

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

Table of Contents

Security Policy	4
Overview of Security Policy	4
About the Illumio Policy Model	4
The Illumio Policy Model	4
Types of Illumio Policy	7
Static and Staged Policies	8
Security Policy Objects	14
Labels and Label Groups	14
Services	20
Virtual Services	25
IP Lists	30
Export Reports	33
Workloads	33
Workloads in the PCE	34
Workloads and VENs	46
Workload Setup Using PCE Web Console	51
VEN Administration on Workloads	52
Loopback Interfaces	52
Blocked Traffic	53
Create Security Policy	56
Core Services Detector	56
Rulesets	59
Rules	69
FQDN-Based Rules	98
About Provisioning	103
How Provisioning Works Internally	103
Versioning, Restore, and Revert	104
Policy Versions	105
Restore Policy	105
Provision Changes	106
Illumio Policy Enforcement Model	107
Why Use Selective Enforcement?	107
Applying Selective Enforcement	107
How Selective Enforcement Works	108
Enforcement Progression Model	108
Use Cases and Limitations	108
SecureConnect Rules and Visibility-Only State	109
Enhanced Data Collection	111
Policy Exclusions	111
Secure Workload Connections	116
SecureConnect	116
AdminConnect	117
Legal Notice	122

Security Policy

Overview of Security Policy

This section describes the security policies, which are configurable sets of rules that protect network assets from threats and disruptions. Illumio Core relies on security policy to secure communications between workloads.

About the Illumio Policy Model

The Illumio security policy for securing workloads differs from traditional network security policies. Traditional security policies use network constructs, such as VLANs, zones, and IP addresses, to tie security to the underlying network infrastructure.

In contrast, the Illumio security policy uses a multidimensional label system to sort and describe the function of workloads. By describing workload functionally, policy statements are unambiguous. Illumio users assign four-dimensional labels to their workloads to identify their roles, applications, environments, and locations. Additionally, users specify labels in the scopes for rulesets and in the providers and consumers components of rules, allowing their organization's workloads to communicate with each other.

Labeling workloads and creating the corresponding rulesets and rules define the security policies for workloads. The PCE converts these label-based security policies into the appropriate rules for the OS-level firewalls of the workloads.

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.

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 clearly 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.

Security Policy Guidelines

The following guidelines are recommendations on how to create your security policy in Illumio Core. Creating a security policy is an iterative process, so following these recommendations will provide a broad initial policy, which can then be incrementally improved until a sufficiently robust policy has been 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 ruleset, you can preview its potential effects using Illumination's Draft View, which shows what changes will occur once the policy is enforced.

- **Visibility only:** Initially, policies are refined 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 enforcing security rules temporarily 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.

About Rulesets and Rules

Rules form the core of Illumio security policy. A ruleset is a collection of rules defining permitted network traffic. Create the rules using labels that identify your workloads.


Understanding Rulesets and Rules

Illumio's Illumio Core **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 Core) in this example are:

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

In Illumio Core, 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 ruleset for each workload. Use labels to identify your workloads and scopes to apply the rulesets to multiple workloads simultaneously.



NOTE

Illumio recommends creating no more than 500 rules per ruleset, or the PCE web console will not be able to display all of the rules.

If you want to create a ruleset with more than 500 rules, Illumio recommends splitting the rules across multiple rulesets or using the Illumio Core REST API, where there is no limit on the number of rules you can create per ruleset.

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 write a policy for workloads for load balancers to manage traffic.
- **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.

Types of Illumio Policy

This section explains the differences between adaptive and static policy in the Illumio Core

Scopeless Policies

Scopeless policies are used in special cases where a rule needs to be applied broadly across all or large groups of workloads in a network.

You can write rules freely using specific labels without restrictions. If labels are excluded in a category, it defaults to "All" to include workloads without labels.

Scopeless policies require caution, as mistakes can open broader communications than intended. For example, a Default ruleset might permit 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 how broadly rules are applied, limiting the impact of mistakes. However, the restrictive scope also limits how broadly rules can be written.

Single Scope Policies

Single-scope policies are the most commonly used policy type.

They narrow the list of workloads rules apply to and allow cross-communication within the scope.

Single-scope policies are commonly used to write rules for:

- Specific app group
- Broader policy such as a core service policy for an environment
- Broader policy for a location

Advanced Scope Policies

Advanced scope policies can be grouped into:

- Multi-Scope Policies
These policies apply the same rules to many different groups of workloads, each scope as an independent policy.
They apply the same set of rules to multiple groups of workloads, one scope at a time. Then, they proceed to the next scope and repeat the process for the remaining scopes.
- Single-scope Policies
These policies narrow the list of workloads to which the rules apply and allow cross-communication within the scope.

Static and Staged Policies

This section explains the static and staged policy in the Illumio Core.

Static Policy

Static Policy allows administrators to stage policy updates on workloads matching a defined label-based scope. These workloads will receive but not apply new firewall rules until manual approval (via Apply Policy action in the UI or API).

For most of your workloads, adaptive security is the best method for protecting them from the lateral spread of threats. By default, the Illumio Core implements adaptive security for your workloads in all roles, all applications, all environments, and all locations. S

However, in certain scenarios, you might want to control when the VENs apply new or changed OS-level firewall rules to workloads. Using labels, you designate which workloads are impacted by static policy. See [Apply Static Policy \[11\]](#) for the steps to configure static policy using labels.

When you configure the Policy Update Mode for workloads to use static policy, you control when the Illumio VENs running on the workloads apply new OS-level firewall rules that they received from the PCE. The Illumio Core blocks the immediate application of new firewall rules that result from users provisioning policy changes in the PCE and from dynamic updates to firewall rules (adaptive policy) when your environment changes. For example, you add a new rule to a ruleset in the PCE and provision the change, or a change occurs in your environment, such as a workload changing its IP address. In both cases, the VENs for your impacted workloads receive the new OS-level firewall rules from the PCE but they do not apply them until you explicitly select the workloads and click **Apply Policy** in the PCE web console.

See [Staged Policy \[11\]](#) for information about how it uses static policy and stages OS-level firewall updates rather than applying them immediately.

You should view static policy as a Security Setting rather than a type of security policy because configuring workloads to use static policy is a mechanism to control when VENs apply new or updated OS-level firewall rules to affected workloads. You can use the static policy setting to establish an audit trail of which Illumio users apply new OS-level firewall rules to workloads and when they apply them.

Use Cases for Static Policy

By default, the PCE is set to apply security policy updates dynamically through adaptive policy. However, scenarios occur where you want to control when updates to the OS-level firewall rules are applied to workloads.

For example, you might want to control when these updates occur in the following scenarios:

- Corporate policy for business-critical applications requires oversight on when updates to the OS-level firewall rules are applied to workloads.
For example, a financial institution requires that its security team explicitly control security updates to its transaction processing application. The security team authorizes the date and time of the update and applies it in the Illumio PCE.
- The corporate IT team has established policies for applying security updates during disparate maintenance windows.
The IT team utilizes distributed maintenance windows to lessen the up-time impact on applications; for example, half the application is upgraded during the first maintenance window and the second part during the second maintenance window to keep the application running and minimize risk.
- The central security team sets the security policy static for certain environments and adaptive for others.
For example, the security policy is adaptive for workloads running in the development environment (using the labels All Applications, Development Environment, and All Locations). However, workloads in the production environment (All Applications, Production Environment, and All Locations) require the static policy.

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

Example: Static Policy Workflow

The security team for an internet retail application has strict requirements for updating their production environment. They require that all updates to the OS-level firewall rules for their Database tier running in production be applied during maintenance windows. For their Illumio-managed workloads, they configure a static policy with the following labels: Role: Database, Applications: All, Environment: Production, Locations: All.

A spike in customer demand occurs, and the production environment automatically scales by adding servers to the Web tier. The Illumio PCE detects the web servers connecting to the Database tier workloads and re-computes their security policy to include rules for the web servers. The PCE re-computes the OS-level firewall rules for those workloads and sends them to the VENs running on the Database workloads. The VENs stage the updates locally but do **not** apply them to OS-level firewalls.

A maintenance window opens, and a security team member filters the Database workloads in the PCE to determine which ones have staged security policy. She selects these workloads and applies the staged changes.

The VENs request the latest OS-level firewall rules from the PCE to ensure that all changes are included. The PCE sends the latest OS-level firewall rules to the VENs, who apply them.

Static Policy Prerequisites, Limitations, and Caveats

Before configuring your workloads to use static policy, review the following prerequisites and limitations and consider the following caveats.

Prerequisites

- You must be a member of the Global Organization Owner or Global Administrator role to manage security settings and add static policy.
- The VENs on affected workloads must be running version 17.2 or later. Earlier versions of VENs cannot stage static policy. They will apply security policy updates immediately to workloads even though you configured them to use static policy.

Limitations

- You should provision label groups before adding them to the static policy.
- In the following situations, a VEN will apply a security update immediately and will not stage it even though the workload on which the VEN is running is configured to use static policy:
 - When you pair a new workload, the VEN immediately applies the policy it receives from the PCE.
 - When a VEN detects tampering, it requests security updates from the PCE and applies them immediately.
 - A VEN is offline when a user applies changes to their workloads. When the VEN returns online, it connects to the PCE and receives updated OS-level firewall rules. The VEN applies the revised rules to the workload even though it is configured to use a static policy.



NOTE

When a VEN goes offline and online, its OS-level firewall rules can become out of sync with other VENs that remain online.

See [Staged Policy \[11\]](#) for an explanation of how the VENs stage policy.

Because a VEN may apply security updates immediately, Illumio recommends that you do not provision security policy updates until they are final. Keep the updates in the Draft state until you complete them.

- To maximize performance, the PCE transmits 5,000 updated OS-level firewalls to the VENs until all updates are sent.

Caveats

Illumio recommends implementing a static policy for special cases, and advanced users should oversee the process.

The Illumio Core is designed to ensure that the latest versions of your security policy across your environment protect your workloads. Users provision policy changes, or the PCE responds dynamically to environmental changes. In both cases, the PCE re-computes new

OS-level firewall rules incorporating the changes and sends them to the VENs to be applied immediately.

However, when configuring workloads to use static policy, you override this design by controlling when the VENs apply the security update to the workloads. As a result, you can have inconsistent security policies across your managed environment, which can cause communication disruptions between workloads.

Troubleshooting communication issues is difficult when the workloads within a scope use different security policy versions.

Illumio recommends keeping the number of workloads in your environment that utilize static policy as low as your business processes allow.

Apply Static Policy

By default, the Illumio Core implements adaptive security for your workloads in all roles, all applications, all environments, and all locations.

However, you might want to add static policy to control when updates to OS-level firewall rules are applied to your workloads.

You designate which workloads use static policy by configuring the Policy Update Mode in the Security Settings. To configure the Policy Update Mode, you specify the role, application, environment, and location labels. Any workloads within the scope of the selected labels will use a static policy. You can add multiple scopes. The overlap between the scopes does not affect how workloads use static policy.

Label groups are currently not supported by static policy. To create scopes using multiple labels from the same type, add them as separate scopes. For example, four Role labels are added to the PCE: Web, Database, API, and Mail. You want to add a static policy for the Web and Database roles only, so you add two scopes.

See [Static Policy Prerequisites, Limitations, and Caveats \[10\]](#) for information before you complete this task.

To add static policy:

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 static policy's scope.
3. Select labels to select workloads for static policy (Role, Application, Environment, Location)
4. Click **OK**.
The static policy appears in the list.
5. Click **Provision** from the PCE web console toolbar.

Staged Policy

When a workload matches a defined Static Policy filter, its state appears as Staged, indicating that policy updates have been delivered to the VEN but are not yet enforced. This policy

requires you to define a policy that matches the workload's Static Policy filter."Workloads that match the Static Policy scope will show the Staged policy state after provisioning. This means the VEN has received updated rules but has not applied them yet. For this to happen, active policy rules (e.g., rulesets or services) must affect those workloads.

