADDITION PENDING Service1 | IPv6, 41 UDP | | 12/01/2020, 12:56:51 | ari@illumio.com |
| <input type="checkbox"/> | test | 22 TCP | | 04/30/2021, 11:37:41 | radi@illumio.com |
| <input type="checkbox"/> | MODIFICATION PENDING testing2 | c:\windows\myprocesses.exe myprocess | | 05/27/2021, 15:09:50 | radi@illumio.com |
| <input type="checkbox"/> | ADDITION PENDING used in VS | 22 TCP | | 04/28/2021, 14:48:48 | am@illumio.com |

Virtual Services

Virtual services (previously known as bound services) allow you to label processes or services on workloads. You can use virtual services directly in rules or use the labels applied to virtual services to write rules.



IMPORTANT

Illumio is deprecating the use of Virtual Services configurations for internal bridge networks. This is effective for versions 23.2.x and later. This aligns with Illumio's continued efforts to streamline platform support and simplify policy configuration.

See the topic on [VEN Support for Standalone Containers](#) to review and update any policies or automation that rely on Virtual Services bridging.

Virtual Services Scenarios

Utilize a virtual service in the following situations:

- **Apply rules to a single service:** Represents a workload service or process using a name or label. This method allows you to enforce a policy that permits only specific communication with that service. If the service moves to a different workload or a new set of workloads, adjustments to workload bindings on the virtual service suffice. The PCE dynamically computes necessary rules on updated workloads to facilitate this service access.
- **Apply rules to multiple services (on the same workload):** Each service or process running on a workload is represented using distinct labels. Rules can be written to enable communication exclusively with individual services. If any service is relocated to a different workload or new workloads, changes solely to workload bindings on the virtual service are needed. The PCE dynamically computes essential rules on the updated workloads for service access.

Policy Generator and Explorer support virtual services. To create rules based on labels, assign labels to a virtual service. A virtual service lacks enforcement and relies on the enforcement of its bound workloads.

Virtual services are provisionable objects, necessitating the creation and provisioning of applications on workloads beforehand. Workload bindings, however, can be altered without the need for provisioning. The most recent changes have moved port overrides from the virtual service to the workload binding.

How Virtual Services Work

Imagine a single workload running an Apache Tomcat and an Apache HTTP server, catering to HRM and ERP applications. To manage these services efficiently, virtual services can be established for each, labeling one for the HRM app and another for ERP. With label-based rules, you can distinctly control the Apache Tomcat serving the HRM app, segregating it from the ERP application.

In this scenario, two distinct virtual services are generated: one for an HRM database and one for an ERP database. Using these configurations, communication between the web and database for each application (HRM or ERP) can be controlled, considering the specific environment (Prod or QA) and location (US or EU).

Virtual Service - HRM

- **Name:** HRM-DB

- **Labels:** DB | HRM | Prod | US
- **Service:** MySQL
- **Bound to:** Workload - Database 1, Port Override: 3308
- **Scope:** HRM | Prod | US
- **Rule:** DB ← From Source ← Web

Virtual Service - ERP

- **Name:** ERP-DB
- **Labels:** DB | ERP | QA | EU
- **Service:** MySQL
- **Bound to:** Workload - Database 1, Port Override: 3309
- **Scope:** ERP | QA | EU
- **Rule:** DB ← From Source ← Web

Virtual Services in Rule Writing

When you create rules for virtual services using the Policy Generator or from Illumination, add the “Uses Virtual Services only” option or “Uses Virtual Services and Workloads” option in the Source or Destination field of the generated rules. You can configure virtual services using a port or a port range.



NOTE

Custom iptables rules and SecureConnect are not supported with virtual services.

When you write a rule in a ruleset, specify these values:

- A service
- Source of the service
- Destination of the service

For example:

The web provides Apache Tomcat service to All Workloads.

When you write rules using virtual services, you do not need to select a service in the rule, because the virtual service is both the service and the source of the service.

For example:

Virtual Service Apache Tomcat is provided to All Workloads

When you want to treat the source as a virtual service, select **Uses Virtual Services** or **Uses Virtual Services and Workloads** from the Source dropdown list as the service.

To write a rule that applies to all virtual services labeled **Database**, write it the same way and select **Uses Virtual Services** or **Uses Virtual Services and Workloads** as the providing service.



NOTE

The above rule does not impact workloads labeled "Database". You need an additional rule listing the specific service applicable to include them.

When you select a specific service, the rule applies only to workloads with the selected label.

For example, for the following virtual service rule:

- DB | MySQL | Web

The rule is only applied to workloads that use the DB label.

However, when the virtual service rule is the following type of rule:

- DB | Uses virtual services or uses virtual services and workloads | Web

The inbound side of the rule is applied to all workloads bound to the virtual service using the DB label.

Advanced Virtual Services Configuration

Consider these advanced configuration options when configuring a virtual service.

- **Optional Configuration: IP Overrides:** This feature permits specifying IP addresses or ranges (CIDR blocks) for virtual service rules, overriding the default IP addresses of bound workloads. Hosts communicating with the virtual service use the specified IP addresses and subnets instead when enabled.

Combining stateless and forwarding rules on the same host and port isn't supported. For instance, if a service on a port has stateless rules, a forwarding rule allowing traffic to a container on the same host and port doesn't function if the destination is identical.

Host-only network

Example of a virtual service rule using host network (default):

| Source | Services | Destination |
|---|-------------|---|
| Virtual Service X | From Source | Workload B |
| Virtual Service X is bound to workload A, with service 80 TCP | | Workload B has IP address 192.168.0.200 |
| Workload A has IP address 192.168.0.100 | | |

This rule programs the following security policy:

- An inbound rule on workload A for 80 TCP with source address 192.168.0.200
- An outbound rule on workload B for 80 TCP with destination address 192.168.0.100

When you add an IP override, the subnet 172.16.0.0/16 on the BPS, this rule programs the following security policy:

- An inbound rule on workload A for 80 TCP with source address 192.168.0.200
- An outbound rule on workload B for 80 TCP with destination subnet 172.16.0.0/16

The IP override dictates that devices allowed to communicate with this virtual service use the addresses/subnets specified in the IP overrides.

Filter the Virtual Services List

You can filter the Virtual Services list by using the properties filter at the top of the list. For example, you can filter and search by label. In the case of DNS-based rules, you can also filter and search by the following objects:

- Service or port
- IP entry or DNS entry (for example, search for *.google.com)

Add and Bind a Virtual Service

When you add a virtual service, enter a name, select the service, and apply labels.

Bind it to the workload where the service is running. This binding instructs the PCE on where to enforce the rules for this virtual service. When you configure two rules with the same service ports, one is stateless, and the other is stateful. The stateless rule takes precedence.

Add a Virtual Service



NOTE

A virtual service must be provisioned before it can be bound to a workload.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Add+a+Virtual+Service.mp4>

1. From the PCE web console menu, choose **Policy Objects > Virtual Services**.
2. Click **Add**.
The Add Virtual Service page appears.
3. Enter a name for the service.

4. Select the service from the Service drop-down list or enter a service name.
5. Select a Role, Application, Environment, and Location label.
6. **Host-only network:** The rules associated with the virtual service are applied over the host network and programmed into the INPUT/OUTPUT chains in Linux iptables.
7. (Optional) In the IP addresses field, you can override the IP address of the workload bound to the virtual service and specify different IP addresses or CIDR blocks that will be used for programming the virtual service rules.
8. Click **Save**.
The virtual service is created and labeled. Next, it is provisioned and bound to a workload.

**NOTE**

SecureConnect is not supported for virtual services.

Bind a Virtual Service to a Workload

Binding a virtual service to a workload enables the PCE to program rules to the VEN on the workload to which the virtual service is bound.

If the workload binding ever changes, the rules of your ruleset are dynamically recalculated for the new binding.

**NOTE**

The virtual service must be provisioned before it can be bound to a workload.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Bind+a+Virtual+Service+To+WL.mp4>

1. From the PCE web console menu, choose **Policy Objects > Virtual Services**.
2. Select the virtual service you want to bind to a workload.
The Virtual Services details page appears.
3. Click the **Workloads** tab.
4. Click **Bind**.

5. In the Workloads drop-down list, select the workload you want to bind this virtual service.
6. Select the Override ports checkbox to allow this virtual service to use a port different from the one specified.

**NOTE**

When you select **All Services** as the virtual service's service, you cannot enable port overrides on the workload bindings.

7. In the Ports/Protocols section, enter this virtual service's TCP and UDP ports.
8. Click **Save**.

IP Lists

IP lists allow you to define an **allowlist** of trusted IP addresses, IP address ranges, or CIDR blocks you want to allow into your data center to access workloads and applications in your network.

Overview of IP Lists

After you define an IP list, you can use it in policies to create rules for workload traffic flows. When you provision the policies, the workload only allows IP addresses in the IP list to access workload services.

The default IP list **Any** represents all IPv6 addresses and IPv4 addresses. Rules that use IP lists are only programmed on one side of the connection. IP lists can be used as a source or a destination.

**NOTE**

To allow outbound access to IP lists, Illumio recommends using an intra-scope rule to prevent the rule from being applied to a broader set of workloads than intended.

Examples of IP List Usage

For example, the following policy (scope + rules):

Scope:

- App: HRM
- Env: Prod
- LOC: US

Rule:

- Source: DB
- Services: SSH
- Destination: Corp-HQ

The rule means “allow outbound from DB in HRM|Prod|US to Corp-HQ IP-list using SSH”

This policy:

Scope:

- App: All
- Env: Prod
- Loc: All

Rule:

- Source: Corp-HQ
- Services: SSH
- Destination: DB

The rule means “allow inbound from Corp-HQ IP-list to any DB in Prod using SSH”

Create an IP List

1. From the PCE web console menu, choose **Policy Objects > IP Lists**.
2. Click **Add**.
3. Enter a name for the IP list.
4. **IP Addresses:** To define the list, add IP addresses, IP address ranges, or CIDR blocks.

**TIP**

You can copy and paste lists of IP addresses from other sources.

5. FQDN: Type or paste in fully qualified names**IP List Exclusions**

In IP lists, you can exclude IP addresses or subnets from a broader IP subnet.

For example, you might want to exclude a list of IP addresses within an IP range that should not access specific workloads. Or, you could open up a set of workloads to any IP address (0.0.0.0/0 and ::/0), but exclude a set of IP addresses that keep attempting unauthorized access to your workloads.

**NOTE**

Any (0.0.0.0/0) refers to IP addresses not associated with workloads, while **All workloads** refers to workloads within a scope.

When you use an IP list with exclusions in a rule, any IP addresses marked as exclusions are not allowed, while all the others in the IP list are allowed.

To create IP list exclusions:

- To add an IP address or subnet exclusion, use an exclamation point followed by the IP address, CIDR block, or IP range. The excluded IP addresses must be within the included IP range.

For example, if you added 192.16.0.0/12 as an allowed IP address and you want to exclude an IP address from this CIDR block, enter the following value:

```
!192.31.43.0-192.31.43.100
```

- To add a CIDR block but exclude a portion of the CIDR block, enter the following values:

```
10.0.0.0/8 !10.1.0.0/24
```

In this example, the first block would be included, and the second block would be excluded.

Filter IP Lists

You can filter the IP list page using the **Select properties for filter view** field at the top. Enter an IP list name, description, IP address, FQDN, and provision status (draft or active).

Virtual Servers

Virtual servers in Illumio Segmentation for Data Centers contain a Virtual IP address (VIP) and port for service exposure, along with local IP addresses used to communicate with backend servers.

Each virtual server is assigned labels and has IP addresses, but does not track traffic for Illumio Segmentation for Data Centers

Each virtual server has a single VIP, while the local IP addresses serve as source IP addresses connecting to pool members (backend servers) in SNAT or Auto mode. These IP addresses might be shared among multiple virtual servers on the server load balancer.

Virtual servers are identifiable by a set of labels. Sources and destinations can have different labels, placing them in the same or different groups within Illumination. Sources may have an incomplete label set, allowing them to be in all groups within a specified location. Therefore, a single virtual server may have sources in any group or number of groups within Illumination.

Virtual Server Members and Labels

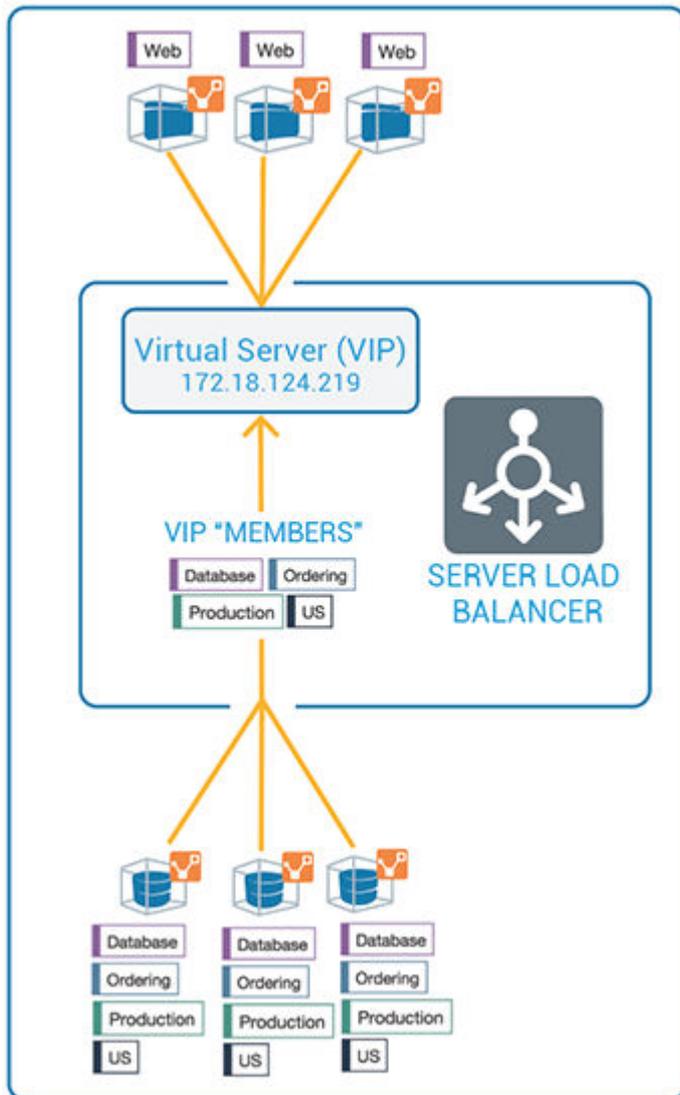
When setting up PCE load balancers, they connect via the REST API. The PCE gathers all load balancer virtual server configurations and displays them in the Discovered Virtual Servers tab on the load balancer's details page. Any virtual servers linked to the load balancer can be transformed into managed virtual servers for PCE utilization.

Labels can be assigned to the virtual server within the PCE web console during configuration. After the virtual server is set up, you can create a rule to enable external clients to interact with it.

The members associated with a virtual server are identified by setting up a series of labels in the virtual server's configuration. Four Illumio labels can be added to the Virtual Server Members tab, mirroring the labels assigned

to workloads within the virtual server’s pool. Suppose any workloads within the virtual server pool share the same four labels specified under the Virtual Server Members tab. In that case, any rules created for the virtual server will also affect the workload members.

The diagram illustrates how workloads part of the virtual server pool exhibit identical labels as those specified on the Virtual Server Members tab.



Policy for Virtual Server

The rule you can write after you label a virtual server and its members:

| | | |
|------------------------------|---------------------|-------------|
| Rule | | |
| Source: Virtual Server (VIP) | Service from Source | Destination |

Configure Virtual Servers

To manage virtual servers once a load balancer is incorporated into the PCE, you can assign each virtual server the complete set of four Illumio labels: Role, Application, Environment, and Location. Including labels on the virtual server allows you to incorporate them into a rule.

These four Illumio labels are added to the Virtual Server's Members tab. If the labels set in the Virtual Server Members correspond to the labels on workloads within the virtual server pool, any rule established for the virtual server extends to the workload members.

The configuration of a load balancer's virtual servers involves three key settings:

- **Enforced or Not Enforced:** Opting for 'Enforced' ensures that rules utilizing the labels linked to the virtual servers and their members are activated. Choosing 'Not Enforced' deactivates the labels, disabling any policies affecting the virtual server or its members.
- **Service:** Select the service necessary for rules permitting virtual server accessibility, such as HTTPD 80 TCP.
- **Labels:** The four Illumio labels—Role, Application, Environment, and Location—must be assigned to the virtual server. Label assignment is essential for integrating the virtual server into rules.



NOTE

Virtual servers are regarded as elements of a security policy. Therefore, any modifications made to a virtual server configuration require provisioning before they become active and take effect.

Virtual Server Limitations

- Illumination does not support location-level and application-level maps.
- The Illumination map does not render correctly if a single SNAT pool is shared between multiple virtual servers.
- SNAT and Auto-map modes of F5 virtual servers are supported. Transparent mode is not supported.



NOTE

You must make your changes before any virtual server configuration takes effect.

Filter the Virtual Servers List

You can filter the Virtual Servers list by using the properties filter at the top of the list. For example, you can filter and search by label. You can also filter and search by the following objects:

- Virtual server mode
- Virtual IP address, the VIP port number, or VIP Protocol
- Server Load Balancer

Configure a Load Balancer's Virtual Servers

1. Choose **Policy Objects > Virtual Servers**.
2. Categories by which you can filter Virtual Servers are:
 - a. Name
 - b. Labels
 - c. No Label
 - d. VIP
 - e. VIP Port number
 - f. VIP Protocol
 - g. Server Load Balancer
 - h. Enforcement

Virtual Server Load Balancers

Illumio Segmentation for Data Centers supports activation of enforcement on F5 BIG-IP Local Traffic Manager (LTM), BIG-IP Advanced Firewall Manager (AFM), and AVI Vantage systems.

Export Reports

Using the Export Reports feature, you can download PCE objects in JSON and CSV formats. These reports are particularly useful when you need to share data with application owners, managers, executives, or auditors who do not have access to the PCE.

CSV is the most common and popular format because it can be easily imported into other tools, such as CMDBs.

You can export the following objects into an export report: Workloads, policies, IP lists, pairing profiles, services, labels, label groups, virtual services, and virtual servers.

Generate an Export Report

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Generate+an+Export+Report.mp4>

1. From the PCE web console menu, select **Troubleshooting > Exports**.
2. On the Export page, click **New Report**.
3. **Containing All:** Select the object for which you want to generate the report:
IP Lists, Deny Rules, Service Accounts, Services, Policies, Labels, Label Groups, Pairing Profiles, Virtual Servers, Virtual Services, and Workloads.
4. **Formatted As:** Select the format, JSON or CSV.
5. **File Name:** Enter a unique name for the report.
6. Click **Export**.

Workloads

This section describes workload attributes, their enforcements, and how to create managed and unmanaged workloads.

Workloads have the following attributes:

- Workload enforcement and visibility state
- Connectivity and policy sync state
- Workload labels
- Attributes

Workload Setup Using PCE Web Console

After you pair workloads, you can view details by clicking a single workload. From the Workload Summary page, you can name the workload, write a description, and change its policy state.

Creating Managed Workloads by Installing VENs

When you install a VEN on a workload and pair it to the PCE, it becomes a managed workload because it can be managed using the PCE. For more information, see [VEN Installation and Upgrade Guide](#).

Unmanaged Workloads

Unmanaged workloads expand rule-writing capabilities to network entities not connected to the PCE and lacking an installed VEN. Integrating unmanaged workloads into the PCE allows you to craft rules enabling communication between workloads connected to the PCE and other entities. The policy between workloads with a VEN and unmanaged workloads is enforced through outbound rules on the workloads running a VEN. In the case of unmanaged workloads, the enforcement display remains blank.

For instance, unmanaged workloads representing the file servers can be added to restrict access to a network file server linked to an HRM application solely to the HRM application's database workloads. Label-based rules can then enforce the communication policy. The PCE employs outbound rules on

the database workloads with a VEN to ensure that only the databases labeled HRM can establish outbound connections to the network file servers.

Adding Unmanaged Workloads

You can add unmanaged workloads from the Workloads list. After assigning labels, write label-based rules that apply to unmanaged workloads.



TIP

You can also create an unmanaged Workload from a blocked traffic IP address.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Security+Workload+Policy+Setup+Demo.mp4>

1. In the Servers & Endpoints category, click **Workloads**.
2. Click **Add > Add Unmanaged Workload**.
3. In the Add Unmanaged Workload details page, enter a name and description for the unmanaged workload.
4. In the Label Assignment section, select the labels you want to be applied to the unmanaged workload.
5. In the Host Attributes section, enter all relevant information about the unmanaged workload, such as its hostname, location, OS Family, Release, and Public IP.
6. (Optional) In the Machine Authentication ID field, enter all or part of the DN string from the Issuer field of the end entity certificate (CA Subject Name). Complete this field when you use this unmanaged workload with the AdminConnect feature, as it involves a Windows or Linux laptop.
7. When using Kerberos for encryption, type a SPN to authenticate VEN.
8. Click **Save**.

Workload Summary

The workload summary displays information about the workload, including the user-specified attributes at the time of pairing and information that Illumio Segmentation for Data Centers has been automatically detected about the workload, specifically:

- The name of the workload

- A description (if provided)
- The workload enforcement state
- The visibility that the VEN uses
- The dates when the policy was revised and last applied
- For the workload’s VEN connectivity status, see “VEN-to-PCE Communication” in VEN Administration Guide.
- For the workload’s VEN policy sync status, see “VEN Policy Sync” in VEN Administration Guide.
- Any labels applied to the workload
- Workload system attributes (such as VEN version number, hostname, and uptime)

[Home](#) > [Servers & Endpoints](#) > [Workloads](#)

backup31

Summary Processes Rules Deny Rules Blocked Traffic Vulnerabilities Ransomware Protection

[Edit](#) [Increase Traffic Update Rate](#)

GENERAL

| | |
|----------------------|---|
| Name | backup31 |
| Description | |
| Enforcement | Visibility Only No traffic is blocked by policy |
| Visibility | Blocked + Allowed VEN logs connection information for allowed, blocked and potentially blocked traffic |
| VEN | backup31 |
| Connectivity | ● Online |
| Policy Sync | ✓ Active |
| Policy Last Received | 07/15/2024 at 00:21:26 |
| Policy Last Applied | 07/15/2024 at 00:21:26 |

RANSOMWARE PROTECTION

| | |
|---------------------------|------------|
| Ransomware Exposure | ● Critical |
| Protection Coverage Score | 0% |

LABEL ASSIGNMENT

Labels

VULNERABILITY

| | |
|-----------------------|------------------------|
| Total V-E Score | 302 |
| Highest V-E Score | 82 |
| Highest Vulnerability | 7.8 |
| Import Time | 09/28/2021 at 10:54:42 |

ATTRIBUTES

| | |
|-------------|-----------------------------------|
| VEN Version | 22.5.0 |
| Hostname | backup31 |
| Location | Amazon EC2 (US West), Oregon, USA |
| OS | ubuntu-x86_64-xenial |

Workload Enforcement States

Policy state determines how the rules affect a workload's network communication.

