



Illumio[®] CloudSecure

User's Guide

July 2024

91000-100-1.0.36

Legal Notices

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

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

Resources

Legal information, 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>

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 Get Started with CloudSecure	7
About the CloudSecure User's Guide	7
Contact the Illumio Documentation Team	8
Welcome to CloudSecure	8
Typical CloudSecure Workflow	9
Supported Browsers	10
Onboard Cloud Accounts	10
About Onboarding Cloud Accounts	10
Signing In	11
Onboard an AWS Cloud Account	12
Onboard an AWS Cloud Organization	18
Onboard an Azure Cloud Subscription	26
Onboard an Azure Cloud Tenant	29
Grant Flow Log Access	32
Prerequisites	33
Review and Grant Flow Log Access	33
Test Flow Log Access	34
Caveats	35
Set up Flow Logs in AWS and Azure	36
AWS	36
Azure	39
Manually Configure CloudSecure to Fetch Flow Logs	39
Overview	39
Add a Destination with a Custom Path to the CloudFormation Template	39
Update the IllumioBucketListPolicy Document	40
What to Do Next	44
1. Visualize Your Cloud Resources	44
2. Define Your Public Clouds in CloudSecure	45
3. Create Policy in CloudSecure	46
About the Illumio Virtual Advisor	47
Use IVA	47
Chapter 2 Visualize Your Cloud in CloudSecure	50
Map	50
Grouping in the Map	51

Map Layout Options	53
How to Read the Map Symbols	53
Filtering the Map	56
Panels in the Map	57
Cloud Map	59
What is the Cloud Map?	59
Supported Resource Types	60
How the Cloud Map is Organized	60
How to Navigate the Cloud Map	63
Display Resource Side Panel	67
Cloud Map Traffic Lines	68
Limitations for Using the Cloud Map	70
Caveats	71
Cloud Map Supported Resources	71
Inventory	75
Supported Resource Types	75
Use Case and Example	75
VPC/VNet Peering Details	76
Security Control Resource Details	77
Details Resource Graph	77
Inventory Supported Resources	78
Traffic	85
Supported Resource Types	85
Exporting Traffic Lists	85
Limitations for Displaying Traffic	86
Generating Risk Reports	86
Search Traffic	88
Risky Services	89
Traffic Supported Resources	90
CloudSecure Dashboard	91
What is the Dashboard?	91
CloudSecure Search	93
What is Context-based Search?	93
What is Filter-based Search?	94
Product Usage	95
Displaying Product Usage	95
Events	96

Displaying Events	96
Reports	97
Reports	97
Chapter 3 Define Your Cloud Resources in CloudSecure	99
Deployments and Applications	99
What is a Deployment in CloudSecure?	99
Relationship between Deployments and Applications	100
CloudSecure Discovers Your Application Environments	102
Example Deployment and Application Definitions	103
Define a Deployment	105
Define an Application	108
View and Approve an Application	115
Cloud Tag to Label Mapping	117
Use Case and Example	117
View System Labels	119
Rule-Based Labeling	119
Before you begin	120
Typical Labeling Rule Workflow	120
Add and Manage Labeling Rules	121
How Label Matching Works	127
Labeling Rule Examples	129
Use AI Labeling	132
Overview	132
Use AI Labeling	133
Chapter 4 Create Policy in CloudSecure	135
CloudSecure Policy Model	135
About the Illumio CloudSecure Policy Model	135
Security Policy Guidelines	136
Understanding Rules	136
Types of CloudSecure Policy	136
Overview of Policy Attributes	136
CloudSecure Policy Attributes Overview	137
IP Lists	137
Labels	139
Services	141
Organization Policy versus Application Policy	141

About CloudSecure Policies	142
What Happens When Org and App Policies Conflict?	143
Writing Organization Policy	143
Writing Application Policy	145
Resources that Support Policy	150
Unified Policy	151
Overview	151
Notices	153
Chapter 5 CloudSecure Administration	154
Connector	154
Use Case and Example	154
Events	155
View your Events	156
Role-Based Access Control	156
Add Roles	156
User Management	157
About Users	157
Add Users	158
Delete Users	158
Add or Remove Roles	159
My Profile	159
My Roles	159
Chapter 6 CloudSecure Reference	160
CloudSecure Requirements	160
AWS Requirements	160
Azure Requirements	168

Chapter 1

Get Started with CloudSecure

This chapter contains the following topics:

About the CloudSecure User's Guide	7
Welcome to CloudSecure	8
Onboard Cloud Accounts	10
Grant Flow Log Access	32
Set up Flow Logs in AWS and Azure	36
Manually Configure CloudSecure to Fetch Flow Logs	39
What to Do Next	44
About the Illumio Virtual Advisor	47

This section helps you get started with using Illumio CloudSecure. In particular, you begin by onboarding your public cloud accounts; then, the section contains a What's Next topic that explains the phases in order for using CloudSecure.

