



Illumio Core[®]

Version 21.1

Security Policy Guide

November 2022

11000-200-21.1

Legal Notices

Copyright © 2020 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.

Product Version

PCE Version: 21.1 (Standard Release)

For the complete list of Illumio Core components compatible with Core PCE, see the Illumio Support portal (login required).

For information on Illumio software support for Standard and LTS releases, see [Versions and Releases](#) on the Illumio Support portal.

Resources

Legal information, see <https://www.illumio.com/legal-information>

Trademarks statements, see <https://www.illumio.com/trademarks>

Patent statements, see <https://www.illumio.com/patents>

License statements, see <https://www.illumio.com/eula>

Open source software utilized by the Illumio Core and their licenses, see [Open Source Licensing Disclosures](#)

Contact Information

To contact Illumio, go to <https://www.illumio.com/contact-us>

To contact the Illumio legal team, email us at legal@illumio.com

To contact the Illumio documentation team, email us at doc-feedback@illumio.com

Contents

Chapter 1 Overview of Security Policy	8
The Illumio Policy Model	8
About the Illumio Policy Model	8
Security Policy Guidelines	9
Rule Creation Approach	9
Policy Refinement	11
Enforcement States	11
Understanding Segmentation Rulesets and Rules	12
Overview of Policy Objects	13
Types of Illumio Policy	14
Adaptive Policy	14
Static Policy	15
Staged Policy	20
Chapter 2 Security Policy Objects	24
Labels and Label Groups	24
Label Types	24
System Default “All” for Labels	25
Filtering Labels and Label Groups	26
Create a Label	26
Label Workloads	27
Edit Labels for Multiple Workloads	27
Label Groups	28
Create a Label Group	29
Use a Label Group in a Scope	29
Use a Label Group in a Rule	30
Services	31
Service Types	31
Windows Process-based Rules	32
IGMP Services	33
ICMP Services	34
Filter the Services List	35
Services in a Rule	36
Create a Service	36
Virtual Services	38

Overview of Virtual Services	38
How Virtual Services Work	38
Virtual Services in Rule Writing	39
Advanced Configuration for Virtual Services	41
Filter the Virtual Services List	42
Add a Virtual Service	43
Bind a Virtual Service to a Workload	44
IP Lists	45
Overview of IP Lists	45
Example of IP List Usage	45
Create an IP List	46
IP List Exclusions	46
Filter IP Lists	47
Load Balancers and Virtual Servers	47
Load Balancers	47
Configure Load Balancers	49
About Virtual Servers	50
Virtual Server Members and Labels	50
Configure Virtual Servers	52
Adaptive User Segmentation	54
Overview of Adaptive User Segmentation	54
Add Active Directory User Groups	55
User Group-Based Rules for AUS	55
Export Reports	56
Overview of Export Reports	56
Generate an Export Report	56
Chapter 3 Workloads	58
Workloads in the PCE	58
Overview of Workload Attributes	58
Workload Summary	59
Workload Processes	60
Workload Enforcement States	61
Visibility Level	62
Workload Rules	63
Workloads Blocked Traffic	63
Filter the Workloads List	63

Enforce a Workload Policy State	65
Set Workload Interfaces to Ignored	65
Workloads and VENS	66
Container Workloads	68
Workload Setup Using PCE Web Console	70
About Creating Managed Workloads by Installing VENS	70
Unmanaged Workloads	71
Add an Unmanaged Workload	71
VEN Administration on Workloads	72
VEN Suspension	73
Blocked Traffic	74
Overview of Blocked Traffic	75
Filter Blocked Traffic	77
Create Unmanaged Workload from Blocked Traffic	78
Reject Connections	79
Chapter 4 Create Security Policy	81
<hr/>	
Segmentation Templates	81
Overview of Segmentation Templates	82
Catalog Retrieved from Support Portal	82
Features of Segmentation Templates	84
Segmentation Template Prerequisites and Limitations	85
About Editing Segmentation Templates	86
Install a Segmentation Template	87
Upload a Segmentation Template	88
Update a Segmentation Template	89
Uninstall a Segmentation Template	90
Policy Generator	90
Overview of Policy Generator	90
Policy Generator Prerequisites and Limitations	92
About Granularity Levels for Rules	92
Ways to Access Policy Generator	93
Create Intra-scope Rules with Policy Generator	94
Create Extra-scope Rules with Policy Generator	97
Create Rules Using IP Lists with Policy Generator	99
Segment Multiple App Groups with Policy Generator	102
Segmentation Rulesets	104

Segmentation Ruleset Scope	104
Multiple Ruleset Scopes	106
Combine Labels in Scopes and Rules	106
Segmentation Ruleset Status	108
Filter the Segmentation Rulesets List	108
Create a Segmentation Ruleset	109
Create a Ruleset with Multiple Scopes	110
Duplicate a Segmentation Ruleset	111
Rules	112
About Rules	112
Enforcement Modes for Rules	113
Intra-scope Rules	114
Extra-scope Rules	114
Custom iptables Rules	115
Permitted Rule Writing Combinations	118
Stateless Rules	118
Segmentation Rule Search	119
Policy Check	120
Rule Writing	122
Create an Intra-Scope Rule	122
Create an Extra-Scope Rule	123
Create a Custom iptables Rule	124
Create a Selective Enforcement Rule	126
Write Multicast Rules	128
Create Service While Creating Rule	129
Tips for Managing Rules	130
Add a Note to a Rule	130
Duplicate a Rule	131
Reverse a Rule	131
Reorder Rules	132
FQDN-Based Rules	133
Benefits of FQDN-Based Rules	133
Features of FQDN-Based Rules	134
FQDN-Based Rule Requirements and Limitations	135
FQDN Visibility	136
Create Policy Objects for FQDNs	136
Write Policies to Allowlist FQDNs	141

Provisioning	143
Items that Require Provisioning	143
Provision All or Selected Items	144
Dependencies for Partial Provisioning	144
Active vs Draft Versions	145
Provisioning Progress Indicator	146
Policy Versions	148
Provision Changes	149
Revert Provisionable Changes	150
Restore Policy	150
Provisioning Note Setting	152
Chapter 5 Secure Workload Connections	153
SecureConnect	153
Our Solution	154
Use Cases	154
Features of SecureConnect	154
SecureConnect Prerequisites, Limitations, Caveats	156
Certificate Setup on Workloads	158
Enable SecureConnect for a Rule	159
AdminConnect	160
Overview of AdminConnect	160
Features of AdminConnect	161
AdminConnect Prerequisites and Limitations	161
Enable AdminConnect for a Rule	162
Secure Laptops with AdminConnect	163

Overview of Security Policy

This chapter contains the following topics:

The Illumio Policy Model	8
Types of Illumio Policy	14

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.

The Illumio Policy Model

Illumio gives you the option to manage your security policies by using either adaptive or static policy. Choosing how to implement security policy is possible because of the Illumio policy model.

About the Illumio Policy Model

The Illumio security policy for securing workloads differs from traditional network securities 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 clear and 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 segmentation rulesets and in the providers and consumers components of rules, which allows the workloads in their organization to communicate with each other.

Together, labeling workloads and creating the corresponding segmentation 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.

See the following related topics:

- [Labels and Label Groups](#) for a description of each label type
- [Workloads in the PCE](#) for information about how workloads use labels and for the steps to assign labels to workloads
- [Segmentation Ruleset Scope](#) and [Rules](#) for information about using labels in segmentation ruleset scopes and rules

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. Create rules following the process below to prevent communication loss and optimize rule coverage.
2. Refine your initial policy to strengthen it by narrowing overly broad access.
3. Use the Visibility Only enforcement to verify and enact your policy.

Rule Creation Approach

For optimal configuration, Illumio recommends using the following approach when creating rules:

1. **Core service policies:** Core services are among the first services that are detected by the PCE and can include DNS, DHCP, backup traffic, and performance management. Use Illumination to explore your network environment and determine what unmanaged workloads are required. Because policy for core services tends to be very granular, rules for these services should be created before using Policy Generator. When core services are defined for a larger entity (like a data-center or an environment), it is best to organize these “global” rules into a dedicated segmentation ruleset. This approach removes all the core services from the application-specific rules and helps keep the policy statements clear and

easily understandable. Additionally, when RBAC is used for user segmentation, the core services are often managed by different users than the application-level rules, so separating them allows separate RBAC control.

Since core infrastructure services, such as NFS are usually available to all workloads, use an IP list-based policy so that those servers accept connections from a CIDR block.

2. **Application ringfencing:** Unless your network has specific compliance or regulatory requirements it must follow, implementing application ringfencing is a good option for your initial policy, because it covers the majority of your traffic. It limits the risk of intrusion to a single application and provides the greatest impact for the least amount of work.
3. **Environment-specific “Outbound Allow-All”:** After creating rules for core services and for ringfencing applications, adding a rule to allow all outbound traffic within an environment (for example, Prod or Dev) usually provides rule coverage for most of the remaining traffic. This approach allows managed workloads to initiate communication with any traffic that is not currently managed by Illumio Core.
4. **Use Policy Generator to build rules for App-to-App traffic:** You can use Policy Generator to write application-specific rules. This approach establishes a base policy for initial enforcement that can be iteratively improved to be more granular, based on security needs.
5. **Managing IP lists:** As a default initial policy, Illumio recommends writing rules for traffic from applications to IP lists using an “All Services” rule. This approach permits traffic from the application group to that subnet on all services.
6. **Managing internet traffic:** In Illumination, any traffic not destined for a workload, unmanaged workload, or IP list is displayed as destined for the cloud icon above the application (“Internet”). Because most of this traffic will likely be destined for a single IP address, manage the traffic in the following ways:
 - a. **Create unmanaged workloads:** Use the hostname to create an unmanaged workload for the destination IP address, then write a rule to block or allow traffic from the application to the unmanaged workload. When dynamic or high-order ports are used, use “All Services” in your rule.
 - b. **Use Illumination to write a workload-specific rule:** Select the traffic flow in Illumination, then click **Add a Rule** and select **All Services**.

- c. **Use Policy Generator to write a default “All Services” rule to each destination:** This option quickly establishes a policy to turn traffic lines green in Illumination, but does not provide the visibility of having unmanaged workloads to represent each external destination.

Policy Refinement

Creating rules using the guidelines above establishes a stable policy where most of the traffic lines in Illumination will be green. It can be adjusted as new traffic is observed, but the recommendations above provide an initial basic policy. This policy can and should be refined incrementally, as it might initially be overly broad. The following guidelines provide the greatest benefits when refining your policy:

1. **Identify critical applications:** Determine your network’s top five or top ten applications and refine the rules for them. Repeat as needed to continuously improve the policy.
2. **Investigate inter-application flows:** In Illumination, inspect the traffic between applications to see how it can be improved. When the traffic is destined for a single port, find out if the port is fixed or dynamic. If it is fixed, update “All Services” to the specified port. If it is dynamic with known ranges, create a new service with that range to narrow the service definition in the rule.
3. **Investigate traffic to IP lists:** Inspect the top-volume flows to IP lists. Optimally, try to reduce IP lists by replacing them with unmanaged workloads, then reduce the service definition to only the required ports.
4. **Identify high-volume unmanaged workloads or applications:** When a single unmanaged workload has many connections to managed applications, it might be a database cluster or critical server. Ideally, you should try to reduce usage of “All Services” where possible to a more narrow definition.
5. **Restrict access for scopes and services to unmanaged workloads:** Unmanaged workloads do not have an installed VEN, so access to these types of workloads should be restricted where possible. Determine if you can use scopes to restrict access to specific roles within applications and review access to services to determine if any use of “All Services” can be reduced to a more specific port range.

Enforcement States

After creating a segmentation ruleset, you can preview the effects in Illumination using the Draft View. This view shows you the changes that will be enacted by your policy when it is enforced.

- **Visibility Only:** After refining your initial policy, most of the traffic lines in Illumination should be green. No traffic will be blocked and you can check your policy's accuracy. Any new traffic will be displayed as a red line, which can be addressed using the strategies discussed in [Rule Creation Approach](#) and [Policy Refinement](#).
- **Selective Enforcement:** Enables you to protect applications or processes on workloads while other services and ports function as if the workloads are in the Visibility Only enforcement state. By using selective enforcement, you can gradually expand the enforcement of policy on your workloads. Using the selective enforcement state is useful for temporarily enforcing security for specific ports in case a vulnerability is detected and action must be taken quickly. Using the selective enforcement state enables security enforcement before you are able to create complete allowlists of what traffic is allowed to reach your workloads.
- **Full Enforcement:** It is useful to move workloads to the Full Enforcement state in stages. This action can be done by workload, by application, by environment, or by datacenter. Start with less critical applications or workloads, stabilize them, then move on to more sensitive systems. This approach minimizes issues to a smaller number of affected workloads.

Understanding Segmentation Rulesets and Rules

Rules are an integral component of the Illumio security policy. A set of rules is known as a “segmentation ruleset” and it specifies the allowed traffic in your network. Create the rules using labels that identify your workloads. See [Labels and Label Groups](#) for more information.


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



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.

For example, in the following diagram:



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 segmentation ruleset for each of your workloads. Use labels to identify your workloads and use scopes to apply the segmentation rulesets to multiple workloads at once. For more information, see [Labels and Label Groups](#) and [Segmentation Rulesets](#)



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 use 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 PCE contains the following policy objects that help you write your security policy:

- **Segmentation Templates:** Prepackaged, tested security policies that provide all the segmentation rules needed for common enterprise applications.
- **Labels and Label Groups:** Group similar labels together and use the label groups in rule writing.
- **Services** Allow you to define or discover existing services on your workloads. When a workload is paired with the PCE (has a VEN installed), it is scanned for any running processes, which are then displayed in the Services list.
- **Virtual Services:** Allow you to label processes or services on workloads. Virtual services can either be used directly in rules or the labels applied to virtual services can be used to write rules.
- **IP Lists:** Create IP lists (allowlists) so you can define IP addresses, IP ranges, and CIDR blocks that should be allowed access to your applications.
- **Load Balancers and Virtual Servers:** Add F5 Load Balancer configurations to the PCE so you can write policy for workloads whose traffic is managed by load balancers.
- **Pairing Profiles:** Configurations that allows you to apply certain properties to workloads as they 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 User Segmentation](#).

Types of Illumio Policy

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

Adaptive Policy

Without adaptive security, enterprises face an overwhelming number of firewall rules, manual changes required to policies, and the possibility of errors leading to outages or serious vulnerabilities and breaches. Adaptive security automatically accounts for moves, scale, and changes to the applications and infrastructure that are typical of modern datacenters.

Because Illumio bases workload security on a policy model, it enables adaptive security that continuously adjusts to changes in the environment and to changed workload relationships. When a change occurs, the PCE responds dynamically by re-computing the OS-level firewall rules for the impacted workloads. The PCE alerts the VENs of the

new OS-level firewall rules. The VENs request the new rules and apply them immediately.

The Illumio Core dynamically adapts and updates security policy when events, such as the following ones, occur in the managed environment.

- Workloads are added to or removed from your environment.
- Workloads change their IP addresses.
- Managed workloads come online and go offline.
- The labels on workloads change.

The PCE does not require Illumio users or automated processes to provision these changes for the PCE to re-compute the OS-level firewall rules for the impacted workloads and transmit them to the VENs.

See the following related topics:

- [Pairing in the *VEN Installation and Upgrade Guide*](#) for information about adding workloads to your environment
- [IP Lists](#) for information about using them in security policies
- [Provisioning](#) for information about provisioning, which is a manual process
- [Staged Policy](#) for information about how provisioning differs from adaptive policy

Static Policy

For the large majority 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. See [Adaptive Policy](#) to learn how Illumio provides adaptive security.

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](#) 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 segmentation ruleset in the

PCE and provision the change, or a change occurs in your environment, such as a workload changes 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](#) for information about how the Illumio Core uses static policy and stages OS-level firewall updates rather than apply 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 security updates to its transaction processing application must be explicitly controlled by its security team. 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 up and running and minimize risk.

- The central security team sets the security policy to 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 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 must be applied during maintenance windows. For their Illumio-managed workloads, they configure a static policy that has the following labels: Role: Database, Applications: All, Environment: Production, Locations: All.

A spike in customer demand occurs and their 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-compute 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 they 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 and they 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 role 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 gGroups before adding them to 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 applies the policy it receives from the PCE immediately.
 - 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. The VEN comes back online, connects to the PCE, and receives updated OS-level firewall rules. The VEN applies the updated rules to the workload even though it is configured to use static policy.



NOTE:

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

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

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

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

Caveats

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

The Illumio Core is designed to ensure that your workloads are protected by the latest versions of your security policy across your environment. Users provision policy changes or the PCE responds dynamically to changes in the environment. 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 you configure 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 policy across your managed environment and cause communication disruptions between workloads.

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

Illumio recommends that you keep 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. See [Adaptive Policy](#) to learn how Illumio provides adaptive security.

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

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 labels for the role, application, environment, and location. Any workloads within the scope of the specified labels will use static policy. You can add multiple scopes. Overlap between the scopes does not affect how workloads use static policy.

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

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

To add static policy:

1. From the PCE web console menu, choose **Settings > Security**.
2. Choose **Edit > Manage Policy Update**.

The page refreshes with the settings to configure Static as the Policy Update Mode.

3. Click **Add**.

A dialog box appears in which you set the scope of the static policy.

4. Select labels to select workloads for static policy.

5. Click **OK**.

The static policy appears in the list.

6. Click **Provision** from the PCE web console toolbar.

Staged Policy

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 and they 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](#) for information.



TIP:

The orange badge on the Provision button (top toolbar) indicates the number of changes you need to provision. See [Provision Changes](#) for more information.

In addition to segmentation rulesets and rules, you must provision changes to the Illumio policy objects, such as services, IP lists, and label groups. To make security policies easier to maintain and update, Illumio supports including re-usable policy objects in intra- and extra-scope rules. 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 segmentation 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 on those workloads 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 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 quickly determine which ones have staged OS-level firewall rules.

- **Active (Syncing):** The PCE is in the process of sending new policy to the VEN. Typically, this process takes only a few seconds.



NOTE:

Workloads configured for adaptive policy and static policy can appear in the active (syncing) state while the PCE is sending 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 the *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 that 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:**

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

Apply Staged Policy

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

1. From the PCE web console menu, choose **Workloads**.

The Workloads page appears.

2. (Optional) Use the Workload property filter in the following ways:
 - To find all your workloads that are configured to use static policy, choose **Policy Update Mode > Static Workloads**.
 - To find workloads that have staged policy that needs to be applied, choose **Policy Sync > Staged Workloads**.
3. To apply staged policy to specific workloads, select the workloads and choose **Apply Policy > Update Selected Workloads**.

**NOTE:**

- Choosing **Update Selected Workloads** only applies staged policy. It does not provision pending policy changes for workloads that are configured to use adaptive policy even when you selected them.
- If you applied policy to a subset of workloads with staged policy, the remaining workloads will continue to use the older policy.
- The **Apply Policy** button is enabled only when you have workloads with staged policy waiting to be applied.

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

**NOTE:**

If you filtered workloads by label and chose **Update All Workloads**, the PCE applies the staged updates to all the workloads matching that label scope and not just the workloads appearing in the PCE web console page. See [Filter the Workloads List](#) for information about the Label filter.

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

5. Click **OK**.

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

Security Policy Objects

This chapter contains the following topics:


Labels and Label Groups	24
Services	31
Virtual Services	38
IP Lists	45
Load Balancers and Virtual Servers	47
Adaptive User Segmentation	54
Export Reports	56


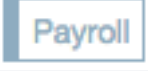


This section describes the policy objects that you can use to write security policies.

Labels and Label Groups

The Illumio Core policy model is a label-based system, which means that the rules you write don't require the use of 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 	This label type allows you to describe the “role” (or function) of a workload. 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 segmentation rulesets for different applications.

Label	Description
	 <p>NOTE: The Role label cannot be used to define the scope.</p>
Application 	<p>This label type allows you 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 role.</p>
Environment 	<p>This label type allows you to describe a workload based upon 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 upon its location. For example, Germany, US, Europe, Asia. Or, Rack #3, Rack #4, Rack #5; or datacenter AWS-east1, AWS-east2, and so on.</p>

Additional Dimensions

A given workload cannot have more than one label per type. It’s possible to allow a workload that used a service or services or across boundaries to communicate; for example, if a server is playing multiple roles, such as a database server used by two different applications, Illumio recommends that you create 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

To find the label or labels groups you are looking for, you can use the property filter at the top of the **Policy Objects > Labels** or **Label Groups** pages.

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



Create a Label

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

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), and their location (Location label). After a workload is labeled, you can write rules using the labels you have applied to the workload.

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. (See "Pairing Profiles" in *VEN Installation and Upgrade Guide*.)
- Add labels to the workload on the [Workload Summary](#) page. In the PCE web console, select **Workloads and VENs > Workloads** from the left navigation menu. Select a workload, and in the details panel click **Edit** to select any or all of the four 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. Previously, if you wanted to delete a Label and it was in use by a Virtual Server, you would not know if it was in use or not. In the Illumio Core 20.1.0 release and higher, on the Labels page, the "In use by" column includes Virtual Servers. The Labels' summary page also displays the "In Use By Virtual Servers Yes/No" field.



NOTE:

Keep in mind that label changes do not require provisioning, so mass label changes can potentially have a major impact on your segmentation rule-sets, rules, and overall security policy.

1. From the PCE web console menu, choose **Workloads and VENs > Workloads**.
2. From the left side of the Workloads list, select the workloads you want to change labels for.
3. From the top of the Workloads list, click **Edit Labels**.

A dialog box appears asking if you are sure you want to edit labels for multiple workloads.