In Illumio Segmentation for Data Centers the workload list page includes four policy states for workloads. If a workload is unmanaged, the Policy State column is not displayed.



NOTE

The PCE representation of the enforcement state is the desired state to be applied to the next policy update. If there is an issue applying the enforcement state, a Policy Sync error will be shown for the workload.

Idle

The Idle state is used to install and activate VENs on workloads without changing the workload's firewalls. In the Idle state, the VEN on the workload does not take control of the workload's host firewall but uses workload network analysis to provide the PCE relevant details about the workload, such as the workload's network interface, operating system, and traffic flows. This information is captured in the following ways and intervals:

- Traffic flows: a snapshot is taken every 10 minutes.
- Operating system: included in the Compatibility Report every four hours.
- Workload network interface: reported to the PCE anytime it changes.

You can use a pairing profile to pair workloads into the Idle state.



NOTE

SecureConnect (IPv6 compatibility) is not supported on workloads in the Idle state. The traffic between these workloads can be impacted when you activate SecureConnect for a rule that applies to workloads in both Idle and Non-idle policy states.

Visibility Only

In the Visibility Only state, the VEN inspects all open ports on a workload and reports traffic flow between it and other workloads to the PCE. In this state, the PCE displays the traffic flow to and from the workload, providing insight into the data center and its applications. No traffic is blocked in this state. This state is useful when firewall policies are not yet known. This state can be used to discover the application traffic flows in the organization and then generate a security policy that governs required communication.

Selective Enforcement

Segmentation rules are enforced only for selected inbound services when a workload is within the scope of a Selective Enforcement Rule.

Full Enforcement

Segmentation Rules are enforced for all inbound and outbound services. Traffic that is not allowed by a Segmentation Rule is blocked.

Visibility Level

You can choose from three levels of visibility for workloads. These modes allow you to specify how much data the VEN collects from a workload when in the Full Enforcement state:

- **Off:** The VEN does not collect any information about traffic connections. This option provides no Illumination detail and demands the least amount of system resources from a workload.
This property is only available for workloads in the Full Enforcement state.
- **Blocked:** The VEN only collects the blocked connection details (source IP, destination IP, protocol, source port, and destination port), including all dropped packets. This option provides less Illumination detail but demands fewer system resources from a workload than high detail.
- **Blocked + Allowed:** The VEN collects connection details (source IP, destination IP, protocol, source port, and destination port) for both allowed and blocked connections. This option provides rich Illumination detail but requires some system resources from a workload.

Workload Processes

The Processes tab of the Workload detail page shows the processes currently running on the workload. For each process running on the workload, the following information is listed:

- V-E Score
- Process name
- Process path
- Ports used by the process
- Protocol (for example, TCP or UDP)



NOTE

On the Workload Processes tab, when you delete the binary for that process while the process is running, the PCE appends the process name with “(deleted).”

The UDP-PCE UI processes tab shows server and client UDP processes and ports.

On the Services tab for a workload, both UDP client and server processes and their port numbers appear. For TCP, only listening ports/processes are presented.

For UDP, only listening ports/processes should be presented. The information is coming from service reports sent by VEN once every 24 hours.

Customers depend on this information to understand the provider processes in their data center and write policies to allow traffic from needed workloads.

Workload Rules

Illumio Segmentation for Data Centers has two types of rules:

- **Inbound Rules:** Show all the services on the workload and the interface endpoints allowed to communicate with these services.

- **Outbound Rules:** Show all the interface endpoints with which the services on that workload can communicate.

To apply rules to a workload, create a policy, and ensure that the policies and workloads share the same labels.



NOTE

The workload rules are listed against individual IP addresses in an ipset. The PCE limits the size of the returned data.

The PCE web console displays an error message whenever the PCE exceeds a certain number of rules, which is the number of peer-to-peer rules calculated for that workload.

Workloads Blocked Traffic

The Blocked Traffic tab shows you all traffic that attempted to communicate with your workload but was blocked due to policy.

For information, see [Blocked Traffic \[62\]](#).

Filter a View

You can filter by workload name, label, hostname, enforcement, etc.

To filter a view, select a category from the list, such as Labels, and then choose an existing element within that category.

Categories you can filter on are:

- Name
- Labels
- No Label
- IP Address
- Description
- OS

- Hostname
- Policy Sync
- Enforcement
- Ransomware Exposure
- Connectivity
- Policy Last Applied
- Policy Last Received
- Policy Update Mode

Enforce a Workload Policy State

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Enforce+a+Workload+Policy+State.mp4>

1. On the left navigation, go to **Servers & Endpoints > Workloads**.
2. Click the link for the workload you want to change the Enforcement state.
3. Click **Edit**.
4. You can select Idle, Visibility Only, Selective, or Full from the Enforcement drop-down list, depending on how you want to allow or block traffic connections.
5. Click **Save**.

Set Workload Interfaces to Ignored

In the PCE web console, you can set interfaces from Managed to Ignored. This option allows you to set the workload to ignore visibility and enforcement on the interconnected interfaces of database clusters, such as Oracle RAC.

During pairing, you can set one or more interfaces to Ignored, which causes the first downloaded firewall to ignore those interfaces.

After you set an interface to 'Ignored,' it is excluded from the policy configuration, and traffic flows uninterrupted through it without any change in latency. You can see which interfaces are marked as Ignored on the Workloads' Summary page.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Set-up+Workloads+Interfaces+to+Ignored.mp4>

1. On the left navigation, go to **Servers & Endpoints > Workloads**.
2. Click a workload to open the details.
3. Click **Edit**.
4. In the Network Interfaces section, change interfaces from Managed to Ignored using the PCE Action drop-down list.

Managed interfaces will be included in policy configuration provided by PCE
i Ignored interfaces will NOT be included in policy configuration provided by the PCE. Traffic will continue to flow through the interface uninterrupted.

| Interface Name | IP/CIDR | PCE Action |
|----------------|-------------------------|------------|
| eth0 | 10.55.55.55./5 10.0.0.5 | Managed ^ |
| eth0.public | 55.111.155.220/32 | ✓ Managed |
| eth0 | fd00::200:a:0:248/64 | Ignored |
| | | Managed v |

WARNING
 DO NOT ignore PCE-generated interfaces such as `eth123.public` for cloud workloads.

If you are editing an unmanaged workload, you will not be able to ignore it using the PCE Action drop-down menu, which does not exist for unmanaged workloads.

However, you can still provide information on the Interface Name and the IP/CIDR address.

Managed interfaces will be included in policy configuration provided by PCE
i Ignored interfaces will NOT be included in policy configuration provided by the PCE. Traffic will continue to flow through the interface uninterrupted.

+ Add - Remove

| <input type="checkbox"/> * Interface Name | * IP/CIDR |
|--|-----------------------------|
| <input type="checkbox"/> • E.g. eth0.public | E.g. 10.0.10.1/24 17.1.0.10 |

5. Click **Save**.

Compare Workload V-E Scores by Enforcement Type

The **Show Vulnerability Exposure (V-E) Score** tool lets you see how the security of your workloads would change if you were to change their current enforcement mode. Columns in the Workload list and details pages provide a side-by-side comparison of the effect different enforcement modes would

have on Vulnerability and Exposure (V-E) scores. A toggle allows you to simulate the switch between Full Enforcement and Visibility Only enforcement modes.



NOTE

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

How it works

- The PCE displays V-E scores in the UI based on ransomware and previously calculated vulnerability statistics stored in a database.
- If the stored data is stale (4 hours or older), the PCE recalculates the statistics and updates the V-E scores in the UI.
- Toggling the Full Enforcement/Visibility Only options allows for a side-by-side comparison of the effects of the different enforcement modes.
- Because the PCE calculates and re-checks for new data periodically, the information in the UI may not immediately reflect the current V-E score.
- API responses include the complete vulnerability data set for the different enforcement modes. V-E data for all modes is pre-processed and stored in a database to eliminate the performance impact of frequent recalculation.
- A V-E score is the calculated value based on the Vulnerability Score and Exposure Score = $\sum f(VS, ES)$. It can be shown for an individual vulnerability on a port for a single workload or as a summation of all the V-E Scores for an App Group, role, or workload.

Workload List pages

On Workload list pages, two adjacent columns show the following:

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

| Connectivity | Full Enforcement V-E Score | Current V-E Score | Enforcement | Visibility | Policy Sync | Ransomware Exposure | Protection Coverage Score | Name |
|--------------|----------------------------|-------------------|-----------------|-------------------|-------------|---------------------|---------------------------|---------------|
| Online | 0 | 3.1 | Visibility Only | Blocked + Allowed | Active | Critical | 0% | 409_vm4.local |
| Online | 0 | 3 | Selective | Blocked + Allowed | Active | Critical | 0% | 409_vm1.local |
| Online | 0 | 0 | Full | Blocked + Allowed | Active | Protected | 82% | 409_vm2.local |
| Online | | | Full | Blocked + Allowed | Active | Protected | 82% | 409_vm3.local |

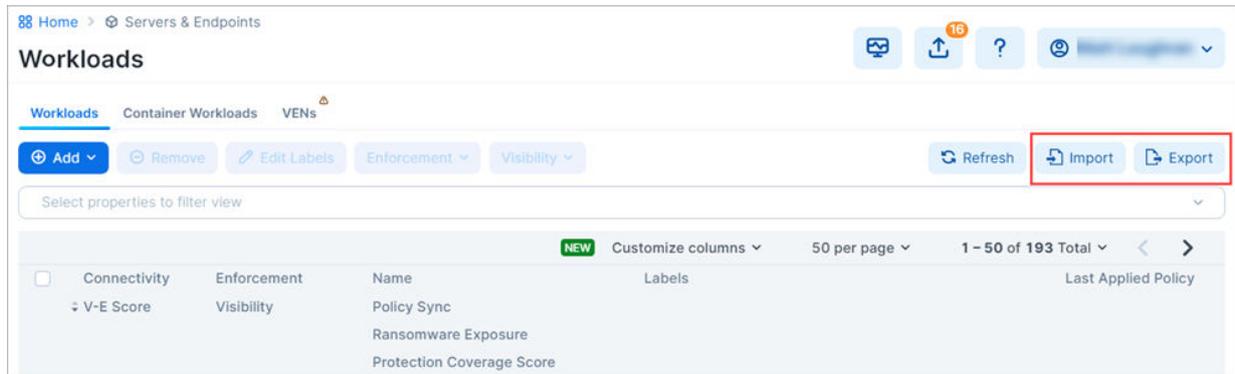
Workload Details pages

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

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

| Full Enforcement V-E Score (Total: 0) | Current V-E Score (Total: 0) | Full Enforcement Exposure | Current Exposure |
|---------------------------------------|------------------------------|---------------------------|------------------|
| 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 |

Update Workload Labels in Bulk



This section describes how to perform bulk operations on labels using the Import / Export feature available on the Workloads List Page. With this feature, you can:

- You can export a CSV or JSON file containing information about the Illumio labels assigned to your workloads. You can also export other information about your workloads.
- Import changes to your workload labels using either the CSV file you exported from the PCE or your own CSV file. You can use the Import feature to do the following:
 - Create new labels of existing label types and assign them to workloads. (Labels you create using **Import** are assigned to the workloads you specify in the CSV file. You can't use **Import** to create an unassigned label.)
 - Change a label assigned to a workload.
 - Unassign a label from a workload.

About the Export File

File format

You can export the file in these formats:

- **CSV**: This format is convenient if you use the same file to import label updates to the PCE. Only CSV files can be imported to the PCE.
- **JSON**: This option exports workload data in a JSON file. It can't be imported to the PCE.

Columns

| | A | B | C | D | E | F |
|----|--|---------------|------------|-----------|-----------|-----------|
| 1 | href | name | label:role | label:app | label:env | label:loc |
| 2 | /orgs/1/workloads/8714b5cb-5779-48cb-a898-8ddc2d856c78 | workload-1223 | Role26511 | App26511 | Env26511 | Loc26511 |
| 3 | /orgs/1/workloads/90b451e3-79dc-4976-a5af-d9034f4b95b4 | workload-1224 | Role26511 | App26511 | Env26511 | Loc26511 |
| 4 | /orgs/1/workloads/24f99794-0915-442c-862a-bf75cf2c322d | workload-1225 | Role26511 | App26511 | Env26511 | Loc26511 |
| 5 | /orgs/1/workloads/49360657-3770-4674-8a0d-0c222c988ef3 | workload-1456 | Role34592 | App34592 | Env34592 | Loc34592 |
| 6 | /orgs/1/workloads/1a80cce8-1a23-4a41-9bbc-7294c75607d5 | workload-1478 | Role94678 | App94678 | Env94678 | Loc94678 |
| 7 | /orgs/1/workloads/aa28b5bb-63f5-41b0-9058-39e40e01d8b2 | workload-1257 | role_7173 | app_7173 | env_7173 | loc_7173 |
| 8 | /orgs/1/workloads/b891fdb0-71bd-467e-86bb-3191a0fa20cc | workload-1259 | role_7173 | app_7173 | env_7173 | loc_7173 |
| 9 | /orgs/1/workloads/04f1bbc8-ec44-4c6c-ba10-0f6d981c63fb | workload-1457 | Role34592 | App34592 | Env34592 | Loc34592 |
| 10 | /orgs/1/workloads/4020a462-2d20-48ce-92f3-28360d80ea6b | workload-1479 | Role94678 | App94678 | Env94678 | Loc94678 |

By default, the exported CSV file has the following columns:



NOTE

The href and hostname columns must occupy the first and second columns from the left, respectively, and column headers should not be changed. Label column headers should not be changed, but the columns can be in any order.

- First column: href
- Second column: hostname
- label: role
- label: app
- label: env
- label: loc

Rows

With the exception of the header row (the top row), each row in the import file corresponds to a workload on the PCE.

| | A | B | C | D | E | F |
|----|--|---------------|------------|-----------|-----------|-----------|
| 1 | href | name | label:role | label:app | label:env | label:loc |
| 2 | /orgs/1/workloads/8714b5cb-5779-48cb-a898-8ddc2d856c78 | workload-1223 | Role26511 | App26511 | Env26511 | Loc26511 |
| 3 | /orgs/1/workloads/90b451e3-79dc-4976-a5af-d9034f4b95b4 | workload-1224 | Role26511 | App26511 | Env26511 | Loc26511 |
| 4 | /orgs/1/workloads/24f99794-0915-442c-862a-bf75cf2c322d | workload-1225 | Role26511 | App26511 | Env26511 | Loc26511 |
| 5 | /orgs/1/workloads/49360657-3770-4674-8a0d-0c222c988ef3 | workload-1456 | Role34592 | App34592 | Env34592 | Loc34592 |
| 6 | /orgs/1/workloads/1a80cce8-1a23-4a41-9bbc-7294c75607d5 | workload-1478 | Role94678 | App94678 | Env94678 | Loc94678 |
| 7 | /orgs/1/workloads/aa28b5bb-63f5-41b0-9058-39e40e01d8b2 | workload-1257 | role_7173 | app_7173 | env_7173 | loc_7173 |
| 8 | /orgs/1/workloads/b891fdb0-71bd-467e-86bb-3191a0fa20cc | workload-1259 | role_7173 | app_7173 | env_7173 | loc_7173 |
| 9 | /orgs/1/workloads/04f1bbc8-ec44-4c6c-ba10-0f6d981c63fb | workload-1457 | Role34592 | App34592 | Env34592 | Loc34592 |
| 10 | /orgs/1/workloads/4020a462-2d20-48ce-92f3-28360d80ea6b | workload-1479 | Role94678 | App94678 | Env94678 | Loc94678 |

CSV file requirements

Whether you're using a file exported from the PCE or your own *.csv file, the file you intend to import to the PCE must meet the following requirements:

- The file must be in a *.csv format.
- The first column header must be href
- The second column header must be hostname.
- The file doesn't need a label column for every label type defined in the PCE Label Settings (**Settings > Label Settings**).
- If you're attempting to create new labels, ensure they don't exist in your Illumio instance. If the label already exists, an error will occur, and an error message will appear.
- You can include label types other than Role, Application, Environment, and Location if they are already defined in the PCE Label Settings.
- Blank cells in the import file are ignored.
- Up to 1000 import rows per CSV are supported.

Customizing the file

If custom label types are defined in **Settings > Label Settings** on the PCE, the exported file will include columns corresponding to those Label Types. For example, if your organization defines custom label types for **OS** and **city**, the exported file will include corresponding columns.

| | Style | Name | Key | Label Type Initial | In use by |
|---------------------|-------------|-------------|------|--------------------|----------------------|
| Default Label Types | Role | Role | role | R | Labels |
| | Application | Application | app | A | Labels, Label Groups |
| | Environment | Environment | env | E | Labels |
| | Location | Location | loc | L | Labels, Label Groups |
| Custom Label Types | os | os | os | o | Labels |
| | city | city | city | c | Labels |

| B | C | F | G | H |
|---------------|------------|-----------|----------|------------|
| name | label:role | label:loc | label:os | label:city |
| workload-1223 | Role26511 | Loc26511 | linux | chicago |
| workload-1224 | Role26511 | Loc26511 | linux | chicago |
| workload-1225 | Role26511 | Loc26511 | linux | chicago |
| workload-1456 | Role34592 | Loc34592 | linux | chicago |
| workload-1478 | Role94678 | Loc94678 | windows | phoenix |
| workload-1257 | role_7173 | loc_7173 | windows | phoenix |
| workload-1259 | role_7173 | loc_7173 | windows | charlotte |
| workload-1457 | Role34592 | Loc34592 | windows | charlotte |

Procedure

STEP 1: Export Workload Information



TIP

You can skip the Export step if you plan to prepare your own CSV file for importation to the PCE. See [Step 1](#)

You can use the Export feature to create and download a file to your local computer for one or both of the following reasons:

- **Prepare for importing bulk updates.** In the exported file, specify the updates you want to make to Workload labels, as described in STEP 2: Prepare the CSV file for import. Then, import the file to the PCE, as described in Step 3.
- **Capture workload information**—export data about your workloads in a text file for informational purposes.

1. In the left navigation, click **Servers & Endpoints > Workloads**.
2. On the Workload list page, click **Export** in the upper right corner.
3. In the Export Workloads dialog box, configure settings:
 - **Export:**
 - **All Workloads:** Select if you want the exported file to include all Workloads. If no filters are applied, only this option is available.
 - **Filtered Workloads:** This option is available only if one or more filters are applied to the list of workloads. Select if you want the exported file to include only the filtered list of Workloads. Otherwise, select **All Workloads**.

- **Columns:**
 - **All Columns:** Select if you want the exported file to include all columns in the Workload List Page, including hidden columns. Note: While the exported file includes all columns, only updates you make to data in the label columns will take effect when you import the file to the PCE. Changes to data in other columns, if any, are ignored.
 - **Labeling Columns:** Select if you want the exported file to include only the label columns in the Workload List Page.
 - **File Format:**
 - **CSV:** Select CSV if you plan to use this file to import label updates to the PCE. Only CSV files can be imported to the PCE.
 - **JSON:** This option is not used for updating labels. It exports workload data in a JSON file, JSON files cannot be imported to the PCE.
4. Click **Export**. The file is sent to your Downloads folder.

STEP 2: Prepare the CSV File for Import

Here's how to prepare the CSV file for creating, assigning, updating, and unassigning labels during import.

1. Open the CSV file located in your Downloads folder and modify it in any of the following ways:
 - **Assign a new or change an existing label.**

In the appropriate label column and workload row, enter a label name or change an existing label name for each workload that you want to have the new or a changed label.
 - **Unassign labels**

In the appropriate label column and workload row, replace the name you want to unassign with any combination of alphanumeric or special characters. Later, in STEP 3: Update Workload Labels Using Import, you'll enter the exact string in "Remove the existing label" if the imported label matches the string listed below. Also, unassigning a label from a given workload doesn't delete the label for use with other workloads in the PCE.

**NOTE**

Simply deleting the label name from the CSV file and then importing the file to the PCE does not unassign the label from the workload.

As described in the above step, you must replace the label name in the CSV file with a string that you'll also enter in the **Import a CSV to edit workload labels** dialog box, as described in STEP 3: Update Workload Labels Using Import. If the strings don't match when you perform the import, an error occurs, and the label isn't unassigned.

2. Save the CSV file.

STEP 3: Update Workload Labels Using Import

The Import feature sends a CSV file to the PCE to update workload labels on your PCE. You can upload a CSV exported from the PCE (STEP 1: Export Workload Information) or prepare and upload your own CSV file.

1. Prepare the CSV file for import (STEP 2: Prepare the CSV File for Import).
2. If you have not already done so, log in to the PCE.
3. In the left navigation, go to **Servers & Endpoints > Workloads**.
4. On the Workload list page, click **Import** in the upper-right corner.
5. In the **Import a CSV to edit workload labels** dialog box, click **Choose File** and select the CSV file you want to import to the PCE.
6. Select one or both of the following options:
 - **Create labels if they don't already exist**
This option allows you to create new labels of an existing label type and assign them to workloads you specified in the CSV file. Available label types are defined in **Settings > Label Settings**.
 - **Remove the existing label if the imported label matches the string listed below**
This option allows you to unassign a label from workloads you specified in the CSV file in STEP 2: Prepare the CSV File for Import. Enter the exact string in this field that you entered in the CSV file as described in STEP 2. If the strings don't match when you perform the import, an error occurs, and the label isn't unassigned.

**NOTE**

Simply entering a string in this field and importing the CSV file to the PCE does not unassign the label from the workload. You must enter the exact string in this field in the CSV file.

If the strings don't match when you perform the import, an error occurs, and the label isn't unassigned. Also, unassigning a label from a given workload doesn't delete the label for use with other workloads in the PCE.

7. Click **Preview Changes**.
8. Review the proposed changes in the Preview Changes message.
9. (Optional) Click **Review** to see the impact of your changes before you complete the import process. Any new labels you created appear in the New Labels list. A copy button allows you to copy the details into your buffer.
Click **Back** to return to the Preview Changes message.
10. Click **Save**. The file is imported into the PCE.
- .
11. Click **Refresh** to see the label changes reflected in the workloads list.
12. If you entered a string in the CSV file to remove an existing label, delete the string from the file and then save the file. Otherwise, if you import the file again, the PCE will interpret the string as a label you want to add to a workload.

Blocked Traffic

Blocked traffic identifies blocked and potentially blocked traffic among workloads and other entities that the PCE manages.

Overview of Blocked Traffic

The Blocked Traffic option is available for each workload.