About the CloudSecure User's Guide

This PDF includes all the content from the CloudSecure documentation portal and presents it in a downloadable PDF file.

The content in the HTML documentation portal and the PDF *Illumio CloudSecure User's Guide* is synchronized so that they contain the same content.

You can use this PDF when you need to view the documentation offline or want to learn about CloudSecure on a device better suited to displaying PDF documents.

Contact the Illumio Documentation Team

At Illumio, we value the customer experience; therefore, we welcome any feedback regarding this documentation. To contact us, email us at doc-feedback@illumio.com.

We look forward to hearing from you.

Welcome to CloudSecure

Organizations everywhere are realizing the benefits of adopting a public hybrid cloud approach to managing their computing resources. However, these benefits are tempered by real challenges. Despite prevention using a range of security tools, customers struggle with breaches that go undetected and spread through channels. With workload and network flow proliferation, and decentralized application deployment through development and operations (DevOps), the risk of breaches moving laterally is greater in the multi-cloud environment.

Signing up with Illumio CloudSecure allows organizations to adopt the public cloud with confidence. CloudSecure helps address the problem of undetected breaches in the cloud.

CloudSecure is an agentless SaaS offering designed to prevent breaches and ransomware from becoming cyber disasters in the public cloud by providing an understanding of communications in combination with attribute-based micro segmentation.


CloudSecure includes the following key features:

- **Multi-Cloud:** Single-pane-of-glass view into multi-cloud environments for AWS and Azure.
- **Cloud Inventory View:** Visibility, searchability, and ability to gather insights from resources spanning multiple cloud environments.
- **Cloud Map View:** Visualize and search resources, resources relationships, and actual traffic flows on an interactive map of the multi-cloud environment, with drill-down on resources, flows, and metadata with a few clicks.
- **Traffic Visibility:** Visualize and search to gather insights from traffic flows.
- **Application Blueprint:** Build and view a multi-cloud application blueprint using cloud tags and metadata. Auto-discover application deployments and resources. View inventory, resource relationships, traffic flow, interactive cloud maps, and policies in the context of an application and its deployments.
- **Policy Authoring:** Author and provision organization-wide and application-specific policies using labels and IP lists for application segmentation.
- **Change Management System Integration:** Integrate with Slack.

CloudSecure drives segmentation in an organization's multi-cloud environment by learning cloud context, and using your network flow logs.

CloudSecure builds the relationship graph between your cloud resources to visualize resource hierarchy and traffic flows. It also gives you the ability to quickly locate resources and security controls of interest for network security reviews and notify the application owners.

Typical CloudSecure Workflow



In this diagram, the typical workflow consists of these tasks in this order:

1. Onboard a cloud account with CloudSecure. CloudSecure supports onboarding AWS and Azure accounts.

See [Onboard an AWS Cloud Account](#) and [Onboard an Azure Cloud Subscription](#) for information.
2. Using the onboarded accounts, CloudSecure begins the process of discovering and ingesting their resources.

See [CloudSecure Discovers Your Application Environments](#) in the topic *Deployments and Applications* for information.
3. Review information for your cloud account. If you are a visual person and want to analyze a graphical, hierarchical display, go to the Cloud Map and browser your resources. If you are most comfortable working with lists and sorted data, go to the Inventory page and browse your discovered resources.

See [Cloud Map](#) and [Inventory](#) for information.

4. Create blueprints of your applications as they are deployed in your cloud accounts. You begin by defining your deployments in use in your accounts and then define your applications; CloudSecure will discover every cloud environment where the applications are hosted.

See [Define a Deployment](#) and [Define an Application](#) for information.

5. Review the application definition. Before an application definition is available for accepting security policy, the correct stakeholders in your organization are responsible for reviewing it and approving that it is correctly defined. Once you are satisfied with your application blueprint, you can write application policies.

See [CloudSecure Policy Model](#) and [Organization Policy versus Application Policy](#) for information.