4. Click **OK**.

5. In the Edit Labels dialog box, you can add or remove labels assigned to the selected workloads. 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 exact 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 of that type. 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.
6. 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 segmentation 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 residing in datacenters 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 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)
- When it is currently in use by a segmentation ruleset, label group, and static policy
- Last modified date and time
- User who last modified the label group

Name	Type	In Use By R...	In Use By L...	In Use By St...	Last Modified On	Last Modified By
	Role				02/09/2018, 17:33:01	umio.com
	Role				02/09/2018, 17:38:45	umio.com
	Role				02/09/2018, 17:30:43	umio.com

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 segmentation rulesets.

Create a Label Group

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

1. From the PCE web console menu, choose **Policy Objects > Label Groups**.
2. On the Label Groups page, click **Add**.
3. In the Add Label Group page, choose the label type and enter a name for the label.
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.

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 segmentation ruleset (scope + rule):

	App	Env	Loc
Scope	HRM	Non-Prod	US
	Providers	Services	Consumers
Rule	DB	MySQL	DB

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 segmentation ruleset (scope + rule):

	App	Env	Loc
Scope	HRM	All	US
	Providers	Services	Consumers
Rule	Non-Prod DB	MySQL	Non-Prod DB

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 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 qsidtype "WebSAT Plant."` You can write rules with unrestricted service IDs (SIDs). When there is a restricted SID, you should write rules without the SID. Including the SID might cause a disconnect between the Reported view and Draft view because the Draft view assumes that the SID works.

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, 500 UDP, 1000-2000 TCP. 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. The Windows Process Path and the Windows Service Name must be surrounded by quotation marks:
 - **Port/Protocol, Windows Process Path, and Windows Service Name**
443 TCP "c:\windows\myprocess.exe" "myservice"
 - **Port/Protocol and Windows Process Path**
443 TCP "c:\windows\myprocess.exe"
 - **Port/Protocol and Windows Service Name**
443 TCP "myservice"
 - **Windows Port/Protocol**
"22" "514 UDP" "ICMP"

- **Windows Process Path**
"c:\windows\myprocess.exe"
- **Windows Service Name**
"myservice"

Windows Process-based Rules

Rules to Allow System Created Processes

Rules can be created to allow all system-initiated processes in Windows. This approach allows all traffic related to drivers and other operating system modules. You can create a service of type Windows—process or service-based—with 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 definition for the service you are adding.
 - To add a service definition, from the *Operating System* drop-down, select either **All Operating Systems:Port Based** or **Windows 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 "22" or "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 "22" or "512 TCP", process, "C:\windows\myprocess.exe," and Windows service, "myprocess".
 - To remove a service definition, from the *Operating System* drop-down, select either **All Operating Systems:Port Based** or **Windows Process/Service-Based**:
 - a. Click the check box next to the *Port and/or Protocol*. You may select a single or multiple entries.
 - b. Click **Remove**.
4. Click **Save**.

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

IGMP can be added as a service and used 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 that use the IGMP protocol by using the Illumio Core REST API. See [Services](#) in the *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	Protocol and code	All ICMP traffic is type 3 (Destination Unreachable) allowed, regardless of the code used in the rule
3/6 ICMP	Protocol name, type, and code	Only type 3 and code 6 ICMP traffic is allowed

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 the *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 systems 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

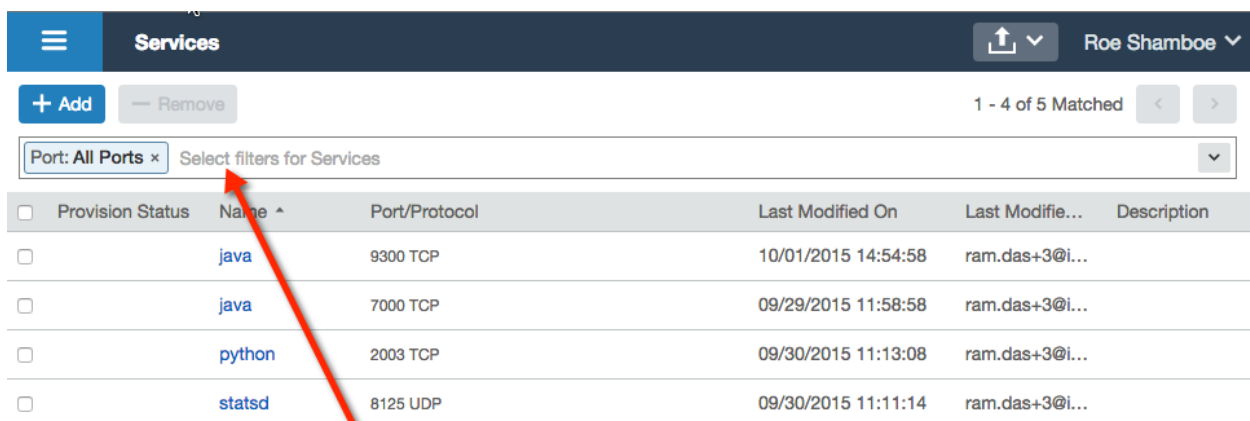
Upgrading from Illumio Core Version 17.1

If the ICMP Echo option was allowed in your PCE prior to upgrade, the PCE automatically adds and provisions a rule during the upgrade to allow ICMP Echo on all workloads. During the upgrade, the PCE checks the current organization settings and takes the following actions:

1. Creates a new service named “ICMP.”
2. Creates a new rule in the default segmentation ruleset to allow outbound ICMP for all workloads.
3. When the ICMP Echo setting was enabled, creates a new service named “ICMP ECHO” to allow echo requests and a new rule to allow all “ICMP ECHO” on all workloads.
4. Adds the rules to the active version of the policy.

Filter the Services List

You can filter the Services list using the property filter at the top of the list. You can filter list by entering a service name, description, port, protocol, and provision status (draft or active).

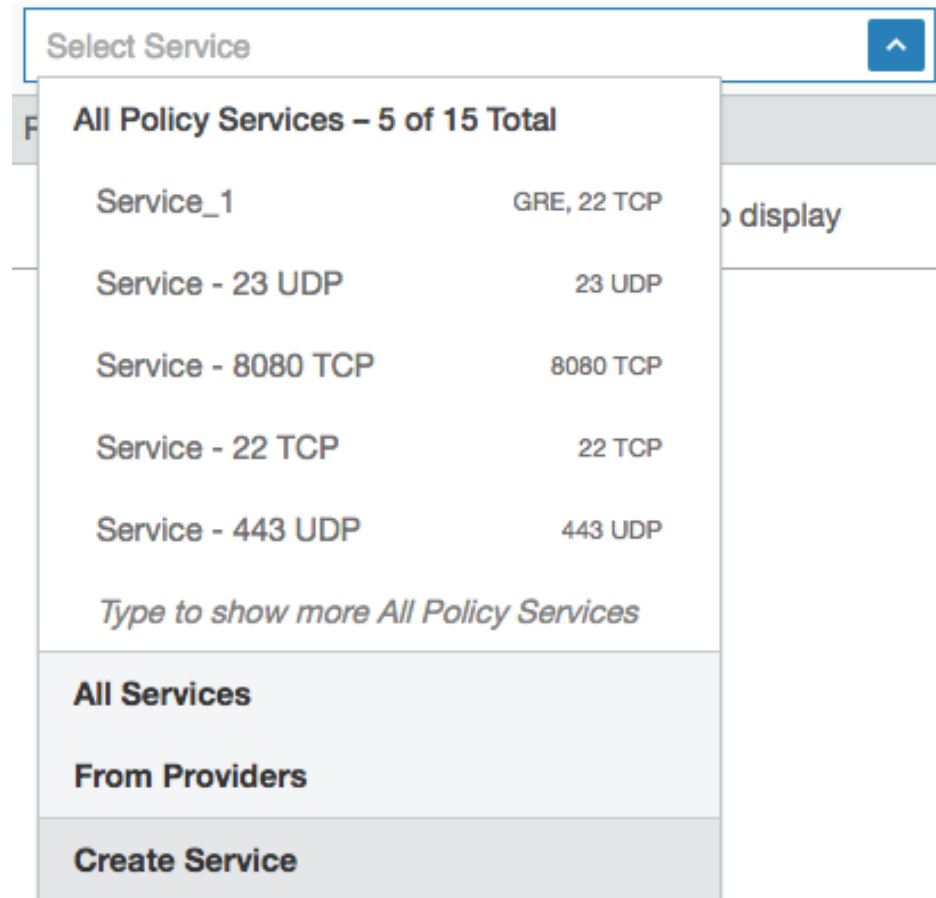


Provision Status	Name ^	Port/Protocol	Last Modified On	Last Modified By	Description
<input type="checkbox"/>	java	9300 TCP	10/01/2015 14:54:58	ram.das+3@i...	
<input type="checkbox"/>	java	7000 TCP	09/29/2015 11:58:58	ram.das+3@i...	
<input type="checkbox"/>	python	2003 TCP	09/30/2015 11:13:08	ram.das+3@i...	
<input type="checkbox"/>	statsd	8125 UDP	09/30/2015 11:11:14	ram.das+3@i...	

Type Service properties to filter list

Services in a Rule

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



Create a Service

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](#).

To create a service from the **Services** page:

1. From the PCE web console menu, choose **Policy Objects > Services**.
2. 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. If you want to enter a range, separate the port numbers by a hyphen. You can also copy and paste lists of services here from another source.
6. When the service uses any UDP ports, enter them as well.
7. Click **Save**.

To create a service from the Ruleset page:

To make rule writing easier, you can create a new service in a segmentation ruleset as you are writing rules.



NOTE:

The service is not associated with the segmentation ruleset.

1. Create an extra-scope or an intra-scope rule. (See [Rule Writing](#).)
2. In the *Select Service* field, choose **Create Service** at the end of the list.

The screenshot shows a dropdown menu titled "Select Service" with a blue arrow icon on the right. The menu is open, displaying a list of services under the heading "All Policy Services - 5 of 15 Total". The services listed are:

Service Name	Protocol
Service_1	GRE, 22 TCP
Service - 23 UDP	23 UDP
Service - 8080 TCP	8080 TCP
Service - 22 TCP	22 TCP
Service - 443 UDP	443 UDP

Below the list, there is a prompt: "Type to show more All Policy Services". At the bottom of the dropdown, there are three options: "All Services", "From Providers", and "Create Service".

Virtual Services

Virtual services (previously known as bound services) allow you to label processes or services on workloads. Virtual services can either 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 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 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 on, Illumination, Policy Generator, and Explorer support virtual services. You have to assign labels to a virtual service in order 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. See [Provisioning](#) and [Bind a Virtual Service to a Workload](#) for more information.

How Virtual Services Work

For example, if a single workload is running both an Apache Tomcat and Apache HTTP server, supporting an HRM and ERP application respectively, you can create a virtual service for each service and then label one service as belonging to an HRM

application and one belonging 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 would 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, you need to 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 port or port range.



NOTE:

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

When you write a rule in a segmentation 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.

When you want to write a rule applicable for all virtual services labeled “Database,” you would 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 be impacted by the above rule. To include them, you need an additional rule listing the specific service applicable.

When you want to use a virtual service as a provider, select “Uses Virtual Services” or “Uses Virtual Services and Workloads” from the Provider drop-down list.

When you select a specific service, then the rule applies only to workloads that have 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 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 (rules to internal bridge are ignored by Windows 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 FORWARD chain; instead, 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 rules 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 device that is 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, this rules programs 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 that 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)

Virtual Services					
<input type="button" value="+ Add"/> <input type="button" value="↑ Provision"/> <input type="button" value="↶ Revert"/> <input type="button" value="— Remove"/>					
Select properties to filter view					
<input type="checkbox"/>	Provision Status	Name	Service	Role	Application
<input type="checkbox"/>	MODIFICATION PENDING	New VS	Diverse Service		
<input type="checkbox"/>	MODIFICATION PENDING	ERP - DB	All Services		
<input type="checkbox"/>	ADDITION PENDING	BsTest43	All Services		
<input type="checkbox"/>		hserv-add-rule 2733070	All Services		

Add a Virtual Service

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

Then, you need to 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 and one of the rules is stateless and the other stateful, the stateless rule takes precedence.



NOTE:

A virtual service must be provisioned before binding it to a workload. See [Provisioning](#) for more information.

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. From the *Service* drop-down list, select the service 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, 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 block that will be used for programming the virtual service rules.
8. Click **Save**.

The virtual service is created and labeled; next, provision it and bind it to a workload. See [Provisioning](#) for more information.



NOTE:

[SecureConnect](#) is not supported for virtual services.

Bind a Virtual Service to a Workload

When you bind a virtual service to a workload, it enables the PCE to program rules to the VEN on the workload the virtual service is bound to.

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



NOTE:

Before binding a virtual service to a workload, the virtual service must be provisioned. See [Provisioning](#) for more information.

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. To allow this virtual service to use a different port than the one specified, select the *Override ports* checkbox.



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 to use.
8. Click **Save**.

IP Lists

IP lists allow you to define allowlists of trusted IP address, IP address ranges, or CIDR blocks that you want to allow into your datacenter and to be able to access workloads and applications in your network.

Overview of IP Lists

After you define an IP list, you can use it in segmentation rulesets to create rules for workload traffic flows. When you provision the segmentation 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 programmed on one side of the connection only. IP lists can be used as a provider or a consumer.



NOTE:

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

Example of IP List Usage

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

	App	Env	Loc
Scope	HRM	Prod	US
	Providers	Services	Consumers
Rule	DB	SSH	Corp-HQ

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

This segmentation ruleset:

	App	Env	Loc
Scope	All	Prod	All
	Providers	Services	Consumers
Rule	Corp-HQ	SSH	DB

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

This segmentation ruleset:

	App	Env	Loc
Scope	All	Prod	All
	Providers	Services	Consumers
Rule	Any	Any	Any

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. 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. Click **Save**.

IP List Exclusions

In IP lists, you can exclude certain 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 certain workloads. Or, you might want to 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 that are 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, 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:

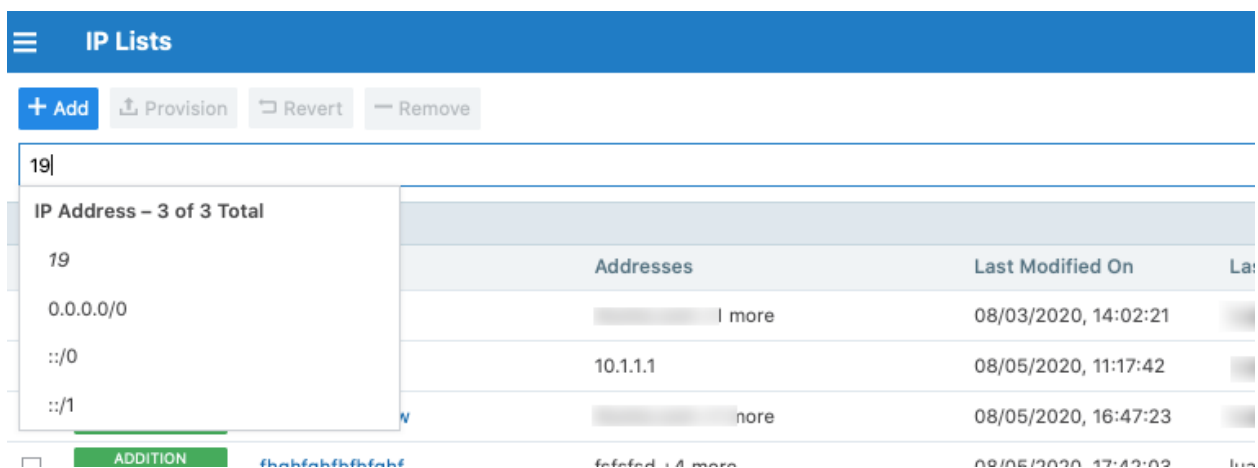
```
10.0.0.0/8
```

```
!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 property filter at the top of the list. You can filter list by entering an IP list name, description, IP address, FQDN, and provision status (draft or active).



Addresses	Last Modified On	La
0.0.0.0/0	08/03/2020, 14:02:21	
10.1.1.1	08/05/2020, 11:17:42	
	08/05/2020, 16:47:23	

Load Balancers and Virtual Servers

Load Balancers

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



IMPORTANT:

From the Illumio Core 19.3.0 release onwards, the Network Function Controller (NFC) is no longer supported. The F5 interface has been moved from the PCE in to the Network Enforcement Node (NEN).

Since the NFC has been discontinued, you need the NEN to interface with Load Balancers. The NEN has both switch and load balancer capabilities.

By applying labels to your load balancer's virtual servers, you can write rules that allows client workloads in front of the load balancer to communicate with the virtual IP address of the load balancer's virtual servers. By adding labels to the pool members behind a virtual sever, you can allow communication from the load balancer to the members of the pool. The source for this communication is determined by the load balancer. The Illumio Core programs policies on the load balancer to enforce security policy. The PCE uses the load balancer's REST APIs to read and write security policies to configure security rules.

The PCE supports configuration of two load balancer units if they are configured in Active/Standby or Active/Active modes. The PCE needs to be configured with the user name and password of an administrative user who has read-write access to all configurations on the load balancer.

The PCE configures new objects on the load balancer and does not alter any existing configurations. When an Illumio-created object in the load balancer configuration is modified, the PCE detects it as tampering and modifies the configuration back to the intended state so that the correct security policy is enforced.

The Illumio Core dynamically adjusts policies on the load balancer based on application and topology changes in the datacenter so that the Illumio Core can enforce consistent security policy on load balancers across the datacenter and cloud environments, as well as show the application traffic in Illumination. The Illumio Core keeps track of the policy it programmed and reconfigures policy if it was altered on the load balancer manually or by other means.

The Illumio Core makes use of the following constructs on load balancers:

- **LTM:** iRules on LTM provide capability to restrict application access. When LTM is used as enforcement mechanism, the Illumio Core uses virtual-server based iRules and Datagroup Lists.
- **AFM:** AFM provides stateful firewalling on BIG-IP. When AFM is used as an enforcement mechanism, the Illumio Core uses Network Firewall policies in the virtual server section and address-lists in the network firewall.
- **AVI:** The Illumio Core uses the Network Security Policy rules to program AVI Vantage.

**NOTE:**

Configuring two F5 units in Active/Standby mode is supported. However, F5 vCMP or F5 clustering is not supported.

F5 BIG-IP Requirements

The Illumio Core uses its REST API to program F5 load balancers, which means that F5 needs to be running a software version that supports REST-API. The requirements include:

- BIG-IP 11.5.x or higher
- Utilize SNAT or Auto-map mode

AVI Vantage Requirements

- AVI Vantage 18.2.3 or higher

Configure Load Balancers

You can add a load balancer using the PCE web console. However, before you add a load balancer, you need to pair the NEN with the load balancer functionality enabled with the PCE.

**NOTE:**

A load balancer does not need to be provisioned to work. However, the virtual servers you associate with this load balancer do need to be provisioned.

1. From the PCE web console menu, choose **Infrastructure > Load Balancers**.
2. Click **Add**.
3. Specify a name for the load balancer and provide a description.
4. From *Device Type*, select appropriate load balancer device type.
5. From number of devices, select **(1) Standard** or **(2) HA Pair**.

The load balancer details are displayed.

6. Specify the following settings to enable the PCE to connect to the load balancer:
 - Management IP address or FQDN of the load balancer
 - Port on which to connect
 - Username
 - Password

7. Select **verify TLS** to verify the trust of the TLS certificate provided by the load balancer before connecting to it.
8. Click **Save**.

About Virtual Servers

Virtual servers in the Illumio Core contain two parts:

- A virtual IP address (VIP) and port through which the service is exposed
- Local IP address(es) used to communicate with backend servers (pool members).

A virtual server is similar to a workload. It can be assigned labels and has IP addresses, but does not report traffic to the Illumio Core. Each virtual server has only one VIP. The local IP addresses are used as a source IP address for connections to the pool members (backend servers) when the virtual server is operating in SNAT mode or Auto mode. These IP addresses are likely to be shared by multiple virtual servers on the server load balancer.

A virtual server is identified by a set of labels. The consumers and providers for a virtual server can be assigned different labels, which could place them in the same group or a different group in Illumination. See **Groups in Illumination** in the *Visualization Guide* for information.

Providers are allowed to have an incomplete label set (for example, only a Location label), so the providers can be in all groups within the specified location. As a result, a single virtual server can have providers in any group or in any number of groups in Illumination.

Virtual Server Members and Labels

The Illumio Core allows you to write rules to allow communication with workloads managed by a load balancer using labels.

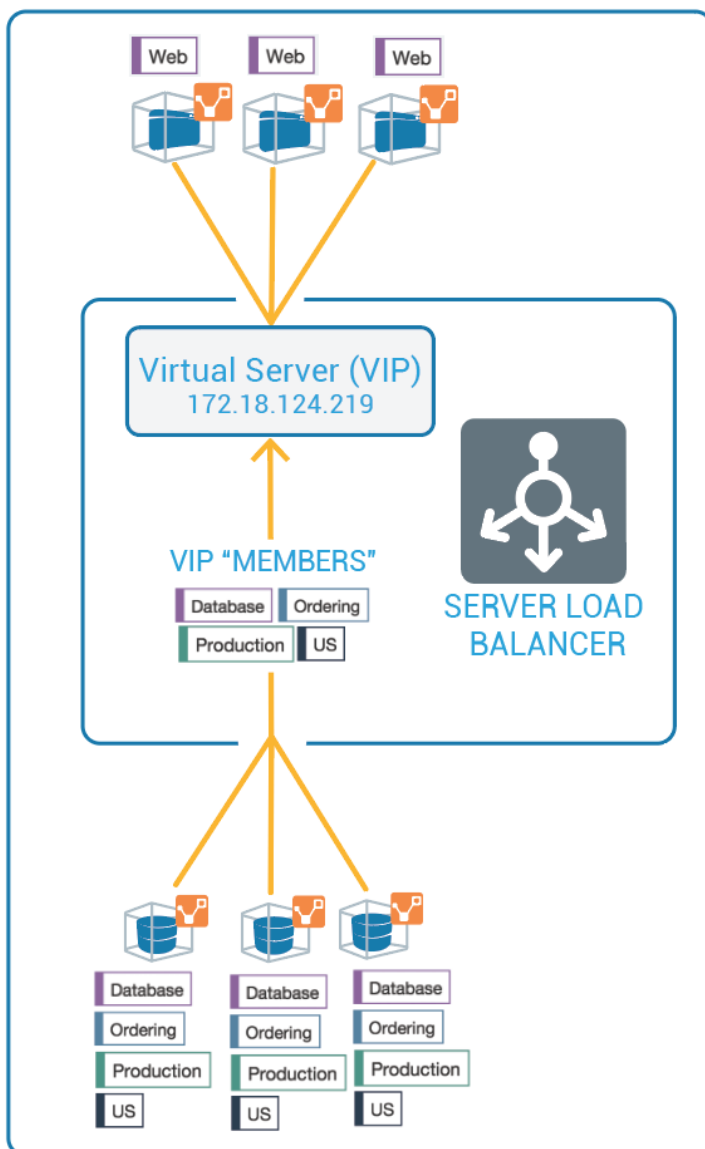
Virtual Server Members

When you configure load balancers in the PCE, it connects to the load balancer using the Illumio Core REST API. The PCE reads all the load balancer virtual servers configurations and populates the Discovered Virtual Servers tab of a load balancer's details page. Any virtual servers associated with the load balancer can be converted to a managed virtual server for use with the PCE. When you configure the virtual server in the PCE web console, you can apply labels to the virtual servers. After con-

figuring a virtual server, you can write a rule that allows other clients to communicate with it.

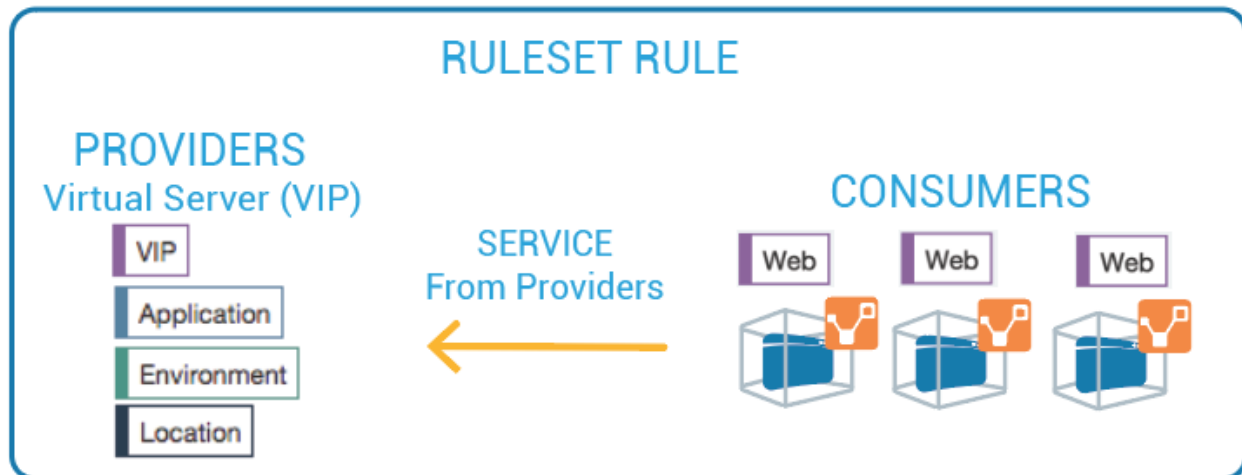
The members behind a virtual server are specified by configuring a set of labels in the virtual server configuration. A set of four Illumio labels can be applied on the Virtual Server Members tab, which is used to match the same set of labels on workloads in the virtual server's pool. If any of the workloads in the virtual server pool share the same set of four labels specified under the Virtual Server Members tab, then any rule you write with the virtual server also applies to the workload members.

In this diagram, you can see how the workloads that belong to the virtual server pool have the same labels specified on the Virtual Server Members tab:



Ruleset Rule for Virtual Server

This diagram illustrates the rule you can write after you label a virtual server and its members:



Configure Virtual Servers

After adding a load balancer to the PCE, you can manage its virtual servers. For each virtual server, you can add a complete set of the four Illumio labels: Role, Application, Environment, and Location. Adding labels to the virtual server enables you to add it in a rule.

You add the four Illumio labels to the Virtual Server's Members tab. When the labels specified under Virtual Server Members match labels of workloads in the virtual server pool, any rule you write with the virtual server are applied to the workload members.

Configuring a load balancer's virtual servers consists of these three settings:

- **Enforced or Not Enforced:** When you select Enforced, any rules you write using the labels associated with the virtual servers and any of its members are enacted. Selecting Not Enforced disables the labels and any policy written that affects the virtual server or its members is disabled.
- **Service:** Select the service to use for the rules that allow access to the virtual server. For example, HTTPD 80 TCP.
- **Labels:** You must apply one each of the four Illumio labels to the virtual server: Role, Application, Environment, and Location. Assigning labels enables the virtual server to be used in rules.



NOTE:

Virtual servers are considered a security policy item, so any changes to a virtual server configuration must be provisioned before any of those changes take effect and become active.

Virtual Server Limitations

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



NOTE:

Before any virtual server configuration can go into effect, you need to provision your changes. See **Provisioning** in the *Security Policy Guide* for information.

Filter the Virtual Server List

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

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

The screenshot shows the 'Virtual Servers' management page. At the top, there are action buttons: Unmanage, Enforce, Stop Enforcement, Provision, and Revert. Below these is a search bar labeled 'Name:'. A filter menu is open on the left, showing options like 'Name - 1 of 1 Total', 'Role Labels', 'Application Labels', 'Environment Labels', 'Location Labels', 'Virtual Server Mode', 'VIP', and 'VIP Port Number'. The main table displays one virtual server with the following details:

Name	Mode	VIP & Port	Management State	SLB	Role	Application
virtual_ip_1	SNAT	192.168.1.1 80 TCP	Unassociated		Mail	Application123 45

Configure a Load Balancer's Virtual Servers

1. From the PCE web console menu, choose **Infrastructure > Load Balancers**.
2. Select the load balancer for which you want to configure virtual servers.
3. Select the **Virtual Servers** tab.
4. Select one of the load balancer's virtual servers and click **Manage**.
5. Select one of the virtual servers and click **Edit**.
6. Enter a name and description for the virtual server.
7. To enable the virtual server's policy, select **Enforced**.
8. Select a service to associate with the virtual server. The service selected enables that service to be used in rules you write for this virtual server.
9. Select one each of the four labels to assign to the virtual server.
10. Click **Save**.
11. Before any virtual servers can go into effect, they must be provisioned.

Adaptive User Segmentation

Illumio's Adaptive User Segmentation (AUS) allows you to leverage Microsoft Active Directory User Groups to control access to computing resources in your organization. With this feature, you can create user groups in the PCE that map directly to your Active Directory Groups.

Overview of Adaptive User Segmentation

You can then write rules with these groups so that you can control outbound access on specific workloads—such as a VDI desktop—based on the group membership of the user logged in to that workload.

For example, you might want to allow only employees in the Sales user group to access the ERP application, but not users in HR. You might want to allow HR users to only access HR applications, but not all internal resources.

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

This type of rule is represented by an icon, where the VDI desktop and AD User Group are added as the consumers of a segmentation ruleset, and entities that these user groups are allowed to access are added as providers.

Add Active Directory User Groups

1. From the PCE web console menu, choose **Policy Objects > User Groups**.
2. In the User Groups page, click **Add**.
3. In the Add User Group page, enter a name, system identifier (SID), and description for the Active Directory Group.
4. Click **Save**.

The new Active Directory Group appears in the User Groups list. You can now use the user group in a segmentation ruleset to control access to specific workloads.

**NOTE:**

A maximum of 100 User Groups can be displayed.

User Group-Based Rules for AUS

1. From the PCE web console menu, choose **Rulesets and Rules > Segmentation Rulesets**.
2. In the Segmentation Rulesets list, click **Add**.
3. Enter a name for the segmentation ruleset.
4. Select an Application, Environment, and Location label to define the segmentation ruleset scope.
5. Click **Save**.

In the *Rules* section, you can start writing identity-based rules.

6. If necessary, expand the *Intra-Scope Rule* section.
7. In the *Consumers* drop-down list, select the user group that you want to provide access to the other workload.
8. From the *Providers* drop-down list, select the workloads or labels that you want to provide access to by a user group.
9. In the *Services* drop-down list, select the service that you want the user groups to be able to access on the providing workloads.
10. Click the **Save** icon at the end of the row.
11. To add additional rules to the segmentation ruleset, Click the **Add (+)** icon.

To enact these changes on the workloads this segmentation ruleset affects, provision your changes. See [Provision Changes](#) for more information.

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
- Segmentation 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 > Export Reports**.
2. Click **New Report**.
3. From the *Containing All* drop-down list, select the object for which you want to generate the report.
4. Select the format, JSON or CSV.

New Report

Containing All Workloads v

Formatted as JSON CSV

File name Workloads,JSON,2019-06-06_15-30-32

Cancel
Generate

Workloads ^

- IP Lists
- Selective Enforcement**
- Services
- Rulesets
- Labels
- Label Groups
- Pairing Profiles
- Virtual Servers
- Virtual Services
- ✓ Workloads
- VEN

5. Click **Generate**.

Workloads

This chapter contains the following topics:

Workloads in the PCE	58
Workloads and VENS	66
Workload Setup Using PCE Web Console	70
VEN Administration on Workloads	72
Blocked Traffic	74

This section describes workload attributes, it's enforcements, and how to create managed and unmanaged workloads.

Workloads in the PCE

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

Overview of Workload Attributes

Workloads have the following attributes:

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

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 States](#)
- The visibility the VEN uses
- The dates when the policy was revised and last applied
- The workload's VEN connectivity status; see [VEN-to-PCE Communication](#) in the *VEN Administration Guide*
- The workload's VEN policy sync status; see [VEN Policy Sync](#) in the *VEN Administration Guide*
- Any labels applied to the workload
- Workload system attributes (such as VEN version number, hostname, and uptime)

☰ 🏠 **Workload – m1.acme.com**

Summary
Processes
Rules
Blocked Traffic

✎ Edit
🗺 App Group Map

General

Name	m1.acme.com
Description	
Enforcement	Selective Segmentation Rules are enforced only for selected inbound services when workload is within scope of a Selective Enforcement Rule AuthService 22 TCP 888 TCP
Visibility	Blocked + Allowed VEN logs connection information for allowed, blocked and potentially blocked traffic
VEN	m1.acme.com
Connectivity	● Online
Policy Sync	✓ Active
Policy Last Received	09/29/2020 at 09:24:04
Policy Last Applied	09/29/2020 at 09:24:04

Labels

Role	
Application	IAM
Environment	Demo
Location	Amazon

Attributes

VEN Version	20.2.0-6956-dev
Hostname	m1.acme.com
Location	Amazon EC2 (US West), Oregon, USA
OS	ubuntu-x86_64-xenial
Release	4.4.0-1107-aws #118-Ubuntu SMP Sun May 3 23:28:51 UTC 2020 (Ubuntu 16.04.3 LTS)
Machine	i-06045b2e9f239cf61
Uptime	106 Days, 2 Minutes
Heartbeat Last Received	09/29/2020, 09:27:04
Public IP Address	34.217.45.24
Interfaces	docker0: 172.17.0.1/16

Workload Processes

In the Workload Processes tab of the Workload detail page, you can view the processes currently running on the workload.

For each process running on the workload, the following information is listed:

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

To organize the listed items, you can select the column headings to sort the processes by that attribute. For example, when you click the Protocol column heading, the processes are grouped by protocol so that all processes using UDP are listed together and all processes using TCP are listed together.

**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 datacenter and write policies to allow traffic from needed workloads.

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 for installing and activating VENs on workloads without changing the workloads’ firewalls. In the Idle state, the VEN on the workload does not take control of the workload’s iptables (Linux) or Windows firewall (Windows), but uses workload network analysis to provides relevant details about the workload to the PCE, such as the workload’s IP address, operating system, and traffic flows. This snapshot is taken every 10 minutes.

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

Selective Enforcement

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

Full Enforcement

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

Visibility Level

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

- **Off:** The VEN does not collect any information about traffic connections. This option provides no Illumination detail and demands the least amount of system resources from a workload.

This property is only available for workloads 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 Rules

In the Rules tab of the Workload detail page, you can view the rules that are currently applied to the workload. The Illumio Core has two types of rules:

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

To apply rules to a workload, create a segmentation ruleset and then make sure that the ruleset and workloads share the same labels. See [Create a Segmentation Ruleset](#) for information.



NOTE:

The workload rules are listed against individual IP addresses in an ipset. The PCE places a limit on 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](#).

Filter the Workloads List

You can filter by one or any combination of workload labels and properties. For example, you can use the Workload filters to see only workloads that have the Role label named Web, that are running Linux Ubuntu, that are in the Build policy state, with a VEN policy sync status of Active, that have a specific IP address, and that have “asset” in the hostname.

- Use the filter at the top of the Workloads and VENs page to do a label-based search. For example, you can filter the list to view all workloads that have the

Application label “Application12345.”

The screenshot shows the 'Workloads and VENS - Workloads' page. At the top, there are tabs for 'Workloads', 'Container Workloads', and 'VENS'. Below the tabs are buttons for '+ Add', '- Remove', 'Edit Labels', 'Enforcement', and 'Visibility'. A search bar is present with the text 'Name:'. The main table has columns: 'Name', 'Role', 'Application', 'Environment', 'Location', and 'Last Applied Policy'. A row is visible with 'solr-s41' as the name, 'solr-s' as the role, 'Application12345' as the application, 'Staging' as the environment, and 'London' as the location. A dropdown menu is open on the left, showing a list of labels: 'Name - 5 of 547 Total', 'db-d70', 'kafka-s1', 'test', 'redis-d1', 'solr-d66', 'Type to show more Names', 'Role Labels', 'Application Labels', 'Environment Labels', 'Location Labels', 'IP Address', 'Description', and 'Hostname'. 'Application Labels' is selected.

- You can filter workloads based on their properties, such as workload name, IP address, description, hostname, OS family, VEN connectivity, and when a policy was last applied to or received by the workload, and when was the last heartbeat received.

Click the **Refresh** button to refresh the content of the page with the latest information without clearing the filters or the results.

Use a Wildcard to Filter Workloads

To help sort and organize large numbers of workloads, the Workloads filter supports a wildcard character for the Name and Hostname properties.

To filter the list of workloads on the Workloads page, select either the Name or Hostname property from the drop-down list and enter the search terms using the asterisk (*) character as a wildcard. The asterisk can represent any number of characters.

For example, you can enter “db-*auto” using the *Name* property to find workloads with names that include “db,” “-auto,” and any number of characters in between (for example, “db-prod-auto,” “db-dev-auto,” or “db-12-auto”).

At least one non-wildcard character must be included before or after the wildcard character. An error message is displayed when you include only the wildcard character in the search field.



NOTE:

The auto-complete feature is disabled when the wildcard character is used.

Use Clone Alerts to Filter Workloads

Workloads can be filtered according to whether a clone has been detected. If there are any workloads that are in the clone detected state, a red banner (similar to workloads in suspension) is displayed at the top of the workload list page.

The VEN communicates with the PCE using HTTPS over Transport Layer Security (TLS). Additionally, a clone token is generated. When an agent token is mistakenly or maliciously reused on another workload, the clone token is used to detect the condition and disambiguate the hosts. The clone token is periodically rotated. The agent token is never rotated.

To filter by clone alerts:

1. In the PCE web console, display the Workloads List page.
2. Look for an alert banner indicating some workloads are in "clone detected" state. This banner will appear only if, for example, you pair one or more VENs and then clone the VEN(s).
3. Click the filter link on the banner. The list now shows only the "clone detected" workloads.
4. Click on one of the "clone detected" workloads. An alert is displayed on the detail page for that workload.
5. If you stop, unpair, or repair the cloned VEN, you can come back and see that the messages and alerts are removed from the Workloads List page.

Enforce a Workload Policy State

1. From the PCE web console menu, choose **Workloads and VENs > Workloads**.
2. Select the workload that you want to change the Enforcement state..
3. From the *Enforcement* drop-down list, select **Idle**, **Visibility Only**, **Selective**, or **Full** depending on how you want to allow or block traffic connections.

A dialog box appears directing you to confirm your change.

4. Click **OK**.

Set Workload Interfaces to Ignored

You can set interfaces from being 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 be 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. From the PCE web console menu, choose **Workloads and VENS > 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.

Network Interfaces

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	<input checked="" type="checkbox"/> Managed <input type="checkbox"/> Ignored
eth0	fd00::200:a:0:248/64	Managed v

In case you are editing an unmanaged workload, you will not have the option to ignore the workload using the PCE Action drop-down. That drop-down menu does not exist for unmanaged workloads. You can still provide information on the Interface Name and the IP/CIDR address.

Network Interfaces

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.

<input type="checkbox"/> Interface Name	<input type="checkbox"/> IP/CIDR
<input type="checkbox"/> <input style="border: 1px solid red;" type="text" value="E.g. eth0.public"/>	<input type="checkbox"/> <input style="border: 1px solid #ccc;" type="text" value="E.g. 10.0.10.1/24 17.1.0.10"/>

5. Click **Save**.

Workloads and VENS

The modified 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 in a new page separate from workloads. The VEN-related actions are not available under the Workloads tab.



NOTE:

Users with the Workload Manager role can manage workloads and VENs.

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

Workloads Container Workloads **VENs**

Unpair 1 Refresh Reports ▾

Select properties to filter view ▾

1 Selected	Customize columns ▾ 50 per page ▾ 1 - 4 of 4 Total ▾ < >									
<input type="checkbox"/>	Connectivity	Health	Name	Version	Role	Application	Environment	Location	OS	Status
<input checked="" type="checkbox"/>	Online	✓	ip-10-0-22-107	19.3.0-6104	Gateway	IOT infrastructure	Development	AWS	ubuntu-x86_64-xenial	Active
<input type="checkbox"/>	Online	✓	rhel7.local	19.1.1-5651	Single Node App	Testbed	Development	VMware	centos-x86_64-7.0	Active
<input type="checkbox"/>	Online	✓	SERVER20 12-B	19.1.1-5651	Single Node App	Testbed	Development	VMware	win-x86_64-server	Active
<input type="checkbox"/>	Online	✓	solaris11-3	19.1.1-5651	Single	Tethered	Development	VMware	SunOS	Active

Unpair 1 Refresh Reports

Select properties to filter view ▾

Unpair VEN ✕

Selected	Pairing Method	Policy Mode	Remove or Unpair Actions	Remove
<input checked="" type="checkbox"/>	Paired using pairing key	Build/Test/Enforce	Uninstalls the selected VEN(s). Removes policy for the associated workloads. Policies are configured into the host firewall based on the option selected below.	✕

Select final firewall status

Remove Illumio policy - Default

- Linux Removes Illumio policy and retains the coexistent firewall rules.
- AIX / Solaris Removes Illumio policy and reverts firewall rules to the pre-pairing state.
- Windows Removes Illumio WFP filters and activates Windows Firewall.

Open all ports

All operating systems Leaves all ports open.

Close all ports except remote management

- Linux / AIX / Solaris Temporarily allows only SSH/22 until system is rebooted.
- Windows Allows only RDP/3389 and WinRM/5985, 5986.

Cancel **OK**

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

Container Workloads