Understanding the distinction between using static policy to stage updates to OS-level firewall rules and provisioning security policy is important because the actions differ in crucial ways.

When you configure workloads to use static policy, the PCE sends the new OS-level firewall rules for Linux iptables or the Windows Filtering Platform (WFP) to the VENs, who stage them locally. The VENs do not apply the new firewall rules immediately. You must select the workloads and explicitly click **Apply Policy** in the Workloads page to activate the staged OS-level firewall rules.

Configuring a set of workloads to use static policy does not eliminate the requirement to provision policy updates for those workloads. Through provisioning, you update the PCE's version of your security policy.

When you provision security policy changes, you trigger the PCE to apply these changes to the workloads. When the workloads are set to use static policy, the VENs on the workloads will stage the changes until you explicitly click **Apply Policy**. However, under certain circumstances, the VENs could apply the latest changes before you explicitly click **Apply Policy**. See [Static Policy Prerequisites, Limitations, and Caveats \[10\]](#) for information.

**TIP**

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

In addition to rulesets and rules, you must provision changes to the Illumio policy objects, such as services, IP lists, and label groups. Illumio supports including re-usable policy objects in intra- and extra-scope rules to make security policies easier to maintain and update. When you update a policy object, all the rules using the object are updated without you needing to change each rule where the object is included.

When you provision changes to rulesets and policy objects, the PCE saves your security policy as a new version. It recomputes the OS-level firewall rules for all the workloads affected by the change and instructs the VENs to download the updated OS-level firewall rules.

See the following topics related to provisioning:

- Overview of Policy Objects, for a description of each type of policy item
- Provisioning for the policy items that require provisioning
- Active vs Draft Versions to learn how provisioning establishes the active version of the policy

Determine When Workloads Have Staged Policy

Workloads Page

The Workloads page displays each VEN's current state in the Policy Sync column. You can filter your workloads by this column to determine which ones have staged OS-level firewall rules.

- **Active (Syncing):** The PCE sends a new policy to the VEN. This process typically takes only a few seconds.



NOTE

Workloads configured for adaptive and static policy can appear in the active (syncing) state while the PCE sends a new policy.

- **Staged:** The VEN has received the latest OS-level firewall rules but has not applied them.
- **Active:** The VEN has received, applied, and confirmed all policies sent from the PCE. (Active workloads have a green dot icon.)

For more information about the VEN Policy Sync states, see “VEN Policy Sync” in VEN Installation and Upgrade Guide.

Workload Details

The Workload details page provides important information about when and how your workloads received staged policy.

- The General section indicates whether the workload is configured to use static policy (Policy Update Mode field) and displays the date and time the VEN staged the policy (Last Policy Staged field).
- The VEN section includes the Policy Sync state, which can be active (syncing), staged, active, error, warning, and suspended.



NOTE

When all your workloads are configured to use adaptive policy, these fields will not appear in the General or VEN sections.

Apply Staged Policy

See [Static Policy Prerequisites, Limitations, and Caveats \[10\]](#) for information before you complete this task.

1. From the PCE web console menu, choose **Workloads**.
The Workloads page appears.
2. (Optional) Filter by Policy Mode:
Use the Workload property filter: **Policy Update Mode > Static Workloads**
3. To apply 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, choose **Apply Policy > Update All Workloads**.

**NOTE**

If you filter workloads by label and choose Update All Workloads, the PCE applies the staged updates to all the workloads matching that label scope, not just the workloads appearing 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.


Security Policy Objects

This section describes the policy objects that you can use to write security policies.

Labels and Label Groups

The Illumio Core policy model is a label-based system, which means that the rules you write don't require using an IP address or subnet, like traditional firewall solutions. You control the range of your policy by using labels. This helps you categorize your workloads more quickly and makes it easier to set up your policy.

Label Types

Label	Description
Role	<p>This label type allows you to describe a workload's role (or function). In a simple two-tier application consisting of a web server and a database server, there would be two roles: Web and Database. You can use the same role as many times as you want in other rulesets for different applications.</p> <div>  <p>NOTE The Role label cannot be used to define the scope.</p> </div>
Application	<p>This label type allows you to describe the application that a workload supports. When two servers in a two-tier application have a relationship with one another because one provides a service (like a database) to another, they likely constitute an application.</p> <p>If an organization has 100 applications, and each application has a separate web role and separate database role, the application role separates each one of the Web and Database roles.</p>
Environment	<p>This label type allows you to describe a workload based on its stage in the product development lifecycle, such as QA, staging, and production.</p>
Location	<p>This label type allows you to describe a workload based on its location. For example, Germany, the US, Europe, and Asia. Or, Rack #3, Rack #4, Rack #5; or data center AWS-east1, AWS-east2, and so on.</p>
Flexible labels	<p>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.</p>

Additional Dimensions

A given workload cannot have more than one label per type. It's possible to allow a workload that uses a service or services or across boundaries to communicate; for example, if a server plays multiple roles, such as a database server used by two different applications, Illumio recommends creating different role labels for that workload.

System Default “All” for Labels

When you log into the PCE for the first time as the organization owner, the following default labels are provided:

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

The built-in (default) Environment, Application, and Location labels are defined as “All,” which enables you to create broad policies to cover All Applications, All Environments, and All Locations.

To avoid confusing policy writers, Illumio recommends not creating labels named “All Applications,” “All Environments,” or “All Locations” (exactly as written in quotes).

When you attempt to create labels of these types with the exact name as the system defaults, for example “All Applications,” an “HTTP 406 Not Acceptable” error will be displayed.

**NOTE**

You can modify or delete these default labels at any time.

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 or Label Groups page. Similarly, you can filter by label name, description, and provision status on the Label Groups page. For example, select Type: Location in the Label property filter if you only want to see location labels.

Create a Label Type

Illumio Core 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. For example, OS.
4. Enter singular and plural versions of the Display Name (Operating System and Operating Systems).
5. Choose an icon, and enter a one- or two-character unique initial to be displayed with the icon (such as OS).
6. Choose foreground and background colors to be used when the label is displayed.
7. 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.

Create a Label

1. From the PCE web console menu, choose **Policy Objects > Labels**.
2. On the Labels page, click **Add**.
3. Name: Enter a label name.

4. Type: Select one of the types such as Environment, Application, Role, Business Unit, or
5. Click **Save**.

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

Label Workloads

You apply labels to workloads to identify their function or purpose in an application (Role label), the application they belong to (Application label), their network environment (Environment label), their location (Location label), and any custom purpose you have defined (flexible labels; for example, OS). After a workload is labeled, you can write rules using the labels you applied.

After you create a label, you can label a workload in two ways:

- Automatically label the workloads when you pair them by adding labels in the pairing profile.
 - Add labels to the workload on the workload summary page.
1. Select **workloads > VENs** in the PCE console.
 2. Select a workload, and click **Edit** to select any or all label types to apply to the workload.

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. In the Illumio Core 20.1.0 release and higher, if you want to delete a Label and it was used by a Virtual Server, you can determine whether it was in use. The "In use by" column on the Labels page includes Virtual Servers. The Labels' summary page also displays the "In Use By Virtual Servers Yes/No" field.



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**.
2. Select the checkboxes next to the workloads you want to change labels.
3. Click **Edit Labels**.
4. 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 the workloads contain multiple labels. You can click in the Label field and add a new label.
 - When you remove a label assignment, that label is removed from all selected workloads.

5. When you are finished, click **OK**.

Label Groups

Label groups help you write your security policy more efficiently when you use the same labels repeatedly in rulesets. When you add those labels to a label group, the label group can be used in a rule or scope as a shortcut or an alias for multiple labels. The Label Groups list pages can contain up to 10,000 label groups, and the individual Label Groups pages can contain up to 10,000 members. You can use filters to find labels or label groups.

For example, you have workloads in data centers in Dallas, New York, and Washington, and you want to apply a rule to all those workloads. Instead of using the labels for Dallas, New York, and Washington in three separate rules, you can define a Location label group named the US, add those three location labels to the label group, and use the US label group.

Label groups are displayed as a list that includes the following details:

- Provision status
- Name of the label group
- Type (Role, Application, Environment, Location, or a custom-defined Flexible Label type)
- When it is currently in use by a ruleset, label group, and static policy
- Last modified date and time
- The user who last modified the label group

Policy Calculation Using Label Groups

Label groups can be nested, so it is important to understand how label groups can affect 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 rulesets.

Create a Label Group

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

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, enter a label name to find labels to add to the group, and then click **Add**. You can add as many labels (or label groups) of the same type to the group as desired.

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.

Use a Label Group in a Scope

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

For example, to create a scope that applies to all environments other than production, first create a Non-Prod label group which consists of the labels for the Dev, QA, and Stage environments. The following ruleset (scope + rule):

Scope:

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

Rule:

- Providers: DB
- Services: MySQL
- Consumers: 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

Use a Label Group in a Rule

When you use a label group in a rule, the label group is expanded into multiple rules. Cross-communication is allowed.

For example, the Non-Prod label group is used again here, but in the rule and not the scope, which allows cross-communication. The following ruleset (scope + rule):

Scope:

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

Rule:

- Providers: Non-prod DB
- Services: MySQL
- Consumers: 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

Services

When workloads are paired with the PCE, the VEN discovers all running processes and services on a workload and makes those services available for use when writing rules. You can see those discovered services when you view the Processes tab on the Workload’s details page.

However, you can also create your own to services to specify the service type, as well as the ports and protocols the services use to communicate.

**NOTE**

Service names can be unrestricted, for example, `sc.exe qsidtypemyservice`. You can write rules with unrestricted service IDs (SIDs). When there is a restricted SID, you should write rules without the SID. Including the service with a restricted SID type causes the traffic to be dropped and might cause traffic between the Reported view and Draft view to be reported inaccurately.

Service Types

When you create a service, you can choose one of two general types:

- **All OS: Port-Based:** This type of service can be used for writing rules for any workloads and is defined by specifying a port and protocol, a port range, or in some cases, only the protocol. For example: `80 TCP`, `1000–2000 TCP`, `500 UDP`. For GRE or IPIP, you only need to specify the protocol.
- **Windows: Process/Service-Based:** This type of service can be used for writing rules for Windows Workloads only, and is defined by specifying one of the following combinations or scenarios:
 - **Port and/or Protocol, Windows Process, and Windows Service**
`443 TCP c:\windows\myprocess.exe myservice`
 - **Port and/or Protocol and Windows Process**
`443 TCP c:\windows\myprocess.exe`

- **Port and/or Protocol and Windows Service**

443 TCP myservice

- **Windows Port and/or Protocol**

514 UDP

- **Windows Process**

c:\windows\myprocess.exe

- **Windows Service**

myservice

Windows Process-based Rules

Rules to Allow System Created Processes

You can create rules 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 word “system” (case-insensitive) in the Port/Protocol text input field. Once you create this service, you can use it in rules.

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

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

Operating System

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

Service Definitions

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

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

Click **Remove**.

Service Using Windows Environmental Variables

The Windows environmental variable can be used to specify the full path. This can be done by creating a Service of 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 to allow all system-initiated processes in Windows. This will allow 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.

To create a service that uses Windows environmental variables:

1. From the PCE web console menu, choose **Policy Objects > Services**.
2. Click **Add**.
3. In the Name field, enter **system** (case-insensitive).
4. From the Operating System drop-down list, select **Windows: Process/Service-based**.
5. In Ports & Protocols, specify the port/protocol, separating the port and protocol with a space. For example:
`%systemroot%\myprocess.exe`
6. Click **Save**.

IGMP Services

You can add Internet Group Management Protocol (IGMP) as a service for use in rules to write granular inbound or outbound policy for IGMP, which is typically used for multicast. No range is required for IGMP.