Blocked traffic alerts provide information such as the service's port and protocol, the destination's IP address, the total number of flows, and the time the last detection occurred.

| Connectivity | V-E Score | Enforcement | Visibility | Policy Sync | Ransomware Exposure | Protection Coverage Score | Name | Labels | Last Applied Policy |
|--------------|-----------|-----------------|-------------------|-------------|---------------------|---------------------------|--------------|--|----------------------|
| Online | 440 | Selective | Blocked + Allowed | Active | Protected | 17% | es-d2 | E: Development, Cn: US, R: Leading Space, Label Examples | 07/15/2024, 00:23:13 |
| Online | 366 | Visibility Only | Blocked + Allowed | Active | Critical | 0% | redisjob-d34 | E: PRD, Cn: AMER-IV2-CORE | 07/15/2024, 00:21:17 |
| Online | 351 | Visibility Only | Blocked + Allowed | Active | Critical | 0% | solr-d66 | app1, E: Staging, Cn: loc2, R: role_4 | 07/15/2024, 00:21:45 |

Under the following conditions, traffic is marked as potentially blocked or blocked based on the active policy at the PCE when the latest flow was recorded:

- Traffic is blocked when a workload is in the enforced state, and the PCE doesn't have rules in the active policy to allow that traffic.
- Traffic is potentially blocked when a workload is in a Visibility Only state, and the PCE doesn't have rules in the active policy to allow that traffic.

During the pairing process, existing connections are reported as static connections. These connections display as blocked or potentially blocked until new traffic for them is detected.

When you select the blocked connection, the Detail view provides more information on when the connection was last reported (when available).

The Blocked Traffic page allows you to verify that only unauthorized traffic is blocked and that permitted communication between workloads is not unintentionally blocked before moving workloads to the enforced state.

You can use the page buttons in the upper left to navigate the listings.

You can also use the **Refresh** button to refresh the page's content with the latest information without clearing the filters or the results.



NOTE

Only the latest 500 blocked traffic entries are displayed.

For each traffic record, the following information is displayed:

- **Traffic Type:** Specifies whether the traffic is blocked or potentially blocked and whether it is blocked by the Destination or by the source.
- **Source:** Displays the source's workload name and IP address.
- **Source Labels:** Displays labels assigned to the source.
- **Service:** Displays the process name, port, and protocol information of the reported traffic, along with an indication of whether the destination or the source reported the record.

**NOTE**

For optimal scale and performance, when the PCE has 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.

- **Destination:** Displays the workload name and IP address of the Destination.
- **Destination Labels:** Displays labels assigned to the Destination.
- **Total Flows:** Displays the total number of traffic flows for that connection.
- **Last Detected:** Displays a timestamp for the most recent recorded connection.

**NOTE**

When the source reports the record, the information in the Destination column is grayed out.

Create an Unmanaged Workload from Blocked Traffic

Sometimes, your policy might be blocked from the host's IP address you want to allow to communicate with one of your managed workloads. You can achieve this by converting the IP address to an unmanaged workload, which enables the PCE to permit its use in policy.

Click the IP address in the blocked traffic event and fill out the Unmanaged Workload page. Once you have converted the IP address into an unmanaged workload, you can use it in policies to allow other managed workloads to communicate with it, or later convert it into a managed workload by pairing it.

For more information about unmanaged workloads, see [Workload Setup Using PCE Web Console \[45\]](#).

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Workload+Blocked+Traffic.mp4>

1. From the PCE web console menu, choose **Servers & Endpoints > Workloads**.
2. Double-click on a workload.
3. You can manage the details by selecting any tab: Summary, Processes, Rules, Deny Rules, Blocked Traffic, Vulnerabilities, or Ransomware Protection.

Reject Connections

You can configure Workloads to reject traffic that does not meet the required policy instead of blocking it in the Enforced state.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Reject+Connections.mp4>

1. Select **Settings > Security >** and then the tab **Reject Connections**.
2. A new firewall security setting provides two options:
 - Reject blocked inbound traffic: When this setting is applied, the firewall is configured to send:
 - TCP RST for TCP connections
 - ICMP port unreachable for UDP connections
 - ICMP protocol unreachable for other connections
 - Drop disallowed traffic (default).

The setting acts at the VEN level, not at the interface level, and is selected by the Label set.

Loopback Interfaces

(Works with Linux VENs) VENs can report loopback interfaces and enforce policy on them.

The VEN reports all interfaces, including loopback interfaces. If the VEN detects a loopback interface but is not in the standard-defined IP block meant for loopback interfaces (127.0.0.0/8), the VEN reports this as a loopback interface to the PCE. If the workload is in the scope where loopback interfaces are to participate in policy enforcement, the workload distributes the IP address to peers and enforces policy on that interface.

The PCE web console defines the scope where loopback interfaces are to participate in policy enforcement.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Loopback.mp4>

1. Log in to the web console as a Global Ruleset Provisioner or a Global Org Owner.
2. Choose **Settings > Security**.
3. Click the Loopback Interfaces tab.
4. Choose labels to define the scope.

VEN Administration on Workloads

You can monitor the connectivity, policy sync, and health status of the VEN from the PCE web console. To view VEN health status, see the VEN list page for your managed environment. From the PCE web console menu, choose Workloads and VENs > VENs. The VEN list page appears.

VEN Suspension

You can mark a workload as suspended by using the PCE web console.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/VEN+Admin+on+Workloads+Demo.mp4>

1. Choose **Workloads > VENs** from the PCE web console.
2. Click on the VEN link to get to the VEN details page.
3. Click **Mark as Suspended**.

Workloads and VENs

The Workloads navigation menu includes Workloads, Container Workloads, and VENs. On separate tabs, you can see all your workloads, container workloads, and VENs. You can view their configuration, perform workload—or VEN-specific actions, and find the related VENs and workloads.

An idle workload does not program a firewall, so its Rules page does not show its rules.

The VENs are listed on a new page separate from workloads. The VEN-related actions are not available under the Workloads tab.

Manage Workloads and VENs



NOTE

Users with the Workload Manager role can manage workloads and Virtual Enterprise Networks (VENs).

You can select VENs to unpair, refresh, and generate support reports. The Container Workloads tab displays container workloads (if any).

Unpair a workload

Click **Unpair** to unpair a VEN.

On the Unpair VEN page, select the appropriate radio button to define the Final Firewall Status:

| Firewall Status | Description |
|--|--|
| Remove Illumio Policy | <p>This is the default option.</p> <p>Linux: Removes Illumio policy and retains the coexistent firewall rules.</p> <p>AIX/Solaris: Removes Illumio policy and reverts firewall rules to the pre-pairing state.</p> <p>Windows: Removes firewall WFP filters and activates Windows firewall</p> |
| Open all ports | All OS systems: leave all ports open |
| Close all ports except remote management | <p>Linux/AIX/Solaris: temporarily allows only SSH/22 until the system is re-booted</p> <p>Windows: allows only RDP/3389 and WinRM/5985, 5986</p> |

Proceed with unpairing as follows:

| Pairing Method | Policy Mode | Unpair Action |
|-----------------------------|--------------------------|--|
| Pairing Key | Visibility only/Enforced | <ul style="list-style-type: none"> Uninstalls the selected VEN(s). Removes policy for the associated workloads. Policies are configured into the host firewall based on options selected in "Select final firewall status". |
| Pairing Key | Idle | <ul style="list-style-type: none"> Uninstall the selected VEN(s). Removes policy for the associated workloads. No changes to the host firewall. |
| PKI Certificate or Kerberos | Visibility only/Enforced | <ul style="list-style-type: none"> Uninstall the selected VEN(s). Associated workloads become unmanaged but retain labels and IP addresses. Policies are configured into the host firewall based on options selected in "Select final firewall status". |
| PKI Certificate or Kerberos | Idle | <ul style="list-style-type: none"> Uninstall the selected VEN(s). Associated workloads become unmanaged but retain labels and IP addresses. No changes to the host firewall. |

Delete a workload from the PCE

You cannot directly delete workloads from the PCE, as the workload represents an entity that the PCE does not control. You can unpair the VEN on that workload from the VENs tab on the Servers & Endpoints/Workloads menu, removing the workload from the workloads table.

Enhanced Data Collection



NOTE

Keep these considerations in mind:

- When Enhanced Data Collection is enabled, CPU and RAM usage increases marginally on the machine hosting the VEN.
- To show byte information for VEN flows in the Bytes column on the [Label Insights](#) page, you first need to enable Enhanced Data Collection as described in this section.

When enhanced data collection is enabled, the PCE reports the amount of data transferred in and out of workloads and applications in a data center.

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

**NOTE**

In **pre-24.4.x** releases, a license is required to enable Enhanced Data Collection. For information about obtaining the license, contact Illumio Customer Support.

In **24.4 and later releases**, no license is required to enable Enhanced Data Collection.

- Select **Visibility -> Enhanced Data Collection**.

You can also enable Enhanced Data Collection as a Visibility option on the Pairing Profile page by selecting the radio button **Enhanced Data Collection**.

After you set the VEN's visibility level to enhanced data collection, it reports the number of bytes transferred over the connections. The PCE collects this data, adds relevant information, such as labels, and sends the traffic flow summaries out of the PCE.

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

- Destination Total Bytes Out (dst_tbo): Number of bytes transferred out of the source (Connection Responder)
- Destination Total Bytes In (dst_tbi): Number of bytes transferred into the source (Connection Responder)

The number of bytes includes:

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

The bytes transferred in a connection's packets are included in the measurement. This is similar to various networking products, such as firewalls,

span-port measurement tools, and other network traffic measurement tools.

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

| Conne- tion State | Description |
|-------------------------|---|
| A | Active: The connection was still active when the record was posted. This is typically observed with long-lived flows on the source and destination sides of communication. |
| T | Timed Out: Flow no longer exists. It has timed out. This is typically observed on the destination side of communication. |
| C | Closed: Flow does not exist anymore. It has been closed. This is typically observed on the source side of communication. |
| S | Snapshot: The connection was active when VEN sampled the flow. This is typically observed when the VEN is in an Idle state. |

Container Workloads

The Container Workloads page lists the containers that exist on the PCE.

The page contains this information:

| Column | Description |
|------------|--|
| Summary | <p>General: Information about the container's Name, namespace/project, policy state, etc.</p> <p>Labels: Information such as Role, Application, Environment, Location</p> <p>Attributes: Information about Interfaces and Workloads</p> |
| Containers | Information about a specific container. |
| Rules | Information about rules. |

Pairing Profiles

Pairing Profiles allow you to apply specific properties to workloads as the key pair with the PCE, such as applying labels and setting workload enforcement.

See "[Pairing Profiles and Scripts](#)" in VEN Installation and Upgrade Guide for more details.

Create Security Policy

This section describes how to create a security policy in Illumio Segmentation for Data Centers.

Creating a security policy is an iterative process. Illumio recommends creating a broad initial policy, which you can incrementally improve until you establish a sufficiently robust policy.

Policies

You can write policies that enable the workloads in your application to communicate effectively.

A policy consists of rules and scopes:

- Rules define which workloads are allowed to communicate.
- Scopes define to which workloads the rules are applied.

Basic versus Scoped Policies

You can create a basic policy or a scoped policy. When you create new policies, you can choose whether to include scopes.

When the PCE is configured to create scopeless policies, you can create simple rules that do not apply to specific environments, locations, applications, or other categories you may have defined using flexible label types. Such rules are scopeless because they do not belong to a policy that uses scopes.

When you are new to using and creating your first security policy rules, consider creating basic rules, such as a simple rule to control SSH traffic for all your workloads. As you become more familiar with Illumio Segmentation for Data Centers or need to create more complex rules, you can create scoped rules: intra-scope, extra-scope, and custom iptables rules.

Creating scoped rules enables you to define policies and rules tailored to specific environments, locations, applications (typically larger systems), or other categories you specify using flexible label types.

When the PCE is configured to create scopeless policies, you can add a scope to a policy after saving the policy.

Go to the Policies page, select a specific policy, and click **Add Scope**.

Scopeless policies in the PCE web console

The following details apply to scopeless policies in the PCE web console.

- An option in the Policy Settings page determines whether new policies are created with or without scopes. However, every user's permission to create policies is always based on the scopes they can access, even when the PCE is configured to create scopeless policies.

Disabling scopes in policies does not invalidate the Policy Manager or Policy Provisioner roles used for user authentication or Role-Based Access Control (RBAC). For more information about these roles, see [PCE Organization and Users](#).

- When the PCE is configured to create scopeless rules, the Policy details page for a policy displays a single Rules tab where you add basic rules, including container hosts as the Destination.
- When you add a scope to a scopeless policy after creating the policy, the page refreshes and displays the Intra-scope Rules and Extra-scope Rules tabs. If any rules include container hosts for Destinations, they are moved to the Extra-scope Rules tab.
- Custom iptables rules cannot be added to scopeless policies. To create custom iptables rules, you must add a scope to the policy.
- When you remove all scopes from a policy, the PCE merges the rules in the Intra-scope Rules and Extra-scope Rules tabs into a single Rules tab. However, any custom iptables rules created in the policy remain in the Custom iptables Rules tab.

Policy Scope

The scope of a policy determines which workloads receive its rules and enables the rules to apply to workloads in a group (one scope).

When workloads share the same set of labels defined in a policy's scope, those workloads receive all the rules from the policy. When you add a second scope, all the workloads within both scopes receive the rules from the policy.

A single scope is defined by using labels that identify the workload:

- **Application:** To which application (for example, ERP or HRM) do these workloads belong?
- **Environment:** Which type of environment (for example, development, production, or testing) describes these workloads?
- **Location:** Where are these workloads physically located (for example, rack server or AWS) or geographically (for example, US, EU, or CA)?
- **Flexible labels:** If you have defined custom label types, you can use them to define a scope.

A scope (or collection of workloads that the rules are applied to) is defined as ERP | Prod | US, which means that the rules apply to any workload that meets the following three requirements:

- Workloads in the ERP application
- Workloads in the Prod (Production) environment
- Workloads in the US location

That example is relatively simple, but combining rules and scopes can create complex security policies.

For example, the following policy (scope + rules):

| Scope | | |
|------------|-------------|----------|
| App | Environment | Location |
| HRM | Prod | US |
| Rules | | |
| Source | Destination | Service |
| Processing | DB | MySQL |
| Web | Processing | Tomcat |
| Corp-HQ | Web | Apache |

Allows the following communication:

- Processing | HRM | Prod | US → DB | HRM | Prod | US
- Web | HRM | Prod | US → Processing | HRM | Pod | US
- Corp-HQ | HRM | Prod | US → Web | HRM | Prod | US

Single policy scopes

Using a single scope in a policy narrows the list of workloads to which the rules apply, enabling workload cross-communication.

When you are defining rules, you have the option of using the “All” label in the scope. The “All” label applies to all instances of that label type (Application, Environment, Location, or a flexible label type you have defined). For example, creating a rule with a scope of “All | All | All” means that the rule applies to all workloads.

When you create a rule with a scope of “HRM | All | US,” this rule applies only to workloads using the HRM and US labels, regardless of Environment (“All”). For example, the following policy:

| Scope | | |
|------------|----------------|----------|
| App | Environment | Location |
| HRM | (un-specified) | US |
| Rule | | |
| Source | Destination | Service |
| Processing | DB | MySQL |

Means “The HRM application in the US can initiate communications between Processing and DB in any environment” and allows the following communication:

**NOTE**

(1) Assume below that “Dev” and “Prod” are types of Environment labels.

(2) When no label is specified in the scope for a given dimension, any label for that dimension is within the scope.

- Processing | HRM | (Env label unspecified) | US | → DB | HRM | Anything | US

- or -

- Processing | HRM | Dev | US | → DB | HRM | Dev | US
- Processing | HRM | Prod | US | → DB | HRM | Dev | US
- Processing | HRM | Dev | US | → DB | HRM | Prod | US
- Processing | HRM | Prod | US | → DB | HRM | Prod | US

Multiple policy scopes

Using multiple scopes in a policy applies the rules to each scope in isolation, preventing workload cross-communication.

For example, consider the following policy:

| Scope | | |
|------------|-------------|----------|
| App | Environment | Location |
| HRM | Prod | US |
| HRM | DEV | US |
| Rule | | |
| Source | Destination | Service |
| Processing | DB | MySQL |

This rule and scope state:

“Workloads using the HRM application in the Prod environment in the US can initiate communications between Processing and the DB.”

And

“Workloads using the HRM application in the Dev environment in the US can initiate communications between the Processing and the DB.”

The rule and scope **do not** state:

“Workloads using the HRM application in the Prod and Dev environments in the US can initiate communications between the Processing and the DB.”

This example **does** allow the following communication:

- Processing | HRM | Prod | US → DB | HRM | Prod | US

And

- Processing | HRM | Dev | US → DB | HRM | Dev | US

But **not**

- Processing | HRM | Prod | US → DB | HRM | Dev | US

Labels in scopes and rules

When the same label is used multiple times in a rule, it is expanded to multiple rules, with one label for each rule.

The following examples further demonstrate how scopes work with rules.

The following policy:

| Scope | | |
|--------|----------------|----------|
| App | Environment | Location |
| HRM | (un-specified) | US |
| Rules | | |
| Source | Destination | Service |
| Dev | Prod | MySQL |
| DB | DB | MySQL |



IMPORTANT

When no label is specified in the scope for a given dimension, any label for that dimension is within the scope.

Means:

“Allow the database used by the HRM application in the Dev environment to communicate with the database used by the HRM application in the Prod environment”

and allows the following communication:

DB | HRM | Dev | US → DB | HRM | Prod | US

The following policy:

| Scope | | |
|----------------|----------------|----------|
| App | Environment | Location |
| (un-specified) | (un-specified) | US |
| Rules | | |
| Source | Destination | Service |
| ERP | HRM | MySQL |
| Dev | Prod | MySQL |
| DB | DB | MySQL |



IMPORTANT

When no label is specified in the scope for a given dimension, any label for that dimension is within the scope.

Means:

“Allow the database used by the ERP application in the Dev environment located in the US to communicate with the database used by the HRM application in the Dev environment located in the US”

And allows the following communication:

DB | ERP | Dev | US → DB | HRM | Dev | US

The following policy:

| Scope | | |
|----------------|-------------|----------|
| App | Environment | Location |
| (un-specified) | Dev | US |
| (un-specified) | Prod | EU |
| Rules | | |
| Source | Destination | Service |
| ERP | HRM | MySQL |
| DB | DB | MySQL |



IMPORTANT

When no label is specified in the scope for a given dimension, any label for that dimension is within the scope.

Allows the following communication:

- ERP | (App label unspecified) | Dev | US → HRM | All | Dev | US
- ERP | (App label unspecified) | Prod | US → HRM | All | Prod | US
- DB | (App label unspecified) | Dev | US → DB | All | Dev | US
- DB | (App label unspecified) | Prod | US → DB | All | Prod | US



NOTE

When the service in a rule is DNS, the Destination must be in IP Lists.

Manage Policies

In this section, you will learn how to enable or disable scopes for policies, view policy status, and create policies.

Create a Policy

You can create a policy to write rules defining the allowed communication between workloads in a single or multiple groups.

When you write a rule for a Windows workload, you can add a Windows service name without specifying a port or protocol. The rule will allow communication for that service over any port and protocol.



NOTE

Illumio recommends creating no more than 500 rules per policy; otherwise, the PCE web console will not be able to display all the rules.

If you want to create a policy with more than 500 rules, split the rules across multiple policies or use the REST API, which allows you to create unlimited rules per policy.

The following task creates a single scope, which means the policy's rules apply to a single group. Add a second scope indicated by the group's labels to apply the rules to another group.

You can use a **template** or create a policy **from scratch**.

Create a Policy from Scratch

1. Choose **Policies > Add**.
2. In the Add Policy dropdown list, choose **Add from Scratch**.
3. In the Add Policy dialog, type in the new policy's name and description
4. In the **Scope** dropdown menu, select:
 - Labels and Label groups:** Select labels and label groups one by one from the list, or
 - Labels and Label Groups Except:** This option allows you to remove only the labels you want excluded.

These labels define the scope of the policy, which is its range or boundary. The scope defines the workloads affected by this policy or all workloads that share the same labels in the scope.

**NOTE**

The Scope field only appears when the PCE is configured to display it.

Add a Policy from a Template

1. Choose **Policies > Add**.
2. To create a policy from a template, you have the following choices:
 - a. **Ransomware**: This creates a set of deny rules for services and ports frequently used by Ransomware to spread across the environment.
 - b. **Inbound Admin Access**: This creates a set of rules for inbound traffic using SSH and RDP services and ports (including Jump boxes).
 - c. **Outbound Admin Access**: This creates a set of rules for outbound traffic using SSH and RDP services and ports.
 - d. **Block Internet Access**: This creates a deny rule that restricts all outbound traffic to the internet.
 - e. **Active Directory**: This creates a set of rules for default services and ports for domain controllers in your environment.
 - f. **ICMP**: Internet. Control Message Protocol is used for network maintenance and troubleshooting.
3. Select one of the templates and click **Next**.

Add a Policy for Ransomware

When you select the Ransomware template, a list of the existing deny rules is displayed.

You can confirm the selection and save or edit the Sources, Destinations, or Destination Services for any Deny rules.

- To edit the Source, click on the specific Source link. The next page will show whether the source can be edited. For example, a default IP List cannot be edited or removed.
- To edit a Destination, click on the specific Destination link.
 - Click **Add** to add new members to the label group.
 - Select as many new members from the dropdown list as you wish.
 - Click **Ok**.
 - The Label Groups page now includes the newly added members.

- Click **Provision** to get this provisioned.
- You can use this same page to remove any existing label groups.
- To edit the Destination Service, click on the specific link in that group.
 - On the Services page, you can edit the service by clicking **Edit**.
 - Change the Description, Protection Severity, or Attributes.
 - RANSOMWARE PROTECTION: Choose one of the severity levels: None, Low, Medium, High, or Critical
 - ATTRIBUTES: Use the option Service Definitions to add or remove ports and/or protocols

Add a Policy for Inbound Admin Access

When you select the Inbound Admin Access template, a list of the existing policies and deny rules is displayed.

Policy 1

- You can edit the name or scope for each policy on the policy page.
- Scope displays whether the policy contains extra-scope or intra-scope rules.
- Edit Sources, Destinations, and/or Destination Services for any existing extra- or intra-scope rules (when allowed).

DENY RULES

- Names of the Deny rules are not editable.
- Sources and Destinations of the Deny rules are not editable as well.
- The Destination Services page shows general information and attributes. To edit the service, click **Edit**.
 - GENERAL: You can edit both the name and the description
 - RANSOMWARE PROTECTION: Choose one of the severity levels: None, Low, Medium, High, or Critical
 - ATTRIBUTES: Use the option Service Definitions to add or remove ports and/or protocols.

Add a Policy for Outbound Admin Access

For the outbound admin access, there are only Deny rules.

DENY RULES

- Names of the Deny rules are not editable.
- Sources are editable, and you can add new members to the label groups using the dropdown list.