Supported Browsers

The following web browsers, and any operating systems on which they run, are officially supported for Illumio CloudSecure:

- Chrome and corresponding Chromium-based browsers (MS Edge, Opera, etc.): 117+
- Firefox: 115+
- Safari: 16.6+

Onboard Cloud Accounts

This chapter contains the following topics:

This section describes how to onboard your cloud accounts

About Onboarding Cloud Accounts

Illumio supports onboarding your AWS and Azure cloud accounts with CloudSecure.

For the steps to onboard the supported cloud accounts, see [Onboard an AWS Cloud Account](#) and [Onboard an Azure Cloud Account](#).

The Onboarding Wizard

To make the onboarding process easier, CloudSecure provides a wizard to assist you with onboarding your cloud accounts. The wizard uses the following workflow:

Connect to the cloud account → Set up Access → Confirm and Save

The overall onboarding process for both AWS and Azure clouds follows this workflow with some minor differences; such as when onboarding AWS, customers can have CloudSecure send the authentication token from AWS back to CloudSecure. When onboarding Azure, retrieving the authentication token is a manual process.

The first time that you log in, the page displays a message that you need to add your cloud accounts to CloudSecure.

After you onboard a new cloud account, the Cloud Onboarding page appears with a new row for the account.

After Onboarding Cloud Accounts

After onboarding a cloud account, you can perform the following tasks from the Onboarding page.

- **View the details about an account**

Click an account name link to view the details about the account.

- **Delete an account or multiple accounts**

To delete multiple account, select them from the Onboarding list and click **Remove**. Or, open an account details page and click **Remove**.

- **Change an account name in CloudSecure**

Click an account name link to view the details about the account and click **Edit**.

- **Search for accounts**

From the Onboarding page, using the Search filter, search by name, account ID, and by account type – AWS vs Azure.

For information about how cloud account onboarding fits into the overall recommended workflow for successful CloudSecure usage, see [What to Do Next](#).

Signing In

This topic explains how to log in and authenticate.

Illumio is set up for multi-factor authentication for sign-in. When a new user signs in for the first time, there are Okta verification steps. Illumio keeps your Okta verification for a limited time, during which you do not need to re-login to Okta. Your Okta session ends if you log out, requiring you to re-login.

Activate your Account and Sign In

After you sign up for a free trial, Illumio sends you an onboarding email with a link to activate your account.

The first user to onboard an account becomes the owner for your organization.

To activate your account and sign in for the first time:

1. Click the link in your onboarding email. A page appears that contains a field with your email address.
2. Click **Continue**. An Okta page appears.
3. In the Okta page, activate your account by setting a password.

After setting your password in Okta, Illumio sends you an email with the following URL to sign-in:

`https://console.illum.io`

4. Click the console URL. The **Log In** page opens in your browser.
5. Enter your email address and click **Continue**. A second sign in page appears with your email address in the **Username** field with an option to stay signed in.
6. If you want to stay signed in, select that option. Click **Next**. An Okta page with various options appears.

For example, if you select to use a push notification, Illumio sends the notification to your phone.

7. Select the option that you want to use to verify your identity.
8. Verify your identity using the method you selected.
9. If prompted, re-enter your password and click **Verify**.

The *Connecting to* page refreshes with a message that you are being signed in.

If you have not onboarded your cloud accounts or paired your servers and endpoints yet, CloudSecure prompts you to do so. See [Onboard Cloud Accounts](#) for Illumio CloudSecure or Core server pairing documentation for Servers and Endpoints.

What Happens Next?

Once signed in, you can use Illumio CloudSecure. You may wish to view your account settings before onboarding accounts. To view your profile, see [My Profile](#). To view your roles, see [My Roles](#).

Onboard an AWS Cloud Account

This topic explains how to onboard an AWS account.

Prerequisites

- Before onboarding an AWS account, copy the account ID for the account you want to onboard so that you can specify it in the first step of the onboarding wizard
- Determine the method to use for onboarding the account — either by using CloudSecure to run the CloudFormation stack or by using an Illumio provided YAML file as a template to manually create the stack
- Prior to onboarding an AWS account, ensure that CloudSecure has the required AWS permissions. See [AWS Requirements](#) for information.
- You need to login to the account in which you will run the stack, so that account will need at least administrator permissions

Ways to Onboard AWS

IMPORTANT:

The wizard for onboarding an AWS account contains the option to onboard a single AWS account or an AWS organization (which is a collection of accounts).

When onboarding an AWS account, you have the option to use CloudSecure to create the stack in the AWS console or by downloading a YAML file and completing the settings outside of the AWS console.