You can export IGMP traffic in JSON, CEF, or LEEF format.

You can also create and update services using the IGMP protocol using the Illumio Core REST API.

See "[Services](#)" in REST API Developer Guide for information about using the REST API to create services.

Caveats

- When IGMP service is used in a rule, all IGMP types are allowed; however, granular control and specific multicast addresses are not supported.
- IGMP is not supported in the Illumination map.

ICMP Services

ICMP can be added as a service and used in rules to write granular inbound or outbound policy for ICMP. ICMP is usually used for traceroute and path MTU discovery.

You can export ICMP traffic in JSON, CEF, or LEEF format.




NOTE

When these services are blocked, they do not appear in the Blocked Traffic list and the connection is dropped silently.

ICMP types/codes (such as 0 ICMP or 3/2 ICMP) are supported. The ICMP range is from 0 to 255.

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. From the 19.1.0 release on, 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 Illumio Core REST API. See [Services in REST API Developer Guide](#) for information about using the REST API to create services.

Caveats

- ICMP is not supported for virtual services.
- When an ICMP service is used in a rule, all ICMP types are allowed; however, granular control and specific multicast addresses are not supported.
- When you enable IPv6 on Windows VENs, IPv6 system rules are not propagated to those VENs. You need to write security rules to ensure robust IPv6 functionality. 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

View or Edit a Service

To view or edit an existing service:

1. Click the name of the desired service. You can filter the list by various attributes. See [Filter the Services List \[24\]](#) 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 the 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

Filter the Services List

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

Services					
<div> + Add Provision Revert Remove Refresh Reports </div>					
Select properties to filter view					
<div> Customize columns 50 per page 1 – 6 of 6 Total </div>					
<input type="checkbox"/>	Provision Status	Name	Port/Protocol	Last Modified On Last Modified By	Description
		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	

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.

For a list of the types of services you can create, see [Service Types \[20\]](#).

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 a 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**.

Virtual Services

Virtual services (previously known as bound services) allow 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.

Overview of Virtual Services

A virtual service can be used in the following scenarios:

- **Apply rules to a single service:** Represents a service or process on a workload using a name or label. This approach allows you to write a policy to allow other entities to communicate only with that service. The policy does not need to change when the service is moved to a different workload or a new set of workloads. Only the workload bindings on the virtual service need to be changed. The PCE dynamically calculates the required rules on the updated workloads to allow this service.
- **Apply rules to multiple services (on the same workload):** Represents each service or process running on a workload with a different set of labels. You can write rules to allow other entities to communicate only with that service. The policy does not need to change when this service is moved to a different workload or a new set of workloads. Only the workload bindings on the virtual service need to be changed. The PCE dynamically calculates the required rules on the updated workloads to allow the service.

From the 18.3.1 release, Illumination, Policy Generator, and Explorer support virtual services. You have to assign labels to a virtual service to write label-based rules. A virtual service does not have an enforcement, so you need to refer to the enforcement of its bound workloads.

Virtual services are provisionable objects, which means they must be created and provisioned before they can be applied to workloads. However, the bindings are not provisionable

objects, so the bindings can be changed without having to provision the changes. Additionally, port overrides have been moved from the virtual service to the workload binding.

How Virtual Services Work

If a single workload runs both an Apache Tomcat and Apache HTTP server, supporting an HRM and ERP application, you can create a virtual service for each service and label one service as belonging to an HRM application and one to an ERP application. You can then write a set of label-based rules that apply only to the Apache Tomcat process serving the HRM application, effectively isolating it from the ERP application.

In the following example, two different virtual services are created: one for an HRM database and one for an ERP database. The following configurations allow the web to communicate with the database for each application (HRM or ERP) in the specified environment (Prod or QA) in the specified 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 Providers ← 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 Providers ← 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 Providers or Consumers field of the generated rules. You can configure virtual services using a port or port range.



NOTE

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

When you write a rule in a ruleset, you need to specify the following values:

- A service

- Providers of the service
- Consumers of the service

For example:

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 provider of the service.

For example:

Virtual Service Apache Tomcat is provided to All Workloads

When you want to treat the providers as a virtual service, select **Uses Virtual Services** or **Uses Virtual Services and Workloads** from the Providers drop-down list as the service.

To write a rule applicable for 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

Workloads labeled **Database** are not impacted by the above rule. 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 Configuration for Virtual Services

You have two advanced configuration options to consider when configuring a virtual service:

- **Apply To: Host Network or Internal Bridge Network:** This optional setting allows you to determine if the rules associated with the virtual service are applied over an internal bridged network or the host network. If you choose Internal Bridge Network, the rules associated with the virtual service are programmed into the FORWARD chain on Linux iptables (Windows ignores rules to internal bridge in this current implementation). Or, you can specify that a virtual service's rules are applied over the host network, programmed into the INPUT/OUTPUT chains in Linux iptables. Stateless rules are not supported when associated with the FORWARD chain; stateful rules are programmed.
- **Optional Configuration: IP Overrides:** Allows you to specify IP addresses or ranges (CIDR blocks) to be used for programming the rules associated with the virtual service instead of using the IP address of the bound workload. When IP overrides are specified on a virtual service and the virtual service is used in a rule, the IP addresses programmed on other hosts communicating with the virtual service are the IP addresses and subnets specified in the IP overrides rather than the IP addresses of the workloads bound to the virtual service.

A combination of stateless and forwarding rules on the same host, port, and consumer is not supported. For example, when a workload has a service running on a port with stateless rules, a forwarding rule to allow traffic to a container running on the same host using the same port does not work when the consumer is the same.

Host-only network

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

Providers	Services	Consumers
Virtual Service X	From Providers	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 for devices that are allowed to communicate with this virtual service, use the addresses/subnets specified in the IP overrides.

Internal Bridge Network

When you remove the IP override and change to Internal Bridge Network, you have the following security policy:

- An inbound rule on workload A for 80 TCP with source address 192.168.0.200 on the FORWARD chain of the firewall.

This means the rule applies to traffic destined for somewhere other than the host network namespace that hits the host firewall.

- An outbound rule on workload B for 80 TCP with destination address 192.168.0.100.

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 a Virtual Service

When adding a virtual service, you need to give it a name, select the service, and apply labels to it.

Bind it to the workload where the service is running. This binding instructs the PCE 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.



NOTE

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

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. (Optional) Choose whether you want the rules associated with the virtual service to be applied over an internal bridged network instead of a host network (default behavior).
 - **Internal Bridge Network:** The rules associated with the virtual service are programmed into the FORWARD chain on Linux iptables.
 - **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 binding a virtual service to a workload.

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 to which 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 service for the virtual service, you cannot enable port overrides on the workload bindings.

7. In the Ports/Protocols section, enter the TCP and UDP ports for this virtual service.
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 rulesets to create rules for workload traffic flows. When you provision the rulesets, the workload only allows IP addresses in the IP list to access workload services.

The default IP list **Any** represents all IPv6 addresses as well as all 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 application of the rule to a broader set of workloads than intended.

Example of IP List Usage

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

Scope:

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

Rule:

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

This means “allow SSH from Corp-HQ to the database.”

This ruleset:

Scope:

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

Rule:

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

This means “allow SSH from the database to Corp-HQ.”

This ruleset:

Scope:

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

Rule:

- Source: Any
- Services: Any
- Destination: Any

This means “do not apply Any IP list to anything.”

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:** Add IP addresses, IP address ranges, or CIDR blocks to define the list.

**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).

Export Reports

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

Overview of Export Reports

CSV is the most common and popular format because you can import it into other tools like CMDBs. You can export the following objects into an export report:

- Workloads
- Rulesets
- IP lists
- Pairing profiles
- Services
- Labels
- Label groups
- Virtual services
- Virtual servers

Generate an Export Report

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

Workloads

This section describes workload attributes, it's 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

Workloads in the PCE

This section describes how to manage workload using the Workload pages in the PCE web console.

Workload Summary

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

- The name of the workload
- A description (if provided)
- The workload enforcement state
- The visibility 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. The Illumio Core includes four policy states for workloads. If a workload is unmanaged, the Policy State column is not displayed on the workload list page.

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.

A pairing profile can be used to pair workloads in the idle state.

**NOTE**

SecureConnect (IPv6 compatibility) is not supported on workloads in the Idle state. When you activate SecureConnect for a rule that applies to workloads that are in both Idle and Non-idle policy states, it can impact the traffic between these workloads.

Visibility Only

In the Visibility Only state, the VEN takes control of the host firewall, attempts to load kernel modules if required, and reports traffic flows to the PCE. The VEN will never block traffic in Visibility Mode (see note below). In this state, the PCE displays the flow of traffic to and from the workload, providing insight into the datacenter and the applications running in it. Visibility Only is useful when firewall policies are not yet known, allowing you to discover the application traffic flows in the organization and then generate a security policy that governs required communication.

**NOTE**

Depending on the workload's operating system, it may be possible for existing configurations or third-party tooling to be already interacting with the host firewall prior to VEN deployment. Therefore, Illumio recommends that you activate the VEN in Idle mode and then run the [Compatibility Report](#) to help you determine if [Firewall Coexistence](#) mode is required.

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 that are in the Full Enforcement state.
- **Blocked:** The VEN only collects the blocked connection details (source IP, destination IP, protocol and source port and destination port), including all packets that were dropped. This option provides less Illumination detail but also demands fewer system resources from a workload than high detail.

- **Blocked + Allowed:** The VEN collects connection details (source IP, destination IP, protocol and 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.

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 both server and client UDP processes and ports.

On the Services tab for a workload, both UDP client and server processes show up along with their port numbers. 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

The Illumio Core 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 ruleset and ensure that the ruleset 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, and that count 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 \[53\]](#).

Filter a View

You can filter based on workload name, label, hostname, enforcement, etc.

To filter a view, select from the list of categories, such as Labels, and then from the existing elements in 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

1. On the left navigation, go to **Servers & Endpoints > Workloads**.
2. Click the link for a workload you want to change the Enforcement state.
3. Click **Edit**.
4. 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

You can set interfaces from Managed to Ignored in the PCE web console. You can use this option when you want 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, that interface is not included in 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.

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 have the option to ignore it using the PCE Action drop-down menu. That drop-down menu 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

* Interface Name	* IP/CIDR
<input type="checkbox"/> E.g. eth0.public	<input type="checkbox"/> 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 vulnerability statistics it previously calculated and stored in a database.
- If the stored data is stale (4 hours or older), the PCE recalculates the statistics and updates the V-E scores in the UI.
- Toggling the Full Enforcement/Visibility Only options provides a side-by-side comparison of the effect of the different enforcement modes.
- Because the PCE calculates and re-checks for new data periodically, the information in the UI may not immediately reflect the current V-E score.
- API responses include the complete vulnerability data set for the different enforcement modes. V-E data for all modes is pre-processed and stored in a database to eliminate the performance impact 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.

Home > Servers & Endpoints

Workloads

Workloads Container Workloads VENS

[Add](#)
[Remove](#)
[Edit Labels](#)
[Enforcement](#)
[Visibility](#)

Select properties to filter view

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

Connectivity	Full Enforcement V-E Score	Current V-E Score	Enforcement	Visibility	Policy Sync	Ransomware Exposure	Protection Coverage Score	Name
<input type="checkbox"/> Online	0 .	3.1 .	Visibility Only	Blocked + Allowed	✓ Active	Critical	0%	409_vm4.local
<input type="checkbox"/> Online	0 .	3 .	Selective	Blocked + Allowed	✓ Active	Critical	0%	409_vm1.local
<input type="checkbox"/> Online	0 .	0 .	Full	Blocked + Allowed	✓ Active	Protected	82%	409_vm2.local
<input type="checkbox"/> 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 was changed.
- **Current Exposure:** The current exposure score of the workload.