- You can also remove any of the existing members of the label group.

Add a Policy that Blocks Internet Access

You can add a deny rule restricting all outbound traffic to the internet.

DENY RULES

- Names of the Deny rules are not editable.
- Sources (applications) can be edited by adding new label group members from the dropdown list.
- Destinations (list of IP addresses) can be edited by removing any existing IKP addresses using a trash icon that appears after you double-click on the address. To add an FQDN, type or paste a fully qualified name or FQDN inside the FQDN window.
- Once the changes are in, click **Confirm and Save**.

Add a Policy for Active Directory

You can add a policy for default services and ports for domain controllers.

In the Rules for Active Directory page:

- The Name of the policy is editable.
- The policy scope is editable: add any existing label groups using the dropdown list.
- Intra-scope rules in the policy
 - Sources:
 - If denoted by  (all), rules are not editable.
 - If denoted by  (any), rules Destinations and Destination Services are editable.
 - Once the changes are in, click **Confirm and Save**.

Add a Policy for ICMP

You can add a policy for ICMP (Internet Control Message Protocol).

POLICY 1

- The Name of the policy is editable.
- The scope of the policy is editable in the following instances:
 - If denoted by  (all), the rule is not editable.
 - If denoted by  (any), the Destination Services rules are editable.

- Once the changes are in, click **Confirm and Save**.

Core Services Detector

Core services like DNS, Domain Controller, NTP, and LDAP are fundamental to your IT environment and operate on one or more workloads. The Core Service Detector tool in Illumio PCE aids in recognizing these crucial services and recommends suitable labels. With the capability to detect 51 core services, identifying and labeling these workloads is crucial since they are central connections vital for other applications.

Application owners may lack expertise in identifying core services, which different teams can manage. Collaboration between application owners and core service teams is essential for application security. By leveraging the Core Services Detector to label and set policies for core services, time spent on application policies is reduced, expediting policy implementation.



NOTE

The Core Services Detector is available only on the leader PCE in the Supercluster.

For information about using the REST API to manage core services, see "[Core Services Detection](#)" in REST API Developer Guide.

Enabling Core Services Detection

The Core Services Detector is not enabled by default because it is optional. Organizations already working extensively with labeling their core services might not be interested in this feature.



IMPORTANT

To enable Core Services detection, you must be an Illumio Org Administrator.

To enable this feature, follow these steps:

1. To allow access to the Core Services feature in the PCE, update the value for the following parameter in the PCE `runtime_env.yml` file: `core_services_enabled: true`.
2. To enable access to the Core Services in the Web console, go to **Settings > Core Services Settings**.
3. Check **Enabled** in the "Core Services Detection".
4. The Core Services menu option will now appear in the PCE web console main menu under Infrastructure. You can manage Core Services using the Illumio REST API.

Managing Core Services

Core Services Detector uses a three-step process to identify and manage core services:

1. **Detect:** The detection tool runs in the backend to recommend potential core services (workloads running core services).
2. **Review:** Review recommendations provided by the detection tool and accept or reject them.
3. **Label:** Label accepted recommendations for core services.

Detection Methods

The PCE uses three methods to detect core services:

- **Port Matching:** Rule-based model based on connections to specific ports.
- **Port-based ML:** Machine learning model based on connections to specific ports.
- **Process-based ML:** A Machine learning model based on processes running on the server.



NOTE

- The PCE's method to detect a core service is not configurable.
- All three algorithms run continuously.
- The core services detection for Microsoft Active Directory uses a machine learning (ML) model.

Detection methods can include **Port-based ML, 93% confidence**

Identifying and Reviewing Core Services

1. From the PCE web console main menu, choose **Infrastructure > Core Services**.

The landing page for core services displays all services detected by the detection tool during the most recent run.

It also tabulates the workloads recommended for running that particular core service and the ones previously accepted or rejected for that service.

2. Click the link for any of the listed core services. The page will refresh and display the detailed status of that service.

The details page for a core service provides the following information:

- **Status:** This shows whether the recommendation is new.
- **Detection Model:** Indicates the method the PCE used to detect the service.
- **Server:** Displays the IP addresses and workloads recommended for that core service. The column includes either a defined workload or an unknown IP address.
- **Labels:** For a defined workload, displays the existing labels.

To view the service's details, click either the detection method or the value in the Server column.

3. Accept or reject the core service by clicking the buttons on the right.

Accept: If the core service is from an unknown IP address, clicking **Accept** creates an unmanaged workload, such as 35.251.68.112.



NOTE

Illumio encourages customers to create unmanaged workloads, install VENS on them to make them managed, and then label them to allow enforcement.

Reject: When you reject the recommendation, that IP address is no longer recommended as a Destination of the detected core service.

Follow Up: If you are unsure whether to accept the recommendation, note your reasons to help in later decision-making.

Labeling the detected Core Services

1. Once you have accepted a recommendation to label a service, select the Accepted tab on the Core Services page.

Each service type has its own recommended label.

2. Click **Edit Labels** to see the current labels. The screen shows the current labels on the left and the recommended labels on the right. The labels

shown include All, Role, Application, Environment, Location, and any custom label types you have defined using flexible labels.

3. Click **Accept** to accept the recommended labeling.

The page refreshes and displays the labels added for the core service.

4. When required for your network environment, change the default labels by selecting **Edit Default Settings** and modifying the labels as necessary.



IMPORTANT

You must be an Illumio Org Administrator to change the default label assignments.

Default Settings



These are default label assignments for workloads providing the Microsoft-Global-Catalog Service. Editing the default setting does not affect previously edited workload Labels.

Role

R-GlobalCatalog ✕



Application

A-ActiveDirectory ✕



Cancel

OK



NOTE

Changing the default label assignment does not change any previously edited workload labels.

Scanner Detection

Scanners running in a network can be automatically detected, just as services are detected.



IMPORTANT

Scanner detection is not enabled by default. You must manually enable it on the Core Services page. After being enabled, scanner detection runs every 24 hours to detect scanner traffic.

After a scanner is detected, the `src_port` can be used to create a collector-side traffic filter so that traffic originating from that `src_port` will be dropped and not stored in the PCE.

About Rules

Rules enable communication between multiple applications or entities across different scopes or within the same scope.

Illumio supports the delegation of rule writing using role-based access control (RBAC). Application administrators can only edit rules where the policy's scope matches the scopes for which they have administrator privileges. They cannot create or manage policies if the scope includes "All."

Rule types enable the application administrator to create rules that allow other applications to communicate with the managed applications without requiring global administrator privileges. This feature enables users to group rules required for inter-application and intra-application communication within a specific application into a single policy.

You can combine multiple types of rules (intra-scope, extra-scope, and custom iptables) in a single policy.

You can use multiple services, ports, and protocols in a rule. This approach helps reduce the number of rules in your PCEs, thereby improving PCE performance.



NOTE

You cannot provision drop actions from the PCE in a NAT table for custom IP tables. Doing so results in a firewall generation failure.

Types of Rules

Rules are created based on application groups, which define the scope of the rules.

To write a rule, you need to define three things: A service, a source of the service, and a destination of the service. You also need to select the type of rule:

- **Intra-scope rule:** Allow communication within a group.
- **Extra-scope rules:** Allow communication between groups.
- **Custom iptables rules:** This policy allows custom iptables configurations. The PCE manages these rules and applies them to each managed Linux workload VEN that matches the labels for the scope and receivers.

Intra-scope Rules



NOTE

The ability to create intra-scope rules is only enabled when the PCE is configured to display it.

Intra-scope rules enable authorized users to create rules that facilitate communication between sources and destinations within a specific scope. This rule type is typically used to allow communication between workloads that belong to the same application. For intra-scope rules, the labels used in the scope must match the labels used for both the source and the destination. If you don't specify a Label, "All" is used by default.

For example, you can create a rule allowing all Database workloads with the labels HRM | US | Dev to accept MySQL connections from all Web workloads with the same labels.

For intra-scope rules, both source and destination labels must match the labels defined in the scope.

Extra-scope Rules



NOTE

The ability to create extra-scope rules is only enabled when the PCE is configured to display it.

Extra-scope rules are created for communications coming from outside the scope.

Extra-scope rules enable authorized users to create rules that facilitate communication between the scoped application and external entities. Specifically, you can write rules that allow Destinations within a scope to be accessed by Sources in or outside the specified scope. For extra-scope rules, the labels used in the scope must match the labels used by the source. If you don't specify a label, "All" is used by default.

Label Matching: Destination labels must match those used in the scope, while source labels can be outside. For example, use "All Workloads" and "All Services" between application groups to learn about the roles and services that must be allowed without prematurely restricting communications.

Once you know the exact services and workload roles, switch to "Specific Workloads" and "Specific Services" to tighten security and allow only necessary communications.

In extra-scope rules, destination labels must match those used in the scope, while source labels can be outside the scope.

MySQL may not be part of the HRM application (for example, the destinations are "Global" and are not restricted by the labels in the scope).



NOTE

If the RBAC user's scope coverage type is "Sources and Destinations," the user cannot select an IP list as the Destination. To select an IP list as a Destination in a rule, the scope coverage type must be "Sources Only." For more information, see "[Role-Based Access Control](#)" in PCE Administration Guide.

**NOTE**

Understanding and correctly applying Intra-Scope and Extra-Scope rules in Illumio is crucial for effective application micro-segmentation and security.

Begin with broad rules and refine them over time to strike a balance between security and operational efficiency.

Custom iptables Rules

**NOTE**

The ability to create iptables rules is only enabled when the PCE is configured to display it.

Illumio supports custom iptables rules to keep configurations from native Linux host setups. This allows users to include custom iptables rules within the policy, ensuring new rules do not override existing Linux configurations.

Suppose you can configure iptables directly on your Linux workloads for your application workloads as part of your host configuration. However, when you pair a workload and put a policy into the Visibility Only or Full enforcement mode, the VEN assumes control of the iptables to enact the policy and does not apply any pre-programmed iptables to the policy.

Custom iptables rules enable you to program the custom iptables rules needed for your applications as part of the rules managed by the PCE. Custom iptables rules help preserve any configured iptables from native Linux host configurations by allowing you to include them with the rules for your policy.

- **iptables** refer to a Linux host configuration before the VEN is installed.
- **Rules** refer to statements the PCE writes to determine permitted traffic, typically by assuming control of iptables and programming the new rules.
- **iptables rules** refer to iptables inserted as rules onto the VENs and managed by the PCE.

Custom rules follow the iptables `-A` (append) command pattern:

```
-t<table>-A<chain> <rule>
```

Example:

```
-t filter -A INPUT -p tcp -s 10.10.10.10 --sport 8888 -j ACCEPT
```

Custom iptables rules consist of a list of iptables statements and the entities to which the rules are applied. Each rule can consist of a list of iptables rules, which allows users to group a sequence of rules for a specific function. The custom iptables rules are programmed after the Illumio PCE generates the iptables rules, but before the last default rule.

Before they are sent to the VEN, the custom iptables rules are checked for unsupported tokens (such as names of firewall chains already in use by Illumio, matches against IP sets, and semicolons). The rule cannot be saved or provisioned if an unsupported token is included.

If the VEN fails to apply a custom iptables rule because of a missing package or an incorrectly formatted rule:

- The error is reported to the PCE and logged in the organization's event logs.
- The error is displayed in the VEN policy sync status.
- The new policy is not used; the last successful policy is used instead.

For policy distribution and enforcement, the VEN creates a custom chain that contains the rules for each table or chain in the iptables. Each custom chain is appended to the end of its corresponding chain in the correct table. When the VEN requests the policy, the iptables command is sent, including the chain where it should be placed.

For security reasons, custom iptables rules only support rules in the `mangle`, `nat`, and `filter` tables.

The following table describes the permitted actions for each iptables type:

| Table Name | Chain Names | Custom Rules Support |
|------------|--|----------------------|
| raw | prerouting, output | No |
| mangle | prerouting, input, output, forward, post routing | Yes |
| nat | prerouting, output, post-routing | Yes |
| filter | input, output, forward | Yes |
| security | input, output, forward | No |



NOTE

If the RBAC user's scope coverage type is "Sources and Destinations," the user cannot manage custom iptables rules. To allow access to custom iptables rules, the scope coverage type must be "Sources Only." For more information, see "[Role-Based Access Control](#)" in PCE Administration Guide.

For more information about APIs, see "[Custom IP Tables Rules Reference](#)" in REST API Developer Guide.

Rule Writing

This section explains how to write various rules.

Permitted Rule Writing Combinations

The following table explains the valid rule combinations between sources and destinations.

| If the Source is | And Service is | The Destination can be |
|---|--|--|
| Workload, All workloads, label, label group | Any service | Workload, IP list (including Any (0.0.0.0/0 and ::/0), label, label group, user groups, All workloads |
| IP list | Any service | Workload, label, label group, user groups, all workloads |
| Uses virtual services | Not applicable (the service is derived from the virtual service) | Workload, label, label group, IP lists, all workloads, uses virtual service, uses virtual services, and workloads. |
| Uses virtual services and workloads | Any service | Workload, label, label group, IP lists, all workloads, uses virtual service, uses virtual services, and workloads. |
| Workload, all workloads, label, or label groups | Any service | User groups and one or more of the following: workload, all workloads, label, label groups. |

Rules for Application Policies

Illumio allows or denies traffic between applications using policies that you write. To write application policies, you must create rules for the policy.

You can define and manage rules to control and secure communication within and between application groups.

Illumio has the following types of rules for application policies:

Application Policy Rule Types

There are three types of Application Policy rules:

- [Override Deny Rules \[98\]](#)
- [Allow Rules \[97\]](#)
- [Deny Rules \[97\]](#)

**NOTE**

From the release 25.2.10, all rule types (Allow, Deny, and Override Deny rules) support label exclusion.

The ability to use an "all labels except. . ." approach when selecting labels for your rules was previously available only for Allow rules.

Override Deny Rules

**NOTE**

Override Deny rules require VEN release 22.3.0 or later.

These rules block all traffic, regardless of any other rules listed below them in the policy.

Because they have the **highest** precedence, they can't be overridden by another rule, such as any implemented **Allow** rules. If an administrator creates an Allow rule by mistake, the Override Deny Rule, which denies such communication, acts as a safeguard.

They are used to stop traffic completely, especially during a security breach.

Create an Override Deny rule:

1. Go to Policies and click **Add**.
2. Select **Override Deny Rule** and then click **Add Rule**.
3. In Sources, select one or more sources.
4. In Destinations, select one or more destinations.
5. In Destination Services, select one or more services.
6. Click **Save**.

Override Deny rule implementation.

There are various implementations for Override Deny rules, such as:

- Blocking all traffic between your Production and Development environments except over `splunk-data (007 TCP)`
- Additionally, blocking all traffic between all workloads over SSH with no possible exceptions (highest precedence)

To satisfy these requirements, proceed as follows:

1. Add a Deny rule specifying 'Production' as the source and 'Development' as the destination, blocking all services.
2. Add an Allow rule specifying the same source and destination, permitting traffic over `splunk-data (9997TCP)`.
3. Add an Override Deny rule blocking all traffic between all workloads over SSH. Because this rule has the highest precedence, it cannot be overridden by an Allow rule.

Allow Rules

Allow rules have the second-highest priority, after Override Deny rules.

They allow traffic to and from specific workloads. They act like security guards, permitting only registered or authorized traffic, and are used to define explicitly permitted traffic.

Deny Rules

Deny rules temporarily block specific traffic, often during initial setup. They are useful for blocking known problematic traffic while determining what should be allowed.

In the Allow List model transition, Deny rules are gradually replaced with Allow rules, which specify precisely which traffic is permitted.

Implementing Deny Rules During the Transition to Allow Rules

Start with deny rules to block risky traffic.

- Monitor traffic patterns to understand what needs to be allowed.
- Create the Allow rules for essential, trusted traffic.
- Gradually remove deny rules as the Allow rules are established.

Once the Allow rules are fully enforced, all traffic is denied by default unless explicitly allowed by an Allow rule. Full enforcement of Allow rules ensures a secure and controlled network environment.

Conflicted Rules panel

You are now alerted when rules in the same or another policy in your organization conflict with one or more other rules.

Click the yellow icon to display a panel with the conflict details. Use the information to perform housekeeping on your policy or troubleshoot unexpected policy behavior.

Rules conflict when:

- Traffic allowed by an Allow rule in your policy is overridden by an Override Deny rule in the same or another policy in your organization.
Result: Traffic is denied, which you may or may not have intended.
- Traffic denied by a Deny rule in your policy is overridden by an Allow rule in the same or another policy in your organization.
Result: Traffic is allowed, which you may or may not have intended.

Policy Check and Rule Search

The Policy Check feature enables you to verify whether a rule allowing communication between workloads or between a workload and another IP address already exists. On the Policy Check page, you select two workloads or IP addresses to determine if a rule exists to allow communication between them. Policy checks can utilize a network profile to account for rules that affect outbound traffic to non-corporate interfaces on endpoints. Servers cannot have non-corporate interfaces.



NOTE

You can do a policy check between two workloads, a single workload, and an IP address.

For example, you have created several rule sets for your workloads and applications, and you want to know whether your organization has an existing rule for that traffic before you start writing new rules that duplicate those existing rules.

Perform a Policy Check

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Perform+a+Policy+Check.mp4>

1. From the PCE web console menu, choose **Troubleshoot > Policy Check**.
2. In the Source field, type or select a workload or IP address.
3. In the Destination field, type or select a workload or IP address.
4. In the Destination Port and Protocol field, enter a port and protocol when the connection runs over TCP or UDP, or just a protocol when it runs over GRE or IPIP.
5. Choose **Corporate, Non-Corporate Networks (Endpoints Only)**, or **Any** in the Network Profile field.

If an IP address is specified in the Destination and Source fields, the Network Profile value must be set to Corporate, which means searching within the internal corporate network only.

6. Click **Check Rules**.

If a connection between the selected two workloads or IP addresses is allowed, the page will display at least one rule that allows the connection. When a rule does not exist, the page displays “No Rules exist to allow that connection.”

Rule Search

You can't easily search for rules across rule sets when you have many rules organized in rule sets. The segmentation rule search solves this issue by making it simple to search for specific rules.

For example, it is time-consuming to narrow down the search without using this feature when determining the number of rules for SNMP (UDP 161), which has approximately 200,000 rules organized across 700 rule sets.

You can search for and analyze rules that allow communication over a specific port and protocol.

- Segmentation Rule Search lets you quickly find rules that apply to sources and destinations.

- A workload, an IP address, or a set of labels can represent sources and destinations.
- This feature helps you identify rules being applied to your workloads due to unnecessarily broad rule sets or human errors.

To search for rules

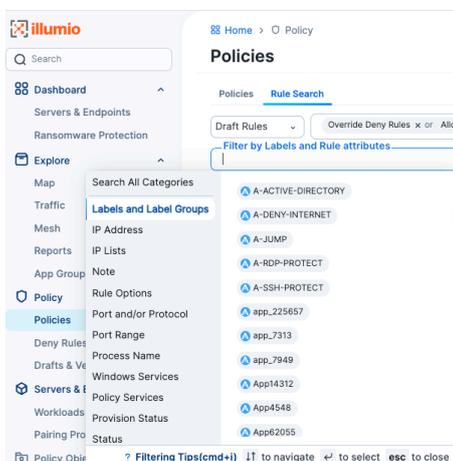
<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Search+for+Policy+Rules.mp4>

1. From the PCE web console menu, select **Policies**.
2. Choose the **Rule Search** tab.
3. Search for Active or Draft rules.
 - a. In release 25.2.10, an additional dropdown list was added to the Rule Search:

All rules: This includes `Override Deny Rules`, `Allow Rules`, `Deny Rules`, and `IP table rules`. When no type category is selected, all rule types are shown.

Subset of rules: Select which rules you want to search by deselecting the other rule types.
 - b. Choose an **Exact Match** of the selected search filters displayed or a match to any of the selected filters (**All Results**).
 - c. Perform a **Basic search** for all attributes or an **Advanced search** by destination, source, or both.
 - d. Filter by Labels and Rule Attributes. Use these options to narrow your search results. You can search all categories or select only the ones you want to use, such as labels and label groups, IP address, IP lists, Rule options, Port and/or Protocol, Port range, Process name, Windows services, Policy services, Provision status, and Status.

Previous releases limited the number of labels per rule search using the UI to eight. This limitation was removed in release 25.4, and users can search for more than eight labels using the UI.



4. Click **Run**.

5. Click **Export** to export the search result in JSON format.

Stateful vs. Stateless Rules

By default, all rules you write in the PCE are stateful, meaning the host's firewall keeps track of a connection for the entire session duration.

Stateless Rules

For workloads, you can specify stateless packet filtering for a rule ("stateless": true). This means the VEN instructs the host's firewall not to maintain persistent connections for all sessions. You can create this stateless rule for data center core services like DNS and NTP.

Caveats

In a stateless rule, you can add the following policy objects as destinations:

- An individual workload
- A label (one each of a specific type, up to four total)
- Any IP list plus all workloads

If you attempt to add any other destinations, you receive an error.

The limit ensures that the number of stateless rules is capped at 100, allowing both stateful and stateless rules to coexist on the host in a way that optimizes system and network performance. If you require more than 100 stateless rules in your Illumio policy, please contact your Illumio Professional Services Representative for further information.



WARNING

Existing active connections on workloads allowed by a stateless rule (for example, an SSH session) are terminated when workloads receive new rules from the PCE. Clients need to reestablish those connections. For this reason, Illumio recommends using stateless rules for services that utilize high-frequency, short-lived connections, such as DNS and SNMP.