The Container Workloads page, lists the containers that exist on the PCE. The Status, Name, Container ID, and the Labels are displayed. You can click on a container to view it's details.

☰ ↶ Container Workloads - g

Summary Containers Rules

🗺 App Group Map

General

Name	guestl
Namespace/Project	guestbook
Policy State	Build Build Rules without events
Policy Sync	Active
Firewall Mode	Exclusive
Container Cluster	KS
Created At	08/12/2019 at 14:17:09

Labels

Role	Web
Application	C_1.0
Environment	Development
Location	1

Attributes

Interfaces	eth0: 19... 1.1
Workload	ip-...internal

Container Workloads - gu	
Summary	Containers
Container	
Name	/k8s_POD_guestbook-ssrp4_guestbook_87e52d05-bd46-11e9-b7c9-0e854af88744_0
Container ID	cc1e1fd6db23151b7e813e3c2227c7bd07ab52280b5b826d904e1d086a44dcb9
Started At	08/12/2019, 14:17:04
Stopped At	Never
annotation.com.illumio.role	Web
annotation.kubernetes.io/config.seen	2019-08-12T21:17:03.691628325Z
annotation.kubernetes.io/config.source	api
annotation.kubernetes.io/psp	eks.privileged
app	guestbook
io.kubernetes.container.name	POD
io.kubernetes.docker.type	podsandbox
io.kubernetes.pod.name	guestbook-ssrp4
io.kubernetes.pod.namespace	guestbook
io.kubernetes.pod.uid	87e52d05-bd46-11e9-b7c9-0e854af88744
Container	
Name	/k8s_guestbook_guestbook-ssrp4_guestbook_87e52d05-bd46-11e9-b7c9-0e854af88744_0
Container ID	34f3a1fa6db71ee754c8d11c977ac0eaa5d8cf2e5b7a5ffa9c18ed9150be5ada
Started At	08/12/2019, 14:17:04
Stopped At	Never
annotation.io.kubernetes.container.hash	5745b73b
annotation.io.kubernetes.container.ports	[{"name":"http-server","containerPort":3000,"protocol":"TCP"}]
annotation.io.kubernetes.container.restartCount	0
annotation.io.kubernetes.container.terminationMessagePath	/dev/termination-log
annotation.io.kubernetes.container.terminationMessagePolicy	File
annotation.io.kubernetes.pod.terminationGracePeriod	30
io.kubernetes.container.logpath	/var/log/pods/87e52d05-bd46-11e9-b7c9-0e854af88744/guestbook/0.log
io.kubernetes.container.name	guestbook
io.kubernetes.docker.type	container
io.kubernetes.pod.name	guestbook-ssrp4
io.kubernetes.pod.namespace	guestbook
io.kubernetes.pod.uid	87e52d05-bd46-11e9-b7c9-0e854af88744
io.kubernetes.sandbox.id	cc1e1fd6db23151b7e813e3c2227c7bd07ab52280b5b826d904e1d086a44dcb9

Workload Setup Using PCE Web Console

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

About 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 Using PCE Web Console](#) in the *VEN Installation and Upgrade Guide*.

Unmanaged Workloads

Unmanaged workloads extend rule-writing capabilities to network entities that are 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 that are 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.

Add an Unmanaged Workload

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](#) for information.

1. From the PCE web console menu, choose **Workloads and VENs > Workloads**.
2. Click **Add > Add Unmanaged Workload**.
3. In the Add Unmanaged Workload details page, enter a name and description for the unmanaged workload.
4. In the *Label* section, select the labels you want to be applied to the unmanaged workload.
5. In the *Attributes* section, enter all the relevant information about the unmanaged workload, such as its hostname, IP addresses, location, and OS.
6. In the *Processes* section, enter the name and the port and protocol for the process.
7. (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. See [Secure Laptops with AdminConnect](#) for information.

8. Click **Save**.

VEN Administration on Workloads

The connectivity, policy sync, and health status of the VEN can be monitored 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.

For more information about managing VENs on workloads, see [About VEN Administration on Workloads](#) in the *VEN Administration Guide*.

VEN Details for a Workload

VEN – analytics-s9

Edit Unpair Upgrade Generate Support Report Mark as Suspended

Node

Name	analytics-s9
Description	
Hostname	analytics-s9
Enforcement Node Type	Virtual Enforcement Node (VEN)
Version	20.2.0
Activation Type	Pairing Key

Status

Status	✔ Active
Health	✔ Healthy
Last Heartbeat Received	10/12/2020 at 13:29:51

Host

Location	Amazon EC2 (US West), Oregon, USA
OS	ubuntu-x86_64-xenial
Release	4.4.0-97-generic #120-Ubuntu SMP Tue Sep 19 17:28:18 UTC 2017 (Ubuntu 16.04.1 LTS)

Workload


Name	analytics-s9
Enforcement	Full Segmentation Rules are enforced for all inbound and outbound services. Traffic not allowed by a Segmentation Rule is blocked
Visibility	Blocked + Allowed VEN logs connection information for allowed, blocked and potentially blocked traffic
Policy Sync	Active
Policy Last Received	10/12/2020 at 10:24:56
Interfaces	eth0: 10.28.36.62/8 10.0.0.1 eth0: fd00::200:a:0:fc/64 eth0,public: 66.151.147.220/32
Public IP Address	66.151.147.220
Role	
Application	
Environment	
Location	

VEN Suspension

You can mark a workload as suspended by using the PCE web console. To suspend a VEN, choose **Workloads and VENs > VENs** from the PCE web console menu. Select

your VEN to open its details page and click **Mark as Suspended**.

For more information about suspending and unsuspending VENs, see [VEN Suspension Using PCE Web Console](#) in the *VEN Administration Guide*.



Blocked Traffic

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

**IMPORTANT:**

In the 19.1.0 release, blocked traffic was marked for deprecation and will be turned off by default in a future release. When a large number of traffic summaries are reported to the PCE, the blocked traffic functionality consumes more memory, which can cause side-effects such as:

- Illumination dropping some traffic flows
- PCE slowing down due to extra processing

When upgrading to 19.3.0, Illumio recommends that you turn off blocked-traffic by setting the appropriate value in the PCE `runtime_env` file.

The functionality provided by blocked traffic is available in Explorer. In 18.3.1 and later, when the Explorer feature is configured, the Blocked Traffic page was updated using the Explorer data. The Blocked Traffic page will continue to work using the data from Explorer.

Overview of Blocked Traffic

To view the Blocked Traffic page, choose **Troubleshooting > Blocked Traffic** from the PCE web console menu. The Blocked Traffic tab shows you all traffic that attempted to communicate with your workload but was blocked due to policy. Blocked traffic alerts provide information such as the port and protocol of the service, as well the IP address of the consumer, the total number of flows, and the time last detected.

Traffic Type	Provider	Service	Consumer	Total Flows	Last Detected
Blocked By the Consumer	Internet 172.31.47.255	← unknown 138 UDP	WIN-F7MNO 172.31.47.3	1201	03/06/2018, 15:09:15
Blocked By the Consumer	Internet 172.31.47.255	← unknown 137 UDP	WIN-F7MNO 172.31.47.213	43	03/06/2018, 15:02:12

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.

Traffic that is blocked in the following ways is reported as blocked traffic in the Illumination map, regardless of the workload enforcement:

- Firewalls on the workload not managed by Illumio Core
- WFP policies not managed by Illumio Core

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 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 content of the page 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 consumer or by the provider.
- **Provider:** Displays the workload name and IP address of the provider.
- **Provider Labels:** Displays labels assigned to the provider.
- **Service:** Displays the process name, port, and protocol information of the traffic that was reported along with an indication of whether the record was reported by the consumer or the provider.



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.

- **Consumer:** Displays the workload name and IP address of the consumer.
- **Consumer Labels:** Displays labels assigned to the consumer.
- **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 provider reports the record, the information in the consumer column is grayed out. When the consumer reports the record, the information in the provider column is grayed out.



From the 18.3.1 release on, the traffic entries displayed on the blocked traffic page cannot be removed via the PCE web console.

Filter Blocked Traffic

The Blocked Traffic page displays the 500 most recent entries from all workloads managed by the PCE. When you are monitoring or writing rules for a specific set of workloads, use Blocked Traffic filters to display up to 500 of the most relevant entries based on the 10,000 entries in the PCE.

The PCE web console allows you to use filters to display only the blocked traffic entries of interest. You can filter based on workload name, label, traffic type (blocked or potentially blocked), or any combination of these attributes. When you apply the filter by clicking **Go**, the 500 most recent entries that match the search criteria are displayed.

To filter blocked traffic, type the keywords for the filter in the *Select properties to filter view* field at the top of the Blocked Traffic page.

Role	Provider	Service	Consumer	Total Flows	Last Detected
Application	ordering-web3	 unknown 5678 TCP	Internet Minus Blacklist	7	01/24/2016 20:13:18
Environment					
Location					
Traffic Status	ordering-web2	 unknown 5678 TCP	Internet Minus Blacklist	7	01/24/2016 15:45:55
Name					



NOTE:

You can filter blocked traffic using multiple properties at the same time. Only entries that match all the entered criteria are displayed.

To specify the type of results, click the arrow at the end of the text entry field and select one or more of the available properties:

- Role
- Application
- Environment
- Location
- Traffic status
- Workload name

After entering your keywords, click **Go** to the right of the text entry field. The results display below the text entry field. The following information is included:

- **Traffic Type:** A link to additional information about that entry
- **Provider:** The provider of the service
- **Service:** The service type
- **Consumer:** The consumer of the service
- **Total Flows:** The total number of times this blocked traffic flow occurred
- **Last Detected:** A timestamp (in hh:mm:ss format) of the last time this flow occurred

Create Unmanaged Workload from Blocked Traffic

In some cases, your policy might be blocked from the IP address of a host 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 [Unmanaged Workloads](#).

1. From the PCE web console menu, choose **Troubleshooting > Blocked Traffic**.
2. From the list of blocked traffic events, under the *Consumer* column, click any of the linked IP addresses.

Traffic Type	Provider	Service	Consumer
<input type="checkbox"/> Blocked By the Provider	support-s11-zones 10.6.255.255	! ← unknown 137 UDP	Internet Minus Blacklist 10.6.4.50
<input type="checkbox"/> Blocked By the Provider	support-s11-zones 10.6.255.255	! ← unknown 138 UDP	Internet Minus Blacklist 10.6.3.15

The Unmanaged Workload page appears.

3. Complete all the fields and click **Save**.

You can now use the unmanaged workload in your policy. For example, you can configure rules to allow incoming traffic from this unmanaged workload to other managed workloads.

Reject Connections

You can configure Workloads to reject traffic that does not meet the required policy, instead of blocking it in the *Enforced* state. You can edit *Reject Connections* from the **Settings > Security** menu option.

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 and not at the interface level. It is selected by a Label set.
- It is visible on the Workload detail page.

☰ ↶ Workload – solr-s66

Summary
Processes
Rules
Blocked Traffic
Vulnerabilities

✎ Edit
✖ Remove
🗺 App Group Map
🗺 Vulnerability Map

General

Name	solr-s66
Description	
Policy State	Build <small>Build Rules without events</small>
Policy Last Received	10/07/2019, 00:19:23
Policy Last Applied	10/07/2019, 00:19:23
Policy Update Mode	Adaptive
Firewall Mode	Exclusive
Blocked Connection Action	Reject
Container Inherit Host Policy	No

Vulnerability

Total V-E Score	453
Highest V-E Score	59
Highest Vulnerability	7.8
Import Time	08/23/2018 at 15:27:06

Labels

Role	🏠 Load Balancer
Application	🏠 normal name
Environment	🌱 Staging
Location	📍 London

Create Security Policy

This chapter contains the following topics:

Segmentation Templates	81
Policy Generator	90
Segmentation Rulesets	104
Rules	112
Rule Writing	122
FQDN-Based Rules	133
Provisioning	143

This section describes how to create 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.

See also [Rule Creation Approach](#) in the “Overview of Security Policy” section.

Segmentation Templates

Applications can be a complex set of services and processes that have different components which communicate with other applications. For example, you might find an application in your Illumination map that has many processes communicating through several ports to connect to and receive connections from Active Directory. Some of these processes, such as Netlogon, can use 10,000 or more dynamic ports as it’s communicating with Active Directory. The ports that are used at any given time can be

unpredictable. Creating security policy for these types of applications is a complex and time consuming endeavor.

Overview of Segmentation Templates

To deliver Segmentation Templates, Illumio leveraged our knowledge of enterprise applications, such as Active Directory, Exchange, and SharePoint, because we know the services and the different processes that these applications use.

Illumio Segmentation Templates provide prepackaged, tested security policies that provide all the segmentation rules needed for common enterprise applications. They can be deployed in minutes; thereby reducing the time it takes to protect key computing assets. They simplify the definition and implementation of security policy while reducing errors and preventing security gaps for widely-used, business critical applications.

Each Segmentation Template serves two purposes. Illumio customers can see an example of how to add the security policies required to protect the application in question. Secondly, customers can use the Segmentation Template as designed to secure the application quickly in their organization.

When you install a Segmentation Template, the PCE web console automatically adds the necessary policy objects (such as services, segmentation rulesets, and labels) to allow the communication required for that application.

Catalog Retrieved from Support Portal

When you go to the Segmentation Templates page, the PCE web console automatically retrieves the latest Segmentation Templates catalog from the Illumio Support portal and displays it in the web console.

The screenshot displays the 'Segmentation Templates' page in the Illumio web console. The page title is 'Segmentation Templates'. Below the title, there is a section titled 'For Windows' containing seven template cards. Each card includes an icon, the template name, an 'Install' button, and details about its content.

Template Name	Version	Content
Exchange 2013	1	3 rulesets, 33 Services, 8 labels, 1 IP list
Active Directory 2008+	1	6 rulesets, 25 Services, 13 labels, 1 label group, 3 IP lists
Microsoft System Center 2012	1	5 rulesets, 14 Services, 8 labels, 3 IP lists
SharePoint 2013	1	5 rulesets, 19 Services, 10 labels, 1 IP list
Microsoft SQL Server 2014	1	1 ruleset, 19 Services, 3 labels
Windows Server Update Service	1	1 ruleset, 3 Services, 1 label, 1 IP list
Windows	1	35 Services

To manually locate the catalog on the Illumio Support portal:

1. From the PCE web console menu, choose **Policy Objects** > **Segmentation Templates**.

A dialog box appears prompting you to log into the Illumio Support portal. (While you are logged into the PCE web console, you only have to log into the Illumio Support portal once.)

2. Click **Log In** and, if prompted, enter your Illumio Support portal username and password. (Illumio Secure Cloud customers are automatically logged into the Illumio Support portal.)

**NOTE:**

Internet connectivity is not required to use the Segmentation Templates. When you are connecting to the PCE web console from a device that does not have internet connectivity, you must access the Illumio Support portal from another device that has internet connectivity and download the templates locally to that device before you can use them. See [Upload a Segmentation Template](#).

The Illumio Support portal automatically redirects you back to the Segmentation Templates page and the templates appear in the page. The templates are organized by operating system.

3. To view the contents of a Segmentation Template, click its name or icon.

The Segmentation Template details page describes the template and lists all the policy objects that belong to the template. Policy objects appear as hyperlinks when they have already been installed by another template. (Templates can share policy objects.)

Features of Segmentation Templates

Segmentation Templates share the following key features.

Template Contents

Each Segmentation Template adds an associated group of unique, non-overlapping, predefined services, and can contain any of the following policy objects:

- Labels
- Label groups
- IP lists
- Segmentation Rulesets

Some templates contain all the segmentation rulesets, services, and labels needed to secure a given application. Other templates contain port-based service definitions only.

Dynamic Processes and Ports

Using Segmentation Templates is especially useful in Microsoft environments, which must accommodate a range of dynamically used ports for RPC. Other Microsoft applications (such as Active Directory) require opening dynamic port ranges. Rather than opening only the ports in use, network-based solutions leave open an entire range of ports, effectively leaving the security environment wide open.

The Illumio PCE is service and process aware. Because of this, installing Segmentation Templates can protect against dynamic processes (like Netlogon) and add the correct policy to open only the ports that are active at a time.

Segmentation Templates are designed to use the specific processes and path used by the server rather than dynamic ports and apply the exact set of fine-grained segmentation rules required for protection.

Sharing Policy Objects

Services, labels, label groups, and IP lists can be used by more than one Segmentation Template. A segmentation ruleset, however, is never used by multiple templates.

Identifying Policy Objects Added by Templates

You can identify all objects added to the PCE that are part of Segmentation Templates. In the External Data Set field of the object's details page, the PCE identifies these policy objects by labeling them using the following convention:

IST - type_of_object

(Where IST stands for Illumio Segmentation Template)

Additionally, the PCE provides full names to increase readability. For example, "IST - [AD] - Client to Domain Controller" appears as "IST - Active Directory Client to Domain Controller."

Segmentation Template Prerequisites and Limitations

Segmentation Templates are bound by the following prerequisites and limitations.

Internet Connectivity

Internet connectivity is not required to use the Segmentation Templates. For example, you might be connecting to the PCE web console from a device that does not have internet connectivity.

Illumio stores the Segmentation Templates on the Illumio Support portal. When the device from which you are connecting to the PCE web console does not have internet connectivity, you can connect to the Illumio Support portal over the internet using another device and download the Segmentation Templates locally, then upload them to the PCE web console from that device.

When you choose **Policy Objects > Segmentation Templates** from the PCE web console, you are prompted to log into the Illumio Support portal to download the templates. When you do not have internet connectivity from your device and have already downloaded the templates to another device, you can skip this step.

See [Catalog Retrieved from Support Portal](#) for information.

Upgrade Policy Object Installed by Segmentation Templates

The PCE recognizes when policy objects are installed by Segmentation Templates from the values in the External Data Reference field. Therefore, if you installed a Segmentation Template prior to 17.2 or you modified the contents of this field for an object, the PCE cannot recognize that a template installed the object and you cannot update it while updating the template.

Unique Names for Labels, Label Groups, and IP Lists

In the PCE web console, the names of policy objects must be unique. For example, when you have an existing label, label group, or IP list that has the same name as a label, label group, or IP list in a template, the template installation will end and prompt you to change the name of the policy object in your organization.



NOTE:

In Segmentation Templates, policy objects are named using the following convention: *IST - type_of_object*

Delete Labels Associated with Segmentation Templates

When you have provisioned a segmentation ruleset or label group associated with a template, the labels associated with the template cannot be removed until the segmentation rulesets and label groups are removed and the removal is provisioned.

About Editing Segmentation Templates

Installing a Segmentation Template adds a predefined set of services and can add labels, label groups, IP lists, and segmentation rulesets.

Editing a policy object associated with a Segmentation Template is no different from editing any other policy object in the PCE web console. Also, the display and designation of a Segmentation Template does not change in the PCE web console after you edit the policy objects associated with it.

However, before you edit the policy objects installed by a Segmentation Template, you should be aware of the following caveats.

Edit the Names or IDs of Policy Objects

The PCE assigns each policy object associated with a template an ID number, which the PCE web console displays in the Description and External Data Reference fields of the object details or Summary pages.

The PCE tracks all objects associated with Segmentation Templates by their names. In Segmentation Templates, these policy objects are named using the following convention:

IST - *type_of_object*

Changing the policy object name does not affect the PCE validation that it is installed; however, using the Illumio API to edit the External Data Reference field does affect the PCE validation that it is installed.



NOTE:

Illumio strongly recommends you do not change the IDs in the *External Data Reference* fields.

Delete Policy Objects or Editing Their Attributes

Deleting policy objects associated with templates or editing their attributes is subject to the following caveats:

- When you remove a policy object associated with a template after the template is installed, the PCE will re-add the object when the template is updated.

For example, you remove the common LDAP service, which is associated with a Segmentation Template. When Illumio releases an update for the template, installing that update will re-add the common LDAP ports to the PCE.

- When you edit the attributes of policy objects associated with a template (for example, edit the ports or protocols of a service, or the scope or rules of a segmentation ruleset), the PCE web console will prompt you to specify whether to preserve or overwrite your changes when you update the template to the next version.

Install a Segmentation Template

1. [Retrieve the Segmentation Template Catalog.](#)

When a template has not been installed, an **Install** button appears on the page.

2. Click **Install**.

The end user licensing agreement (EULA) appears.

3. Accept the EULA and click **Continue**.

Before the PCE installs the template, it checks that the policy objects required by the template don't conflict with any existing policy objects in your

organization. The time that the check takes depends on the number of policy objects in your organization. When the PCE finds any conflicts during the check, it cancels the installation and doesn't install any policy objects. You are prompted to rename the conflicting objects.

When the check is successful, the PCE adds the included policy objects in Draft mode so that you can review or edit them before provisioning them. See [Provisioning](#) for more information.

As the policy objects are added, links to the objects appear in the template details page.



NOTE:

Global policy objects—such as All Services and Any (0.0.0.0/0 and ::/0)—don't include links in the Segmentation Template details page to the objects.

Upload a Segmentation Template

Internet connectivity is not required to use the Segmentation Templates. However, Illumio stores them on the Illumio Support portal. When you are connecting to the PCE web console from a device that does not have internet connectivity, you can retrieve the templates using another device (which has internet connectivity), then manually upload them to the PCE web console so that you can install or update them.

When you download a Segmentation Template from the Illumio Support portal, you save the template locally as a JSON file.

1. Log into the Illumio Support portal with your Illumio Support username and password.
2. Click **Tools > Illumio Segmentation Templates**.
3. Navigate to the Segmentation Templates and download them locally.
4. Log into the PCE web console and choose **Policy Objects > Segmentation Templates**.

The Segmentation Templates dialog box appears.

5. Click **Load File**.

A dialog box appears prompting you to specify the Segmentation Template file to upload.

6. Click **Choose File**.

A file explorer appears.

7. Navigate to the file and click **Open**.

The Segmentation Templates dialog box reappears.

8. Click **Load**.

The page refreshes and a tile for the Segmentation Template appears in the page.

Update a Segmentation Template