Home > Servers & Endpoints > Workloads


409_vm2.local

Summary Processes Rules Deny Rules Blocked Traffic **Vulnerabilities** Ransomware Protection

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

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:

- Export a CSV or JSON file containing information about the Illumio labels assigned to your workloads. There's also an option to 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.
 - Un-assign 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 simply 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/1a80cceb-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/8714b5c3-5779-48c3-a898-8dd3a854c78	workload-1223	Role26511	App26511	Env26511	Loc26511
3	/orgs/1/workloads/90b451e3-79db-4976-a5af-d9034f4b95b4	workload-1224	Role26511	App26511	Env26511	Loc26511
4	/orgs/1/workloads/34999794-0915-442c-862a-b775c72c322d	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/1a80cc0b-1a23-4a43-9bdc-7294c75d07d5	workload-1478	Role94678	App94678	Env94678	Loc94678
7	/orgs/1/workloads/aa28c53b-63f5-41b0-9058-39e40e01d8b2	workload-1257	role_7173	app_7173	env_7173	loc_7173
8	/orgs/1/workloads/b891d8d0-71bd-467e-86db-3191a0fa30cc	workload-1258	role_7173	app_7173	env_7173	loc_7173
9	/orgs/1/workloads/0471b6c8-ec44-4cfc-ba10-0f6d981c61fb	workload-1457	Role34592	App34592	Env34592	Loc34592
10	/orgs/1/workloads/8020a462-3d20-48ce-92f3-2836d88cead6	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 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 already 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 is 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	Label In Use	Label Group In Use
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, you'll specify the updates you want to make to Workload labels as described in STEP 2: Prepare the CSV file for import. You'll 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 that 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:** Not used for updating labels. This option exports workload data in a JSON file. JSON files can't 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 to create, assign, update, and unassign 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 un-assign 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, un-assigning 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 workloads 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 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 un-assign 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, un-assigning 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** if you want 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 to 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.

Workloads and VENs

The Workloads navigation menu includes Workloads, Container Workloads, and VENs. You can see all your workloads, container workloads, and VENs on separate tabs. You can view

their configuration, do workload or VEN-specific actions, and find the related VENs and workloads.

An idle workload does not program a firewall, therefore the Rules page of an idle workload 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 VENs.

You can select VENs to unpair, refresh, and generate support reports. Container workloads (if any) are displayed under the Container Workloads tab.

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 rebooted</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 in to 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

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 the VEN's visibility level is set 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 source (Connection Responder)
- Destination Total Bytes In (dst_tbi): Number of bytes transferred into 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 the packets of a connection 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 that measure network traffic.

Term	Description
dst_tbi	<p>Destination Total Bytes In</p> <p>In total bytes received till now by the destination over the flows, they are included in this flow summary in the latest sampled interval. This is the same as bytes sent by the source. Present in 'A', 'C', and 'T' flow-summaries. source = client = connection initiator, destination = server = connection responder.</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 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 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 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>

Connec- tion State	Description
A	Active: The connection was still active when the record was posted. Typically observed with long-lived flows on the source and destination side of communication.
T	Timed Out: Flow does not exist anymore. It has timed out. Typically observed on the destination side of communication.
C	Closed: Flow does not exist anymore. It has been closed. Typically observed on the source side of communication.
S	Snapshot: The connection was active at the time VEN sampled the flow. 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, and so on.</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.

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.

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**.

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.

Workload Setup Using PCE Web Console

After you pair workloads, you can view details by clicking a single workload. You can name the workload from the Workload Summary page, write a description, and change the workload's 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 extend rule-writing capabilities to network entities not paired with the PCE and do not have an installed VEN. Adding unmanaged workloads to the PCE allows you to write rules so that workloads paired with the PCE can communicate with those other entities. The policy between workloads with a VEN and unmanaged workloads is enforced using the outbound rules on the workloads where the VEN is running. For unmanaged workloads, enforcement is displayed blank.

For example, when you want to ensure that a network file server belonging to an HRM application is only accessible from the database workloads of the HRM application, you can add unmanaged workloads for the file servers and use label-based rules to enforce the policy. The PCE uses the outbound rules on the database workloads running the VEN to ensure that only the databases labeled HRM are allowed to make 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. See [Create Unmanaged Workload from Blocked Traffic \[54\]](#) for information.

1. In the Servers & Endpoints category, click **Workloads**.
2. Click **Add > Add Unmanaged Workload**.
3. **GENERAL**: In the Add Unmanaged Workload details page, enter a name and description for the unmanaged workload.
4. **LABEL ASSIGNMENT**: In the Label Assignment section, select the labels you want to be applied to the unmanaged workload.
5. **HOST ATTRIBUTES**: In the Host Attributes section, enter all the relevant information about the unmanaged workload, such as its hostname, location, OS Family, Release, and Public IP.
6. **MACHINE AUTHENTICATION**: (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 plan to use this unmanaged workload with the

AdminConnect feature because the unmanaged workload is a laptop running Windows or Linux.

7. VEN TO PCE AUTHENTICATION: When using Kerberos for encryption, type a SPN to authenticate VEN
8. Click **Save**.

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.

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**.

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.

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 an interface that is a loopback interface, but is not in the standard defined IP block that is 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 scope where loopback interfaces are to participate in policy enforcement is defined through the PCE web console.

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.

Blocked Traffic

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

Overview of Blocked Traffic

The Blocked Traffic option is available under each workload.

Blocked traffic alerts provide information such as the port and protocol of the service, as well as the IP address of the Destination, the total number of flows, and the time last detected.

Workloads

Container Workloads

VENs

Add

Remove

Edit Labels

Enforcement

Visibility

Apply Policy

1 Suspended

Refresh

Import

Export

Select properties to filter view

NEW

Customize columns

50 per page

1 - 50 of 658 Total

<

>

<input type="checkbox"/>	Connectivity	V-E Score	Enforcement	Visibility	Policy Sync	Ransomware Exposure	Protection Coverage Score	Name	Labels	Last Applied Policy
<input type="checkbox"/>	<div>Online</div>	440	Selective	Blocked + Allowed	Active	Protected	17%	es-d2	<div>E: DevelopmentCn: US</div> <div>R: Leading Space Label Examples</div>	07/15/2024, 00:23:13
<input type="checkbox"/>	<div>Online</div>	366	Visibility Only	Blocked + Allowed	Active	Critical	0%	redisjob-d34	<div>E: PRDCn: AMER-IV2-CORE</div>	07/15/2024, 00:21:17
<input type="checkbox"/>	<div>Online</div>	351	Visibility Only	Blocked + Allowed	Active	Critical	0%	sofr-d66	<div>E: app1E: StagingCn: loc2R: role_4</div>	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.

Existing connections are reported as static connections during pairing. These connections display as blocked or potentially blocked until new traffic for the connections 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. When the Destination reports the record, the information in the source column is grayed out.

Create Unmanaged Workload from Blocked Traffic

In some cases, your policy might be blocked from the host's IP address that you want to allow to communicate with one of your managed workloads. You can do this by converting the IP address to an unmanaged workload, which enables the PCE to permit it to be used 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 rulesets to allow other managed workloads to communicate with it, or you can later convert it into a managed workload by pairing it.

For more information about unmanaged workloads, see [Workload Setup Using PCE Web Console \[51\]](#).

1. From the PCE web console menu, choose **Servers & Endpoints > Workloads**.
2. Double-click on a workload.

The Workload details page appears.

Figure 1. Workload Detail Page

Home > Servers & Endpoints > Workloads

solr-s35

SummaryProcessesRulesDeny RulesBlocked TrafficVulnerabilitiesRansomware Protection

ConnectionsGroup byR: Roles x

Allow Selected Connections...Edit LabelsResolve Unknown FQDNs14Export

Timestamp: 03/14/2025, 1

Reported Policy Decision

Potentially Blocked

no Rule

Source Labels

Source

Source Port/Process

User

Destination Labels

Destination

Destination Port Process

User

Flows/Bytes

First Detected

Last Detected

1 Source IP

Visibility Only

solr-s35

ldap

root

1 Destination IP

10.21.240.235

389 TCP

Service - 389 TCP

Service - 389 TCP

Service - 389 TCP

+4 more

1 Connection

1 Flow

45.1 KB →

45.1 KB ←

Corporate

03/14/2025, 00:03:02

03/14/2025, 00:03:02

6 Source IPs

Visibility Only

redislob-d2643

lob-d12

solr-s35

1 Destination IP

Visibility Only

56552 UDP

broad udp

6 Connections

6 Flows

Corporate

03/13/2025, 16:36:09

03/13/2025, 23:57:08

3. Manage the details by selecting any tab: Summary, Processes, Rules, Deny Rules, Blocked Traffic, Vulnerabilities, or Ransomware Protection.
- option.

Home > Settings

Reject Connections

GeneralStatic PolicyFirewall CoexistenceReject ConnectionsLoopback InterfacesContainers PolicyIP ForwardingSecure Connect

Edit

Blocked Connection Action

The default blocked connection action is drop. Workloads that match these labels will reject blocked inbound connections.

Scopes

Select properties to filter view

50 per page

Scope

No data to display

Reject Connections

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

1. Select **Settings > Security >** and then the tab **Reject Connections**.

Figure 2. Reject Connections

Security – Reject Connections

General Static Policy Firewall Coexistence **Reject Connections** Secure Connect Containers Policy

Edit

Blocked Connection Action The default blocked connection action is drop. Workloads that match these labels will reject blocked inbound connections.

Scope using Labels and Label Groups Select properties to filter view

Role	Application	Environment	Location
No data to display			

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.

Create Security Policy

This section describes how to create a security policy in the Illumio Core. 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.

Core Services Detector

Core services (DNS, Domain Controller, NTP, and LDAP) are essential to your computing environment and run on one or multiple workloads. The Core Service Detector feature helps you identify these core services and suggests an appropriate label. The Illumio PCE can detect 51 core services. Identifying and labeling these workloads is important because they are centrally connected, and other applications depend on them.

Application owners sometimes don't know enough about the core services or how to identify them. In addition, different teams could be managing core services, and application owners must coordinate with these teams to secure their applications. When you use the Core Services Detector to label and write policies for core services, you can save time on application policies and progress to policy enforcement faster.



NOTE

The Core Services Detector is available only on the leader PCE in the Super-cluster.

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 access 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. Log in to the PCE web console and choose **Settings > Core Services**.
The setting for the Core Services feature appears.
3. Select **Enabled**.
The Core Services menu option will now appear in the PCE web console main menu under Infrastructure, and you can use the Illumio REST API to manage Core Services.

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**: 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 all the time.
- The core services detection for Microsoft Active Directory uses the machine learning (ML) model.

Detection methods can be such as **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 shows all services detected by the detection tool during the last run.

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

2. Click the link for any of the listed core services. The page refreshes and displays 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 with 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 so they become 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 by default is not enabled. You must manually enable scanner detection 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.

Rulesets

You can use rulesets to write policies so the workloads in your application can communicate. A ruleset consists of rules and scopes:

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

Basic versus Scoped Rulesets

You have the option to create basic or scoped rulesets. Choose whether you want to include scopes when creating new rulesets.

When the PCE is configured to create scopeless rulesets, you 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 ruleset that uses scopes.

You might want to create basic rules when you are new to using Illumio Core and creating your first security policy rules, such as creating a simple rule to control SSH traffic for all your workloads. As you become more familiar with Illumio Core or you need to create more complicated rules, you can create scoped rules: intra-scope, extra-scope, and custom iptables rules.

Creating scoped rules allows you to create rulesets and rules that are defined for specific environments, locations, applications (typically larger environments), or other categories you define in flexible label types.

When the PCE is configured to create scopeless rulesets, you can still add a scope to a ruleset after saving the ruleset. Select a rule on the Rulesets & Rules page and click **Add Scope**.