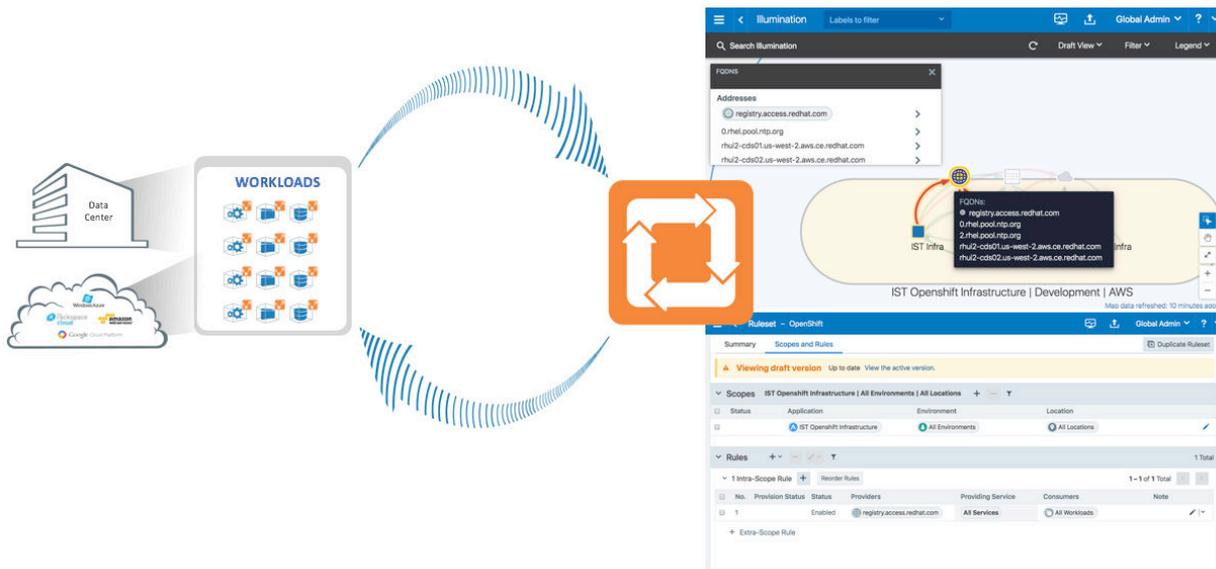
FQDN-Based Rules

Applications in data centers and cloud environments generate significant east-west traffic due to communication between various workloads, such as bare-metal, virtual machines, and containers. Additionally, these applications often need to interact with external services like SaaS, PaaS, or registries, which are accessed via frequently changing IP addresses or URLs. This creates a security challenge as traditional security policies rely on IP addresses or subnets. While administrators might allow broad outbound communication to mitigate this issue, it poses a security risk. To address this, Illumio has introduced FQDN-based visibility and enforcement in Illumio Segmentation for Data Centers

Benefits of FQDN-Based Rules

Implementing FQDN-based rules has the following benefits:

- **Deeper visibility:** Delivers visibility into communications from workloads to any workload reachable via a URL. For example, when a workload needs to pull an image from an unmanaged repository or use Amazon RDS for database services, Illumio provides visibility to those FQDNs, not just the IP addresses behind them.
- **Natural language policy:** Automatically generate or write allowlist policies that allow workloads to consume services from FQDNs rather than IP addresses or subnets.
- **Adaptive security:** Using distributed DNS snooping at the workload, PCE dynamically conforms policy to any changes, such as a domain name resolving to a new IP address.
- **Lock-down outbound communications and reduce risk:** With FQDN-based enforcement, you decide which outbound services should be allow-listed for your application rather than allowing all outbound communications. This ability mitigates the risk of applications potentially communicating with a malicious IP address or domain name.
- **Wildcard support:** Enables you to write FQDN-based policy using wildcards, such as *.redhat.com.



Features of FQDN-Based Rules

Distributed DNS Snooping

The VEN snoops DNS responses each time a workload sends a DNS request, gathering data and storing it in its DNS cache without generating DNS requests. It tracks DNS responses, avoiding repeated requests.

DNS Visibility

The VEN reports flow data such as IP addresses, ports, and protocols to the PCE. It also maps FQDNs to outbound flow data and reports DNS-based traffic flows in near real-time and for historical data retention.

DNS Enforcement

Security teams can create allowlist policies for FQDNs, specifying which DNS hostnames or FQDN workloads can communicate with them without knowing the associated IP addresses.

Wildcards

Illumio Segmentation for Data Centers supports wildcards in FQDNs, such as *.google.com, to simplify rule creation. For optimal performance, Illumio recommends limiting it to around 100 entries.

FQDN-Based Rule Requirements and Limitations

FQDN-based visibility and enforcement are subject to the following requirements and limitations:

- FQDN-based rules are not supported for any type of container technology with the VEN.
- Supported for any Linux OS supported with the Illumio VEN 19.1.0 release.
- Supported for any Windows OS supported with the Illumio VEN 19.1.0 release.
- Supported for any Mac OS supported with the Illumio endpoint VEN 23.2.0 release.
- Solaris and AIX workloads are not supported.
- Visibility and enforcement for DNS-based traffic when the source is a DNS hostname are not supported.
- FQDNs can be described in IP lists or virtual services, but not in an unmanaged workload interface.
- Only one FQDN (wildcard supported) can be specified when using virtual services. IP lists can support a list or a group of FQDNs.
- A mix of virtual services and IP lists is supported.
- A period character is not supported in a wildcard. For example, **www.server*.mycorp.com** matches **www.server1.mycorp.com** but not **www.server1.farm2.mycorp.com**.
- A wildcard-only entry (specifying only “*”) is not allowed.



IMPORTANT

A wildcard will not cover subdomains. For example, ***.mycorp.com** will not match **host1.downloads.mycorp.com**



WARNING

A workload can receive more than 100+ FQDN entries. However, problems can occur once there are more than the maximum of 100 FQDN rules in a single IP list.

For better performance, when you write FQDN-based rules, limit the number of rules to around 100 entries.

FQDN Visibility

Illumio requires no new configuration to gain visibility into outbound traffic towards FQDNs. However, you can create Illumio policy objects representing an FQDN or a list of FQDNs. Illumination presents outbound FQDN flows in the following example when no policy objects have been created. A web server is fetching updates from `us-west-1.ec2.archive.ubuntu.com`.

You can create an Illumio policy object, such as an IP list or a virtual service representing the FQDN.

Create Policy Objects for FQDNs

IP List

By default, IP lists can describe IP ranges, groups, and subnets. From the 19.1.0 release on, IP lists can also describe FQDNs.

You can use the previous example (`us-west-1.ec2.archive.ubuntu.com`) to create an IP list for FQDNs:

1. From the PCE web console menu, choose **Policy Objects > IP Lists**.
2. Click **Add**.
3. Enter a name (can be a custom name).
4. In the IP Addresses and FQDNs field, enter one or multiple FQDNs (wild-cards are supported).
5. Click **Save**.
6. Provision the changes.



IMPORTANT

The provided checkbox can be selected to "Disable validation of IP addresses and FQDNs." When working with large sets of IP Addresses and FQDNs, it is recommended that you disable real-time IP address and FQDN validation for performance reasons.

The following methods of describing the specific FQDN are supported:

Supported examples

- us-west-1.ec2.archive.ubuntu.com
- *.ec2.archive.ubuntu.com
- *.*.archive.ubuntu.com
- *.*.*.ubuntu.com

You can use a wildcard in the IP list, such as ***.ec2.archive.ubuntu.com**.

Virtual Service

When you have created an IP list to describe the FQDN, you do not need to create a virtual service to describe the same FQDN.

You should only create a virtual service for an FQDN when you do not want to create an IP list:

1. From the PCE web console menu, choose **Policy Objects > Virtual Services**.
2. Click **Add**.
 - Enter a name.
 - Enter a service or port.
 - Enter your R-A-E-L labels for the FQDN.
 - Click **Add FQDN** and enter an FQDN.
3. Click **Save**.
4. Provision the changes.

Based on the example above, these methods of describing the specific FQDN are supported or unsupported.

Supported

- us-west-1.ec2.archive.ubuntu.com
- us-west-1.ec2.*.ubuntu.com
- *.ec2.*.ubuntu.com
- us-*.ec2.archive.ubuntu.com

The syntax below is supported, but does not describe the FQDN in the example.

- ubuntu.com
- *.ubuntu.com

Write Policies to Allowlist FQDNs

IP List

The syntax and ruleset structure for IP list policies does not change for FQDNs.

| Ruleset Scope Example | | | |
|---|-------------------|---------------|---|
| Application | Environment | Location | |
| HRM | Production | All Locations | |
| Intra-Scope Rule Example | | | |
| Destination | Providing Service | Source | Note |
| *.ec2.archive.ubuntu.com (IP List object) | All Services | Web | You can use 80 TCP as the providing service |

Virtual Service

Writing a policy against a virtual service for an FQDN is the same as writing a policy for an IP-based virtual service.

See the following example that uses the Ubuntu Repo (*.ec2.archive.ubuntu.com):

| Ruleset Scope Example | | | |
|--|--|---------------|---|
| Application | Environment | Location | |
| HRM | Production | All Locations | |
| Intra-Scope Rule Example | | | |
| Destination | Providing Service | Source | Note |
| Ubuntu repo (Virtual Service role label for *.ec2.archive.ubuntu.com + Uses Virtual Services Only) | Derived from Destination Virtual Service | Web | There are two objects selected in the Destination column; one is for the Role label, and the other is called "Uses Virtual Services Only" |

Windows Process-Based Rules

Rules can be created to allow all system-initiated processes in Windows. This approach allows all traffic related to drivers and other operating system modules.

You can create a service of type Windows—process or service-based—with the word “system” (case-insensitive) in the Port/Protocol text input field. Once you create this service, you can use it in the rules.

Creating Services with System-Initiated Processes

To create a service that allows for all system-initiated processes:

1. From the PCE web console menu, choose **Policy Objects > Services** .
2. Click **Add**.
3. Enter a name and definition for the service you are adding.

To add a service definition, from the Operating System drop-down, select either All Operating Systems: Port-Based or Windows Inbound: Process/Service-Based:

- If you select All Operating Systems: Port-Based , you can only indicate a port, a protocol, or both, separating the port and protocol with a space. For example, port **512 TCP**.
- If you select Windows Inbound: Process/Service-Based from the Port and/or Protocol drop-down, specify a port/protocol, a process or service, or a port/protocol with a process or service, separating the port and protocol with a space. For example, port **512 TCP**, process **C:\windows\myprocess.exe**, and Windows service, **myprocess**.

Select **All Operating Systems: Port-Based** or Windows **Inbound: Process/Service-Based** to remove a service definition from the Operating System drop-down.

- Click the check box next to the Port and/or Protocol. You may select a single or multiple entries.
- Click **Remove**.

Windows Environmental Variables

The Windows environmental variable can be used to specify a full path.

This can be done by creating a service type Windows: Process- or Service-based with the environment variables in the Port Protocol text input field.



NOTE

Currently, only the Windows System variable is supported for use in the process path.

For example, `%systemroot%\myprocess.exe`.

Rules can be created for all system-initiated processes in Windows, allowing all traffic related to drivers and other operating system modules.

This can be done by placing the word `system` (case-insensitive) in the text input field.

Creating a service with Windows environmental variables

To create a service that uses Windows environmental variables, do the following:

1. Choose **Policy Objects > Services**.
2. Click **Add**.
3. In the Name field, enter **system** (case-insensitive).
4. Select Windows Inbound: Process/Service-Based from the Operating System drop-down list.
5. In Port and/or Protocol, specify the port/protocol, separating the port and protocol with a space.
For example: `%systemroot%\myprocess.exe`
6. Click **Save**.

Rule-Based Labeling

Rule-based labeling allows you to assign labels to one or more workloads when their attributes match the conditions you specify in easily configurable rules. This simplifies the task of labeling multiple workloads.

Before you begin Rule-Based Labeling

- **Rule limit:** You can create up to 600 rules per Illumio org.
- **Label assignment:**
 - You can assign system default and user-defined labels to matching workloads.
 - You can assign only one label of a given type to a workload.
 - Beginning in release 25.21, Rule-Based Labeling can replace existing labels already assigned to workloads if the **Overwrite** option is selected. Otherwise, existing labels already assigned to workloads can't be overwritten. For example, if Overwrite is selected and a matching workload has an existing Location label of New York and your labeling rule specifies a Location label of London, the existing New York Location label is replaced with the London Location label. If Overwrite is not selected, the London Location label is bypassed, and the New York label remains.
 - Depending on how many workloads match labeling rules, it may take a few minutes for the labels to be assigned to all of them. You can navigate to other areas of the PCE UI while the load process continues in the background. When matching and loading have finished, a notification appears wherever you are in the PCE user interface.
- **Events:** An event is created when a rule-based label is assigned to a workload. The name format of the event differs depending on how the label is assigned. When a label is assigned from the PCE UI, the name format is `label_mapping_rules_run.assign_labels`. Also, the `generated_by` field displays the user's email address. For system jobs, the **generated_by** field displays `system`.
- **Removal restriction:** It's impossible to remove a label from the list of labels (**Policy Objects > Labels**) if used in a labeling rule.

Typical Labeling Rule Workflow

Here is a typical workflow for adding rules, launching a search for matching workloads, and assigning labels.

Step 1: Add a Labeling Rule

Labeling rules work by identifying workloads in your environment that match certain conditions you specify and then assigning one or more labels to those workloads. See [Add a Labeling Rule \[113\]](#).

Step 2: Find and review matching workloads

After adding labeling rules, let the Rule Labeling feature search your environment for workloads that match the rule conditions. Then, review the generated list of workloads. [Find and Review Matching Workloads \[115\]](#).

Step 3: Assign labels to matching workloads

Once the feature finds matching workloads, you can assign the labels specified in Step 1: Add a Labeling Rule. See [Assign labels to matching workloads immediately \[116\]](#) or [Schedule Label Assignments \[117\]](#).

Work with Labeling Rules

This section describes how to add, remove, reorder, edit, and enable/disable labeling rules. It also includes procedures for finding and matching workloads and exporting a list of labeling rules to a CSV file.

Add a Labeling Rule

Labeling rules work by identifying workloads in your environment that match conditions you specify and then assigning one or more labels to those workloads.



NOTE

You can create up to 600 labeling rules per Illumio org.

1. (Optional) To determine the workload attributes you want your labeling rule to match, it may help to go to **Servers & Endpoints > Workloads** and examine the workloads in your environment.

2. Go to **Policy Objects > Labels**.
3. Click the **Labeling Rules** tab.
4. Click **Add**.
5. Specify the matching condition.
 - a. Select an attribute.
 - b. Select an operator.

About the regex match operator

(Not available in all Illumio Core releases)

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

Regular Expression operators (regex) allow you to define complex patterns to precisely match workloads in your environment. This precision is particularly useful when you're trying to find and label workloads that have multiple attributes.

As with all operators, **regex match** can be used singly or in combination with other operators to search for attributes.



NOTE

When using the asterisk (*) wildcard character, you must precede it with a period.



TIP

While Rule Based Labeling performs some validation of the regular expressions you enter in the Values field, it may help when crafting complex patterns to use an online regex validator such as <https://regex101.com/>.

Table 2. Regular Expression (Regex) Examples

| Goal | Attribute | Operator | Values | Resulting Condition |
|---|---------------|----------------------|---------------------------|--|
| Hostname match case insensitive for 'Example' | Host-name | regex match | Example.* OR example.* | Hostname regex match Example.* OR example.* |
| Port/Protocol match for TCP 22 | Port/Protocol | regex match | 22 tcp | Port/Protocol regex match 22 tcp |
| Process path match on /usr/bin | Process | regex match | /usr/bin.* | Process regex match /usr/bin.* |
| Match for hostnames formatted as word-word-digit-digit (west-prod-2-0) | Host-name | regex match | (\w+)-(\w+)-(\d+)-(\d+) | Hostname regex match (\w+)-(\w+)-(\d+)-(\d+) |
| Return results that do not match hostnames formatted as word-word-digit-digit (west-prod-2-0) | Host-name | does not match regex | (\w+)-(\w+)-(\d+)-(\d+) | Hostname does not match regex (\w+)-(\w+)-(\d+)-(\d+) |

c. Specify one or more values.

6. Select one or more labels in the **Label** field.
7. (Not available in all releases) Select **Overwrite** if you want to replace existing labels of the same type. For example, if a labeling rule is set to assign a Location label to matching workloads, any Location label(s) that may have been assigned previously to these workloads will be overwritten by the new Location label if Overwrite is selected. Otherwise, the existing label is preserved. This behavior applies to labels of any type.



CAUTION

Label changes are likely to affect your security policy. Make sure you've considered potential policy changes before you select the Overwrite option.

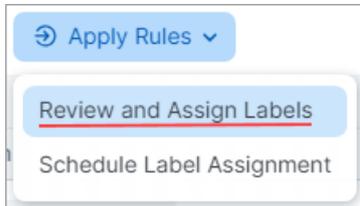
8. Click **Save**.

Find and Review Matching Workloads

This procedure describes how to search your environment for workloads that match the rule conditions.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Find+and+Review+Matching+Workloads.mp4>

1. Go to **Policy Objects > Labels**.
2. Click **Apply Rules** and then choose **Review and Assign Labels**.



The Workloads that match criteria side panel opens, showing the workloads in your environment that match your rules (if any).



NOTE

Depending on the number of workloads that match labeling rules, it may take several minutes for the PCE to load the workloads that match your rules. You can close the **Workloads that match criteria** side panel while the load process continues in the background. A progress message appears on the main page while the operation is underway. When matching and loading have finished, a notification appears wherever you are in the PCE user interface.

3. Review the list to ensure it includes the workloads you want your rules to match. If the list doesn't include the workloads you intended, click **Close**, recheck the condition(s) you specified in the rule(s), and then modify the rules if necessary. You may need to return to the Workloads page and re-examine the workloads to make sure you've specified the correct workload attributes in your rule(s).
4. If the list of matching workloads meets your expectations, assign the specified labels.

Assign labels to matching workloads immediately

To immediately assign labels to the workloads that match your labeling rules, perform these steps.

**NOTE**

In certain use cases, it may be preferable to assign labels immediately as described in this procedure rather than using the [Apply Rules when triggered \[117\]](#) option.

1. Go to **Policy Objects > Labels**.
2. Make sure the **Workloads that match criteria** side panel is open (see Find and Review Matching Workloads).
3. From the **Workloads that match criteria** side panel, click **Assign**. The message **Labels have been assigned to _ workloads** appears.

To assign labels to workloads programmatically, see [Schedule Label Assignments \[117\]](#).

Schedule Label Assignments

If you aren't assigning labels immediately, perform these steps to specify when you want to assign labels.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Schedule+Labeling+Assignments.mp4>

1. Click **Apply Rules** and then select **Schedule Label Assignment**.
2. In the **Recurring Rule Application** dialog box, move the slider(s) to **On** to enable one or both of the following options:
 - **Apply rules when triggered**. Enable this option if you want labels to be assigned automatically to the matching workload(s) whenever a [VEN is activated](#). Note the following about using this option.

**NOTE**

- **Four-hour pause between searches.** Every four hours, Rule-Based Labeling searches for VENs in your environment that were activated within the past four hours. If the search finds such VENs, labels are assigned to the VEN's host workloads if the workloads' conditions match any of your labeling rules. Labels are not reassigned to previously labeled workloads because the search ignores VENs activated more than four hours previously.
- **Activating multiple VENs over a brief period of time.** If your organization uses a tool to automate VEN activation for multiple VENs within a brief time period and you've enabled the **Apply rules when triggered** option, be aware of the following:
 - a. Your tool activates VENs according to the cadence you configured.
 - b. Activation of the first VEN triggers Rule-Based Labeling to search your environment for matching workloads.
 - c. After Rule-Based Labeling finds the first matching workload and assigns labels to it, further search for matching workloads and label assignment is halted for four hours, which you may not have expected.
 - d. When the four-hour pause has ended, Rule-Based Labeling resumes its search for matching workloads and assigns labels to them according to your labeling rules.

To avoid waiting four hours as described above, you can assign labels to the remaining matching workloads immediately by performing the steps in [Assign labels to matching workloads immediately \[116\]](#). The subsequent search that occurs after four hours still runs, but ignores the workloads to which labels were already assigned. Labels are not overwritten.

- **Apply rules regularly.** Enable this option if you want Rule-Based Labeling to assign labels automatically according to a schedule. Click through the Date and Time options to configure a schedule.

3. Click **Done**.

Edit a Labeling Rule

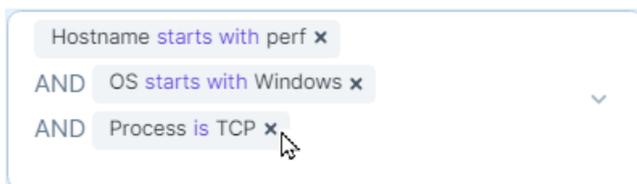
You can edit a rule's condition and label(s).

To add a statement to an existing rule:

1. Go to **Policy Objects > Labels**.
2. Click the **Labeling Rules** tab.
3. Click the **Edit** icon for the rule you want to edit.
4. Click the down arrow to activate the Condition selectors.
5. Specify the statement you want to add.
6. If needed, add or remove label(s) in the **Label** field.
7. Click **Save**.

To delete a value from an existing rule:

1. Go to **Policy Objects > Labels**.
2. Click the **Labeling Rules** tab.
3. Click the **Edit** icon for the rule you want to edit.
4. On the condition you want to delete, click the **X** to delete it.



5. If needed, edit label(s) in the **Label** field.
6. Click **Save**.

To edit a value in an existing condition:

**NOTE**

To change a value in an existing condition, you must delete the original condition and then re-add it, specifying the value you want. You can't directly edit a value in an existing condition and preserve it.

For example, if you want to change the IP range

10.13.0.26-10.13.8.26

to ...

10.13.0.26-10.92.8.26

... you must add the new range as a new condition and also delete the original condition.

1. Click the **Edit** icon for the rule you want to edit.
2. Click the down arrow to activate the Condition selectors.
3. Add the new statement.
4. Delete the original value.
5. If needed, edit label(s) in the **Label** field.
6. Click **Save**.

Enable/Disable Labeling Rules

The Enable/Disable options allow you to generate different matching results by excluding or including one or more labeling rules from the workload matching process.

https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Enable_Disable+Labeling+Rules.mp4

1. Go to **Policy Objects > Labels**.
2. Click the **Labeling Rules** tab.
3. Select one or more labeling rules in the list of rules.
4. Click **Enable** or **Disable**.

- To see the effect of the enable/disable option you selected, re-run the workload matching process.

Reorder Labeling Rules

When labeling rules are assigned, evaluation begins from the top of the list in ascending order (Rule 1, then Rule 2, etc), with Rule 1 having the highest precedence.

To change the precedence of a rule, change its rule number in the list of rules. Note that this will also reorder other rules in the list and change their precedence accordingly.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Reorder+Labeling+Rules.mp4>

- Go to **Policy Objects > Labels**.
- Click the **Labeling Rules** tab.
- Click the **Edit** icon for the rule you want to move. The rule number becomes an editable field.
- Enter the new rule number in the field.
- Click **Save**.

The screenshot shows the 'Labeling Rules' tab in the Illumio Security Policy interface. The interface displays a list of rules with their numbers and conditions. Rule 1 is highlighted in blue, and its number '1' is enclosed in a red box. The conditions for Rule 1 are: 'Hostname starts with perf- x', 'AND Hostname ends with 3345 x OR 3346 x', 'AND Process is /usr/bin/914c-g x OR /usr/bin/3com-njack-1 x', and 'AND Port/Protocol is 211 UDP x OR 5264 UDP x'. The interface also shows buttons for 'Add', 'Remove', and 'Apply Rules'.

| No. | Condition |
|-----|--|
| 1 | Hostname starts with perf AND OS starts with Windows |
| 2 | Hostname starts with perf-workload-3390 |
| 1 | Hostname starts with perf- x AND Hostname ends with 3345 x OR 3346 x AND Process is /usr/bin/914c-g x OR /usr/bin/3com-njack-1 x AND Port/Protocol is 211 UDP x OR 5264 UDP x |
| 4 | OS is Solaris |

**NOTE**

Note that reordering rules changes the precedence of other rules.