Updating a Segmentation Template to a later version can edit or add services, segmentation rulesets, labels, label groups, or IP lists. However, updating a template never removes policy objects added by a previous version.



NOTE:

Later versions of templates are fully backwards compatible with previous versions.

1. [Retrieve the Segmentation Template Catalog](#).

When a new version of a Segmentation Template is available for a template that you have installed, the template has an **Update** button.

2. Click **Update**.

If you edited the Segmentation Template after installing it, a dialog box appears prompting you to specify how to install the new version. For example, you added a new port and protocol to a service added by the template. You can revert the template to the Illumio list of ports and protocols for that service or keep your changes.

3. If necessary, choose how to handle template changes:
 - **Overwrite:** The PCE replaces the policy objects that you edited with the version in the new template and removes the word “edited” after the ID number in the *External Data Reference* field.
 - **Preserve Changes:** Your changes to the policy objects added by the template are kept.

**NOTE:**

If you have edited multiple policy objects associated with a template, you must choose whether to overwrite or preserve all your changes. You cannot overwrite some and preserve some.

The PCE updates the version numbers of all policy objects associated with the template even when the new template changes only a subset of the objects associated with the template.

**NOTE:**

Segmentation Templates can share policy objects; therefore, a policy object can have a later version than a template it's associated with because the object was updated by another template. For example, you can have version 1 of a template installed and it includes version 2 of some policy objects.

Uninstall a Segmentation Template

1. [Retrieve the Segmentation Template Catalog.](#)

After you install a Segmentation template, an **Uninstall** button appears on the page.

2. Click **Uninstall**.

When you uninstall a Segmentation Template, the PCE removes all the policy objects that are associated with that template except when an object is in use. Policy objects that are shared with other installed templates are not removed. Policy objects that are added to other policy objects are not removed. For example, you added a service associated with a template to a segmentation rule-set.

Policy Generator

The Policy Generator simplifies the Illumio policy creation process by recommending the optimal security policy for your App Groups. You can use it to accelerate security workflows and reduce the risk of human error while creating security policy.

Overview of Policy Generator

The Policy Generator uses network traffic to recommend and generate micro-segmentation policies for every workload and application, regardless of its location. It

can generate rules for applications running on physical devices, virtualized platforms, and behind network devices on-premises or deployed in the cloud.

Policy Generator supports the creation of DNS-based rules under all the wizards (intra-scope, extra-scope, and IP lists). You can edit the proposed virtual services and add wildcards.

Application owners use the Policy Generator to write the following types of rules for the applications they manage:

- Intra-scope rules
- Extra-scope rules
- Rules using IP lists.

For more information about each rule type, see [Segmentation Rulesets](#), [Rules](#), and [IP Lists](#).

For a selected App Group, the Policy Generator provides:

- A workflow to create a ruleset that controls internal and external traffic.
- A way to assess your current rule coverage, which represents the number of detected connections that are controlled by rules divided by the total number of connections.

You can increase your rule coverage by creating rules for detected connections that are not controlled by rules. The Policy Generator proposes rules for connections that are not allowed currently by rules and displays the consolidated flow count for each new proposed rule to help ensure the maximum impact on rule coverage.



NOTE:

The Policy Generator calculates rule coverage automatically every 24 hours or after creating a draft ruleset.

You can rewrite rules as your datacenter needs change and the Policy Generator will show you the before and after effect of those rules.

- A way to assess your current rule coverage, which represents the number of detected connections that are controlled by rules divided by the total number of connections.
- Visualization of the traffic between roles associated with a specific application, as represented by App Groups.

- Options to select the level of granularity for new rules; see [About Granularity Levels for Rules](#) for information.

The first time you use the Policy Generator for an App Group, it creates a new draft ruleset with the title of the selected App Group. When you use Policy Generator again to create additional rules, it adds them to the existing ruleset that was created by the Policy Generator. You review the proposed rules and can customize them before you save them into a draft ruleset. For Windows, the Policy Generator detects and suggests Windows process- and service-based rules accordingly. You can edit the service before saving it.

**NOTE:**

You must provision the rules to apply them to workloads. See [Provisioning](#) for more information.

When an App Group has several consumers communicating with a specific provider, the Policy Generator merges all the consumers into one rule for easy readability and better scalability.

On the Summary tab of the Ruleset page, any rulesets created with Policy Generator have the default description “Automatically generated using the Illumio Policy Generator” and the value of `illumio_policy_generator` for the External Data Set field. The value for the External Data Reference is the App Group name.

Policy Generator Prerequisites and Limitations

The Policy Generator is bound by the following prerequisites and limitations:

- You cannot add Role-level rules until Role labels have been added to all workloads in the App Group.
When some workloads in an App Group do not have Role labels, you can still write an App Group level rule using Policy Generator to allow all the workloads to communicate with each other.
- Rule coverage is updated one App Group at a time.

About Granularity Levels for Rules

The following options allow you to select the restrictiveness of your security policy.

Intra-Scope Rules

Granularity Level	Description
App Group Level	(Also known as micro-segmentation or Ringfencing) All workloads in the App Group can communicate with each other across all services. This option is best for creating a broad initial policy that can be further refined later if needed.
Role Level - All Services	All workloads with a specific Role label can communicate with all workloads with Role labels matching the observed flows across all services. This option is useful for restricting role-to-role traffic between workloads when you have many core services that need to communicate with these workloads. When a workload is missing the Role label, the Policy Generator excludes that connection from the wizard.
Role Level - Specified Services	(Also known as nano-segmentation) All workloads with a specific Role label can communicate with all workloads with Role labels matching the observed flows across specified services, based on the collected traffic flow summaries. When a matching port/protocol cannot be located in an existing service, a new service with the necessary port/-protocol is created when the proposed rules are saved. Use this option to create the most restrictive policy.


Extra-Scope and IP List Rules

Granularity Level	Description
All Services	Workloads can communicate over all services. This service policy type provides less restriction for workload communication.
Specified Services	Workloads can communicate over specified services. This service policy type provides more restriction for workload communication.
Auto Level	

Ways to Access Policy Generator

You can access Policy Generator from the following locations in the PCE web console:

Entry Point	Description
Policy Generator	Launches the Policy Generator. You must select an App Group to begin.
App Group Map > App Group panel	Launches the Policy Generator for the App Group selected. When you've opened a Consuming App Group and selected the App

Entry Point	Description
> Start Policy Generator	Group, the Policy Generator creates an extra-scope rule. You can proceed or add more Consuming App Groups.
Illumination > Group panel > Start Policy Generator	<p>Launches the Policy Generator for the App Group selected.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  NOTE: App Groups must be configured to use three labels to start the Policy Generator from the Illumination map. </div>
Rulesets and Rules > Start Policy Generator	Launches the Policy Generator. You must select an App Group to begin.
Rulesets and Rules > Ruleset Details > Start Policy Generator	When a ruleset was created using the Policy Generator, the rule-set scopes and Rules tab includes a Start Policy Generator button. Clicking the Start Policy Generator button, launches the Policy Generator with the App Group selected.
Troubleshooting > Blocked Traffic > Start Policy Generator	Launches the Policy Generator. You must select an App Group to begin.
App Group Map > App Group panel > Mitigate Vulnerabilities	Launches the Policy Generator. You can update your policy to minimize the risks due to the vulnerabilities.

Create Intra-scope Rules with Policy Generator

1. From the PCE web console menu, choose **Policy Generator**.

The Select App Group page appears. The page displays when the Policy Generator last calculated the coverage for each type of rule. Click the refresh icon to recalculate Rule coverage.

2. Select an App Group.

See [Segment Multiple App Groups with Policy Generator](#) for information about adding App Group level rules for multiple App Groups.

3. Click the **Start with Intra-Scope** button.

The Intra-Scope Rule Configuration page appears.

4. In the *Choose Intra-Scope Rule Configuration* section, select a granularity level for the rules.

See [Create Intra-scope Rules with Policy Generator](#) for a description of these rule granularity levels.

The detected connections (including details such as provider, port/protocol, and consumer) appear in the Review All Connections section.

Rule Configuration	Connections Displayed
App Group Level	Labels, ports, and protocols in a single row
Role Level - All Services	Number of connections and associated labels
Role Level - Specified Services	Associated labels and ports/protocols



NOTE:

The Policy Generator displays a truncated list of ports and protocols when the App Group has more than four types of ports or protocols. To display the remaining ports or protocols in a modal window, click the **+ More** link.

5. (Optional for Role level) To exclude a connection from the proposed rules, click **Exclude**. The row is grayed out to indicate that no rules will be proposed for this connection and the amount of rule coverage decreases. To include an excluded connection, click **Include**.



NOTE:

At least one connection must be included to continue.

6. Click **Next**.

The proposed rules appear in the Preview page.

Policy Generator – Intra-Scope Rule Preview



App Group: ApplicationXYZ2 | Production

Ruleset Scope		
Application	Environment	Location
ApplicationXYZ2	Production	All Locations

Rule Construction:

Rule Merging

- Merge rules with common Provider and Consumer
- Merge rules with common Provider and Service

Providing Service

- Use the port/protocol in a rule if a Service does not exist
- Create a new Service if one does not exist for a port/protocol

5 New Intra-Scope Rules

Providers	Providing Service	Consumers
API	Service - 3E CP	Web
API	Service - 3 CP	API
API	Service - 3i CP	analytics-s
analytics-s	Service - 3 y	app-s
API	Service - 3 TCP	api

134 Mitigated Vulnerabilities


Role	Vulnerability	Before	After
Web	27 Vulnerabilities	1.6K	1.6K
API	23 Vulnerabilities	1.1K	1.2K
analytics-s	9 Vulnerabilities	238	240
app-s	15 Vulnerabilities	563	567
discovered	15 Vulnerabilities	703	708
analytics-s	10 Vulnerabilities	542	546
api	35 Vulnerabilities	1.8K	1.8K

Adding to Ruleset: ApplicationXYZ2 | Production Rule Builder
86% New Intra-Scope Rule Coverage

Adding New Objects to Ruleset
5 New Rules

Existing Objects in Ruleset
1 Scope

Cancel Save

7. (Optional) To edit the service for a rule, click the pencil icon  beside a service. The Edit Service dialog box appears.

Select a service from the drop-down list or create a new one. You can select services that have broader ranges of ports. The list includes every service that matches that port and protocol. When you've added a service that has multiple ports and protocols or ranges, they all appear in the list.

Select **Apply Changes to all matching ports** to allow the service to be used in other rules that match that service. You are prompted to allow the Policy Generator to merge rules. To cancel the merge, reload the page and start over.

When you create a process-based service, the connection appears like it's not covered.

For information about creating a service, see [Create a Service](#).

8. To accept the proposed rules, click **Save** and **OK**.

The Policy Generator Successful message appears which displays the number of new rules and services. The rules are added to a draft ruleset. Click **Continue with App Group** to add extra-scope rules or rules using IP lists for the same App Group. On the last step of the Policy Generator, you can return to the App Group to add or append to the rules.



NOTE:

You must provision the rules to apply them to workloads. See [Provisioning](#) for more information.

Create Extra-scope Rules with Policy Generator

When you create extra-scope rules, the Policy Generator displays all traffic that originates from a different App Group and is targeted at the selected App Groups. The Policy Generator displays all App Groups that the selected App Groups communicate with. You can choose which connects to cover with rules.

1. From the PCE web console menu, choose **Policy Generator**.

The Select App Group page appears. The page displays when the Policy Generator last calculated the coverage for each type of rule. Click the refresh icon to recalculate rule coverage.

2. Select an App Group.

See [Segment Multiple App Groups with Policy Generator](#) for information about adding App Group level rules for multiple App Groups.

3. Click the **Start with Extra-Scope** button.

The Extra-Scope App Group Selection page appears.

4. Select one or more Consuming App Groups and click **Next**.

For each App Group, the Policy Generator displays the current number of connections and connections covered by a rule.



NOTE:

Consuming App Groups with 100% rule coverage are not displayed in the page.

The Configure Extra-Scope page appears.

5. Select whether to configure rules by App Group or by role:
 - **App Group Level:** All workloads in the specified App Group can communicate with all workloads in the other App Groups
 - **Role Level:** Specified workloads in the App Group can communicate with specified workloads in the other App Groups

6. Select the permitted services for the rules:

- **All Services:** Workloads can communicate over all services
- **Specified Services:** Workloads can communicate over specified services

See [Extra-scope Rules](#) and [IP Lists](#) for more information.

7. Review the connections selected for the proposed rules.

App Groups are separated by a thick line and the organization of the connections differs depending on the selected configuration. Connection details are organized based on your selection:


Selected Configuration	Details Organized by:
App Level + Specified Services	App Group and port/protocol
App Level + All Services	App Group
Role Level + All Services	Role and App Group
Role Level + Specified Services	Role and port/protocol

8. (Optional for any configuration except App Level + All Services) To exclude a connection from the proposed rules, click **Exclude**.

The row is grayed out to indicate that no rules will be proposed for this connection and the amount of rule coverage decreases. To include an excluded connection, click **Include**.

9. To preview the rules proposed by Policy Generator, click **Next**.

The Extra-Scope Rule Preview page appears.

10. (Optional) To edit the service for a rule, click the pencil icon  beside a service. The Edit Service dialog box appears.

Select a service from the drop-down list or create a new one. You can select services that have broader ranges of ports. The list includes every service that matches that port and protocol. When you've added a service that has multiple ports and protocols or ranges, they all appear in the list.

Select **Apply Changes to all matching ports** to allow the service to be used in other rules that match that service. You are prompted to allow the Policy Generator to merge rules. To cancel the merge, reload the page and start over.

When you create a process-based service, the connection appears like it's not covered.

For information about creating a service, see [Create a Service](#).

11. To accept the proposed rules, click **Save** and **OK**.

The Policy Generator Successful message appears, which displays the number of new rules and services. The rules are added to a draft ruleset. Click **Continue with App Group** to add intra-scope rules or rules using IP lists for the same App Group.



NOTE:

You must provision the rules to apply them to workloads. See [Provisioning](#) for more information.

Create Rules Using IP Lists with Policy Generator

Policy Generator creates rules that use IP lists as intra-scope rules.

When using IP lists to create rules, the Policy Generator defines a connection as a role on a port and protocol to an IP address. For example, when you have five IP addresses that are included in an IP list, the Policy Generator displays five connections.

1. From the PCE web console menu, choose **Policy Generator**.

The Select App Group page appears. The page displays when the Policy Generator last calculated the coverage for each type of rule. Click the refresh icon to recalculate rule coverage.

2. Select an App Group.

See [Segment Multiple App Groups with Policy Generator](#) for information about adding App Group level rules for multiple App Groups.

3. Click the **Start with IP Lists** button.

The IP List Selection page appears.

4. Select the IP lists for which you want to write rules and click **Next**.

The Configure IP List page appears.



TIP:

- To view the IP addresses configured in a list (not the IP addresses in the traffic), expand an IP list by clicking the arrow icon in the Name column.
- To write rules covering all connections, select the Any IP list. This list covers all connections because it includes all the IP addresses.
- Each IP address can be part of more than one IP list and you can choose which list to write your rules to.
- When you choose overlapping IP lists, you can write overlapping rules.

When an IP address is in more than one IP lists, the rule is going to be in all those IP lists.

- You can write rules for inbound and outbound connections, or both. For example, you can write permissive rules for outbound traffic, and specific rules for inbound traffic.

5. Select whether to configure rules by App Group or by role:
 - **App Group Level:** All workloads in the specified App Group can communicate with all workloads in the other App Groups
 - **Role Level:** Specified workloads in the App Group can communicate with specified workloads in the other App Groups

6. Select the permitted services for the rules:

- **All Services:** Workloads can communicate over all services
- **Specified Services:** Workloads can communicate over specified services

It writes a rule for anything that those IP lists applied to.



TIP:

- To display the IP addresses of the traffic for each port and protocol, hover over the info (i) icon in the Consumer column.
- To filter connections by the IP address of the traffic, port number, protocol, role, and label, use the search field above the list of connections. You can use the search field to find and exclude specific traffic.
- To quickly include or exclude all traffic, use the **Include** and **Exclude** buttons by the search field. You can exclude all traffic, then selectively include specific connections.

2. Review All Connections

Rules will be generated for the following connections

Include 2

Exclude 2


125.10.15.45 ×

Find

Ruleset Inclusion	Provider	Port/Protocol	Consumer
1 Connection - 10 Flows <div style="display: flex; justify-content: space-between; margin-top: 5px;"> Include Exclude </div>	<div style="border: 1px solid #ccc; padding: 2px 5px; display: inline-block;">Role2</div>	<div style="color: red; font-size: 2em; margin-bottom: 5px;">←</div> 22 TCP	<div style="border: 1px solid #ccc; padding: 2px 5px; display: inline-block; margin-bottom: 5px;">IPL_1 ⓘ</div> <div style="border: 1px solid #ccc; padding: 2px 5px; display: inline-block;">Any (0.0.0.0/0 and ::/0) ⓘ</div>
1 Connection - 10 Flows <div style="display: flex; justify-content: space-between; margin-top: 5px;"> Include Exclude </div>	<div style="border: 1px solid #ccc; padding: 2px 5px; display: inline-block;">IPL_1 ⓘ</div>	<div style="color: red; font-size: 2em; margin-bottom: 5px;">←</div> 443 TCP	<div style="border: 1px solid #ccc; padding: 2px 5px; display: inline-block;">Role3</div>

7. To preview the rules proposed by Policy Generator, click **Next**.

The IP List Rule Preview page appears.

8. (Optional) To edit the service for a rule, click the pencil icon  beside a service. The Edit Service dialog box appears.

Select a service from the drop-down list or create a new one. You can select services that have broader ranges of ports. The list includes every service that matches that port and protocol. When you've added a service that has multiple ports and protocols or ranges, they all appear in the list.

Select **Apply Changes to all matching ports** to allow the service to be used in other rules that match that service. You are prompted to allow the Policy Generator to merge rules. To cancel the merge, reload the page and start over.

When you create a process-based service, the connection will appear like it's not covered.

For information about creating a service, see [Create a Service](#).

9. To accept the proposed rules, click **Save** and **OK**.

The Policy Generator Successful message appears, which displays the number of new rules and services. The rules are added to a draft ruleset.

Segment Multiple App Groups with Policy Generator

You can apply nano-segmentation (also known as ringfencing) on multiple App Groups using the Policy Generator. Nano-segmenting App Groups allows all workloads to communicate across all services within each App Group.

When segmenting App Groups, the Policy Generator creates one ruleset per App Group. The ruleset includes a rule that covers traffic for all workloads to all workloads on all services.

1. From the PCE web console menu, choose **Policy Generator**.

The Select App Group page appears. The page displays when the Policy Generator last calculated the coverage for each type of Rule. Click the refresh icon to recalculate rule coverage.

2. In the Select App Group down-down menu, select **Segment Multiple App Groups** from the bottom of the list.

The Choose App Groups page appears.

3. Select the App Groups to segment and click **Next**.




TIP:

- To recalculate rule coverage for an App Group, hover over the Last Calculated column and click the refresh icon. The column displays the time that the rule coverage was calculated.

The column indicates whether the ruleset for the group has been edited since the last calculation and triggers you to recalculate it.

- To quickly select App Groups using different criteria, click the arrow icon to the right of the Name column:



- The Choose App Groups page displays all your App Groups regardless of their percentage of rule coverage or whether they have connections. For example, the page displays App Groups that have 100% rule coverage and groups with zero connections.

4. To accept the proposed rules, click **Save** and **OK**.

The Policy Generator Successful message appears, which displays the number of new rules. The rules are added to a draft ruleset.

Segmentation Rulesets

You can use segmentation rulesets to write policy so the workloads in your application can communicate with each other. A segmentation ruleset consists of rules and scopes:

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

Segmentation Ruleset Scope

The scope of a segmentation 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 segmentation 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 segmentation ruleset.

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

- **Application:** To what 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 located—either physically (for example, rack server or AWS) or geographically (for example, US, EU, or CA)?



NOTE:

The Role label should not be used in the scope.

For example, 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 be used to create complex security policies.

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

Scope		
App	Env	Loc
HRM	Prod	US
Rules		
Providers	Services	Consumers
DB	MySQL	App
App	Tomcat	Web
Web	Apache	Corp-HQ

Allows the following communication:

- HRM | Prod | US | DB ← HRM | Prod | US | App
- HRM | Prod | US | App ← HRM | Prod | US | Web
- HRM | Prod | US | Web ← HRM | Prod | US | Corp-HQ

Single Ruleset Scopes

Using a single scope in a segmentation 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, or Location). 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	Env	Loc
HRM	All	US
Rule		
Providers	Services	Consumers
DB	MySQL	App

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

- HRM | Anything | US | DB ← HRM | Anything | US | App
- Or