Scopeless rulesets in PCE web console

The following details apply to scopeless rulesets in the PCE web console:

- An option in the Policy Settings page determines whether new rulesets are created with or without scopes. However, the permission every Illumio Core user has to create rulesets is always based on the scopes they can access, even when the PCE is configured to create scopeless rulesets.
Disabling scopes in rulesets does not invalidate the Ruleset Manager or Ruleset Provisioner roles used for user authentication or Role-Based Access Control (RBAC). For more information about these roles, see "[PCE Organization and Users](#)" in PCE Administration Guide .
- When the PCE is configured to create scopeless rules, the Ruleset details page for a ruleset displays a single Rules tab where you add basic rules, including container hosts as Destination.
- When you add a scope to a scopeless ruleset after creating the ruleset, the page refreshes and displays Intra-scope Rules and Extra-scope Rules tabs. If any rules include container hosts for Destinations, those rules are moved to the Extra-scope Rules tab.
- Adding custom iptables rules is not available for scopeless rulesets. To create custom iptables rules, you must add a scope to the ruleset.
- When you remove all scopes from a ruleset, 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 ruleset remain in the Custom iptable Rules tab.

Ruleset Scope

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

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

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 ruleset (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 | Prod | US
- Corp-HQ | HRM | Prod | US → Web | HRM | Prod | US

Single ruleset scopes

Using a single scope in a ruleset narrows the list of workloads that the rules apply to and allows 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 ruleset:

Scope		
App	Environment	Location
HRM	(unspecified)	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 ruleset scopes

Using multiple scopes in a ruleset applies the rules to each scope in isolation and does not allow workload cross-communication.

For example, consider the following ruleset:

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 ruleset:

Scope		
App	Environment	Location
HRM	(unspecified)	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 ruleset:

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 ruleset:

Scope		
App	Environment	Location
(unspecified)	Dev	US
(unspecified)	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 Rulesets

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

Create a Ruleset

You can create a ruleset to write rules that define the allowed communication between workloads in a single group 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 ruleset, or the PCE web console will not be able to display all of the rules.

If you want to create a ruleset with more than 500 rules, split the rules across multiple rulesets or use REST API, where there is no limit on the number of rules you can create per ruleset.

The following task creates a single scope, which means the rules in the ruleset 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 ruleset from scratch.

Create a Ruleset from Scratch

1. Choose **Policies > Add**.
2. In the Add Ruleset dialog, enter the name and description of the ruleset.
3. In **Scope**, select the labels for the ruleset: Application, Environment, Location, or any custom label types you have defined using Flexible Labels.

These labels define the scope of the ruleset, which is its range or boundary. The scope defines the workloads affected by this ruleset 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 Ruleset from a Template

1. Choose **Policies > Add**.
2. To create a ruleset 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, used for network maintenance and troubleshooting.

3. Select one of the templates and click **Next**.

Add a Ruleset 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, and 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 added new members.
 - Click **Provision** to get this addition 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, click **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 Ruleset for Inbound Admin Access

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

RULESET x

- You can edit the name or scope for each ruleset on the rulesets page.
- Scope displays whether the ruleset 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 Ruleset 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 drop-down list.
- You can also remove any of the existing members of the label group.

Add a Ruleset 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 drop-down list.
- Destinations (list of IP addresses) can be edited by removing any of the existing IKP addresses using a trash icon that shows after you double-click on the address. To add 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 Ruleset for Active Directory

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



In the Rules for Active Directory page:

- The Name of the ruleset is editable.
- The scope of the ruleset is editable: add any existing label groups using the dropdown list.
- Intra-scope rules in the ruleset:
 - 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 Ruleset for ICMP

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

RULESET x

- The Name of the ruleset is editable.
- The scope of the ruleset is editable in the following instances:
 - If denoted by  (all), the rule is not editable.
 - If denoted by  (any), rules Destination Services are editable.
 - Once the changes are in, click **Confirm and Save**.

Rules

Rules can allow communication between multiple applications or entities in different scopes or the same scope. 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. The ruleset scope applies to both Sources and Destinations.
- **Extra-scope rules:** Allow communication between groups.
- **Custom Iptables rules:** Allows custom iptables configurations in a ruleset. These rules are managed by the PCE and applied on each managed Linux workload VEN that matches the labels for the scope and receivers.

About Rules

Illumio supports the delegation of rule writing using role-based access control (RBAC). Application administrators can only edit rules where the scope of the ruleset matches the scopes where they have administrator privileges. They cannot create or manage rulesets if the scope includes “All.”

Rule types allow the application administrator to write rules that allow other applications to communicate with the applications that they manage without requiring global administrator privileges. This feature allows users to group rules required for inter-application and intra-application communication for a specific application into one ruleset.

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

From 18.2.0 version on, you can use multiple services or ports and protocols in a rule. This approach helps reduce the number of rules in your PCEs, which helps improve the 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.

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 allow authorized users to write rules that allow communication between providers and consumers 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 provider and the consumer. If you don't specify a Label, “All” is used by default.

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

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 allow authorized users to write rules that allow communication between applications. Specifically, you can write rules that allow providers within a scope to be accessed by consumers that can be in or outside the specified scope. For extra-scope rules, the labels used in the scope must match the labels used by the provider. If you don't specify a label, "All" is used by default.

For example, you can create a rule where all Database workloads with the labels HRM | US | Dev can accept connections on MySQL from all workloads with the label Web, irrespective of other labels.

The MySQL might not belong to the application HRM (for example, the consumers are "Global" and are not restricted by the labels in the scope).



NOTE

If the RBAC user's scope coverage type is "Providers and Consumers," the user cannot select an IP list as the consumer. To select an IP list as a consumer in a rule, the scope coverage type must be "Providers Only."

Custom iptables Rules



NOTE

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

You might have configured iptables directly on your Linux workloads as needed 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 in Illumio Core provide the ability for 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.

To clarify:

- **Iptables** refer to a Linux host configuration before the VEN is installed.
- **Rules** refer to statements written by the PCE to determine permitted traffic, typically by assuming control of iptables and programming the new rules.
- **Iptables rules** refer to iptables that are 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 that receive the rules. 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 prior to the last default rule.

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

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 is logged in the organization events
- The error is displayed in the VEN policy sync status
- The new policy is not used and the last known 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, postrouting	Yes
nat	prerouting, output, postrouting	Yes
filter	input, output, forward	Yes
security	input, output, forward	No

**NOTE**

If the RBAC user's scope coverage type is "Providers and Consumers," the user cannot manage custom iptables rules. To allow access to custom iptables rules, the scope coverage type must be "Providers Only."

Permitted Rule Writing Combinations

The following table explains the valid rule combinations between providers and consumers.

If Provider is	And Service is	Consumer 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

Stateless Rules

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

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

Caveats

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

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

Be aware also of the following when enabling stateless rules:

- Linux traffic does not get logged in the PCE
- Windows traffic gets logged in the PCE if connections are established
- Traffic is allowed in the opposite direction

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

The Illumio Core limits the number of stateless rules to 100, to ensure that both stateful and stateless rules coexist on the host in a way that optimizes system and network performance. If you need more than 100 stateless rules in your Illumio policy, contact your Illumio Professional Services Representative for more 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. Those connections need to be reestablished by the clients. For this reason, Illumio recommends that you use stateless rules for services that use high-frequency short-lived connections, such as DNS and SNMP.

Rule Search

When you have a large number of rules organized in rulesets, you can't easily search for rules across rulesets. Segmentation rule search solves this issue by making it simple to search for specific rules.

For example, when you want to know how many rules there are for SNMP (UDP 161) and you have around 200,000 rules organized across 700 rulesets, it is time-consuming to narrow down that search without using this feature.

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

- Segmentation Rule Search allows you to quickly find rules that apply to a set of providers and consumers.
- Providers and consumers can be represented by a workload, an IP address, or a set of labels.
- Using this feature helps you identify rules that are getting applied to your workloads due to unnecessarily broadly applicable rulesets or human errors.

To search for rules:

1. From the PCE web console menu, choose **Policy > Rulesets & Rules**.
The Rulesets and Rules page appears.
2. Choose the **Rule Search** tab.
3. Search for Active or Draft rules.
4. Perform a Basic or Advanced search of your rules:
 - Basic: Searches all attributes
 - Advanced: Searches by source, destination, or both.

**NOTE**

When you perform an advanced search by workload name, the search results do not display the IP list rules when the iplist contains workload IP addresses because the Illumio Core does not resolve CIDRs and ranges within an IP list.

5. From the Results drop-down list, choose to either have the exact match of the selected search filters to be displayed or a match to any of the selected filters (All Results).
6. Click the Customize columns drop-down list to select the attributes you want to be displayed in the search results.
7. Filter options to further narrow your search.
8. Click **Run**.
9. In the Ruleset column, you can click a ruleset name and make changes to the rules.
- 10 Click **Download** to download the results of your search in JSON format.

Rule Search by Port

The following guidelines and uses cases are provided to clarify how Rule Search works when you search for rules by the port(s) they specify.

General Guidelines

- Single-port searches generally work as expected. See **Row 1** in the Use Case table.
- When searching for a port range, the port ranges in the search and in the rule must match exactly. See **Row 3** in the Use Case table.
- When searching for rules that specify multiple ports, only rules that specify all of the ports are found. See **Row 5** in the Use Case table.

Use Cases: Search for Rules by Port

Row	Use case	Examples		Is the rule found?
		(A) Search specifies port(s)	(B) Rule specifies port(s)	
1	(A) Search for rules that specify only a single port and (B) There's a rule that specifies the same single port	80	80	Yes
2	(A) Search for rules that specify only a single port and (B) There's a rule that specifies a port range that encompasses the searched-for port	80	50-100	No
3	(A) Search for rules that specify a port range and (B) There's a rule that specifies the same port range	50-100	50-100	Yes
4	(A) Search for rules that specify a port range and (B) There's a rule that specifies only a single port within the searched-for range	50-100	80	No
5	(A) Search for rules that specify multiple ports and (B) There's a rule that specifies the same multiple ports	50, 100	50, 100	Yes
6	(A) Search for rules that specify only a single port and (B) There's a rule that specifies multiple ports, including the searched-for port	50	50, 100	Yes
7	(A) Search for rules that specify multiple ports and	50, 80	50, 100	No

Row	Use case	Examples	Examples	Is the rule found?
		(A) Search specifies port(s)	(B) Rule specifies port(s)	
	(B) There's a rule that specifies some, but not all, of the searched-for ports			

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.

Illumio has the following types of rules for application policies:

Override Deny Rules

- This rule type is typically used to deny communication between sources and destinations that might inadvertently be given to allow rules by another administrator. Override Deny rules take precedence over all other types of rules.

Allow Rules

- You can write rules that allow communication between sources and destinations.

Deny Rules

- You can write rules that deny communication between sources and destinations.

Custom IPtables Rules

- You can write rules for Linux workloads.

Policy Check

The Policy Check feature allows you to determine if 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 check can use a network profile to account for rules affecting outbound traffic to non-corporate interfaces on endpoints. Servers cannot have non-corporate interfaces.



NOTE

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

For example, you have created several rulesets 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.

To perform a policy check:

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 is running over TCP or UDP, or just a protocol when the connection is running over GRE or IPIP.
5. In the *Network Profile* field, choose either *Corporate*, *Non-Corporate Networks (Endpoints Only)*, or *Any*.
If an IP address is specified in both *Consumer* and *Provider* fields, the *Network Profile* value must be *Corporate* -- that is, searching within the internal corporate network only.
6. Click **Check Rules**.
If a connection is allowed between the selected two workloads or IP addresses, 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."

About Rules

Rules allow communication between multiple applications or entities in different scopes or the same scope.

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

Rule types allow the application administrator to write rules allowing other applications to communicate with the applications they manage without requiring global administrator privileges. This feature allows users to group rules required for inter-application and intra-application communication for a specific application into one ruleset.