- The former Rule 3 becomes Rule 1 with the highest precedence.
- The former Rule 1 moves to become Rule 2.
- The former Rule 2 moves to become Rule 3.

The screenshot shows the 'Labeling Rules' configuration page. At the top, there are buttons for 'Add', 'Remove', and 'Apply Rules'. Below is a table of rules:

| No. | Condition |
|-----|--|
| 1 | Hostname starts with perf- AND Hostname ends with 3345 OR 3346 AND Process is /usr/bin/914c-g OR /usr/bin/3com-njack-1 AND Port/Protocol is 211 UDP OR 5264 UDP |
| 2 | Hostname starts with perf AND OS starts with Windows |
| 3 | Hostname starts with perf-workload-3390 |
| 4 | OS is Solaris |

Remove Labeling Rules

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Remove+Labeling+Rules.mp4>

1. Go to **Policy Objects > Labels**.
2. Click the **Labeling Rules** tab.
3. Select one or more labeling rules in the list of rules.
4. Click **Remove**.

Export a Workload-Label-Review List

You can export a CSV file showing the workloads that match your rules and the label(s) assigned to those workloads. This is helpful when you have a large number of rules and workloads.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Export+a+workload+label+review+list.mp4>

1. Go to **Policy Objects > Labels**.
2. Click the **Labeling Rules** tab.
3. Click **Apply Rules** and then click **Review and Assign Labels**.
4. On the **Workloads** that match criteria side panel, click **Export**.

The generated CSV file is downloaded to your Downloads folder with a filename similar to `Workload_Label_Review_(month_day_year)`.

5. Open and review the CSV file.

| | A | B | C | D | E | F | G | H |
|----|--------------------|-----------------------|---|---|---|---|---|---|
| 1 | Workload Hostname | Labels to be Assigned | Existing Labels | | | | | |
| 2 | perf-workload-3717 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 3 | perf-workload-3718 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 4 | perf-workload-3719 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 5 | perf-workload-3720 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 6 | perf-workload-3721 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 7 | perf-workload-3722 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 8 | perf-workload-3723 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 9 | perf-workload-3724 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 10 | perf-workload-3725 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 11 | perf-workload-3726 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 12 | perf-workload-3727 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 13 | perf-workload-3728 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 14 | perf-workload-3729 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 15 | perf-workload-3730 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 16 | perf-workload-3731 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 17 | perf-workload-3732 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 18 | perf-workload-3733 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 19 | perf-workload-3734 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 20 | perf-workload-3735 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 21 | perf-workload-3736 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 22 | perf-workload-3737 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 23 | perf-workload-3738 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 24 | perf-workload-3739 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 25 | perf-workload-3740 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |
| 26 | perf-workload-3741 | OS:Linux | app:App15665 env:Env15665 loc:Loc15665 role:Role15665 | | | | | |

How Label Matching Works

This section provides a detailed example of the Rule-Based Labeling feature's label matching logic and presents a brief list of terms used throughout this document.

When you click Review and Assign Labels to generate a list of workloads that match your labeling rules, workloads are evaluated against the conditions defined in the rules.

A match occurs if all statements in a rule's condition match a workload's attributes.

Labeling Rule Terminology

- **Rule:** Rules consist of a condition and one or more labels.
- **Condition:** Conditions are the user-defined criteria that workloads must match to be eligible for label assignment. A condition consists of one or more statements connected by AND, ensuring that workloads must satisfy all statements of the condition to match the rule.
- **Statement:** Statements define the specific workload attributes, operators, and values that are evaluated. Multiple values within a statement are considered using OR, allowing you to specify match criteria flexibly.
- **Precedence:** Rules are numbered, with Rule 1 having the highest precedence. A workload is evaluated against the rules in order, ensuring that rules with the labeling criteria most important to you are considered first.

Labeling Rule Matching Logic

Example: Workload and Rule Evaluation

| Workload attributes and existing labels | Rules you created | Rule conditions and label(s) | Overwrite option selected? | Match outcome | Were labels assigned? | Assigned labels |
|--|-------------------|--|----------------------------|--|--|---|
| <p>This is the example workload we want to assign labels to. It has the following attributes and existing labels:</p> <ul style="list-style-type: none"> • Hostname: <code>job-d8cc</code> • OS: <code>Windows</code> • IP address: <code>10.10.10.30</code> • Existing labels: <code>App88, Env22, Loc99</code> | Rule 1 | <ul style="list-style-type: none"> • Hostname is: <code>job-d8cc</code> • OS: <code>Windows</code> • IP range: <code>10.10.10.20 - 10.10.10.99</code> • Assign label: <code>Env44</code> | Yes | Match All statements in the rule's condition match the workload's attributes. | Yes The workload's existing Environment label <code>Env22</code> will be overwritten because the Overwrite option <code>is</code> selected. Result: The <code>Env44</code> Environment label specified in the rule condition will be assigned. | Assigned by Rule Based Labeling <ul style="list-style-type: none"> • <code>Env44</code> • <code>Loc22</code> • <code>Role11</code> Existing label already assigned and not overwritten by your rules: <ul style="list-style-type: none"> • <code>App88</code> |
| | Rule 2 | <ul style="list-style-type: none"> • Hostname Contains: <code>d8cc</code> • OS: <code>Windows</code> • Assign label: <code>Loc22</code> | Yes | Match <ul style="list-style-type: none"> • Hostname matches • OS matches | Yes The workload's existing Location label <code>Loc99</code> will be overwritten because the Overwrite option <code>is</code> selected. Result: The <code>Loc22</code> Location label specified in the rule condition will be assigned. | |
| | Rule 3 | <ul style="list-style-type: none"> • Hostname Ends with: <code>-d8cc</code> • Assign label: <code>App66, Role11</code> | No | Match <ul style="list-style-type: none"> • Hostname matches | 1 of 2 (1) The workload already has Application label <code>App88</code> but it will not be overwritten because the Overwrite option is <code>not</code> selected. Result: Application label <code>App66</code> will not be assigned. (2) The workload doesn't already have a Role label, so <code>Role11</code> will be assigned. | |
| | Rule 4 | <ul style="list-style-type: none"> • Hostname starts with: <code>job</code> • OS: <code>Windows</code> • Assign label: <code>Env99, Loc33, App66</code> | No | Match <ul style="list-style-type: none"> • Hostname matches • OS matches | 0 of 3 <ul style="list-style-type: none"> • Environment label <code>Env22</code> is already assigned by Rule 1, which has precedence, so <code>Env99</code> will not be assigned. • Location label <code>Loc22</code> is already assigned by Rule 2, which has precedence, so <code>Loc33</code> will not be assigned. • Application label <code>App88</code> is already assigned and the Overwrite option isn't selected. Result: Application label <code>App66</code> label will be not assigned. | |
| | Rule 5 | <ul style="list-style-type: none"> • OS: <code>Linux</code> • Assign label: User-Defined | N/A | No Match None of the statements in the rule's condition match the workload's attributes. | 0 | |

Labeling Rule Examples

This section provides several detailed examples of crafting labeling rules.

Keep in mind the following as you craft labeling rules:

- The **operator** you select and the particular values you enter in the **Values** field allow you to control the granularity of the labeling rule.
- Rule-Based Labeling automatically inserts an AND between the statements when you include multiple statements in a condition.
- When you specify multiple values in a statement, Rule-Based Labeling automatically inserts an OR between the values.

Example 1. Hostname Rule to match workloads that contain part of a specified host name

1. Select **Hostname** in the **Attribute** field.
2. Select the contents in the **Operator** field.
3. Enter **AWS** in the **Values** field.

4. Click **Close**.
5. Select one or more labels in the **Label** field.
6. Click **Save**.

Example 2. OS Rule to match workloads running a specific operating system



NOTE

Match on OS version or release

You can configure OS labeling rules to match all or part of the workload's OS version or release by selecting operators and entering the details. To find details, go to **Servers & Endpoints > Workloads** and click the workload. On the **Summary** tab, go to the **Attributes** section of the workload's details page.

| ATTRIBUTES | |
|-------------------------|--|
| VEN Version | 23.3.0 |
| Hostname | perf-workload-3724 |
| Location | Unnamed Datacenter, Unknown Location |
| <u>OS</u> | ubuntu-x86_64-xenial |
| <u>Release</u> | 4.4.0-97-generic #120-Ubuntu SMP Tue Sep 19 17:28:18 UTC 2017 (Ubuntu 16.04.1 LTS) |
| Uptime | 2 Days, 18 Hours, 41 Minutes |
| Heartbeat Last Received | 05/14/2024, 17:10:20 |
| Interfaces | eth0: 10.0.14.140/8 10.0.0.1 (Corporate) eth0: fd00::200:a:0:e8c/64 (Corporate) |

1. Select **OS** in the **Attribute** field.
2. Select an **Operator**.
3. Select **Linux** in the **Value** field.
4. Click **Close**.
5. Select one or more labels in the **Label** field.
6. Click **Save**.

Example 3. IP Address Rule to match workloads within a specific IP address range:

1. Select the IP Address in the **Attribute** field.
2. Select is in the **Operator** field.
3. In the **Value** field, enter a narrow range such as 10.2.0.0 – 10.2.200.0.
4. Click **Close**.
5. Select one or more labels in the **Label** field.
6. Click **Save**.

Example 4. CIDR Block Rule to match workloads within a specific CIDR block:

1. Select the IP Address in the **Attribute** field.
2. Select is in the **Operator** field.
3. In the **Value** field, enter a CIDR block. For example: 10.2.20.0/24
4. Click **Close**.
5. Select one or more labels in the **Label** field.
6. Click **Save**.

Example 5. Rule with multiple attributes, each with a single value:

1. Specify a hostname:
 - Select **Hostname** in the **Attribute** field.
 - Select the contents in the **Operator** field.
 - Enter details in the **Values** field.
2. Specify an operating system:
 - Select **OS** in the **Attribute** field.
 - Select the contents in the **Operator** field.
 - Select an operating system in the **Values** field.
3. Specify an IP address:
 - Select **IP Address** in the **Attribute** field.
 - Select **is** in the **Operator** field.
 - In the **Values** field, enter an IP range or CIDR block.
4. Specify a listening port and/or protocol:
 - Select **Port/Protocol** in the **Attribute** field.
 - In the **Operator** field, select is for a specific port/protocol; select is in to specify a range.
 - In the **Values** field, enter a specific port/protocol or a range as appropriate.

5. Specify a process path:
 - Select **Process** in the **Attribute** field.
 - In the **Operator** field, select an appropriate operator.
 - In the **Values** field, enter all or part of a process path according to your selected operator.
6. Click **Close**.
7. Select one or more labels in the **Label** field.
8. Click **Save**.

Illumio Policy Enforcement Model

Illumio employs an allowlist security model. By default, workload-to-workload communication is blocked unless explicitly permitted by defined Illumio policy rules. Administrators create these explicit rules to allow only necessary traffic, significantly enhancing security.

Why Use Selective Enforcement?

Deploying the allowlist model universally and simultaneously can be challenging or disruptive. Illumio addresses this by providing selective enforcement, an intermediate enforcement state that allows a gradual security rollout.

Selective Enforcement provides:

- Gradual Security Implementation: Smooth transition from open ("Idle" or "Visibility-only") states to full enforcement ("Full Enforcement").
- Targeted Visibility: Enforcement focused on selected services and ports via labels or groups, while other services remain in visibility mode.
- Rapid Threat Response: Immediate enforcement on vulnerable or critical ports and services without impacting entire workloads.

Applying Selective Enforcement

The Selective Enforcement mode is configured per workload using labels or groups of labels.

When Selective Enforcement is activated:

- Enforced Ports and Services: Security rules are actively enforced. Only inbound traffic is explicitly permitted.
- Visibility-Only Ports and Services: No active blocking, but communication is monitored and logged.

The Workload Behavior under Selective Enforcement:

- Enforced Ports: According to defined policy rules, only explicitly allowed inbound traffic is permitted; all other traffic is blocked.

- Visibility-Only Ports: Traffic remains unblocked but is actively monitored and logged

How Selective Enforcement Works

Selective enforcement is applied individually per workload through labels or label groups.

When enabled:

- Enforced Ports/Services: Security rules are actively enforced; only explicitly permitted traffic passes.
- Other Ports/Services: Remain in visibility-only mode; traffic is monitored but not blocked.

Workload Behavior under Selective Enforcement:

- Selective Enforcement (Enforced Ports): Only explicitly permitted inbound traffic is allowed. All other inbound traffic to these ports is blocked.
- Visibility-only (Other Ports): Traffic continues normally but is monitored and logged.

Enforcement Progression Model

Selective Enforcement is a crucial step in Illumio's structured enforcement progression:

Idle (Visibility-only) → Selective Enforcement → Full Enforcement

where

- Idle: Visibility and monitoring are in place, but there is no enforcement.
- Selective Enforcement: Partial enforcement on chosen ports/services.
- Full Enforcement: Complete allowlist enforcement on all ports and services.

This structured approach simplifies the implementation of secure policies, offering flexibility in managing risk and operational complexity.

Use Cases and Limitations

Basic use cases for Selective enforcement are:

- Incremental Policy Rollout: Enables the gradual introduction of policies, reducing risks to critical systems.
- Rapid Security Response: Quickly enforce specific, critical, or vulnerable port and service policies.

Selective enforcement only applies to inbound (source-side/ingress) traffic. It controls incoming requests to protected workloads, but it does not control outbound traffic from workloads.

Selective Enforcement Mode Limitations

Limitations of Selective Enforcement are grouped as follows:

- Directional enforcement: Selective enforcement operates only on inbound traffic.
 - Inbound Policy (Destination-centric): Manages incoming traffic to workloads.
 - Outbound Policy (Source-centric): Manages outgoing traffic from workloads.
- Support for Managed Workloads is available only because selective enforcement is available for workloads managed directly by Illumio.
 - Managed workloads are supported.
 - Unmanaged workloads or workloads managed via Network Enforcement Nodes (NEN) cannot utilize selective enforcement.
- Impact on Virtual Services: Selective enforcement does not apply directly to virtual services as a single entity.

Instead, policies must target individual workloads within virtual services. Enforcement is applied at the workload level within virtual services.

Virtual services themselves are not directly enforced.

Workload Enforcement States

Workload policy modes determine how Illumio rules impact workload network communications. Illumio provides four policy modes.

The enforcement state displayed in the Policy Compute Engine (PCE) indicates the desired state for the next policy update. Failure to apply this state successfully will result in a Policy Sync error.

Idle Enforcement State

This state is typically used during initial VEN installation or activation. Its characteristics are:

- No firewall rule enforcement.
- Collects and reports network traffic data every 10 minutes.
- Report OS compatibility every four hours.
- Immediately reports network interface configuration changes.

Policy Exclusions

In Illumio Segmentation for Data Centers the PCE supports including policy exclusions in policy scopes and rules.

This topic explains what they are, how they are supported, and how to add them to your security policy.

Policy Exclusions Overview

Using policy exclusions, a policy can significantly simplify the rule-writing process. Specifically, using a policy exclusion in a scope or rules allows you to replace the inclusion of a large number of required labels with the exclusion of a small number of unwanted labels. Security policies written with exclusions can be easier to read and maintain.

Using policy exclusions allows you to state in your security policy that you want a policy or rule to apply to “all except X,” where X can be both labels and label groups. To state this another way, “all except X” means “All labeled workloads except X” or “All label group objects of a dimension except X.”

You can include policy exclusions in policy scopes and rules, specifically in the destination and source actors.

Use Cases

The following examples demonstrate a few common use cases for using policy exclusions:

- All environments except Production should be able to pull updates directly from Red Hat.
- The standard jump boxes should be able to connect to all environments except PCI.
- All applications, except Quarantine, should be able to connect to Core Services.

Policy Exclusions Support

Policy exclusions are supported by the features and in the PCE web console in the following ways:

| Illumio Segmentation for Data Centers Features | Details |
|--|--|
| Policy scopes and rules | <p>In policies and rules, excluding a label creates an “all-but” rule or boundary that applies to all workloads that don’t have that excluded label but do have another label of the same label type as the excluded label.</p> <p>For example, your data center supports Production, QA, and Development environments.</p> <p>Adding an exclusion for “All environments except Production” means that the rules apply to all workloads with Environment labels, except those labeled as Production. It does not translate to “All workloads except those with the Production label,” which would include workloads that don’t have an Environment label. When you create a rule that applies to “All environments but Production,” this rule achieves the same effect as creating one that applies to the QA and Development environments only.</p> |
| Labels | <p>Fully supports except for the restrictions below. See Requirements and Restrictions [135].</p> |
| Label Groups | <p>Label groups are supported for policy exclusions in the same way as labels. For example, you want to create a boundary between Finance and all other applications. You create a label group named “Finance Apps” and use it as a policy exclusion.</p> <p>Individual workloads, virtual services, virtual servers, “All Workloads,” the “Uses virtual service only” option, the “Uses virtual service and workloads” option, and container hosts do not support label group exclusions.</p> <p>Additionally, you cannot specify exclusions out of label groups. For example, you have created a label group for the environment “Non-production.” You want to use the label group, except you don’t want it to apply to the “Development” environment. You want to create a policy exclusion for the “Development” environment label from the “Non-production” label group. This action is not supported. Selecting to exclude a label group excludes all labels within that group.</p> |
| Rule Search and filters | <p>You cannot search by policy exclusions; however, any rules that contain policy exclusions appear in the results of your rule search.</p> <p>In label filters and rule searches, entering a label name displays both included and excluded labels with that name.</p> |
| App Groups | <p>In the App Groups > App Group List, select the Rules tab.</p> <p>Rules with policy exclusions appear in the Rules tab.</p> |
| Policy Check | <p>Rules with policy exclusions are displayed on the Policy Check page.</p> |

| Illumio Segmentation for Data Centers Features | Details |
|--|--|
| Policy Generator | <p>The PCE does not propose policy exclusions when using Policy Generator to create policy.</p> <p>When using Policy Generator to calculate V-E scores for vulnerabilities (you have the Vulnerability maps feature enabled), Policy Generator won't work for rules that contain policy exclusions because they are not supported in Policy Generator.</p> |
| Access Management | <p>Access management (Role-Based Access Control, or RBAC) detects policy exclusions when determining user access in the Policy Control Engine (PCE). However, you cannot exclude a policy exclusion from an RBAC role.</p> <p>Policy exclusions are only supported in policies and rules. If a scope includes a policy exclusion based on labels outside the scopes you have permission for, you cannot view or manage those policies and rules.</p> <p>For example, suppose a policy excludes "All environments except Production," and you have permission for the Production environment but not for the Staging environment. In that case, you will be unable to view or manage that policy.</p> |
| Explorer | <p>When writing rules using Explorer, you can select policies that contain policy exclusions. You can edit the rules in the policy that have exclusions. You can add new proposed rules, taking the exclusion scopes into account.</p> <p>However, you cannot add a new policy exclusion to an existing proposed rule or an exclusion to a new one.</p> |
| PCE web console maps | <p>Policy exclusions apply to rules; they are not properties of the traffic links (the lines between the workloads) in the Illumio maps (Illumination Map, App Group Map, and Vulnerability Map).</p> <p>When you click View Rule for any traffic link, you can view the policy exclusions in the panel.</p> |
| Enforcement Boundaries | <p>Policy exclusions are not supported in Enforcement Boundaries.</p> <p>However, you can view policy exclusion rules in the Rules tab of an Enforcement Boundary details page.</p> |

Requirements and Restrictions

Requirements

When specifying a policy exclusion, it must be the same label type as the group it's being excluded from; the following examples are allowed:

- All Locations except the New Jersey location
- All Applications except Billing

However, this example is not allowed because it specifies different label types - Location vs Environment:

- All Locations except those with Development systems

Restrictions

- You can specify an included or excluded label for each dimension, but not both. The following examples show valid combinations:
 - App: Swift
 - App: All but Swift
 - Env: Prod, App: All but Swift
 - Loc: EU, Env: All but Prod
- You cannot specify both included and excluded labels within the same label type. The following examples are invalid combinations:
 - Env: Prod, Env: Dev, Env: All but UAT
 - Env: Prod, App: HRM, App: CRM, App: All but Swift
 - App: HRM + App: CRM - App:Swift
 - Loc: EU - Loc: Switzerland
- You cannot use policy exclusions with the following objects in the PCE:
 - Individual (named) workloads
 - Virtual servers
 - Virtual services
 - Container hosts

Adaptive User Segmentation

Illumio's Adaptive User Segmentation (AUS) allows you to leverage user groups from two different Microsoft directory applications to manage network traffic across your organization:

- Microsoft Active Directory
- Microsoft Entra ID (formerly Azure Active Directory)

Support for these directory applications allows you to include user groups in the PCE that map directly to Active Directory or Entra ID Groups. You can then create rules using these groups to control outbound access on specific workloads, such as a VDI desktop, based on the user's group membership logged in to that workload.

For example, you may want to restrict access to the ERP application to only employees in the Sales user group and not to users in the HR department. You may also wish to allow HR users to access only HR applications, but not all internal resources.

If you have a Windows workload that controls access to other resources in your network, such as a VDI desktop with the VEN installed, you can add the VDI desktop workload and Active Directory or Entra ID User Groups to the rule. Writing this type of rule allows user access only to the resources explicitly allowed by the rules.

Configure a Microsoft directory application for use with Illumio AUS

Configure one of the following directory applications.

Microsoft Entra ID

**NOTE**

Feature Enablement: Using Microsoft Entra ID with Illumio AUS requires enablement by Illumio. For assistance, contact your Illumio Account team.

To use Microsoft Entra ID User Groups in Illumio policy, configure the Entra ID Enterprise Application as described in [Configure Microsoft Entra ID \(Azure AD\) for use with Illumio AUS \[138\]](#). Once configured and provisioned, Entra ID periodically pushes group and user information to the PCE automatically. The information eventually appears in **Policy Objects > User Groups**. This automated operation differs from the manual steps required for integrating AUS with Microsoft Active Directory.

Microsoft Active Directory

To use Microsoft Active Directory User Groups in Illumio policy, perform the procedures provided in [Configure Microsoft Active Directory for use with Illumio AUS \[142\]](#).

Configure Microsoft Entra ID (Azure AD) for use with Illumio AUS

To use Microsoft Entra ID User Groups with Illumio Adaptive User Segmentation (AUS) in Illumio policy, you must first configure Microsoft Entra ID Enterprise Application as described in this topic.



NOTE

Feature Enablement: Using Microsoft Entra ID with Illumio AUS **requires enablement by Illumio**. For assistance, contact your Illumio Account team.



NOTE

Similar support is also available for using AUS with Microsoft Active Directory. For details, see [Configure Microsoft Active Directory for use with Illumio AUS \[142\]](#).