- HRM | Dev | US | DB ← HRM | Dev | US | App

- HRM | Dev | US | DB ← HRM | Prod | US | App
- HRM | Prod | US | DB ← HRM | Dev | US | App
- HRM | Prod | US | DB ← HRM | Prod | US | App

(Assuming that “Dev” and “Prod” are types of Environment labels.)

Multiple Ruleset Scopes

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

For example, consider the following segmentation ruleset:

Scopes		
App	Env	Loc
HRM	Prod	US
HRM	Dev	US
Rule		
Providers	Services	Consumers
DB	MySQL	App

This rule and scope states: “Workloads using the HRM application in the Prod environment in the US can initiate communications between the DB and the App and workloads using the HRM application in the Dev environment in the US can initiate communications between the DB and the App,”

The rule and scope do *not* state: “Workloads using the HRM application in the Prod and Dev environment in the US can initiate communications between the DB and the App”

This example allows the following communication:

- HRM | Prod | US | DB ← HRM | Prod | US | App

And

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

But *not*

- HRM | Prod | US | DB ← HRM | Dev | US | App

Combine Labels in Scopes and Rules

When the same type of label is used multiple times in a rule, they are expanded as multiple rules with one label for each rule.

The following examples further demonstrate how scopes work with rules.

The following segmentation ruleset:

Scope		
App	Env	Loc
HRM	All	US
Rules		
Providers	Services	Consumers
Prod	MySQL	Dev
DB	MySQL	DB

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:

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

The following segmentation ruleset:

Scope		
App	Env	Loc
All	All	US
Rules		
Providers	Services	Consumers
HRM	MySQL	ERP
Prod	MySQL	Dev
DB	MySQL	DB

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 | HRM | Prod | US ← DB | ERP | Dev | US

The following ruleset:

Scopes		
App	Env	Loc
All	Dev	US
All	Prod	EU
Rules		
Providers	Services	Consumers
HRM	MySQL	ERP

Scopes		
App	Env	Loc
All	Dev	US
All	Prod	EU
Rules		
Providers	Services	Consumers
DB	MySQL	DB

Allows the following communication:

- All | HRM | Dev | US ← All | ERP | Dev | US
- All | HRM | Prod | EU ← All | ERP | Prod | EU
- All | Dev | US | DB ← All | Dev | US | DB
- All | Prod | EU | DB ← All | Prod | EU | DB



NOTE:

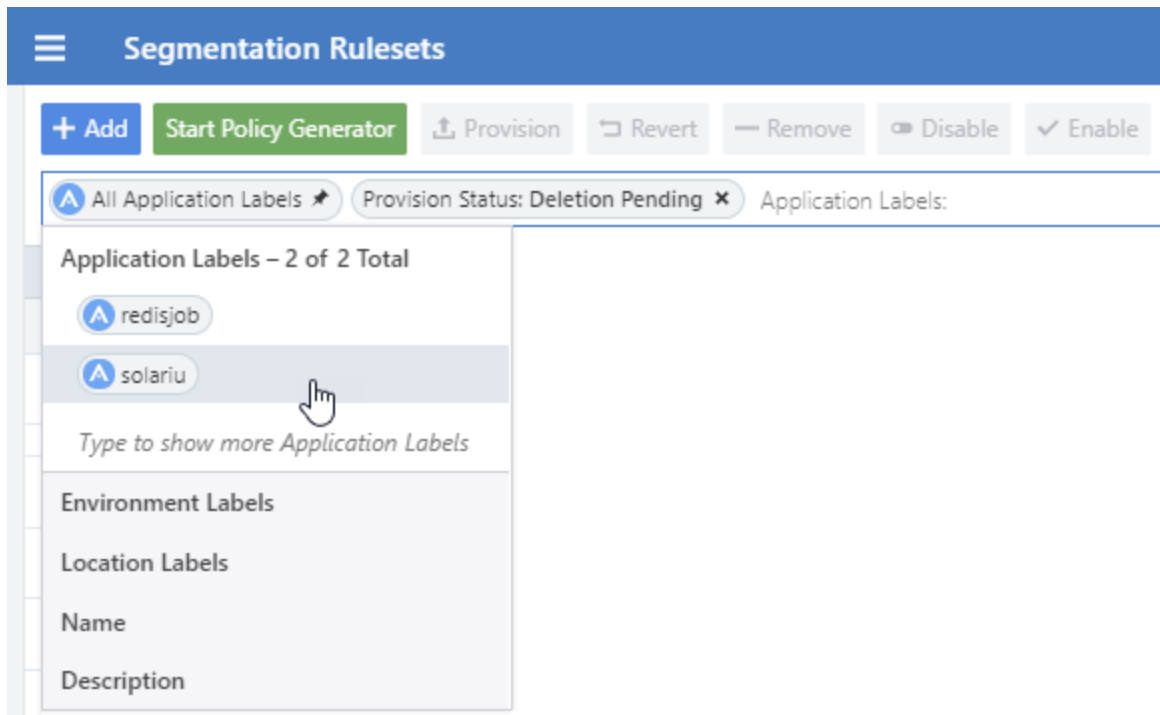
When the service in a rule is DNS, the consumer must be an [IP List](#).

Segmentation Ruleset Status

You can view the ruleset status on the Segmentation Ruleset page. The current status of each segmentation ruleset (enabled or disabled) is displayed in the Status column. When you change a segmentation ruleset but haven't provisioned the change yet, the type of change (addition, deletion, or modification) appears in the Provision Status column with the word "Pending" to indicate that these changes must be provisioned to be applied.

Filter the Segmentation Rulesets List

You can filter the rulesets list using the label and property filter at the top of the list. You can filter the list by entering a label type to show only those rulesets that use the selected labels. You can further filter the list by selecting specific properties of the rulesets. For example, you can filter the list by provision status, such as rulesets that are in draft state and have not yet been provisioned.



Create a Segmentation Ruleset

You can create a segmentation ruleset to write rules that define the allowed communication between workloads in a single group or multiple groups. See [Groups in Illumination](#) in the Visualization Guide for information.

When you write a rule for a Windows workload, you can add a Windows service name without specifying a port or protocol and 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, Illumio recommends splitting the rules across multiple rulesets, or use the Illumio Core;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. To apply the rules to another group, add a second scope, which is indicated by the group's Application, Environment, and Location labels.

To create a segmentation ruleset:

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

The Segmentation Rulesets page appears.

2. Click **Add**.
3. Enter a name for the segmentation ruleset.
4. Select Application, Environment, and Location labels for the ruleset.

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

5. Click **Save**.

Now that the ruleset is created, you can add rules to define your security policy. See [Rules](#) for information about the types of rules you can add.

Create a Ruleset with Multiple Scopes

You can create rulesets with multiple scopes to define the allowed communication between workloads in one or more groups. See [Groups in Illumination](#) in the Visualization Guide for information.

How you define the scope in a segmentation ruleset enables you to write rules for workloads in multiple groups (two or more scopes). Each scope corresponds to one group. The scope defines the boundaries of the rules in the segmentation ruleset.

To create a multi-scope segmentation ruleset:

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


The Segmentation Rulesets list page appears.

2. Click **Add**.
3. Enter a name for the ruleset.
4. In the *Scope* section, set the Application, Environment, and Location label by selecting the them from the drop-down lists.
5. After you select the three labels, click **Save**.

The page refreshes and the Scopes and Rules tab appears.



NOTE:

To edit the Scope, click the Edit button .

- To add another scope, click the Add icon (+).
A new row appears in the scopes section.
- Set the Application, Environment, and Location labels for the new scope and click the Save icon at the end of the row.

The green Addition Pending icon that this addition is pending, so you need to provision the new segmentation ruleset in order for the rule to take effect. See [Provisioning](#) for more information.



NOTE:

This task contains the steps to define multiple-scopes in the segmentation ruleset. For information about rules to the ruleset, see [Rules](#).

Duplicate a Segmentation Ruleset

When you have a ruleset that you want to use to create other new rulesets, you can duplicate an existing ruleset.

- From the PCE web console menu, choose **Rulesets and Rules > Segmentation Rulesets**.
The Segmentation Ruleset list page appears.
- Click the **Scopes and Rules** tab, and then click **Duplicate Ruleset**.
The Duplicate Ruleset page appears.
- Rename the copy of the ruleset.



NOTE:

The default name is “Copy of [Ruleset Name]” (where [Ruleset Name] is the name of the original Ruleset).

- Click **Save**.

After saving the new duplicate ruleset, make any needed scope or rule changes and then provision to apply them. See [Provisioning](#) for more information.

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 provider of the service, and a consumer of the service. You also need to select the type of rule:

- **Intra-scope rule:** Allow communication within a group. The segmentation ruleset scope applies to both providers and consumers.
- **Extra-scope rules:** Allow communication between groups. The segmentation ruleset scope applies only to the providers.
- **Custom iptables rules:** Allows custom iptables configurations in a segmentation 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 segmentation ruleset matches the scopes where they have administrator privileges. They cannot create or manage segmentation 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 segmentation ruleset.

You can combine multiple types of rules (intra-scope, extra-scope, and custom iptables) in a single segmentation ruleset. They can be used to either [Create a Segmentation Ruleset](#) or [Create a Ruleset with Multiple Scopes](#).

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.

Enforcement Modes for Rules

The Illumio policy model follows an allowlist model. Basically, all communication between workloads is denied unless explicitly allowed by Illumio security policy. Users create segmentation rules to allow traffic between their workloads. For information about the allowlist model, see [Enforcement States](#) and [Understanding Segmentation Rulesets and Rules](#).

This method of controlling traffic ensures secure communication between your workloads. However, as you work toward applying the allowlist model for security policy, you might choose a more targeted approach to applying security policy. In addition to creating segmentation rules for your workloads, you can create selective enforcement rules.

Selective enforcement rules can work alongside segmentation rules. They separate policy enforcement and visibility states per workload. Selective enforcement policies are based on one or more labels or label groups.

Using selective enforcement of policy, you can protect a subset of your services and ports on your managed workloads. The other ports on the workload remain in visibility-only mode and function as if the entire workload is in visibility-only mode. In addition to gradually expanding your policy enforcement envelope, selective enforcement is useful for temporarily enforcing policy on specific ports in case a vulnerability is detected and you need to take action quickly.

Another way to think of selective enforcement of security policy is as an intermediate enforcement mode on the workload:

IDLE → **VISIBILITY** → **SELECTIVE** → **FULL**

In this intermediate enforcement mode, label-based rules designate the workloads and the services/ports that need to be enforced; while other services and ports are in visibility-only mode. Policy enforcement is applied only on the provider side (ingress traffic) of the rules.

For more information about visibility modes, see [Workload Enforcement States](#) and [Visibility Level](#) in this guide and [Set Group Enforcement](#) in the *Visualization Guide*.

Limitations for Selective Enforcement of Policy

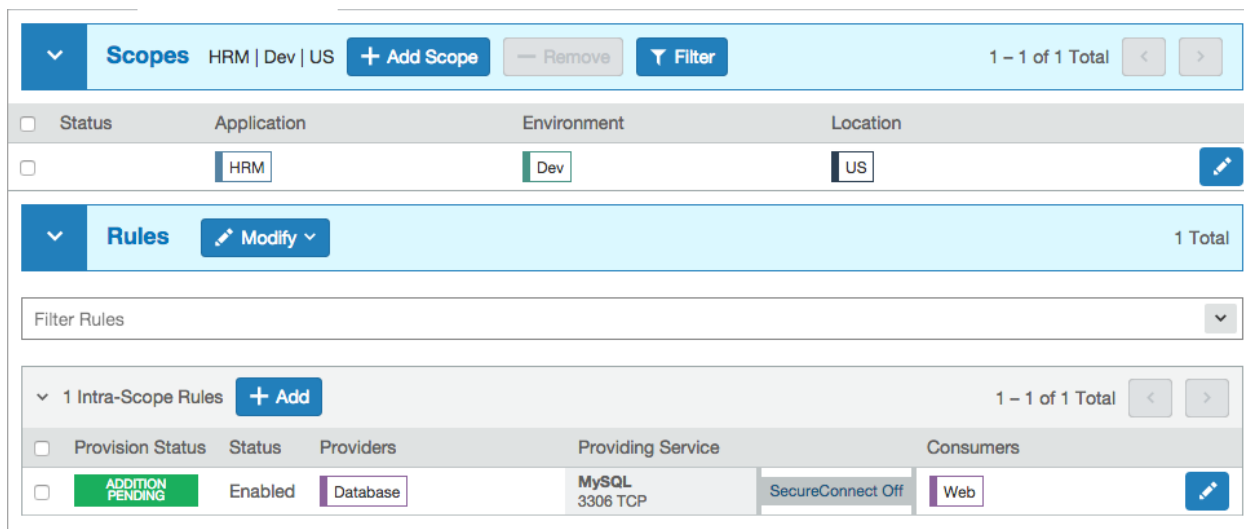
- Selective enforcement is provider-centric and applies to inbound connections only.

- Selective enforcement only applies to managed workloads; it is not supported for NEN-controlled or other unmanaged workloads.
- Virtual Services are enforced at the workload level. As a result, selective enforcement does not affect virtual services directly; instead, selective enforcement affects the workloads they are comprised of.
- Selective enforcement policies control which services/ports are enforced on workloads. Selective enforcement policies do not affect and are completely independent of SecureConnect, AdminConnect (referred to as Machine Authorization in the PCE web console), and stateless rules.

Intra-scope Rules

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.

Example:



The screenshot shows the PCE web console interface. At the top, there is a 'Scopes' section with a dropdown arrow, the text 'Scopes HRM | Dev | US', and buttons for '+ Add Scope', '- Remove', and 'Filter'. Below this is a table with columns: Status, Application, Environment, and Location. The table contains one row with values: [checkbox], HRM, Dev, US, and an edit icon.

Below the scopes section is a 'Rules' section with a dropdown arrow, the text 'Rules', a 'Modify' button, and '1 Total'. Below this is a 'Filter Rules' dropdown menu.

Below the filter is a section for '1 Intra-Scope Rules' with an '+ Add' button and '1 - 1 of 1 Total' with navigation arrows. Below this is a table with columns: Provision Status, Status, Providers, Providing Service, and Consumers. The table contains one row with values: [checkbox], ADDITION PENDING, Enabled, Database, MySQL 3306 TCP, SecureConnect Off, Web, and an edit icon.

In this example, 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

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.

Example:

The screenshot shows the 'Scopes' section of the illumio interface. It features a header with 'Scopes HRM | Dev | US', '+ Add Scope', '- Remove', and 'Filter' buttons. Below the header is a table with columns: Status, Application, Environment, and Location. The table contains one row with values: , HRM, Dev, and US. Below this is the 'Rules' section, which has a 'Modify' button and '2 Total' rules. A 'Filter Rules' dropdown is present. Underneath, there are sections for '1 Intra-Scope Rules' and '1 Extra-Scope Rules'. The 'Extra-Scope Rules' section includes a table with columns: Provision Status, Status, Providers, Providing Service, and Global Consumers. The table contains one row with values: ADDITION PENDING, Enabled, Database, MySQL 3306 TCP, SecureConnect Off, and Web.

In this example, 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." For more information, see [IP Lists](#) and [Role-based Access Control](#) in the *PCE Administration Guide*.

Custom iptables Rules

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

The screenshot shows the Illumio management interface. At the top, there's a 'Scopes' section with a dropdown menu, 'HRM | Dev | US', '+ Add Scope', '- Remove', and 'Filter' buttons. Below this is a table with columns: Status, Application, Environment, and Location. The first row shows 'HRM', 'Dev', and 'US'. Below the table is a 'Rules' section with a 'Modify' button and '2 Total' rules. There are three expandable sections: '1 Intra-Scope Rules', '1 Extra-Scope Rules', and '0 Custom iptables Rules'. The 'Custom iptables Rules' section is expanded, showing a form with a dropdown for 'Web', 'IPv6', and a text input field for the rule. Below the form is a table with columns: Provision Status, Receivers, IP Version, and iptables Rules applied to Scope. The table is currently empty, displaying 'No custom iptables Rules to display'.

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
<code>raw</code>	<code>prerouting</code> , <code>output</code>	No
<code>mangle</code>	<code>prerouting</code> , <code>input</code> , <code>output</code> , <code>forward</code> , <code>postrouting</code>	Yes
<code>nat</code>	<code>prerouting</code> , <code>output</code> , <code>postrouting</code>	Yes
<code>filter</code>	<code>input</code> , <code>output</code> , <code>forward</code>	Yes
<code>security</code>	<code>input</code> , <code>output</code> , <code>forward</code>	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." For more information, see [Role-based Access Control](#) in the *PCE Administration Guide*

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 Linux 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.

If you add a stateless rule to a policy that has both Windows and Linux workloads in its scope, then that rule is configured as a stateful rule on the Windows workload.

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

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.

Segmentation Rule Search

When you have a large number of rules organized in segmentation 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 segmentation 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 **Rulesets and Rules > Segmentation Rule Search**.

The Segmentation Rule Search page appears.

2. Search for Active or Draft rules.

3. Perform a Basic or Advanced search of your rules:
 - Basic: Searches all attributes
 - Advanced: Searches by provider, consumer, 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.

4. 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).
5. Click the *Column* drop-down list to select the attributes you want to be displayed in the search results.
6. Filter options to further narrow your search.
7. Under the Ruleset column, select a ruleset and make changes to the rules.
8. Click **Download** to download the results of your search.

You can download up to 500 rules in the CSV format.

Policy Check

The Policy Check feature allows you to determine if a rule allowing communication between workloads or 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.

**NOTE:**

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

For example, you have created several segmentation 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 **Troubleshooting > Policy Check**.
2. In the *Consumer* field, type or select a workload or IP address.
3. In the *Provider* field, type or select a workload or IP address.

- In the *Provider 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.
- 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.

Policy Check
Jane Doe

Verify if Rules exist that allow connections between Workloads or IP addresses

Provider
Port and Protocol
Consumer

perf-workload-1637 x

67 UDP

perf-workload-1638 x

Check Rules

✓

The Rules below allow this connection

Rulesets	Rules
5	5

Ruleset [unmanaged-workloads@1515293530.25322](#) All Applications All Environments All Locations

Status	Providers	Service	Consumers	Note
	dhcp	dhcpcd 67 UDP, 67 TCP	All Workloads	

Ruleset [unmanaged-workloads@1515293592.66112](#) All Applications All Environments All Locations

Status	Providers	Service	Consumers	Note
	dhcp	dhcpcd 67 UDP, 67 TCP	All Workloads	

Ruleset [unmanaged-workloads@1515293653.17277](#) All Applications All Environments All Locations

- NOTE:**
The status column does not display any values. For more information, you can do a detailed search using the [Segmentation Rule Search](#) feature.

When a rule does not exist, the page displays “No Rules exist to allow that connection.”

Policy Check

Verify if Rules exist that allow connections between Workloads or IP addresses

Provider: perf-workload-1636

Port and Protocol: 22 TCP

Consumer: perf-workload-1645

Check Rules

No Rules exist to allow this connection

Change the criteria or add a Rule to allow this connection

Rule Writing

This topic explains how to create the different types of rules in the Illumio Core. For descriptions of the types of rules, see [Rules](#).



TIP:

You can also use the Illumination map to write rules. For information, see [Write a Group-Level Segmentation Rule](#) in the *Visualization Guide*.

Create an Intra-Scope Rule

Intra-scope rules allow communication within a group. The ruleset scope applies to both providers and consumers. For more information about intra-scope rules, see [Intra-scope Rules](#).

1. If necessary, create a Segmentation Ruleset or open an existing one. See [Segmentation Rulesets](#) for information.
2. In the *Intra-Scope Rules* section, click the **Add** icon (+).
3. From the *Consumers* drop-down list, select or type the name of the consumer of the service.



NOTE:


The consumer must match the Segmentation Ruleset scope.

4. In the *Providers* drop-down list, select or type the name of the provider of the service (for example, Database) . You can select from a range of entity types that match the scope, such as an individual workload, a virtual service, or an unmanaged workload.

5. From the *Providing Service* drop-down list, select a service (for example, PostgreSQL).



NOTE:
Only one service or all services can be selected.

6. (Optional) To enable SecureConnect, select **SecureConnect** from the *Providing Service* drop-down list. For more information, see [SecureConnect](#).
7. After completing your selections, click the **Save** icon () at the end of the row for that rule.



NOTE:
To edit this rule, click the **Edit** icon at the end of the row.

After adding a rule, the Status column displays a green Addition Pending icon. To enforce this rule, you must provision the change. For more information about provisioning, see [Provisioning](#).

Create an Extra-Scope Rule

Intra-scope rules allow communication within a group. The ruleset scope applies to both providers and consumers. For more information, see [Extra-scope Rules](#).

1. If necessary, create a Segmentation Ruleset or open an existing one. See [Segmentation Rulesets](#) for information.
2. In the *Extra-Scope Rules* section, click the **Add** icon (+).
3. From the *Consumers* drop-down list, select the consumer of the service.




NOTE:
The consumer does not need to match the Segmentation Ruleset scope.

4. In the *Providers* drop-down list, select or type the name of the provider of the service (for example, Database). You can select from a range of entity types that match the scope, such as an individual workload, a virtual service, or an unmanaged workload. For a full list of permitted provider and consumer combinations in a rule, see [Permitted Rule Writing Combinations](#).
5. From the *Providing Service* drop-down list, select a service (for example, PostgreSQL).



NOTE:
Only one service or all services can be selected.

- (Optional) To enable SecureConnect, select **SecureConnect** from the *Providing Service* drop-down list. For more information, see [SecureConnect](#).
- After completing your selections, click the **Save** icon () at the end of the row for that rule.



NOTE:
To edit this rule, click the **Edit** icon at the end of the row.

After adding a rule, the Status column displays a green Addition Pending icon. To enforce this rule, you must provision the change. For more information about provisioning, see [Provisioning](#).

Create a Custom iptables Rule

Custom iptables rules allow you to integrate existing iptables into a ruleset. For more information about custom iptables rules, see [Custom iptables Rules](#).



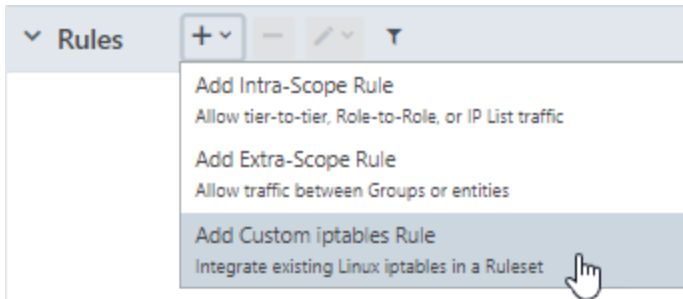
NOTE:
Creating custom chains is not supported.

About Custom iptables Rules

- Receivers column:** Shows the labels representing the resource that receives the custom iptables rule.
- IP Version column:** Specifies the IP version used for this traffic.
- iptables Rules applied to Scope column:** Contains the entire iptables.

To add a custom iptables rule:

- If necessary, create a Segmentation Ruleset or open an existing one. See [Segmentation Rulesets](#) for information.
- From the Rules drop-down menu, select the Custom iptables Rule.



3. In the *Type or select a label by name* drop-down list, select the entity or entities that will receive the iptables rules by selecting or typing a label name.


**NOTE:**

More than one label can be selected. When you select labels as receivers, the custom iptables rules are sent only to workloads matching those labels and not virtual services or virtual servers.

4. From the drop-down *IP Version* drop-down list, select the IP version (IPv4 or IPv6).
5. Type or paste the iptables commands into the *Type or paste a custom iptables rule* field. Supported iptables display in green. Unsupported iptables or iptables with errors display in red.

**NOTE:**

The iptables commands must begin with `-t`. To delete a row, use **Shift+Delete**.

6. After completing your selections, click the **Save** icon () at the end of the row for that rule.

**NOTE:**

To edit this rule, click the **Edit** icon at the end of the row.

After adding a rule, the Status column displays a green Addition Pending icon. To enforce this rule, you must provision the change. For more information about provisioning, see [Provisioning](#).

When you provision a new custom iptables rule, the VEN performs basic validation before applying on the Linux workload host firewall. If this validation test fails, the VEN will log an event and switch to an Error State. If the validation is successful, the VEN installs the custom iptables rules before the last default rule.

**NOTE:**

Ordering is not guaranteed across custom iptables rules. Any iptables rule that is closely tied to or depends on other iptables rules must be written as part of the same rule. For example, when you have three iptables rules to allow ICMP Types 3, 8, 13 and another rule to drop other types of ICMP traffic, all four of these iptables rules must be a part of the same ruleset.

Create a Selective Enforcement Rule

For information about selective enforcement, see [Enforcement Modes for Rules](#) in this guide.

You can add a new selective enforcement rule to your managed environment or you can change the enforcement mode for existing workloads.

To add a new selective enforcement rule:

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

The Selective Enforcement Rules page appears.

2. Click **Add**.
3. Enter a name for the rule.
4. From the Enforced Services drop-down menu, select the services and ports (or port ranges) that you want covered by the rule.
5. Select Role, Application, Environment, and Location labels for the rule.

These labels define the scope of the selective enforcement rule, which is its range or boundary. The labels define the workloads affected by this rule, which is all workloads that share the same labels in their scopes.

6. Click **Save**.

The page refreshes and the rule summary displays which workloads are impacted by the selective enforcement rule.

The screenshot shows the configuration page for a selective enforcement rule named "DB-internal-ports". The page has a blue header with a menu icon, a back arrow, and the title "Selective Enforcement Rules – DB-internal-ports". Below the header are three tabs: "Summary" (selected), "Segmentation Rules", and "Explorer". A yellow warning banner at the top states "You are viewing the draft version (Addition Pending)". Below the banner are "Edit" and "Remove" buttons. The "General" section shows the rule name "DB-internal-ports" and "Service or Ports" as "45-67". The "Labels" section lists four categories: "Role" (All Roles), "Application" (All Applications), "Environment" (All Environments), and "Location" (Amazon). The "Matching Workloads by Enforcement Mode" section shows a table with the following data:

Enforcement Mode	Count
Idle	0
Visibility Only	6
Selective Enforcement	3
Full Enforcement	89

7. To apply the selective enforcement rule to the workloads, provision the new rule. See [Provisioning](#) for information.

To change the enforcement mode for existing workloads:

For example, you want to change the enforcement mode to selective for a specific port, perform the following steps.

1. From the PCE web console menu, choose **Workloads and VENS > Workloads**. The Workloads page appears.
2. Select the workloads that have traffic accessing the port you want enforced.
3. From the Enforcement drop-down list, choose **Selective**.

A confirmation dialog box appears listing the impacted workloads.