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

You can use multiple services or ports and protocols in a rule. This approach helps reduce the number of rules in your PCEs, which helps improve the 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

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:** Allows custom iptables configurations in a ruleset. These rules are managed by the PCE and applied on 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 allow authorized users to write rules that allow 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 labels HRM | US | Dev.

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 allow authorized users to write rules that allow communication between the scoped application and external entities. Specifically, you can write rules that enable Sources within a scope to be accessed by Destinations 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 restricting communications too soon.

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

MySQL might not belong to 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 microsegmentation and security.

Start with broad rules and refine them over time to maintain both 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 in Illumio Core 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 wrote 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 that receive the rules. 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 is logged in the organization's events.
- The error is displayed in the VEN policy sync status.
- The new policy is not used and 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
<code>raw</code>	prerouting, output	No
<code>mangle</code>	prerouting, input, output, forward, post routing	Yes
<code>nat</code>	prerouting, output, post-routing	Yes
<code>filter</code>	input, output, forward	Yes
<code>security</code>	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.

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:

- [Override Deny Rules \[82\]](#)
- [Allow Rules \[83\]](#)
- [Deny Rules \[83\]](#)

Override Deny Rules



NOTE

Override Deny rules require VEN release 22.3.0 or later.

These rules block all traffic, no matter what other rules are 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 that 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.

Policy Check and Rule Search

This section explains how to use the Policy Check feature and search for rules.

Policy Check

The Policy Check feature allows you to determine if 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 use a network profile to account for rules affecting outbound traffic to non-corporate interfaces on endpoints. Servers cannot have non-corporate interfaces.

**NOTE**

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

For example, you have created several rulesets 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.

To perform a policy check:

1. From the PCE web console menu, choose **Troubleshoot > Policy Check**.
2. In the Source field, type or select a Workload, Container Workload, or IP address.
3. In the Destination field, type or select a Workload, Container 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 the connection 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 both the Destination and Source fields, the Network Profile value must be "Corporate, that is, 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 rulesets when you have many rules organized in rulesets. Segmentation rule search solves this issue by making it simple to search for specific rules.

For example, it is time-consuming to narrow that search without using this feature when you want to know how many rules there are for SNMP (UDP 161) and have around 200,000 rules organized across 700 rulesets.

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.
- Using this feature helps you identify rules that are getting applied to your workloads due to unnecessarily broadly applicable rulesets or human errors.

To search for rules:

1. From the PCE web console menu, choose **Policies**.
The Policies page appears with a tab Rule Search.
2. Choose the **Rule Search** tab.

3. Search for Active or Draft rules.

[Home](#) > [Policy](#)

Policies

ALLOW RULES ONLY

Policies Rule Search

Active Rules ▾

Advanced ▾

Exact Results ▾

↓ Download

4. Perform a Basic or Advanced search of your rules:
 - Basic: Searches all attributes
 - Advanced: Searches by source, destination, or both.



NOTE

When you perform an advanced search by workload name, the search results do not display the IP list rules when the iplist contains workload IP addresses because the Illumio Core does not resolve CIDRs and ranges within an IP list.

5. From the Exact Results drop-down list, choose to either have the exact match of the selected search filters displayed or a match to any of the selected filters (All Results).
6. **Filter by Sources**
You can filter rule search by the following categories: Labels and Label Groups, IP Addresses, IP Lists, Virtual Services, Workloads, and User Groups.
7. **Filter by Destinations and Rule Attributes**
You can filter rule search by these categories: Labels and Label Groups, IP Address, IP Lists, Note, Rule Options, Port and/or Protocol, Port Range, Process Name, Windows Services, Policy Services, Status, Created At, Created By, Virtual Servers, Virtual Services, Workloads, and Ruleset Name.
8. Click **Run**.
9. Click **Download** to download the results of your search in JSON format.

Rule Search by Port

The following guidelines and uses cases are provided to clarify how Rule Search works when you search for rules by the port(s) they specify.

General Guidelines

- Single-port searches generally work as expected. See **Row 1** in the Use Case table.
- When searching for a port range, the port ranges in the search and in the rule must match exactly. See **Row 3** in the Use Case table.
- When searching for rules that specify multiple ports, only rules that specify all of the ports are found. See **Row 5** in the Use Case table.

Use Cases: Search for Rules by Port

Row	Use case	Examples		Is the rule found?
		(A) Search specifies port(s)	(B) Rule specifies port(s)	
1	(A) Search for rules that specify only a single port and (B) There's a rule that specifies the same single port	80	80	Yes
2	(A) Search for rules that specify only a single port and (B) There's a rule that specifies a port range that encompasses the searched-for port	80	50-100	No
3	(A) Search for rules that specify a port range and (B) There's a rule that specifies the same port range	50-100	50-100	Yes
4	(A) Search for rules that specify a port range and (B) There's a rule that specifies only a single port within the searched-for range	50-100	80	No
5	(A) Search for rules that specify multiple ports and (B) There's a rule that specifies the same multiple ports	50, 100	50, 100	Yes
6	(A) Search for rules that specify only a single port and (B) There's a rule that specifies multiple ports, including the searched-for port	50	50, 100	Yes
7	(A) Search for rules that specify multiple ports and	50, 80	50, 100	No

Row	Use case	Examples		Is the rule found?
		(A) Search specifies port(s)	(B) Rule specifies port(s)	
	(B) There's a rule that specifies some, but not all, of the searched-for ports			

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 Source is	And Service is	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.

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.

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

Be aware also of the following when enabling stateless rules:

- Linux traffic does not get logged in the PCE.
- Windows traffic gets logged in the PCE if connections are established.
- Traffic is allowed in the opposite direction.

You will receive an error if you attempt to add any other Destinations.

The Illumio Core number of stateless rules is limited to 100 to ensure that both stateful and stateless rules coexist on the host in a way that optimizes system and network performance. If you need more than 100 stateless rules in your Illumio policy, contact your Illumio Professional Services Representative for more 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 re-establish those connections. For this reason, Illumio recommends using stateless rules for services that use high-frequency, short-lived connections, such as DNS and SNMP.

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

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.
- Rule-Based Labeling assigns labels to workloads but doesn't replace existing labels already assigned to workloads. For example, if 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 preserved and the London Location label is bypassed.

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 has finished, a notification appears wherever you are in the PCE user interface.

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 assigned from the PCE UI: `label_mapping_rules_run.assign_labels`
- You can see the difference between a system job and an assignment from the PCE UI in the `generated_by` field.

It displays either `system` for the system, or the user's e-mail when assigned from the PCE UI.

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 \[89\]](#).

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, and then review the generated list of workloads. See [Find and Review Matching Workloads \[89\]](#).

Step 3: Assign labels to matching workloads

Once the feature finds matching workloads, you can assign the labels you specified in Step 1: Add a Labeling Rule. See [Assign labels to matching workloads \[90\]](#).

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.


1. Identify the workloads you want to label by examining the workloads on the Workloads page and then take note of the attributes you'll need to specify in later steps.
2. Go to **Policy Objects > Labels**.
3. Click the **Labeling Rules** tab.
4. Click **Add**.
5. Specify the matching condition. (For terminology and matching logic, see [How Label Matching Works \[95\]](#).)
 - Add an attribute.
 - Add an operator.
 - Add one or more values.
6. Select one or more labels in the **Label** field.
7. Click **Save**.

Find and Review Matching Workloads

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

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 has 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 \[90\]](#).

Assign labels to matching workloads immediately

Perform these steps to immediately assign labels to the workloads that match your labeling rules.



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 \[91\]](#) 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 \[89\]](#)).
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 \[90\]](#).

Schedule Label Assignments

If you aren't assigning labels immediately as described in the [Assign labels to matching workloads immediately \[90\]](#) procedure, perform these steps to specify when you want to assign labels.

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 re-assigned to previously-labeled workloads because the search ignores VENs that were 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 \[90\]](#). 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 learn more about rule components, see [Terminology \[95\]](#).

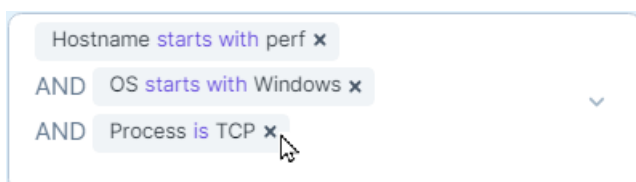
To add a statement to an existing rule:

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.

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**.
5. 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.

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

The screenshot shows the 'Labeling Rules' tab in a web interface. At the top, there are buttons for 'Add', 'Remove', and 'Apply Rules'. Below these is a table with two columns: 'No.' and 'Condition'. The table contains four rules. Rule 3 is selected, and its number '3' is being changed to '1' in an input field. The condition for Rule 3 is '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'. Rule 1 has the condition 'Hostname starts with perf AND OS starts with Windows'. Rule 2 has the condition 'Hostname starts with perf-workload-3390'. Rule 4 has the condition 'OS is Solaris'.

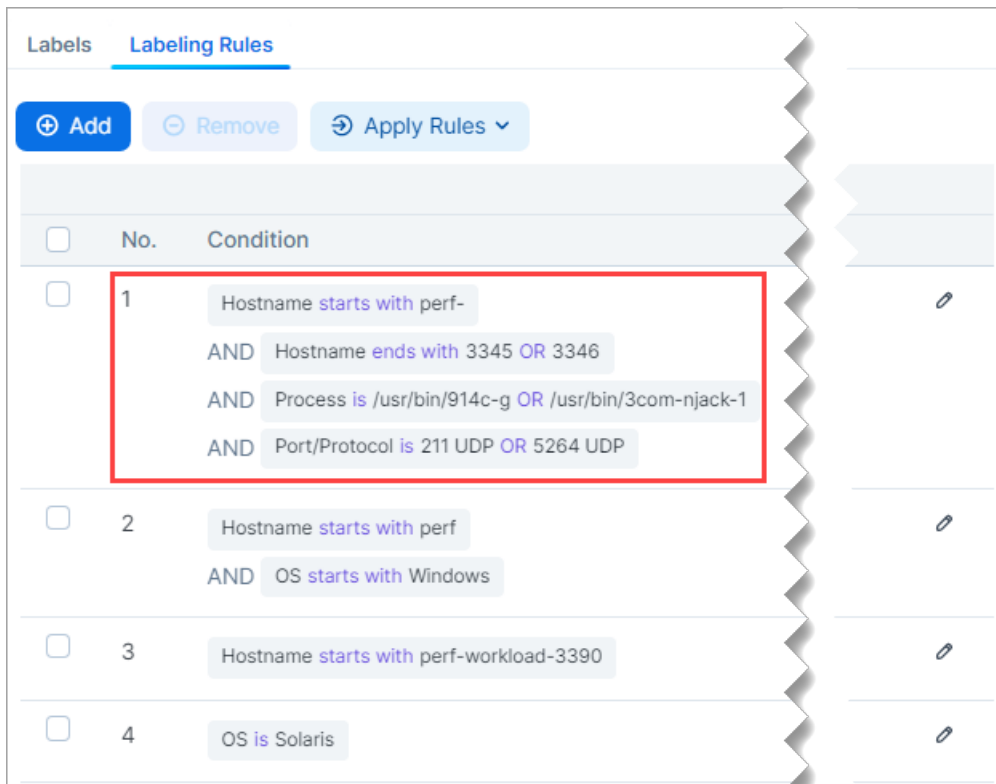
No.	Condition
1	Hostname starts with perf AND OS starts with Windows
2	Hostname starts with perf-workload-3390
1	Hostname starts with perf- x AND Hostname ends with 3345 x OR 3346 x AND Process is /usr/bin/914c-g x OR /usr/bin/3com-njack-1 x AND Port/Protocol is 211 UDP x OR 5264 UDP x
4	OS is Solaris



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.