Key features for Microsoft Entra ID and Illumio AUS integration

This solution includes these features:

- **Entra ID Group Sync:** Support for importing and syncing Entra ID security groups, preserving group names and unique identifiers (such as object IDs or SIDs).
- **Group Visibility in the UI:** Synced Entra ID groups appear in the Illumio Console under the **User Groups** section with searchable metadata such as group name, unique ID, sync source (Entra ID), and last sync status.
- **Use Groups in Rulesets:** You can select imported Entra ID groups in the **Sources** field when creating rulesets, allowing you to apply policies based on the identity of the logged-in users.
- **Full Enforcement:** If traffic originates from a workload where the Entra ID-logged-in user session is enforced by an AUS rule in your policy, then the traffic will be allowed.

Prerequisites

- Using Microsoft Entra ID with Illumio AUS requires **enablement by Illumio**. For assistance, contact your Illumio Account team.
- Your Illumio organization must be in a particular type of SaaS cluster. Verify this with your Account team.
- You must be a user with Entra ID credentials to log in to workloads protected by a Windows VEN.
- You must be running these versions of the VEN and the PCE:
 - VEN version 25.2.30 or later
 - PCE version 25.4 or later

Limitations for Entra ID and User Groups

Note the following limitations for Entra ID and user groups:

- You cannot have a nested group as a member of a group.
- You can only write Adaptive User Segmentation policies for users in an immediate group.

Configure Entra ID for use with Illumio AUS

1. On the Home screen within the Azure portal, click **Microsoft Entra ID**.
2. Navigate to **Manage > Enterprise Application**, and add a new application called **Illumio - AUS**.
3. Navigate to **Illumio - AUS Enterprise Application > Manage > Provisioning**.
4. Click **Admin Credentials** and enter the following information:
 - a. Select **Bearer Authentication** from the **Authentication Method** drop-down list.
 - b. Enter the Illumio PCE URL in the **Tenant URL** field (for example, {Enter the Illumio PCE URL in the **Tenant URL** field (for example, {server1.mycompany.com}/scim/orgs/{orgid}/).
 - c. Create a PCE API key within the Illumio PCE.
 - d. Add the secret token to the **Secret Token** field using the following format: {api_key}:{secret}.

Admin Credentials

Admin Credentials

Microsoft Entra needs the following information to connect to illumio-aus's API and synchronize user data.

Authentication Method ⓘ

Bearer Authentication

Tenant URL * ⓘ

https://[redacted]

Secret Token

[redacted]

Test Connection

5. Navigate to the **Mappings** section.

Mappings

Mappings

Mappings allow you to define how data should flow between Microsoft Entra ID and customappsso.

| Name | Enabled |
|---|---------|
| Provision Microsoft Entra ID Groups | Yes |
| Provision Microsoft Entra ID Users | Yes |

Restore default mappings

6. For **Provision Microsoft Entra ID Groups**, enable the attributes.

Home > illumio-aus | Provisioning >

Attribute Mapping

Save Discard

Name
Provision Microsoft Entra ID Groups

Enabled
 Yes No

Source Object
Group

Source Object Scope
[All records](#)

Target Object
urn:ietf:params:scim:schemas:core:2.0:Group

Target Object Actions
 Create
 Update
 Delete

Attribute Mappings
Attribute mappings define how attributes are synchronized between Microsoft Entra ID and customappsso

| customappsso Attribute | Microsoft Entra ID Attribute | Matching precedence | Edit | Remove |
|------------------------|------------------------------|---------------------|----------------------|------------------------|
| displayName | displayName | 1 | Edit | Delete |
| externalid | objectId | | Edit | Delete |
| members | members | | Edit | Delete |

[Add New Mapping](#)

Show advanced options

7. For **Provision Microsoft Entra ID Users**, enable these attributes:



NOTE

It can take some time for provisioning changes to propagate. This is an Entra ID limitation.

Attribute Mapping

Save Discard

Name
Provision Microsoft Entra ID Users

Enabled
 Yes No

Source Object
User

Source Object Scope
[All records](#)

Target Object
urn:ietf:params:scim:schemas:extension:enterprise:2.0:User

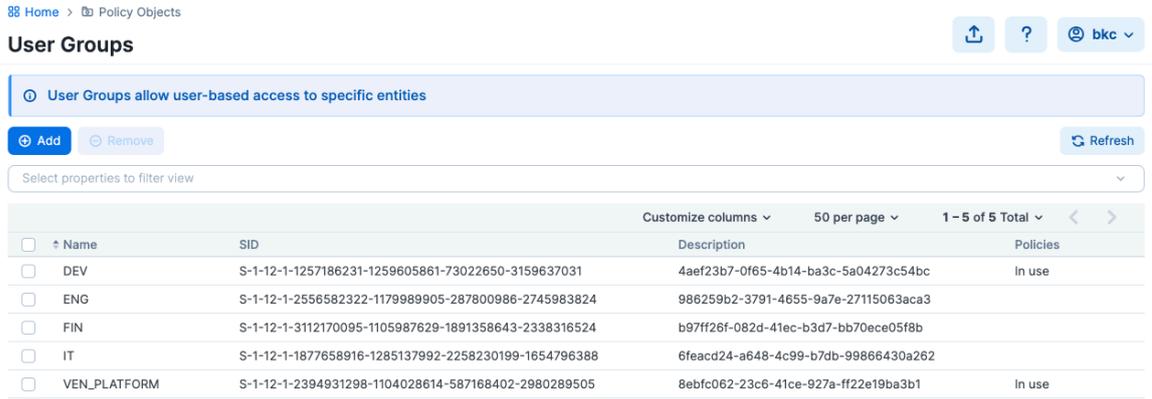
Target Object Actions
 Create
 Update
 Delete

Attribute Mappings
Attribute mappings define how attributes are synchronized between Microsoft Entra ID and customappsso

| customappsso Attribute | Microsoft Entra ID Attribute | Matching precedence | Edit | Remove |
|------------------------|------------------------------|---------------------|----------------------|------------------------|
| username | userPrincipalName | 1 | Edit | Delete |
| externalid | objectId | | Edit | Delete |
| name-formatted | displayName | | Edit | Delete |

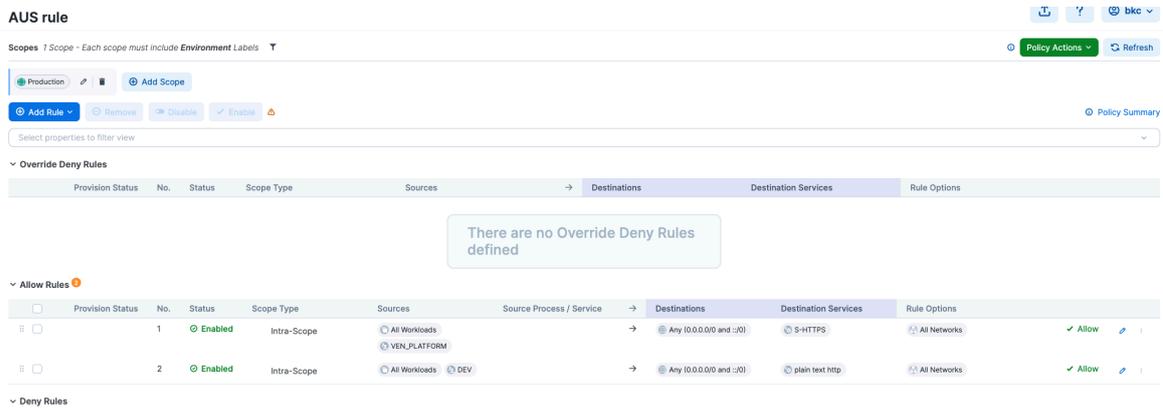
[Add New Mapping](#)

8. Navigate to Settings, and set the scope to **Sync only assigned users and groups**.
9. After Entra ID provisioning starts, Entra ID periodically pushes group and user information to the PCE automatically. The information eventually appears in **Policy Objects > User Groups**. This automated operation differs from the manual steps required for integrating AUS with Microsoft Active Directory (described in [Adaptive User Segmentation \[137\]](#)).



After the **User Groups** list is populated, you can use the user groups to write policies. (See [Adaptive User Segmentation \[137\]](#).)

Here is an example of a policy that allows egress traffic from all workloads with users in the DEV user group to all destinations on port 80/TCP:



Set Up the VEN for Entra ID for AUS

Use Entra ID credentials to log in to the workload where the VEN is installed and the flow is enforced. Entra ID provisioning does not take effect immediately in the VEN, so you may have to log out and log back in for provisioning to occur.

Configure Microsoft Active Directory for use with Illumio AUS

You can use Illumio's [Adaptive User Segmentation \(AUS\) \[137\]](#) with Microsoft Active Directory to manage network traffic across your organization. This feature allows you to include user groups in the PCE that map directly to your Active Directory Groups. You can then create rules using these groups

to control outbound access on specific workloads, such as a VDI desktop, based on the user's group membership logged in to that workload.

**NOTE**

Similar support is also available for using AUS with Microsoft Entra ID. For details, see [Configure Microsoft Entra ID \(Azure AD\) for use with Illumio AUS \[138\]](#).

Key features for using Illumio AUS integration with Microsoft Active Directory

This solution includes these features:

- Group Visibility in the UI: Synced Active Directory groups appear in the Illumio Console under the **User Groups** section with searchable metadata such as group name, unique ID, sync source, and last sync status.
- Use Groups in Illumio policy: You can select imported AD groups in the **Sources** field when creating rules, allowing you to apply policies based on the identity of the logged-in users.
- Full Enforcement: If traffic originates from a workload where the AD-logged-in user session is enforced by an AUS rule in your policy, then the traffic will be allowed.

Add Active Directory Groups to the PCE

**NOTE**

This procedure applies only to Active Directory. Don't perform this procedure if setting up Entra ID for use with AUS.

1. From the PCE web console menu, choose **Policy Objects > User Groups**.
2. On the User Groups page, click **Add**.
3. On the User Group page, enter the name, system identifier (SID), and description of the Active Directory.

4. Click **Save**.

The new group appears in the User Groups list in the PCE. You can now use the user group in an Illumio policy to control access to specific workloads.



NOTE

The PCE can display a maximum of 100 User Groups.

Add User Group-Based Rules for AUS



NOTE

This procedure applies only to Active Directory. Don't perform this procedure if setting up Entra ID for use with AUS.

1. From the PCE web console menu, select **Policies**.
2. In the Policies tab, click **Add** and choose **Add from Scratch**.
3. Enter a name and description for the policy.
4. Select an Application, Environment, and Location label to define the policy scope.
5. Click **Add Rule** and select the rule type:

6. In the Destinations drop-down list, select the user group to which you want to provide access to the other workload.
7. From the Source drop-down list, select the workloads or labels to which you want to grant a user group access.
8. In the Destination Services drop-down list, select the service you want the user groups to access on the provided workloads.
9. Click the **Save** icon at the end of the row.
- 10 Provision the changes.

.

About the Policy Generator

The Policy Generator simplifies the Illumio policy creation process by recommending the optimal security policy for your App Groups. It can accelerate security workflows and reduce errors while creating a security policy.

Overview of Policy Generator

The Policy Generator uses network traffic to recommend and generate micro-segmentation policies for every workload and application, regardless of its location. It can generate rules for applications running on physical devices, virtualized platforms, and behind network devices, both on-premises and in the cloud.

The Policy Generator supports the creation of DNS-based rules across all wizards, including intra-scope, extra-scope, and IP lists. You can edit the proposed virtual services and add wildcards.

Application owners use the Policy Generator to write the following types of rules for the applications they manage:

- Intra-scope rules
- Extra-scope rules
- Rules using IP lists.

For a selected App Group, the Policy Generator provides:

- A workflow to create a policy that controls internal and external traffic.
- A way to assess your current rule coverage is by dividing the number of detected connections controlled by rules by the total number of connections. You can increase your rule coverage by creating rules for detected connections that are not controlled by rules. The Policy Generator proposes rules for connections not currently allowed by the existing rules. It displays the consolidated flow count for each new proposed rule to help ensure the maximum impact on rule coverage.

**NOTE**

The Policy Generator calculates rule coverage automatically every 24 hours or after creating a draft policy.

You can rewrite rules as your datacenter needs change, and the Policy Generator will show you the before and after effects of those rules.

- A way to assess your current rule coverage is by dividing the number of detected connections controlled by rules by the total number of connections.

The first time you use the Policy Generator for an App Group, it creates a new draft policy with the title of the selected App Group. When you use the Policy Generator to create additional rules, they are added to the policy it created. You can review and customize the proposed rules before you save them as a draft policy. The Policy Generator detects and suggests rules based on Windows processes and services for Windows. You can edit the service before saving it.

When an App Group has multiple destinations communicating with a specific source, the Policy Generator consolidates all the destinations into a single rule for improved readability and scalability.

On the Summary tab of the Policy page, any policy created with Policy Generator has the default description “Automatically generated using the Illumio Policy Generator” and the value of `illumio_policy_generator` for the External Data Set field. The value for the External Data Reference is the App Group name.

Policy Generator Prerequisites and Limitations

The following prerequisites and limitations bind the Policy Generator:

- You cannot add Role-level rules until Role labels are added to all App Group workloads.

When some workloads in an App Group do not have Role labels, you can still write an App Group-level rule using Policy Generator to allow all the workloads to communicate with each other.

- Rule coverage is updated one App Group at a time.

Create Intra-scope Rules with the Policy Generator

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Create+Intra-Scope+Rule.mp4>

1. From the PCE web console menu, choose **Policies > Start Policy Generator**.

In the dropdown menu, **Select an App Group to build Rules** and choose one of the App groups.

The Policy Generator displays the latest calculated coverage for each type of rule (Intra-scope, Extra-scope, and IP Lists).

Click the refresh icon to recalculate Rule coverage.

2. Select a workload you want to work with, such as Intra-scope.
3. Click the **Start with Intra-Scope** button.

The Intra-Scope Rule Configuration page appears.

4. Select a granularity level for the rules in the Choose Intra-Scope Rule Configuration section.

The detected connections, including source, port/protocol, and destination details, appear in the Review All Connections section.

| Rule Configuration | Connections Displayed |
|---------------------------------|--|
| App Group Level | Microsegmentation: Allow all Workloads to talk across all Services |
| Role Level - All Services | Divide Workloads by Role and allow them to talk on all Services. |
| Role Level - Specified Services | Nanosegmentation: Divide Workloads by Role and specific services. |
| Auto Level | Vulnerability Mitigation: Eliminate or reduce the exposure of vulnerable ports |



NOTE

When the App Group has more than four types of ports or protocols, the Policy Generator displays a truncated list of ports and protocols. To display the remaining ports or protocols in a modal window, click the **+ More** link.

5. (Optional for Role level) To exclude a connection from the proposed rules, click **Exclude**. The row is grayed out to indicate that no rules will be proposed for this connection, and the amount of rule coverage decreases. To include an excluded connection, click **Include**.



NOTE

At least one connection must be included to continue.

6. Click **Next.**

The proposed rules appear on the Preview page.

7. (Optional) To edit a service for a rule, click the pencil icon beside the service. The Edit Service dialog box appears.

Select a service from the drop-down list or create a new one. You can select services that have a broader range of ports. The list includes every service that matches that port and protocol. When you've added a service with multiple ports and protocols or ranges, they all appear in the list.

Select **Apply Changes to all matching ports** to allow the service to be used in other rules that match this service. You are prompted to allow the Policy Generator to merge rules. To cancel the merge, reload the page and start over.

When you create a process-based service, the connection appears as if it's not covered.

8. To accept the proposed rules, click **Save and **OK**.**

The Policy Generator's Successful message displays the number of new rules and services.

The rules are added to a draft policy. Click **Continue with App Group** to add extra-scope rules or rules using IP lists for the same App Group. On the last step of the Policy Generator, you can return to the App Group to add or append to the rules.

**NOTE**

You must provision the rules to apply them to workloads.

Create Extra-scope Rules with the Policy Generator

When you create extra-scope rules, the Policy Generator displays all traffic from a different App Group and is targeted at the selected App Groups. The Policy Generator displays all App Groups with which the selected App Groups communicate. You can choose which connections to cover with rules.

Follow the steps as explained in [Create Intra-scope Rules with Policy Generator \[147\]](#).

Create Rules Using IP Lists with the Policy Generator

Policy Generator creates rules that use IP lists as intra-scope rules.

When using IP lists to create rules, the Policy Generator defines a connection as a role on a port and protocol to an IP address. For example, the Policy Generator displays five connections when there are five IP addresses in an IP list.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Create+Rules+using+IPLists.mp4>

- 1. From the PCE web console menu, choose **Policies > Start Policy Generator**.**
The Select App Group page appears. The page displays the date and time when the Policy Generator last calculated the coverage for each type of rule. Click the refresh icon to recalculate rule coverage.
- 2. Click the refresh icon to recalculate Rule coverage.**
The Policy Generator displays the latest calculated coverage for each type of rule (Intra-scope, Extra-scope, and IP Lists).
Select an App Group to build Rules , and choose one of the App groups.
- 3. Select an App Group to build Rules** , and choose one of the App groups.
- 4. Click the **Start with IP Lists** button.**
The IP List Selection page appears.
- 5. Select the IP lists you want to write rules and click **Next**.**
The Configure IP List page appears.

**TIP**

- To view the IP addresses configured in a list (not the IP addresses in the traffic), expand an IP list by clicking the arrow icon in the Name column.
- Select the "Any IP" list to write rules that cover all connections. This list includes all IP addresses.
- Each IP address can be part of multiple IP lists, and you can choose which list to write your rules to.
- When you choose overlapping IP lists, you can write overlapping rules.

When an IP address appears in more than one IP list, the rule will be included in all those IP lists.

- You can write rules for inbound and outbound connections, or both. For example, you can write permissive rules for outbound traffic and specific rules for inbound traffic.

6. Select whether to configure rules by App Group or by role:

- **App Group Level:** All workloads in the specified App Group can communicate with all workloads in the other App Groups
- **Role Level:** Specified workloads in the App Group can communicate with specified workloads in the other App Groups

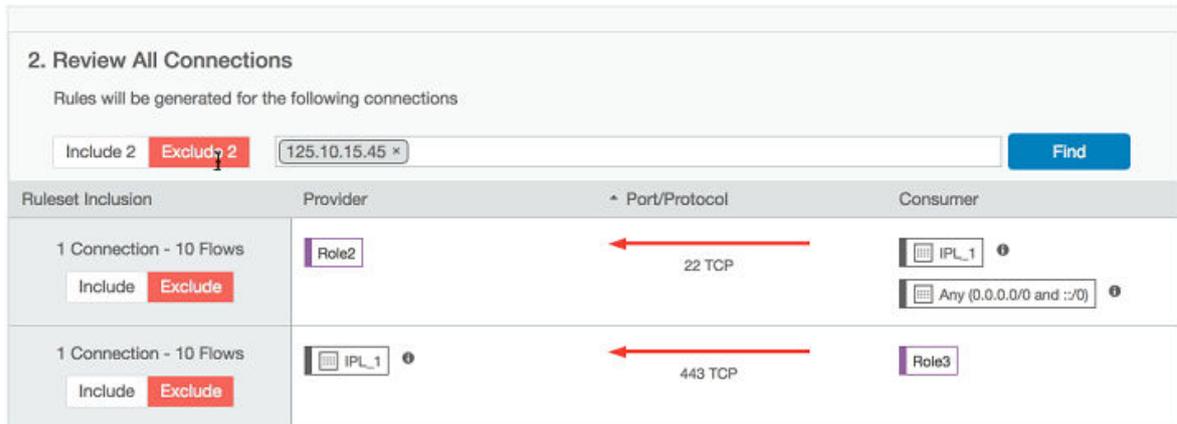
7. Select the permitted services for the rules:

- **All Services:** Workloads can communicate over all services
- **Specified Services:** Workloads can communicate over specified services

It creates a rule for any device to which those IP lists apply.

**TIP**

- To display the IP addresses of the traffic for each port and protocol, hover over the info (i) icon in the Destination column.
- Use the search field above the list of connections to filter connections by IP address, port number, protocol, role, or label. This allows you to find and exclude specific traffic.
- To quickly include or exclude all traffic, use the **Include** and **Exclude** buttons by the search field. You can exclude all traffic, then selectively include specific connections.



8. To preview the rules proposed by Policy Generator, click **Next**. The IP List Rule Preview page appears.
9. (Optional) To edit the service for a rule, click the pencil icon. The Edit Service dialog box appears. Select a service from the drop-down list or create a new one. You can select services that have a broader range of ports. The list includes every service that matches that port and protocol. When you've added a service that has multiple ports and protocols or ranges, they all appear in the list. Select **Apply Changes to all matching ports** to allow the service to be used in other rules that match that service. You are prompted to allow the Policy Generator to merge rules. To cancel the merge, reload the page and start over. When you create a process-based service, the connection will appear as if it's not covered.
- 10 To accept the proposed rules, click **Save** and **OK**.
 - The Policy Generator's Successful message appears, which displays the number of new rules and services. The rules are added to a draft policy.

Segment Multiple App Groups with the Policy Generator

Using the Policy Generator, you can apply nano-segmentation (also known as ringfencing) to multiple App Groups. Nano-segmenting app groups enables all workloads to communicate across all services within each App Group.

When segmenting App Groups, the Policy Generator creates one policy per App Group. The policy includes a rule that covers traffic for all workloads to all workloads on all services.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Segment+Multiple+App+Groups.mp4>

1. From the PCE web console menu, choose **Policy Generator**.
The Select App Group page appears. The page displays the date and time when the Policy Generator last calculated the coverage for each type of Rule. Click the refresh icon to recalculate rule coverage.
2. In the Select App Group drop-down menu, select **Segment Multiple App Groups** from the bottom of the list.
The Choose App Groups page appears.
3. Select the App Groups to segment and click **Next**.

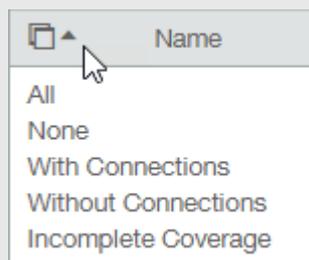


TIP

- To recalculate rule coverage for an App Group, hover over the Last Calculated column and click the refresh icon. The column displays the time at which the rule coverage was calculated.

The column indicates whether the policy for the group has been edited since the last calculation, prompting you to recalculate it.

- To quickly select App Groups using different criteria, click the arrow icon to the right of the Name column:



- The Choose App Groups page displays all your App Groups regardless of their rule coverage percentage or whether they have connections. For example, the page displays App Groups with 100% rule coverage and groups with zero connections.

4. To accept the proposed rules, click **Save** and **OK**.
The Policy Generator's Successful message appears, which displays the number of new rules. The rules are added to a draft policy.

Policy Generator Wizard

The new Policy Generator wizard replaces the earlier Policy Generator.