4. Click **OK**.
5. To apply the enforcement mode change to these workloads, provision the mode change. See [Provisioning](#) for information.

Write Multicast Rules

You can write rules to allow multicast traffic between workloads by writing two rules that follow a very specific workflow.

Multicast Use Case 1

For example, you want some database workloads labeled DB to have multicast for data replication and you want to allow the multicast traffic.

To do this:

1. Create an unmanaged workload or an IP list to represent the multicast group IP address (for example, mDNS Group: 239.0.0.251).
2. Create a service with port (for example, mDNS: UDP 5353).
3. In the ruleset, create these two rules:

Rule	Provider	Service	Consumer
1	mDNS Group	mDNS Group	DB
2	DB	mDNS Group	DB

In Rule 1, the consuming entity DB allows outbound packets from DB *to* 239.0.0.251.

In Rule 2, the mDNS group allows inbound packets *from* DB.

Multicast Use Case 2

For example, you want to ensure that the DB workloads receive a multicast feed on 224.5.5.5:5800 (multicast source).

To do this:

1. Create an unmanaged workload or an IP list to represent the multicast source (for example, Stock-Feed-Group: 224.5.5.5).
2. Create a service with the correct port (for example, Stock-Feed-Service: UDP 5800).
3. In the ruleset, you create these two rules:

Rule	Provider	Service	Consumer
1	Stock-Feed-Group	Stock-Feed-Service	Multicast-Source
2	DB	Stock-Feed-Service	Multicast-Source

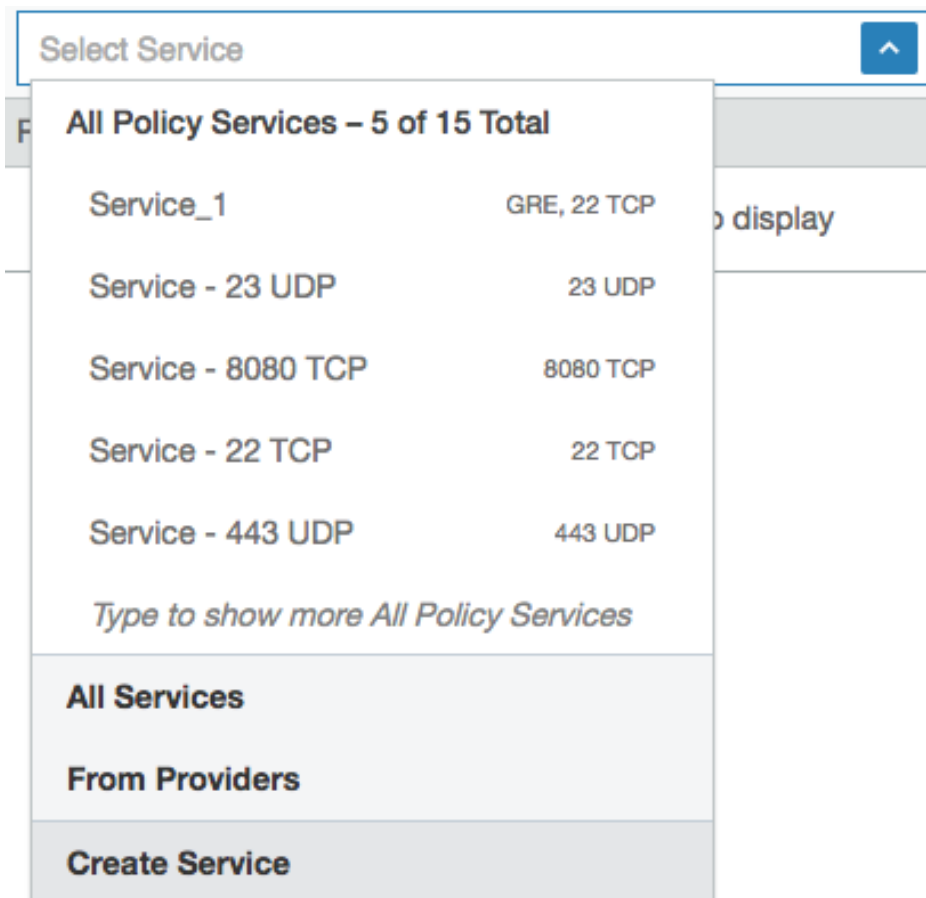
In Rule 1, the Stock-Feed-Service allows outbound packets from Multicast-Source to 224.5.5.5.

In Rule 2, the provider DB allows inbound packets *from* Multicast-Source.

Create Service While Creating Rule

To make rule writing easier, you can create a new service in a ruleset as you are writing rules.


1. Create an [Extra-Scope](#) or an [Intra-Scope Rule](#).
2. In the *Providing Service* drop-down list, select **Create Service** at the end of the list.






3. In the Create Service pop-up that appears, enter a name for the service in the *Name* field and optionally a description in the *Description* field.

4. In the Attributes section, choose whether you want to create a [Port-Based](#) or Windows [Service-Based](#) service.
5. In the Ports section, enter the ports (including any UDP ports) used by the service. To enter a range, separate the port numbers by a hyphen (-). You can also copy and paste lists of services. To delete a row, use **Shift+Delete**.
6. Click **OK**.

Tips for Managing Rules

- To modify an existing rule, click the edit icon () at the end of the rule row.
- To modify or remove an existing rule, open the Edit menu for that rule at the end of the row.



- To remove multiple rules, select their checkboxes and click the remove (-) icon () in the **Rules** header row at the top of the page.
- To enable or disable multiple rules, select their checkboxes and click the edit icon in the **Rules** header row at the top of the page.
- To filter your existing rules, click the Filter icon () in the **Rules** header row at the top of the page. The filter drop-down menu appears. Click the drop-down list and select an option to filter rules by label, IP lists, label groups, virtual services, virtual servers, workloads, user groups, services, All workloads, or Any (0.0.0.0/0 and ::/0). If there are no rules that match the selected criteria, a message appears indicating that no rules match your filter criteria.
- After creating or modifying a rule, an icon appears in the Provision Status column indicating the current provisioning status of the rule (for example, Addition Pending or Removal Pending).

Add a Note to a Rule

You can add a note to a rule to document more information about that rule for context. The note is visible to all users in the organization, but can only be edited by users


with Ruleset Manager privileges for the ruleset.

**NOTE:**

You must provision the changes after adding a note to a rule.

1. Select a rule on the Segmentation Rulesets page.
2. Click the **Edit** button and select **Add Note**.
3. Enter the note in the drop-down entry field that appears. You can enter up to 255 characters.
4. Click **Save**. You must provision the changes to confirm the note.

Details:

- To indicate the rule contains a note, the following icon is displayed in the Note column: 
- To edit an existing note, select the note icon. The entry field displays the existing text. Make any needed changes, then click the Save icon in the lower-right to save the changes to the note.

Duplicate a Rule

1. Select the ruleset on the Segmentation Rulesets page.
2. Select the drop-down list next to the **Edit** button of the rule to be duplicated.
3. Select **Duplicate**. The rule is duplicated in Edit mode, allowing you to make any needed changes.
4. Click **Save**.

After saving the duplicate rule, you must [provision](#) the ruleset changes to apply them.

Reverse a Rule

To expedite the rule writing process, you can duplicate and reverse existing rules. The entity selected as the provider in the original rule will be the consumer in the reversed rule and the entity selected as the consumer in the original rule will be the provider.

Caveats:

- Only intra-scope rules are supported. Extra-scope and custom iptables rules cannot be reversed.
- Only rules that use the following resources are supported: Labels, label groups, workloads, IP lists, All workloads, and Any.

- When you do not have sufficient privileges due to RBAC, an error message displays.
- Only one rule can be reversed at a time.
- When the original rule is disabled, the reversed rule is disabled as well.

To reverse (swap Providers and Consumers) in a rule:

1. Select the ruleset on the Segmentation Rulesets page.
2. Select the drop-down list next to the **Edit** button of the rule to be reversed.
3. Select **Reverse**. The rule is reversed in Edit mode allowing you to make any needed changes.
4. Click **Save**.

After saving the reversed rule, you must [provision](#) the ruleset changes to apply them.

Reorder Rules

Ruleset owners have the ability to rearrange rules in a specific order to improve readability on the Segmentation Rulesets details page. Different rule types can be reordered independently.

After reordering the rules, you must provision the changes for them to take effect. Rearranging rules does not affect the order in which they are enforced in the policy.



NOTE:

You can only reorder rules in rulesets that you own. For more information, see [Role-Based Access Control](#) in the *PCE Administration Guide*

To customize the arrangement of the rules, click **Reorder Rules** on the Scopes and Rules tab of the Segmentation Ruleset details page.

When you hover over a rule, it is highlighted in the list. To move it, drag and drop the rule to its new location in the list. The other rules are rearranged to accommodate the move. When you place the rule in its new location, the numbers of the rules are reassigned to reflect the new order. If you delete a rule, it remains in place but is appended with “Deletion Pending.” When you have finished rearranging the rules, click **Save Order** to confirm the new order for the rules.



NOTE:

If more than one user is reordering the rules at the same time, the most recent changes are saved.


FQDN-Based Rules

Applications across datacenters 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 as well as 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-zA-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 a 100 entries.

FQDN-Based Rule Requirements and Limitations

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

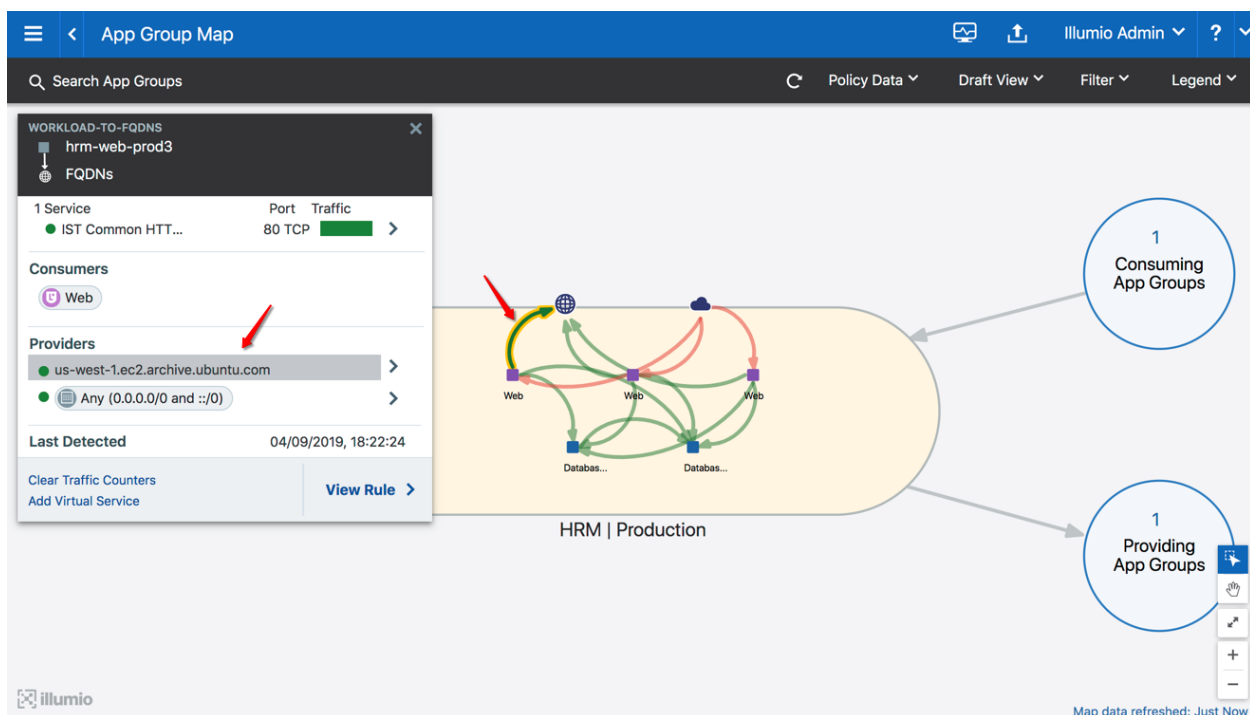
- Requires Illumio Core 19.1.0 or later.
- Supported for any Linux OS that is supported with the Illumio VEN 19.1.0 release.
- Supported for any Windows OS that is supported with the Illumio VEN 19.1.0 release.
- Solaris and AIX workloads are not supported.
- Visibility and enforcement for DNS-based traffic where the source is a DNS hostname is not supported.
- FQDNs can be described in IP lists or virtual services, but not in an unmanaged workload interface.
- When using virtual services, only one FQDN (wildcard supported) can be specified. 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 (namely, specifying only `“*”`) is not allowed.

FQDN Visibility

Illumio does not require any new configuration to gain visibility into outbound traffic towards FQDNs. However, you can create Illumio policy objects to represent an FQDN or a list of FQDNs. In the following example, Illumination presents outbound FQDN flows when there are no policy objects 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 to represent the FQDN shown in this example.



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.

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

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

IP List - *.ec2.archive.ubuntu.com

Viewing draft version You are viewing the draft version. [View the active version.](#)

[Edit](#) [Remove](#)

General

Name *.ec2.archive.ubuntu.com

Description

Created 04/10/2019 at 10:21:09 by jason.williams@illumio.com

Last Modified 04/10/2019 at 10:21:09 by jason.williams@illumio.com

IP Addresses and FQDNs *.ec2.archive.ubuntu.com

Example of the traffic in Illumination:

App Group Map

Search App Groups

Policy Data Draft View Filter Legend

WORKLOAD-TO-FQDNS

hrm-web-prod3

FQDNs

Service	Port	Traffic
IST Common HTT...	80 TCP	<div style="width: 100%;"></div>

Consumers

- Web

Providers

- Any (0.0.0.0/0 and ::/0)
- *.ec2.archive.ubuntu.com**

Last Detected 04/09/2019, 18:22:24

[Clear Traffic Counters](#) [View Rule](#)

HRM | Production

1 Consuming App Groups

1 Providing App Groups

Map data refreshed: Just Now

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

This example of a virtual service represents *.ec2.archive.ubuntu.com.





General

Name	Ubuntu Repo
Description	
Created	04/10/2019 at 10:34:26 by Illumio Admin
Last Modified	04/10/2019 at 10:34:26 by Illumio Admin

Connection Service Or Ports

Service or Ports	All Services
-------------------------	------------------------------

Labels

Role	 Ubuntu Repo
Application	 HRM
Environment	 Production
Location	 CA

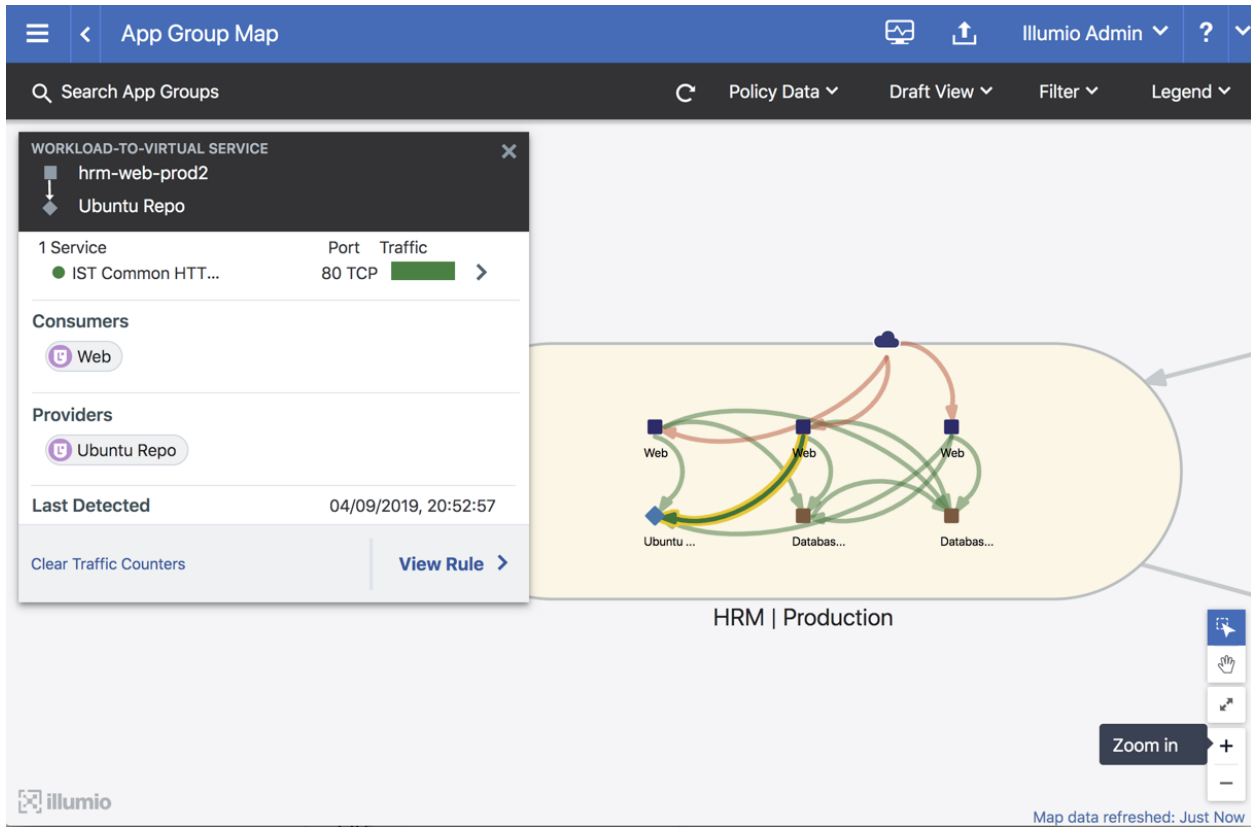
Address Pool

IP Addresses and FQDNs	*.ec2.archive.ubuntu.com
-------------------------------	--------------------------

Advanced

Pool Target	Host Only (Default)
--------------------	---------------------

Example of traffic to the virtual service in Illumination:



Write Policies to Allowlist FQDNs

IP List

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

See the following example:

Ruleset Scope Example			
Application	Environment	Location	
HRM	Production	All Locations	
Intra-Scope Rule Example			
Provider	Providing Service	Consumer	Note
*.ec2.archive.ubuntu.com (IP list object)	All Services	Web	You may also 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			
Provider	Providing Service	Consumer	Note
Ubuntu Repo (Virtual Service Role label for *.ec2.archive.ubuntu.com) + Uses Virtual Services only	Derived from Provider Virtual Service	Web	There are 2 objects selected in the Provider column - one is for the role label and the second is called "Uses Virtual Services only"

Ruleset - HRM | Production

Summary | **Scopes and Rules** | Start Policy Generator | Duplicate Ruleset

Viewing draft version Up to date | View the active version.

Scopes: HRM | Production | All Locations

Status	Application	Environment	Location
	HRM	Production	All Locations

Rules: 2 Total

2 Intra-Scope Rules | Reorder Rules | 1-2 of 2 Total

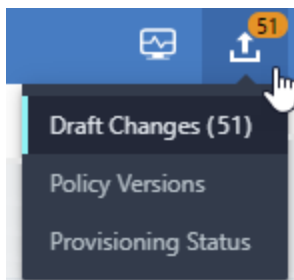
No.	Provision Status	Status	Providers	Providing Service	Consumers	Note
1	Enabled	Enabled	Ubuntu Repo Uses Virtual Services only	Derived from Pro...	Web	
2	Enabled	Enabled	All Workloads	All Services	All Workloads	

+ Extra-Scope Rule

Provisioning

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.



Items that Require Provisioning

The following security policy items must be provisioned before they can take effect:

- Rulesets
- Rule notes

- IP lists
- Services
- Label groups
- Security settings
- Virtual services
- Virtual servers

Provision All or Selected Items

When you create or make a change to security policy items (such as rulesets, IP lists, services, label groups, and security settings), you can provision the item immediately from the item page after you save the change.

You can click **Provision** button on the top PCE web console toolbar, which allows you to see all of the security policy changes that require provisioning. The list shows any items that have been modified (gray) or deleted (red), or added (green).

In the list of changes requiring provisioning, you can select all items, or select items individually to provision.

Change	Name	Item	Last Modified By	Last Modified On
<input checked="" type="checkbox"/>	DELETION PENDING -IPList3	IP List	illumio@illumio.com	06/25/2020, 09:50:29
<input checked="" type="checkbox"/>	MODIFICATION PENDING *.google.com	Virtual Service	illumio@illumio.com	09/24/2020, 16:31:17
<input checked="" type="checkbox"/>	ADDITION PENDING 9300 - 9301	Service	illumio@illumio.com	07/25/2020, 09:08:19
<input type="checkbox"/>	DELETION PENDING 90 TCP	Service	illumio@illumio.com	06/10/2020, 14:32:55

Dependencies for Partial Provisioning

When you select only some items to provision (rather than provisioning all policy item changes), some of those items might have dependencies that are also provisioned. Before you commit to the provision, the PCE shows you the items that are dependent and will also be provisioned.



NOTE:

You cannot partially provision resources with more than 500 dependencies. All changes must be provisioned at the same time.

Active vs Draft Versions

Any changes you make to security items, such as rulesets, services, IP lists, label groups, and security settings, need to be provisioned. All the changes you make to those items are considered to be in a “draft” state (un-versioned) until you provision them. After you provision your changes, those changes become the “active” version.

When you edit a security item that has been published at least once, and new changes have occurred since the last provisioning, you see a note at the top of the page that indicates the item is currently in draft state.

If you want to view the active version, click the **View the active version** link.

☰ 🏠 Virtual Services – *.google.coms

⚠️ You are viewing the draft version of Virtual Service [View the active version.](#)

Summary Workloads Container Workloads

✎ Edit ➖ Remove

General

Name	*.google.coms
Description	eeee
Created	01/31/2019 at 14:13:09 by <i>Jay Scott</i>
Last Modified	09/24/2020 at 16:31:17 by <i>Radhika</i>

Connection Service Or Ports

Service or Ports	<div style="display: flex; gap: 10px;"> <div style="border: 1px solid green; border-radius: 10px; padding: 2px 10px;">78 UDP</div> <div style="border: 1px solid green; border-radius: 10px; padding: 2px 10px;">67 TCP</div> <div style="border: 2px solid red; border-radius: 10px; padding: 2px 10px; color: red;"> ⚙️ TCP-UDP 80-TCP, 500-UDP, 1000-2000-TCP </div> </div>
------------------	--

Labels

Role	🌐 Web1
Application	📄 kafka1
Environment	🌲 Production
Location	📍 test

Address Pool

IP Addresses and FQDNs	*.google.com + 1.1.1.1
------------------------	---


Advanced

Pool Target	Host Only (Default)
-------------	---------------------

Provisioning Progress Indicator

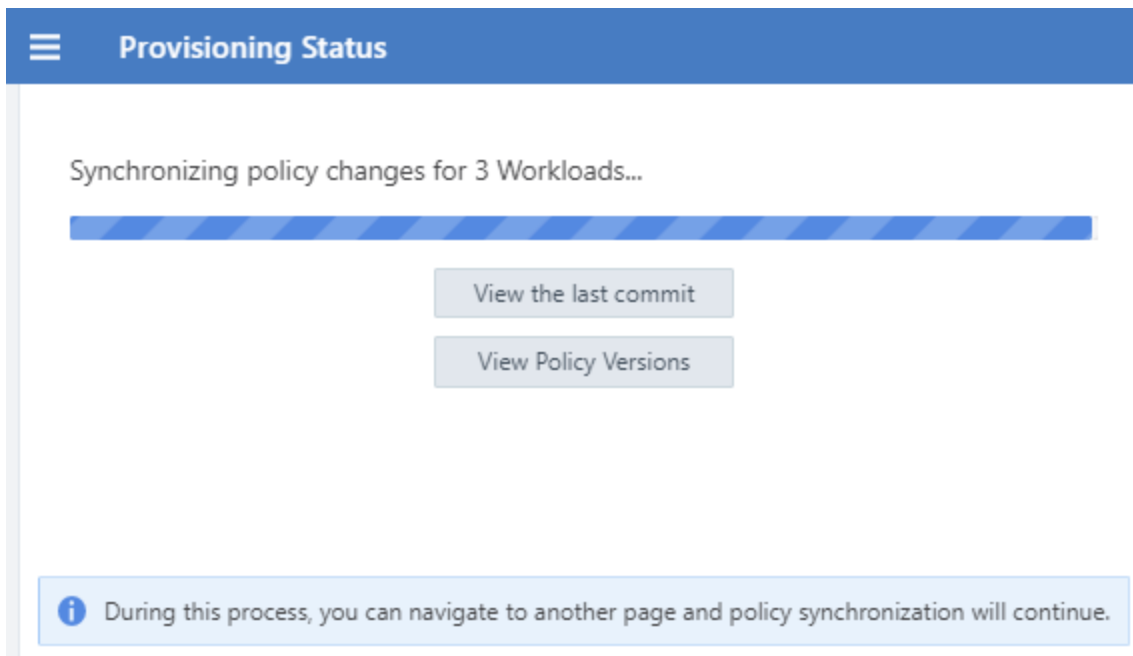
When you confirm provisioning by clicking **Confirm & Provision**, the Provisioning progress indicator displays the number of workloads that need to be synchronized with

the latest provisioned policy changes and the progress for applying the policy changes to those workloads.



On the Provisioning page, you can:


- View the previous policy change by clicking **View the last commit**
- View a list of provisioned changes by clicking **View Provision History**



**NOTE:**

If multiple subsequent policy changes have been provisioned, the number is the total number of workloads that have not yet received all provisioned policy changes, not just the most recently provisioned changes.

During this process, if you navigate to another page, the policy synchronization will continue and a window in the lower-right displays the number of workloads pending synchronization with the latest policy.



To return to the Provisioning page, click the window in the lower right corner or select **Provisioning** from the drop-down *Provisioning* list.

When the provisioning completes successfully, a confirmation message displays.

**NOTE:**

If multiple users simultaneously provision changes, the Provisioning progress indicator is updated to show the new changes, so all users will see the same Provisioning progress indicator.

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. From the PCE web console toolbar, click the **Provision** button and choose **Policy Versions**.

The Provision History page appears, which displays the history of the last provisions in your organization.

Version	Services	Virtual Servers	Selective Enforcement Rules	Restore
Workloads Affected	Label Groups	Secure Connect Gateways	Provisioned By	
Rulesets	Settings	Virtual Services	Provisioned On	
IP Lists			Note	
DRAFT	29 Services 10 Label Groups	1 Virtual Server 8 Secure Connect Gateways 11 Virtual Services	14 Selective Enforcement Rules Pending Provision	
48 Rulesets 29 IP Lists			This is the draft policy. Provision to make it the active version.	
ACTIVE 1530	304 Services 46 Label Groups	1 Virtual Server 22 Secure Connect Gateways 69 Virtual Services	14 Selective Enforcement Rules @illumio.com	
4 Workloads Affected 194 Rulesets 122 IP Lists	1 Setting		10/14/2020, 06:01:14	
1529	304 Services 46 Label Groups	1 Virtual Server 22 Secure Connect Gateways 69 Virtual Services	14 Selective Enforcement Rules @illumio.com	Restore
194 Rulesets 122 IP Lists	1 Setting		10/13/2020, 18:28:53	

- To view details about the changes, click one of the items. For the selected item, you can see the changes that were provisioned in this version.

Policy Versions

Version	Workloads Affected	Rulesets
DRAFT		48
ACTIVE 1530	4	194
1529		194

Policy Versions - Version #1529 vs. Version #1528

Restore 1 Settings

Provision Status	Type	Provisioned On	Provisioned By
MODIFIED	Settings	10/13/2020, 18:28:53	@illumio.com
	Security Settings		@illumio.com

Provision Changes

If you have made any changes to provisionable objects, such as rulesets, IP lists, services, label groups, and security settings, you need to provision those changes before they can take effect.

1. From the PCE web console toolbar, click the **Provision** button > **Draft Changes**.
The Draft Changes page appears, which displays a list of all policy items that have been added, modified, or removed. The top of the page shows a summary of changes based on item type.
2. Select one, several, or all the items you want to provision.
3. Click **Provision** to see a preview of the changes that will occur when you provision them.

**NOTE:**

When you selectively choose items to provision, some of those items might have dependencies that are also published. Any object dependencies are also be provisioned.

4. You can add a note to the provision. If a note is mandatory, the **Confirm & Provision** button is grayed out until you enter text in the field. After you enter appropriate text in the field the button is enabled.

For information about making provisioning notes mandatory, see [Provisioning Note Setting](#).

5. Click **Confirm & 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 you provision them.

1. From the PCE web console toolbar, click the **Provision** button > **Draft Changes**.
The Draft Changes page appears, which lists all security policy items have been added, modified, or removed. You also see a summary of changes based on item type.
2. Select individual items to revert or you can revert all changes.
3. Click **Revert**.

Restore Policy

With the policy restore feature, you can revert to an older version of the policy when the newly provisioned policy did not work as expected.



NOTE:

You need to be a Global Administrator or Global Organization Owner to use this feature.

The older version of the policy is copied to the current working draft version. You can immediately provision it to replace the version that is not working.

When there are pending changes, you cannot restore to a previous version. 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. Choose **Provision > Policy Versions** from the PCE web console menu or from the top-right provision menu.

The policy versions are displayed under the **Version** column.

2. Click **Restore** for the policy version that you want to revert to.


Version	Workloads Affected	Rulesets	Services	Settings	Secure Connect Gateways	Provisioned By	Note	Restore
DRAFT	3	3 Rulesets	2 Services	1 Setting	1 Virtual Service	Pending Provision	This is the draft policy. Provision to make it the active version.	Restore
ACTIVE	55 Workloads Affected	1 IP List	2 Services	1 Setting		01/25/2019, 11:13:25		Restore
1		1 Ruleset	2 Services	1 Setting	1 IP List	01/17/2019, 14:25:21	System created default	Restore

3. Click **Save as Draft** to restore the policy to the selected version.

Version	Workloads Affected	Rulesets	Services	Settings	Secure Connect Gateways	Provisioned By	Note
ACTIVE	3	3 IP Lists	2 Services	1 Setting		01/25/2019, 11:13:21	
SELECTED	55 Workloads Affected	1 IP List	2 Services	1 Setting			

4. Review the draft changes and click **Provision** to restore the policy to the selected version.

ted version or click **Revert** to return to the Policy page.



Provisioning Note Setting

You have the option to 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 it is enabled, you have to populate the note field before provisioning changes.

You might want your users to 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.

When enabled, you cannot provision updates until you enter text in the Provision Note field. The **Confirm & Provision** button is grayed out. After you enter appropriate text in the field the **Confirm & Provision** button is enabled and you can provision the update.



NOTE:

You must have the correct role and permissions to access this feature. If necessary, contact your Illumio administrator for assistance.

To make the provision note mandatory:

1. From the PCE web console menu, choose **Settings > Policy Settings**.
The Policy Settings page appears. By default, this option is set to No.
2. Click **Edit**.
3. Change the *Require Provision Note* option to Yes.
4. Click **Confirm**.
5. Click **Save**.

Secure Workload Connections

This chapter contains the following topics:

SecureConnect	153
AdminConnect	160

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 datacenters. 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 datacenter? 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.

Our Solution

SecureConnect leverages the built-in encryption libraries of host operating systems. On Windows hosts, SecureConnect utilizes Windows IPsec. On Linux hosts, SecureConnect utilizes StrongSwan and Linux kernel IPsec for traffic encryption.

With SecureConnect, Illumio delivers a feature that configures the Security Associations (SAs) necessary to enable traffic encryption between workloads. Once authenticated, encryption and cryptographic suites provide confidentiality and data integrity to network traffic flowing between workloads.

The PCE centrally manages all traffic encryption 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 is 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.

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 datacenters.
- Secure communications across applications and application tiers for regulatory compliance and tighter security.
- Enable secure data migration across different public cloud providers.

Features of SecureConnect

SecureConnect has the following key features.

Platforms Supported by SecureConnect

SecureConnect works for connections between Linux workloads, between Windows workloads, and between Linux and Windows workloads.

IPsec Implementation

SecureConnect implements a subset of the IPsec protocol called Encapsulating Security Payload (ESP), which provides confidentiality, data-origin authentication, connectionless integrity, an anti-replay service, and limited traffic-flow confidentiality.

In its implementation of ESP, SecureConnect uses IPsec transport mode. Using transport mode, only the original payload is encrypted between the workloads. The original IP header information is unchanged so all network routing remains the same. However, the protocol being used will be changed to reflect the transport mode (ESP).

Making this change causes no underlying interfaces to change or be created or any other underlying networking infrastructure changes. Using this approach simply obfuscates the data between endpoint workloads by encrypting the data between them.

If SecureConnect is unable to secure traffic between two workloads with IPsec, it will block unencrypted traffic when the policy was configured to encrypt that traffic.

IKE Versions Used for SecureConnect

SecureConnect connections between workloads use the following versions of Internet Key Exchange (IKE) based on workload operating system:

- Linux ↔ Linux: IKEv2
- Windows ↔ Windows: IKEv1
- Windows ↔ Linux: IKEv1

For a list of supported operating systems for managed workloads, see the [VEN OS Support and Package Dependencies](#) on the Illumio Support portal (login required).

Using Pre-Shared Keys with SecureConnect

SecureConnect includes the option of using pre-shared keys (generated by the PCE) or client-side PKI certificates for IKE authentication.

You can configure SecureConnect to use pre-shared keys (PSKs) to build IPsec tunnels that are automatically generated by the PCE. SecureConnect uses one key per organization. All the workloads in that organization share the one PSK. SecureConnect uses a randomly generated 64-character alpha-numeric string, for example:

```
c4aeb6230c508063db3e3e1fac185bea9c4d17b4642a87e091d11c9564fbd075
```

When SecureConnect is enabled for a workload, you can extract the PSK from a file in the `/opt/illumio` directory, where the VEN stores it. You cannot force the PCE to regenerate and apply a new PSK. If you feel the PSK has been compromised, contact [Illumio Customer Support](#) (login required).

**NOTE:**

Illumio customers accessing the PCE from the Illumio cloud can have multiple Organizations. However, the Illumio PCE does not support multiple Organizations when you have installed the PCE in your datacenter.

Using PKI Certificates with SecureConnect

SecureConnect allows you to use client-side PKI certificates for IKE authentication and IPsec communication between managed workloads. If you have a certificate management infrastructure in place, you can leverage it for IKE authentication between workloads because it provides higher security compared to using pre-shared keys (PSKs).

Certificate-based SecureConnect works for connections between Linux workloads, between Windows workloads, and between Linux and Windows workloads.

The IPsec configuration uses the certificate with the distinguished name from the issuer field that you specify during PCE configuration for IKE peer authentication.

Existing IPsec Configuration on Windows Systems

Installing a VEN on a Windows system does not change the existing Windows IPsec configuration, even though SecureConnect is not enabled. The VEN still captures all logging events (`event.log`, `platform.log`) from the Windows system that relate to IPsec thereby tracking all IPsec activity.

Performance

The CPU processing power that a workload uses determines the capacity of the encryption. The packet size and throughput determine the amount of power that is required to process the encrypted traffic using this feature.

In practice, enabling SecureConnect for a workload is unlikely to cause a big spike in CPU processing or a decrease in network throughput. However, Illumio recommends benchmarking performance before enabling SecureConnect and comparing results after enabling it.

SecureConnect Prerequisites, Limitations, Caveats

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

PKI Certificates with SecureConnect

The following prerequisites and limitations apply when configuring SecureConnect to use certificates:

- You must have a PKI infrastructure to distribute, manage, and revoke certificates for your workloads. The PCE does not manage certificates or deliver them to your workloads.
- The PCE supports configuring only one global CA ID for your organization.
- The VEN on a workload uses a Certificate Authority ID (CA ID) to authenticate and establish a secure connection with a peer workload.

Connected workloads must have CA identity certificates signed by the same root certificate authority. When workloads on either end of a connection use different CA IDs, the IKE negotiation between the workloads will fail and the workloads will not be able to communicate with each other.

VEN Versions

To use PKI certificates with SecureConnect, your workloads must be running VEN version 17.2 or later.

Maximum Transmission Unit (MTU) Size

IPsec connections cannot assemble fragmented packets. Therefore, a high MTU size can disrupt SecureConnect for the workloads running on that host.

Illumio recommends setting the MTU size at 1400 or lower when enabling SecureConnect for a workload.

Ports

Enabling SecureConnect for a workload routes all traffic for that workload through the SecureConnect connection using ports 500/UDP and 4500/UDP for NAT traversal and for environments where ESP traffic is not allowed on the network (for example, when using Amazon Web Services). You must allow 500/UDP and 4500/UDP to traverse your network for SecureConnect.

Unsupported SecureConnect Usage

SecureConnect is not supported in the following situations:

- SecureConnect cannot be used between a workload and unmanaged entities, such as the label “Any (0.0.0.0/0 and ::/0)” (such as, the internet).
- SecureConnect is not supported on virtual services.
- SecureConnect is not supported on workloads in the Idle state. If you enable it for a rule that applies to workloads that are in both Idle and non-Idle enforce-

ment, you can impact the traffic between these workloads.

- SecureConnect is not supported on AIX and Solaris platforms.

SecureConnect and Visibility Only state

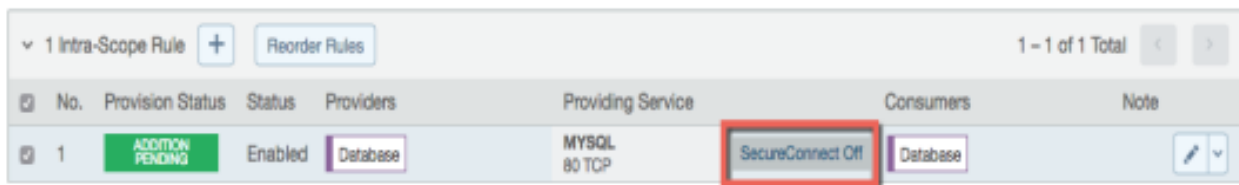
When you configure workloads to use SecureConnect be aware of the following caveat.

SecureConnect encrypts traffic for workloads running in all enforcements except Idle. If misconfigured, you could inadvertently block traffic for workloads running in the Visibility Only enforcement state.

SecureConnect Host-to-Host Encryption

When you configure workloads to use SecureConnect be aware of the following caveat.

SecureConnect encrypts traffic between workloads on a host-to-host basis. Consider the following example.



No.	Provision Status	Status	Providers	Providing Service	Consumers	Note
1	ADDITION PENDING	Enabled	Database	MYSQL 80 TCP	SecureConnect Off Database	

In this example, it appears that enabling SecureConnect will only affect MySQL traffic. However, when you enable SecureConnect for a rule to encrypt traffic between a database workload and a web workload over port 3306, the traffic on all ports between the database and web workloads is protected by IPsec encryption.


Certificate Setup on Workloads

To use PKI certificates with SecureConnect, you must independently set up certificates on your Windows and Linux workloads.

Generate or obtain certificates from a trusted source in your organization. You should only use certificates obtained from trusted sources.

File Requirements

File	Requirements
Issuer's certificate	The global CA certificate, either root or intermediate, in PEM or DER format

File	Requirements
	 NOTE: On Linux, the issuer's certificate must be readable by the Illumio user.
pkcs12 container	Archive containing the public key, private key, and identity certificate generated for the workload host. Sign the identity certificate using the global root certificate. You can password protect the container and private key but do not password protect the public key.

Installation Locations

Windows Store

Use the Windows OS, for example Microsoft Management Console (MMC), to import the files into these locations of the local machine store (not into your user store).

- Root certificate: Trusted Root Certificate Store
- pkcs12 container: Personal ("My") certificate store

Linux Directories

Copy the files into the following Linux directories. (You cannot change these directories.)

- Root certificate: `/opt/illumio_ven/etc/ipsed.d/cacert`
- pkcs12 container: `/opt/illumio_ven/etc/ipsed.d/private`

Enable SecureConnect for a Rule

The following table lists the valid provider and consumer combinations for which you can enable SecureConnect:

Provider	Service	Consumer
Workload	Any service	Workload
All workloads	Any service	All workloads
Label/Label Group	Any service	Label/Label Group

SecureConnect is supported on workloads in Visibility Only and Full enforcement states.

**NOTE:**

When SecureConnect is enabled on a VEN, it is not disabled when the VEN is suspended.

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

The Segmentation Rulesets page appears.

2. Create a new ruleset or open an existing one. See [Segmentation Rulesets](#) for information.
3. In the ruleset, select the **Scopes and Rules** tab.
4. If necessary create an intra-scope or an extra-scope rule. See [Rule Writing](#) for information. To edit an existing rule, click the edit icon at the end of the row.
5. To enable SecureConnect for the rule, select **SecureConnect** from the *Providing Service* drop-down list.



6. Click the **Save** icon at the end of the row.
7. To apply the changes to the applicable workloads, provision the changes. See [Provision Changes](#) for information.

When you enable SecureConnect for a rule, the PCE duplicates symmetrical encryption for both sides of the connection.

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 the *PCE Administration Guide*.
- [Certificate Setup on Workloads](#)
- Configure certificates for AdminConnect. For information, see the *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](#) for the list of rule types that do not support AdminConnect.

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

The Segmentation Rulesets page appears.

2. Create a new ruleset or open an existing one. See [Segmentation Rulesets](#) for information.
3. In the ruleset, select the **Scopes and Rules** tab.

4. If necessary create an intra-scope or an extra-scope rule. See [Rule Writing](#) for information. 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. See [Provision Changes](#) for information.

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 the *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](#) for information.
4. Configure IPsec on a laptop.

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

Illumio does not support installing the VEN on laptops. Therefore, 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. See [Add an Unmanaged Workload](#) for information.

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.
```