Remove Labeling Rules

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) that will be assigned to those workloads. This is helpful when you have a large number of rules and workloads.

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. It also 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 of the statements in a rule's condition match a workload's attributes.

Terminology

- **Rule:** Rules consist of a condition and one or more label(s). If a workload matches the rule's condition, it is assigned the corresponding label(s), provided the workload has not already been assigned a label of the same type.
- **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.

Matching Logic

Example: Workload and Rule Evaluation

Workload Attributes and Existing Label(s)	Rules in order	Rule Condition and Label	Match Outcome	Label Assignment	Assigned Labels
<ul style="list-style-type: none"> • Hostname: job-d8cc • OS: Windows • IP range: 10.10.10.30 • Existing label: App88 	Rule 1	<ul style="list-style-type: none"> • Hostname is: job-d8cc • OS: Windows • IP range: 10.10.10.20 ~ 10.10.10.90 • Assign label: Env44 	Match All statements in the rule's condition match the workload's attributes.	Yes The workload doesn't have an existing Environment label, so label Env44 will be assigned.	Assigned by Rule Based Labeling <ul style="list-style-type: none"> • Env44 • Loc22 • Role11 Existing label already assigned <ul style="list-style-type: none"> • App88
	Rule 2	<ul style="list-style-type: none"> • Hostname Contains: d8c • OS: Windows • Assign label: Loc22 	Match	Yes The workload doesn't have an existing Location label, so label Loc22 will be assigned.	
	Rule 3	<ul style="list-style-type: none"> • Hostname Ends with: -d8cc • Assign label: App66, Role11 	Match	1 of 2 The workload already has an Application label, so label App66 will not be assigned. But the workload doesn't already have a Role label, so Role11 will be assigned.	
	Rule 4	<ul style="list-style-type: none"> • Hostname starts with: job • OS: Windows • Assign label: Env99, Loc33, App66 	Match	0 of 3 <ul style="list-style-type: none"> • An Environment label is already assigned by Rule 1, which has precedence. • A Location label is already assigned by Rule 2, which has precedence. • A pre-existing Application label is already assigned. 	
	Rule 5	<ul style="list-style-type: none"> • OS: Linux • Assign label: User-Defined 	No Match	No	

Labeling Rule Examples

This section provides several detailed examples of adding labeling rules.

Keep in mind the following as you add labels:

- The **operator** you select and the particular values you enter in the **Values** field allow you to control the granularity of the labeling rule.
- When you include multiple statements in a condition, Rule-Based Labeling automatically inserts an AND between the statements.
- 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 contains 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 the **Starts with**, **Contains**, or **Ends with** operator 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 IP Address in the **Attribute** field.
2. Select is in 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 IP Address in the **Attribute** field.
2. Select is in 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 contains in the **Operator** field.
 - Enter details in the **Values** field.
2. Specify an operating system:
 - Select **OS** in the **Attribute** field.
 - Select contains 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 either 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**.

FQDN-Based Rules

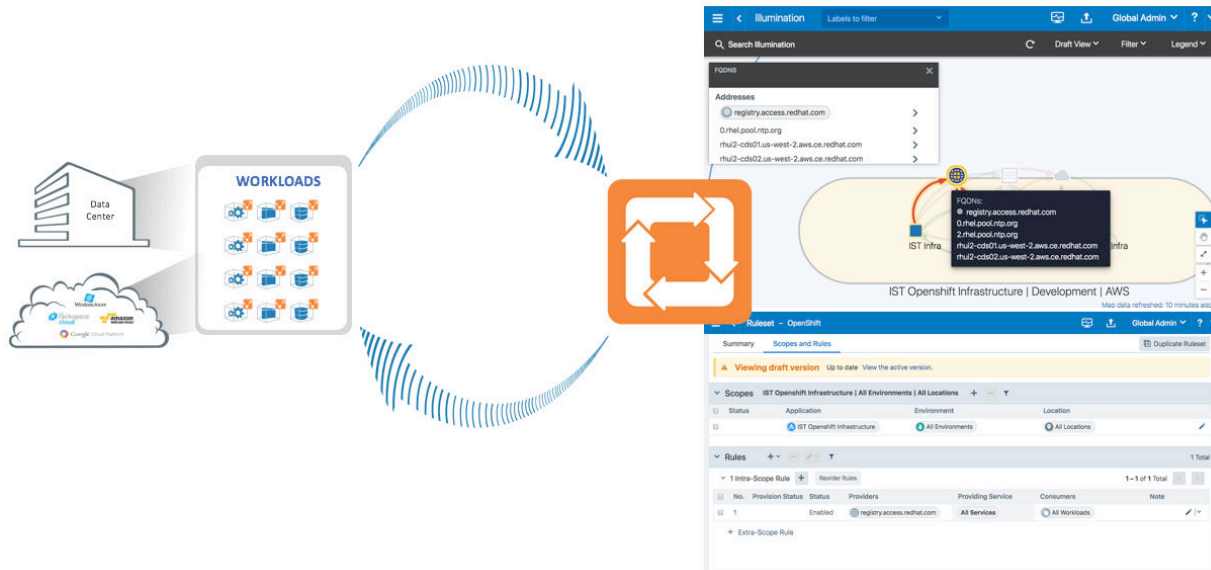
Applications across data centers and cloud environments are responsible for a vast amount of east-west traffic. This traffic is the result of communication between workloads, including bare-metal, virtual machines, and containers. However, many applications might need to communicate with services, such as SaaS, PaaS or external registries. These services are coupled with an IP address but that address might be unknown or the services might only be reachable by a URL because their IP addresses are frequently changing. This situation introduces a challenge to security teams because security policies are based on IP addresses or subnets. Administrators can allow outbound communication to any workload or to a broad range of IP addresses to overcome this challenge; however, this approach opens a security gap. To resolve this challenge, Illumio has added FQDN-based visibility and enforcement to its Illumio Core.

Benefits of FQDN-Based Rules

Implementing FQDN-based rules in the Illumio Core 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 and not just to 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, the Illumio Core 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 performs DNS snooping for both visibility and enforcement purposes. Each time a workload sends out a DNS request, the VEN snoops for a DNS response for that request. The VEN collects data from the DNS response including the CNAMEs, and records and programs it into a DNS cache created by the VEN. The VEN does not generate control plane traffic, for example DNS requests. Additionally, the VEN does DNS-request tracking, which means when the workload receives a DNS response for an FQDN it did not send a request for, the VEN will not add the DNS response data into its cache.

DNS Visibility

One of the core elements of the Illumio Core is visibility into communications between workloads. The VEN periodically reports flow data to the PCE including IP addresses, ports, and protocols. With FQDN-based visibility, the VEN can report outbound communications to FQDNs in addition to IP addresses, ports, and protocols. As the VEN writes flows to its local traffic database, it also checks the VEN DNS cache and maps FQDNs with outbound flow data. When there is a match between the destination IP address in the flow logs and a record in the DNS cache, the VEN adds the FQDN to the outbound flow records. Once the VEN reports flow data to the PCE, the PCE presents the outbound DNS-based traffic flows in Illumination in near real-time and in Explorer for historical data retention.

DNS Enforcement

The Illumio Core allows security teams to write allowlist policies that allow communications across workloads or between workloads and IP addresses. FQDN-based enforcement allows users to control which DNS hostnames or FQDNs that each managed workload can communicate to without the user needing to understand the IP addresses tied to that FQDN. Once an FQDN gets allow-listed by the policy, the PCE sends firewall instructions to the VEN and the VEN creates an FQDN policy table. This policy table tracks the allow-listed FQDN as well as which ports and protocols the workload is allowed to use for outbound communication to the FQDN. The VEN also checks the local DNS cache table for IP listings.

Wildcards

The FQDN-based rules support wildcards such as *.google.com, s3*.aws.amazon.com, and proc1.azure*.com. Wildcards are expanded to zero or more of the characters in [a-z|A-Z|0-9|-]. Wildcards are allowed at the end of the FQDN, for example www.google.*.

Illumio recommends the use of wildcards for the same patterns. This will help reduce rather than increase the number of FQDN-based rules with the same patterns. For better performance, when you write FQDN-based rules, limit the number of rules to around 100 entries.

FQDN-Based Rule Requirements and Limitations

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

- Requires Illumio Core 19.1.0 or later and VEN 19.1.0 or later.
- 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 are 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**

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, you can leverage IP lists to describe IP ranges, groups, and subnets. From the 19.1.0 release on, you can use IP lists to 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 (wildcards are supported).
5. Click **Save**.
6. Provision the changes.



IMPORTANT

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

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; however, it 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"

About Provisioning

Provisioning is the process of applying security policy changes from the Policy Compute Engine (PCE) to the Virtual Enforcement Nodes (VENs) on managed workloads. These changes define the enforcement behavior for workloads using iptables and nftables (Linux) or WFP (Windows), allowing Illumio to maintain a consistent and scalable segmentation policy.

When you provision updates, the PCE recalculates any changes made to rulesets, IP lists, services, label groups, and security settings and then transmits those changes to all VENs installed on your workloads.

When your PCE has changes that need to be provisioned, the orange badge on the Provision button indicates the number of changes that need to be provisioned.

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 rule) and the controlled test deployment.
- 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 as 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 summary of changes.

- Policies can be restored either partially (selectively importing components like rulesets, 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 revert a recent error, and Restore is used to restore a stable baseline.

Policy Versions

Each time you provision changes to policy items (such as rulesets, 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 1000 versions of the policy and automatically removes the older versions for improved 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. To view details about the changes, click one of the items. You can see the changes provisioned in this version 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 try 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 rulesets, IP lists, services, label groups, and security settings. In that case, you need to provision those changes before they can take effect.

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 based on the 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 Provisionable Changes

Any changes you make to policy configuration items (rulesets, 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 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 your internal bug-tracking system or project number for tracking, and the error message they see when they leave the field empty will remind them to do so.

Illumio Core 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**

To access this feature, you must have the correct role and permissions. 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**.**

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/services without impacting entire workloads.

Applying Selective Enforcement

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

When Selective Enforcement is activated:

- Enforced Ports and Services: Active enforcement of security rules; explicitly permitted inbound traffic only.
- Visibility-Only Ports and Services: No active blocking, but communication is monitored and logged.

Workload behavior under Selective Enforcement:

- Enforced Ports: Permits only explicitly allowed inbound traffic according to defined policy rules; 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 only; there is no enforcement.
- Selective Enforcement: Partial enforcement on chosen ports/services.
- Full Enforcement: Complete allowlist enforcement on all ports/services.

This structured approach simplifies secure policy implementation and offers flexibility in managing risk and operational complexity.

Use Cases and Limitations

Basic use cases for Selective enforcement are:

- Incremental Policy Rollout: Enables gradual policy introduction, reducing risks to critical systems.
- Rapid Security Response: Quickly enforce specific critical or vulnerable ports/services policies.

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

Selective Enforcement Mode Limitations

Limitations of Selective Enforcement are grouped as follows:

- Directional Enforcement, where 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.
- Reports OS compatibility every four hours.
- Immediately reports network interface configuration changes.

SecureConnect Rules and Visibility-Only State



NOTE

SecureConnect rules are only applied to workloads where the VEN is in a non-idle enforcement state.

However, unlike other rules, SecureCionnect 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.

Full Enforcement Mode

Full enforcement blocks all non-explicitly allowed traffic, providing the highest security level.

Visibility options mirror Selective Enforcement:

- Off
- Blocked
- Blocked + Allowed
- Enhanced Data Collection – Detailed logs with traffic flow metadata.

Full enforcement is recommended after successful testing and validation of the allowlist model.

Enhanced Data Collection

Enhanced data collection is available only in the Full enforcement mode.

Enhanced Data Collection allows the VEN to log byte counts and connection details for Allowed, Blocked, and Potentially Blocked traffic.



NOTE

SecureConnect rules are only applied to workloads where the VEN is in a non-idle enforcement state.