To use the Policy Generator wizard, go to Policies > Start Policy Generator in the Web console and follow the instructions.

The screenshot shows the 'Policy Generator' page in the Illumio web console. At the top left, there is a 'Home' link and the page title 'Policy Generator'. On the top right, there are utility icons: a mail icon with a '1' notification, an upload icon with a '1,381' notification, a help icon, and an 'Owner' dropdown menu. Below the header is a blue banner with the text: 'The Illumio Policy Generator allows you to write Rules for uncovered connections of traffic in your App Groups' and a network diagram. Underneath the banner is a progress bar with three steps: 'Select App Group' (active), 'Configure Rules', and 'Preview Rules'. Below the progress bar is a section titled 'Select an App Group to build Rules' with a search input field containing 'Search App Group'. At the bottom, there is a three-step process diagram: 1. 'Select the App Group' (yellow), 2. 'Configure Intra-Scope Or/And Extra-Scope Rules' (orange), and 3. 'App Groups are secured' (green).

About Provisioning

Provisioning applies security policy changes from the Policy Compute Engine (PCE) to Virtual Enforcement Nodes (VENs) on managed workloads. These changes define enforcement behavior using iptables/nftables (Linux) and WFP (Windows) to ensure consistent segmentation.

When you provision updates, the PCE recalculates changes to policies, IP lists, services, label groups, and security settings, then transmits them to all VENs on your workloads. An orange badge on the Provision button indicates the number of changes awaiting provisioning.

Provisioning applies changes such as:

- New or modified policy rules
- Label and label group updates
- Virtual service and service definitions
- Policy scope or global settings changes
- Static policy assignments

How Provisioning Works Internally

Provisioning works in stages that follow one another:

1. **Database commit:** where the PCE first records and commits changes as follows:
 - Policy rules (Allow, Deny, Override Deny, Custom iptables)
 - Labels and Label Groups
 - Virtual Services and Services
 - Policy scopes and global settings
2. **Policy Calculation:** where the PCE matches policy objects and rules against workloads using label-based scopes.
 - Virtual Services and Services are resolved to port and traffic definitions.
 - Workload-specific rule sets are generated based on which rules match their labels.
3. **IP Resolution and Rule Compilation:** where Labels are mapped to IP addresses and interfaces.
 - Final iptables (Linux) or WFP (Windows) rules are compiled per workload.

4. **Distribution to VENS:** where VENS receive policy update notifications. Affected VENS securely retrieve and enforce their updated rules.

Full Provisioning

Full provisioning applies to all pending policy changes and impacted workloads.

- It is used to roll out standard policies and for large-scope updates that involve multiple rules or labels.
- With full provisioning, all impacted workloads receive updated rules.

Selective Provisioning (Quick Provision)

This type of provisioning applies changes to a single object (e.g., a specific rule, label, or service).

- It is used for urgent updates (e.g., emergency deny rules) and controlled test deployments.
- To implement Quick Provisioning, use the **Quick Provision** button in the object view.

Provisioning in Static Policy Mode

Static Policy allows provisioning to stage rules on workloads without enforcing them until they are explicitly applied. This enables testing and controlled rollout in sensitive environments.

- Provisioned rules are staged but not enforced, and the Workloads display "Staged" status.
- Administrators must manually click **Apply Policy** to enforce rules.
- Use Static Policy mode when testing policy on production workloads without immediate enforcement, and for manual enforcement approval of workflows.

Versioning, Restore, and Revert

Illumio tracks every provision as a version, allowing administrators to audit, restore, or revert policy states.

Versioning Features

Each provision is saved as a version in the **Changes > History** tab and includes a timestamp, user ID, and a summary of the changes.

- Policies can be restored either partially (selectively importing components like policies, labels, or services) or as a complete restore (replacing all current policy objects with the selected version)
- **Partial restore** allows you to cherry-pick changes without overwriting unrelated policies.
- **Complete restore** is useful when returning to a known good baseline or recovering from major errors.

Restore vs Revert Action

The Restore action loads an older version into the working configuration, while Revert Immediately rolls back and provisions a version.

Revert is used to undo a recent error, and Restore is used to return to a stable baseline.

Policy Versions

Each time you provision changes to policy items (such as policies, services, IP lists, label groups, and security settings), the entire set of changes you provisioned receives a version number. You can view the history of your policies and view their differences.

You can select a previous version to see information about that specific version. By default, the PCE retains only the last 1,000 versions of the policy and automatically removes older versions to improve performance. When a new change is provisioned, the oldest version of the policy is removed.

- 1.** Go to the page **Drafts & Versions > Versions**.

The Policy Versions page displays the history of the last provisions in your organization.

- 2.** Click one of the items to view details about the changes. In this version, you can see the changes that have been provisioned for the selected item.

Restore Policy

With the policy restore feature, you can revert to an older version of the policy when the newly provisioned policy does not work as expected.



NOTE

To use this feature, you must be a Global Administrator or Global Organization Owner.

The older policy version is copied to the current working draft version. You can immediately provision it to replace the non-working version.

You cannot restore to a previous version when there are pending changes. If you attempt to restore to this version, it will result in references to deleted non-versioned objects, such as labels and workloads. The restore will fail, and an error message will be displayed.

To revert to an older policy version:

1. Go to the page **Draft & Versions > Versions**.
2. On the Policy Versions page, click **Restore** for the policy version you want to revert to.

Provision Changes

Suppose you have made any changes to provisionable objects, such as policies, IP lists, services, label groups, and security settings. In that case, you need to provision those changes before they can take effect.

<https://product-docs-repo.illumio.com/Tech-Docs/Animated+GIFs/Provision+Changes.mp4>

1. Go to **Draft & Versions**

The Draft Changes page lists all policy items that have been added, modified, or removed. The top of the page summarizes changes by item type.

2. Select the items you want to provision.
3. Click **Provision** to preview the changes that will occur when you provision them.

**NOTE**

When you selectively choose items to provision, some of those items might also have dependencies that are published. Any object dependencies will also be provisioned.

4. Click **Provision** to push all the policy changes to workloads.

Revert Provisioned Changes

Any changes you make to policy configuration items (policies, IP lists, label groups, services, or security settings) appear as pending provisioning.

You can revert those changes before provisioning them.

1. From the PCE web console toolbar, click **Draft & Versions > Drafts**.
When you selectively choose items to provision, some of those items might also have dependencies that are published. Any object dependencies will also be provisioned.
2. Select individual items or all items to revert changes.
3. Click **Revert**.

Provision Note Setting

You can make a provision note mandatory before you provision the rules. It is disabled by default, but you can enable it to make it mandatory. This feature supports the need to describe context before provisioning and can support your organization's internal workflow.

When enabled, you must populate the note field before provisioning changes.

Users should populate the Provision Note field with a link to their internal bug-tracking system or a project number for tracking purposes. The error message they see when they leave the field empty will remind them to do so.

Illumio Segmentation for Data Centers does not validate the content entered in the Provision Note field.

You cannot provision updates until you enter text in the Provision Note field. The button **Confirm & Provision** is initially grayed out. After you enter the

appropriate text in the field, the button is enabled, and you can provision the update.



NOTE

You must have the correct role and permissions to access this feature. If necessary, contact your Illumio administrator for assistance.

To make the provision note mandatory:

1. Choose Policy Settings

The option in the Policy Settings page for PROVISIONING is set to **No** by default.

2. Click Edit.

3. Change the option Require Provision Note to Yes.

4. Click Confirm and Save.

Segmentation Templates

Illumio's Segmentation Templates offer pretested security policies for intricate applications such as Active Directory. These templates simplify creating and implementing security policies, minimizing errors and bolstering protection for vital assets. Leveraging enterprise application knowledge, Illumio streamlines policy creation, ensuring rapid deployment within organizations. The PCE web console automatically sets up key policy objects after installation to facilitate seamless application communication.

Catalog Retrieved from Support Portal

When you go to the Segmentation Templates page, the PCE web console automatically retrieves the latest Segmentation Templates catalog from the Illumio Support portal and displays it in the web console.



NOTE

You can access the Segmentation templates only directly through the Support Portal.

1. Access the Support portal using your Illumio Support portal username and password. (Illumio Cloud customers are automatically logged into the Illumio Support portal.)
Click **TOOLS > Illumio Segmentation Templates** .
2. To view the contents of a Segmentation Template, click its name or icon. The Segmentation Template details page describes the template and lists all the policy objects that belong to the template. Policy objects appear as hyperlinks when another template has already installed them. (Templates can share policy objects.)

Features of Segmentation Templates

Segmentation Templates share the following key features.

Template Contents

Each Segmentation Template adds an associated group of unique, non-overlapping, predefined services, and can contain any of the following policy objects:

- Labels
- Label groups
- IP lists
- Rulesets

Some templates contain all the necessary rulesets, services, and labels to secure a specific application, while others contain only port-based service definitions.

Dynamic Processes and Ports in Microsoft Environments

Segmentation templates are valuable in Microsoft environments, where dynamically allocated ports are frequently used for Remote Procedure Calls (RPC). Microsoft applications like Active Directory require dynamic port ranges to enhance security. The Illumio PCE, being service and process-aware, secures against dynamic processes such as Netlogon by focusing on specific server processes and paths while implementing precise rules for heightened security.

Sharing Policy Objects

Multiple Segmentation Templates can use services, labels, label groups, and IP lists. However, multiple templates never use a ruleset.

Identifying Policy Objects Added by Templates

You can recognize all objects integrated into the PCE through Segmentation Templates. In the object's details page External Data Set field, these policy objects are labeled with the format:

IST - type_of_object

(IST represents Illumio Segmentation Template). For better readability, the PCE also presents complete names. For instance, "IST - [AD] - Client to Domain Controller" is displayed as "IST - Active Directory Client to Domain Controller."

Segmentation Template Prerequisites and Limitations

The following prerequisites and limitations bind Segmentation Templates.

Internet Connectivity

Internet connectivity is not mandatory to use Segmentation Templates, enabling you to connect to the PCE web console without internet access. If offline, you can download Segmentation Templates from the Illumio Support portal on an internet-connected device and upload them locally.

Upgrade Policy Object Installed by Segmentation Templates

The PCE recognizes when Segmentation Templates install policy objects from the values in the External Data Reference field.

Unique Names for Labels, Label Groups, and IP Lists

Policy object names in the PCE web console must be unique. If duplicates exist, the template installation process prompts users to modify the object names for clarity and consistency.



NOTE

In Segmentation Templates, policy objects are named using the following convention: **IST - *type_of_object***

Delete Labels Associated with Segmentation Templates

Removing labels associated with Segmentation Templates requires removing rulesets and label groups first. Labels cannot be deleted until these prerequisites are met.

Editing Segmentation Templates

When you install a Segmentation Template, it brings in a fixed set of services and allows for the inclusion of labels, label groups, IP lists, and rule sets.

Editing a policy object tied to a Segmentation Template differs from editing other objects in the PCE web console. The appearance and identification of a Segmentation Template remain constant in the PCE web console even after modifying associated policy objects.

Before altering policy objects linked to a Segmentation Template, consider:

Editing Policy Object Names or IDs

The PCE assigns an ID number to each policy object associated with a template, displayed in the Description and External Data Reference fields on object details or Summary pages.

Policy objects tied to Segmentation Templates are identified by their names, structured like:

IST - *type_of_object*

Altering the policy object name doesn't impact PCE validation of its installation, but editing the External Data Reference field through the Illumio API does affect this validation.



NOTE

Illumio strongly recommends not changing the IDs in the External Data Reference fields.

Deleting or Editing Policy Objects

Deleting policy objects linked to templates or modifying their attributes comes with the following considerations:

- If you remove a policy object associated with a template after installation, the object will be re-added when the template is updated.
For instance, if you delete the common LDAP service from a Segmentation Template, an update to the template will re-add the LDAP ports.
- Editing attributes of policy objects tied to a template necessitates a choice in the PCE web console when updating to the next version: whether to maintain or overwrite the changes you made.

Segmentation Templates Installation and Upload

Install a Segmentation Template

1. Retrieve the Segmentation Template Catalog.
When a template has not been installed, an **Install** button appears on the page.

2. Click **Install**.

The End User License Agreement (EULA) appears.

3. Accept the EULA and click **Continue**.

Before the PCE installs the template, it checks that the policy objects required by the template don't conflict with any existing policy objects in your organization. The time it takes to process the check depends on the number of policy objects in your organization. When the PCE detects any conflicts during the check, it cancels the installation and does not install any policy objects. You are prompted to rename the conflicting objects. When the check is successful, the PCE adds the included policy objects to Draft mode, allowing you to review and edit them before provisioning. As the policy objects are added, links to the objects appear in the template details page.



NOTE

Global policy objects—such as All Services and Any (0.0.0.0/0 and ::/0)—don't include links to the objects in the Segmentation Template details page.

Upload a Segmentation Template

When you download a Segmentation Template from the Illumio Support portal, you save the template locally as a JSON file.

1. Log in to the Illumio Support portal with your Illumio Support username and password.
2. Click **Tools** > **Illumio Segmentation Templates**.
3. On the "Illumio Segmentation Templates" page, click the **DOWNLOAD** button.
4. Accept the EULA license agreement and click **Continue**.
5. Name the template and define where to download it on your system. Click **Save**.

Update a Segmentation Template

Updating a Segmentation Template to a later version allows you to edit or add services, rule sets, labels, label groups, or IP lists. However, updating a template does not remove policy objects added by a previous version.



NOTE

Later versions of templates are fully backwards-compatible with previous versions.

1. Retrieve the Segmentation Template Catalog.
When a new version of a Segmentation Template is available for a template that you have installed, the template displays an "Update" button.
2. Click **Update**.
If you edit the Segmentation Template after installing it, a dialog box appears prompting you to specify how to install the new version. For example, you added a new port and protocol to a service that the template created. You can revert the template to the Illumio list of ports and protocols for that service or keep your changes.
3. If necessary, choose how to handle template changes:
 - **Overwrite:** The PCE replaces the policy objects that you edited with the version in the new template and removes the word "edited" after the ID number in the External Data Reference field.
 - **Preserve Changes:** Your changes to the policy objects added by the template are kept.

**NOTE**

If you have edited multiple policy objects associated with a template, you must choose whether to overwrite or preserve all your changes. You cannot overwrite some and preserve some.

The PCE updates the version numbers of all policy objects associated with the template, even when the new template changes only a subset of the objects.

**NOTE**

Segmentation Templates can share policy objects; therefore, a policy object can have a later version than its associated template, because another template updated the object. For example, you can have version 1 of a template installed, and it includes version 2 of some policy objects.

Uninstall a Segmentation Template

1. Retrieve the Segmentation Template Catalog.
After you install a Segmentation template, an **Uninstall** button appears on the page.
2. Click **Uninstall**.
When you uninstall a Segmentation Template, the PCE removes all the policy objects that are associated with that template, except when an object is in use. Policy objects that are shared with other installed templates are not removed. Policy objects that are added to other policy objects are not removed. For example, you added a service associated with a template to a ruleset.

Secure Workload Connections

This section describes SecureConnect and AdminConnect, which are Illumio-provided encryption options.

SecureConnect was developed for host-to-host traffic encryption between paired workloads.

AdminConnect was developed to get control access to network resources based on Public Key Infrastructure (PKI) certificates.

SecureConnect

Enterprises have requirements to encrypt in-transit data in many environments, particularly in PCI and other regulated environments. Encrypting in-transit data is straightforward for an enterprise when the data is moving between data centers. An enterprise can deploy dedicated security appliances (such as VPN concentrators) to implement IPsec-based communication across open, untrusted networks.

However, what if an enterprise needs to encrypt in-transit data within a VLAN, data center, or PCI environment, or from a cloud location to an enterprise data center? Deploying a dedicated security appliance to protect every workload is no longer feasible, especially in public cloud environments. Additionally, configuring and managing IPsec connections becomes more difficult as the number of hosts increases.

SecureConnect Overview

SecureConnect leverages the built-in IPsec subsystem of host operating systems. On Windows hosts, SecureConnect utilizes the Windows IPsec subsystem. On Linux hosts, SecureConnect utilizes StrongSwan and Linux kernel IPsec for traffic encryption.

With SecureConnect, Illumio delivers a feature that configures the Security Policy (SP) necessary to enable traffic encryption between workloads. Once authenticated, encryption and cryptography suites provide confidentiality and data integrity to network traffic between workloads.

The PCE centrally manages all Security Policies (SPs) for workloads, enabling them to be policy-driven. For example, a customer can require that all traffic between their web servers and database servers be encrypted. Selecting the SecureConnect option for these workloads enables the PCE to apply the requisite security policy to your organization, making that happen. SecureConnect simplifies the configuration of IPsec encryption and automatically scales according to your policy definitions.

SecureConnect Use Cases

Employing SecureConnect is especially beneficial in these common scenarios:

- Facilitate PCI compliance by ensuring that confidential data is encrypted over the network.
- Secure off-site backup and recovery of data across geographically distributed data centers.
- Secure communications across applications and application tiers for regulatory compliance and tighter security.
- Enable secure data migration across different public cloud providers.

SecureConnect Features and Enforcement

SecureConnect supports connections between Linux workloads, Windows workloads, and mixed Linux and Windows workloads.



NOTE

SecureConnect rules are only applied to workloads where the VEN is in a non-idle enforcement state.

However, unlike other rules, SecureConnect requires matching rules to be applied to workloads on BOTH sides of any connection. Therefore, SecureConnect traffic is not supported between two workloads where a VEN on either side is in the idle state.

SecureConnect Rules and Visibility-Only State

Illumio employs an allowlist security model. By default, workload-to-workload communication is blocked unless explicitly permitted by defined Illumio pol-

icy rules. Administrators create these explicit rules to allow only necessary traffic, significantly enhancing security.

SecureConnect Rules



NOTE

SecureConnect rules are only applied to workloads where the VEN is in a non-idle enforcement state.

However, unlike other rules, SecureConnect requires matching rules to be applied to workloads on both sides of any connection. Therefore, SecureConnect traffic is **not** supported between two workloads where a VEN on either side is in an idle state.

For SecureConnect rules in visibility-only state, it is essential to keep in mind that these rules are:

- Applicable only to workloads in an enforced state (Visibility-only, Selective, or Full Enforcement).
- Matching rules are required on both source and destination workloads.
- Unsupported for workloads in Idle state.

The visibility-only state offers no enforcement and represents continuous monitoring and reporting of network traffic. It is ideal for initial policy planning and traffic analysis. However, it may disrupt applications dependent on NAT or IP forwarding.

Blocked + Allowed Logging Mode

This mode provides detailed logging of:

- Allowed traffic (explicitly permitted by rules).
- Blocked traffic (explicitly denied or not explicitly permitted).
- Unlocked traffic (permitted without explicit rules).

Visibility Options by Enforcement Mode

These options are available for the selective and full enforcement modes:

Selective Enforcement Mode

Selective enforcement provides:

- Off—There is no logging. The VEN does not collect any information about traffic connections. This option provides no Illumination detail and demands the least amount of system resources from a workload.
- Blocked—Logs only blocked traffic. The VEN collects only the blocked connection details (source IP, destination IP, protocol, and source port and destination port), including all dropped packets. This option provides less Illumination detail but demands fewer system resources from a workload than high detail.
- Blocked + Allowed – Logs both allowed and blocked traffic. The VEN collects connection details (source IP, destination IP, protocol, source port, and destination port). This applies to both allowed and blocked connections. This option provides rich Illumination detail but requires some system resources from a workload.
- Enhanced Data Collection – Detailed logs with traffic flow metadata.

Full Enforcement Mode

Full enforcement blocks all non-explicitly allowed traffic, providing the highest level of security.

Visibility options mirror Selective Enforcement:

- Off
- Blocked
- Blocked + Allowed
- Enhanced Data Collection

Full enforcement is recommended after successful testing and validation of the allowlist model.

Enhanced Data Collection

As of release 25.2.10, Enhanced Data Collection is enabled in all enforcement modes. Before February 25, 2010, it could be enabled only in Full Enforcement mode.

Enhanced Data Collection allows the VEN to log byte counts and connection details for Allowed, Blocked, and Potentially Blocked traffic.

AdminConnect

Relationship-based access control rules often use IP addresses to convey identity. This authentication method can be effective. However, it is not advisable to use IP addresses to establish identity in certain environments.

When you enforce policies on servers for clients that frequently change their IP addresses, the policy enforcement points (PEPs) must continuously update security rules to accommodate IP address changes. These frequent changes can cause performance and scaling challenges, and the IP sets of protected workloads can become unstable.

Additionally, using IP addresses for authentication is vulnerable to IP address spoofing. For example, server A can connect to server B because the PEP uses IP addresses in packets to determine when connections originate from server A. However, in some environments, bad actors can spoof IP addresses and impact the PEP at server B, causing it to mistake a connection from server A.

Illumio designed its AdminConnect (Machine Authentication) feature with these environments in mind. Using AdminConnect, you can control access to network resources based on Public Key Infrastructure (PKI) certificates. Because the feature is based on the cryptographic identity associated with the certificates, not IP addresses, mapping users to IP addresses (common in firewall configurations) is not required.

With AdminConnect, a workload can use a client's certificate-based identity to verify its authenticity before allowing it to connect.

Features of AdminConnect

Cross Platform

Microsoft Windows provides strong support for access control based on PKI certificates assigned to Windows machines. Modern data centers, however, must support heterogeneous environments. Consequently, Illumio designed AdminConnect to support both Windows and Linux servers and Windows laptop clients.

AdminConnect and Data Encryption

When AdminConnect is the only feature enabled, data traffic does not utilize ESP encryption. This ensures that data is in clear text even though it is encapsulated in an ESP packet.

The ESP packets are encrypted when AdminConnect and SecureConnect are enabled for a rule.

Ease of Deployment

Enabling AdminConnect for identity-based authentication is easy because it is a software solution that does not require deploying network choke points, such as firewalls. It also does not require you to deploy expensive solutions such as Virtual Desktop Infrastructure (VDI) or bastion hosts to control access to critical systems in your data centers.

AdminConnect Prerequisites and Limitations

Prerequisites

You must meet the following prerequisites to use AdminConnect:

Limitations

You cannot enable AdminConnect for the following types of rules:

- Rules that use All services
- Rules with virtual services in sources or destinations
- Rules with IP lists as sources or destinations
- Stateless rules

AdminConnect is not supported in these situations:

- AdminConnect does not support “TCP -1” (TCP all ports) and “UDP -1” (UDP all ports) services.
- You cannot use Windows Server 2008 R2 or earlier versions as an AdminConnect server.
- Windows Server does not support more than four IKE/IPsec security associations (SAs) concurrently from the same Linux peer (IP addresses).

Legal Notice

Copyright © 2026 Illumio 920 De Guigne Drive, Sunnyvale, CA 94085. All rights reserved.

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

Resources

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

Contact Information

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