When you use CloudSecure to create and run the CloudFormation stack, CloudSecure populates the required data in AWS to run the stack. When you choose to download and use a YAML file, you must complete the file with the required data.

Illumio recommends that you use the first option to onboard an AWS account and allow CloudSecure to run the stack.

Onboard AWS by Running CloudFormation Stack

This procedure describes the Illumio recommended method for creating the stack. For information about creating the stack by downloading a YAML file, see [Onboard AWS using Stack Template](#).

1. If this is the first time you are logging in, click **+ AWS** to onboard your first account.

If you've already onboarded other accounts, choose **Onboarding** from the left navigation. The Onboarding page appears. Click **+Add AWS** at the top of the page.

The *Add AWS Cloud Account* wizard starts and displays the first step: **Connect to AWS**

2. Provide the following information about your AWS account:

- Name for the account

This name is what will appear in CloudSecure. The name should be descriptive so that you can easily identify it in CloudSecure.

- The AWS account ID of the account you are onboarding into CloudSecure

NOTE:

The page contains a toggle below the Account ID field to specify the type of access CloudSecure will have to your AWS account. Choosing *Yes* grants the Illumio Cross Account Role permission to view your AWS account resources and to apply policy to them. Choosing *No* provides the Illumio Cross Account Role read-only access. To view the permissions you are granting CloudSecure to your AWS account, click **Download Permissions**.

When done completing these settings, click **Next**.

The wizard advances to step two: **Set up Access**

3. Select or create a service account.

NOTE:

During onboarding, you configure a service account for CloudSecure. CloudSecure uses this digital identity to interact with your AWS account. The service account has read/write access, which you granted in the first step of the wizard.

If you haven't onboarded any accounts yet, click **Add a new Service Account** in the Service Account drop-down list and specify a name and description (optional) and click **Create**.

A pop-up dialog box appears displaying information about the credentials created for the service account. You cannot copy information from the dialog box. Click **Download Credentials** to save this information locally, then click **Close**.

IMPORTANT:

Open the downloaded credentials file (`Service-Account-<name>.txt`) for the service account and copy the value in the `serviceAccountToken` field. You will need this value when creating the CloudFormation stack in AWS. CloudSecure only provides these credentials for download during this step of the onboarding wizard.

NOTE:

Alternatively, you can select an existing service account from a previous onboarding. When you use an existing service account, you must still have access to the downloaded credentials file and service account secret. If you do not have access to that file, you must create a new service account.

4. Under *Type of Integration*, select **Create Cloud Formation Stack**. The button **Create IAM Roles on AWS** becomes enabled.
 - a. To create a new stack, click **Create IAM Roles on AWS**. CloudSecure opens the AWS Sign in page in a new browser window. Sign into AWS as a Root or Administrator user. The *Quick create stack* page appears.

The page is pre-populated with the required values, such as the URL for the YAML file, the stack name, the key for the service account you specified, and more. The field for the service account secret is not populated.

NOTE:

The stack name needs to be unique for CloudSecure. If you already have a stack in AWS with the pre-populated name, modify the name so that it is unique.

- b. In the *Quick create stack* page, paste the credential secret that you copied from the downloaded credentials file.
 - c. Select the check box to acknowledge that CloudSecure will create IAM resources in AWS.
 - d. Click **Create stack**.

The script to create the stack runs. When it finishes, your AWS account includes custom IAM roles required by CloudSecure and a temporary

Lambda function named `LambdaExecutionRoleIllumioCloudAPICall`. The Lambda function passes back to CloudSecure two credentials:

- The ARN of the role from the Trusted entities
- The secret key that AWS uses for authentication when CloudSecure accesses account resources

Now, CloudSecure has the required credentials to access your AWS account so that you don't have to repeatedly provide them.

For the complete list of permissions granted to CloudSecure for your account, see [AWS Requirements](#).

- e. Leave the AWS console and return to CloudSecure.
- f. Click **Next**. The final step of the wizard appears.

The wizard displays a summary of the account information you just specified.

5. Review the account information and if everything looks correct, click **Save and Confirm**. If you see issues you need to correct, click **Back** and return to that wizard step.

Your account is successfully onboarded and a row for that account appears in the Onboarding page.

Onboard AWS using Stack Template

NOTE:

Choose this option when you don't have the required permissions in AWS to create a CloudFormation stack or you want to create the CloudFormation stack manually.