However, unlike other rules, SecureCionnect 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 idle state.

Policy Exclusions

In Illumio Core 22.2.0 and later releases, the PCE supports including policy exclusions in ruleset scopes and rules. This topic explains what they are, how they are supported in Illumio Core, and how to add them to your security policy.

Policy Exclusions Described

Using policy exclusions in your Illumio Core policy can greatly simplify the rule writing process. Specifically, using a policy exclusion in a ruleset scope or in 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 policy written with policy exclusions can be easier to read and definitely easier to maintain.

Using a policy exclusions gives you a way to state in your security policy that you want a ruleset 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 ruleset scopes and in rules actors, namely consumers and providers.

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 RedHat
- 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

Support for Policy Exclusions

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

Illumio Core Feature	Details
Ruleset scopes and rules	<p>In rulesets 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 three environments: Production, QA, and Development. Adding an exclusion for “All environments except Production” means that the rules apply to all workloads with Environment labels minus the Production label. 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 affect as creating a rule that applies to the QA and Development environments only.</p>
Labels	Fully supports except for the restrictions below. See Requirements and Restrictions [114] .
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 applications and all other applications. You create a label group named “Finance Apps” and use it as a policy exclusion.</p> <p>Using label group exclusions is not supported with individual workloads, virtual services, virtual servers, “All Workloads,” the “Uses virtual service only” option, the “Uses virtual service and workloads” option, and container hosts.</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 search, entering a label name displays both included and excluded labels with that name.</p>
App Groups	<p>App Groups > App Groups List > select a group > Rules tab</p> <p>Rules with policy exclusions appear in the Rules tab.</p>
Policy Check	Rules with policy exclusions appear in the Policy Check page.
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 aren’t supported in Policy Generator.</p>
Access Management	<p>Access management (also know as Role-based Access Control or RBAC) detects policy exclusions when determining user access in the PCE. However, you cannot add a policy exclusion to an RBAC role.</p> <p>Policy exclusions are only supported in rulesets and rules. If a ruleset scope includes a policy exclusion based on labels outside the scopes you have permission for, you cannot view or manage those rulesets and rules.</p> <p>For example, a ruleset includes a policy exclusion of “All environments except Production” and you have permission for the Production environment but do not have permission for the Staging environment, you could not view or manage that ruleset.</p>

Illumio Core Feature	Details
Explorer	<p>When writing rules using Explorer, you can choose rulesets containing policy exclusions. You can edit the rules in the ruleset 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 add an exclusion to a new proposed rule.</p>
PCE web console maps	<p>Policy exclusions are applicable 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 View Rule 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

- For each label dimension, you can specify an included or excluded label, 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

Create a Policy Exclusion

You can add policy exclusions to the scope of a new ruleset and new rules, or edit existing ruleset scopes and rules. This procedure provides the steps to add policy exclusions to the scope of a new ruleset and in new rules.

1. From the PCE web console main menu, choose **Rulesets and Rules > Rulesets**.

The Rulesets page appears.

2. Click **Add**.

The Add Ruleset dialog box appears.

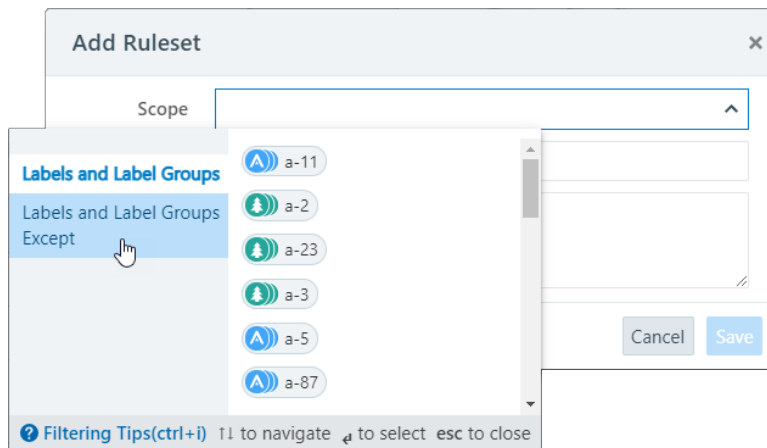


NOTE

The *Scope* field appears in the **Add Ruleset** dialog box only when the PCE is configured to display scopes in rulesets.

If the PCE is configured not to display scopes in rulesets, you can still add a scope with an exclusion after saving the ruleset. From the **Ruleset Actions** menu at the top right corner of the page, select **Add Scope**.

3. To add a policy exclusion to the scope of ruleset, open the Scope drop-down list and select **Labels and Label Groups Except**; then, select labels from the list. When done, click **Save**.



The page refreshes and displays the new ruleset and displays Intra-Scope Rules and Extra-Scope Rules tabs below the scope.

4. Select **Add > Add Intra-Scope Rule** or **Add > Extra-Scope Rule** depending on the type of rule needed.

An empty row for the new rule appears in the page.

5. Configure the values for the row.

To add a policy exclusion for either the Consumer or Provider, or both:

- a. From the *Select...* drop-down list, select the **Advanced Options** checkbox.
A second panel opens displaying your options for adding exclusions.
- b. Select Labels and Label Groups Except and then select labels to exclude from the right-hand list.
- c. When done configuring the rule, click the **Save** icon at the end of the rule row.

**IMPORTANT**

If you unintentionally create a rule that has conflicting elements between added labels or label groups and excluded labels or label groups, the PCE web console will display a warning that the security policy as configured might not apply. Specifically, the rule won't have an actual effect on workloads because the rule conflicts with the ruleset scope or the union of the two is will not have an impactful effect on workloads.

For example, you create a ruleset that has the scope "all but the Production environment" and then you create a rule in the ruleset that specifies the Production environment. This rule ends up having no effect because the rule conflicts with the ruleset scope and the union of the two is nothing.

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 configuring 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 Policy (SP) for workloads so that it can 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 allows the PCE to apply the requisite security policy to your organization to make that happen. SecureConnect reduces the complexity of configuring IPsec encryption and auto-scales per 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 works for connections between Linux workloads, Windows workloads, and 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.

AdminConnect

Relationship-based access control rules often use IP addresses to convey identity. This authentication method can be effective. However, in certain environments, using IP addresses to establish identity is not advisable.

Overview of AdminConnect

When you enforce policy on servers for clients that change their IP addresses frequently, the policy enforcement points (PEPs) continuously need to update security rules for IP address changes. These frequent changes can cause performance and scale challenges, and the ipsets of protected workloads to churn.

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 so that it mistakes a connection as coming from server A.

Illumio designed its AdminConnect (Machine Authentication) feature with these types of environments in mind. Using AdminConnect, you can control access to network resources based on Public Key Infrastructure (PKI) certificates. Because the feature bases identity on cryptographic identity associated with the certificates and not IP addresses, mapping users to IP addresses (common for firewall configuration) is not required.

With AdminConnect, a workload can use the certificates-based identity of a client 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 datacenters, however, must support heterogeneous environments. Consequently, Illumio designed AdminConnect to support Windows and Linux servers and Windows laptop clients.

AdminConnect and Data Encryption

When only AdminConnect is enabled, data traffic does not use ESP encryption. This ensures that data is in cleartext even though it is encapsulated in an ESP packet.

When AdminConnect and SecureConnect are enabled for a rule, the ESP packets are encrypted.

Ease of Deployment

Enabling AdminConnect for identity-based authentication is easy because it is a software solution and it does not require deploying any 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 datacenters.

AdminConnect Prerequisites and Limitations

Prerequisites

You must meet the following prerequisites to use AdminConnect:

- You must configure SecureConnect to use certificate-based authentication because both features rely on the same PKI certificate infrastructure. See the following topics for more information:
 - Configure SecureConnect to Use Certificates. For information, see PCE Administration Guide.
 - Configure certificates for AdminConnect. For information, see PCE Administration Guide.
- AdminConnect must be used with VEN version 17.3 and later.
- AdminConnect supports Linux/Windows IKE v1 (client only) with unmanaged workloads.

Limitations

You cannot enable AdminConnect for the following types of rules:

- Rules that use All services
- Rules with virtual services in providers or consumers
- Rules with IP lists as providers or consumers
- 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).

Enable AdminConnect for a Rule


AdminConnect is supported on workloads in the Visibility Only and Full enforcement . See [AdminConnect Prerequisites and Limitations \[118\]](#) for the list of rule types that do not support AdminConnect.

1. From the PCE web console menu, choose **Rulesets and Rules > Rulesets**.
The Rulesets page appears.
2. Create a new ruleset or open an existing one.
3. In the ruleset, select the **Scopes and Rules** tab.
4. If necessary create an intra-scope or an extra-scope rule. To edit an existing rule, click the edit icon at the end of the row.
5. To enable AdminConnect for the rule, select **Machine Authentication** from the *Providing Service* drop-down list.



NOTE

AdminConnect is displayed as Machine Authentication in the services drop-down lists.

6. Click the **Save** icon  at the end of the row.
The page refreshes and the Providing Service column indicates that AdminConnect is enabled for that Rule.
7. To apply the changes to the applicable workloads, provision the changes.

Secure Laptops with AdminConnect

You can use Illumio to authenticate laptops and grant them access to managed workloads. To manage a laptop with AdminConnect, complete the following tasks:

1. Deploy a PKI certificate on the laptop. See “Certificates for AdminConnect” in PCE Administration Guide.
2. Add the laptop to the PCE by creating an unmanaged workload and assign the appropriate labels to it to be used for rule writing
3. Create rules using those labels to grant access to the managed workloads. See [Enable AdminConnect for a Rule \[119\]](#) for information.
4. Configure IPsec on a laptop.

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

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

1. From the PCE web console menu, choose **Workloads > Add > Add Unmanaged Workload**.

The Workloads - Add Unmanaged Workload page appears.

2. Complete the fields in the General, Labels, Attributes, and Processes sections.
3. In the Machine Authentication ID field, enter all or part of the DN string from the Issuer field of the end entity certificate (CA Subject Name). For example:
CN=win2k12, O=Illumio, OU=Portal, ST=CA, C=US, L=Sunnyvale



TIP

Enter the exact string that you get from the `openssl` command output.

4. Click **Save**.

To configure IPsec on a laptop:

To use the AdminConnect feature with laptops in your organization, you must configure IPsec for these clients.

See the Microsoft Technet article [Netsh Commands for Internet Protocol Security \(IPsec\)](#) for information about using netsh to configure IPsec.

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

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

```
Global Settings:
```

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

StatefulFTP                     Enable
StatefulPPTP                    Enable

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

Categories:
BootTimeRuleCategory            Windows Firewall
```


FirewallRuleCategory	Windows Firewall
StealthRuleCategory	Windows Firewall
ConSecRuleCategory	Windows Firewall

Ok.

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

Rule Name:	telnet

Enabled:	Yes
Profiles:	Domain,Private,Public
Type:	Static
Mode:	Transport
Endpoint1:	Any
Endpoint2:	10.6.3.189/32,10.6.4.35/32,192.168.41.163/32
Port1:	Any
Port2:	23
Protocol:	TCP
Action:	RequireInRequireOut
Auth1:	ComputerKerb,ComputerCert
Auth1CAName:	CN=MACA, O=Company, OU=engineering, S=CA, C=US, L=Sunnyvale, E=user@sample.com
Auth1CertMapping:	No
Auth1ExcludeCAName:	No
Auth1CertType:	Intermediate
Auth1HealthCert:	No
MainModeSecMethods:	ECDHP384-AES256-SHA384
QuickModeSecMethods:	ESP:SHA1-AES256+60min+100256kb
ApplyAuthorization:	No

Ok.

Legal Notice

Copyright © 2025 Illumio 920 De Guigne Drive, Sunnyvale, CA 94085. All rights reserved.

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

Resources

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

Contact Information

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