1. Launch the onboarding wizard in either of the following ways:
 - Click **+ AWS** in the Onboarding page to onboard your first account when you sign in for the first time
 - From the left navigation, choose **Onboarding** and click **+ AWS** at the top of the page.
2. Follow steps [2](#) and [3](#) from the procedure above
3. In step two (**Set up Access**) of the onboarding wizard, select **Download Cloud Formation Stack** and click **Download**.

CloudSecure downloads an AWS Integration YAML file to your local system. This YAML file contains sections for the data required to create and run the CloudFormation stack in AWS. Some sections of the YAML file are pre-populated with default values. In other sections, the default value is empty.

NOTE:

If you wish to share the CloudFormation stack with others so that they can run it, you will need the CloudSecure ID. It will display in the *Add AWS Account* dialog.

4. Complete the missing values as required and save the file.
5. Log into your AWS console with the required permissions to run a CloudFormation stack or provide the file to members of your organization who have the required AWS account access.
6. Use the completed AWS Integration file as an AWS CloudFormation template to run the stack. The CloudSecure YAML file provided by Illumio is a valid stack template file.

For information, see “Creating a stack” in the Amazon AWS online documentation.

7. Click **Next**. The final step of the wizard appears.
8. Review the account information and if everything looks correct, click **Save and Confirm**. If you see issues you need to correct, click **Back** and return to that wizard step.

When the stack command finishes running in AWS and you’ve successfully created the stack, a CloudSecure script will notify CloudSecure that the stack was successfully created and CloudSecure will detect that account was onboarded and begin synchronizing the account resources with CloudSecure. A new row for that account appears in the Onboarding page.

Remove the Integration

You can delete the integration for a given account by selecting the account and clicking **Remove > Remove**. However, you will need to then manually delete the CloudFormation Stack in AWS.

1. Login to the AWS Console and choose **Services > CloudFormation**.
2. Select **Stacks**, and, in the list of stacks, choose the stack name you used while onboarding Cloudsecure and click **Delete**.

Initially the stack deletion will fail. The CloudFormation template provided by CloudSecure creates Lambda-backed custom resources, which AWS does not automatically clear.

3. If it fails, select the stack and click **Delete** again.

A pop-up window appears with the option to retain the resources that are failing to delete.

4. Choose that checkbox option and click **Delete**.

Note: Although you selected the option to retain resources, custom resources are specific to CloudFormation and they will be cleared upon the deletion of the stack. Ref: <https://repost.aws/knowledge-center/cloudformation-lambda-resource-delete>.

The Stack will be deleted, removing all the resources (Role, Lambda, Custom Resource) created when running the stack.

What's Next?

For the next steps after onboarding an account, see [After Onboarding Cloud Accounts](#) and [What to Do Next](#).

Onboard an AWS Cloud Organization

This topic explains how to onboard an AWS organization.

Background

An AWS organization is a service AWS provides that allows you to consolidate multiple accounts into an organization and manage them centrally. The hierarchy of AWS organization is as follows:

- Root - The parent container for all accounts. It consists of Organizational Units (OU) and accounts.
- Organizational Unit (OU) - The container for accounts within root. It can also contain other Organizational Units.
- Account - The standard AWS account that contains the AWS resources

When the AWS root account is onboarded into Cloudsecure, all the accounts under the root will be onboarded (provided the user runs the StackSet). CloudSecure supports onboarding AWS Organization (root account) and AWS accounts. It does not support onboarding AWS Organizational Units.

Onboarding of an AWS organization (root account) is a two-step process.

1. Run a CloudFormation stack on a root account.
2. Run a CloudFormation stackset on a root account, which in turn runs the stack in all accounts under the root account.

Note: If you want to onboard only the accounts under root account, but not the root account itself, then the first step can be skipped.

Prerequisites

- Before onboarding an AWS organization, copy the root account ID for the account you want to onboard so that you can specify it in the first step of the onboarding wizard
- Prior to onboarding an AWS account, ensure that CloudSecure has the required AWS permissions. See [AWS Requirements](#) for information.
- You need to login to the root account in which you will run the stack or stackset

Onboard AWS Organizations in CloudSecure

The following instructions describe how to begin the organization onboarding sequence in CloudSecure, irrespective of whether you will use a CloudFormation stack or stackset. Subsequent instructions will guide you on the specific steps for using either a CloudFormation stack or stackset in the AWS console.

1. Launch the onboarding wizard in either of the following ways:
 - Click **+ AWS** in the Onboarding page to onboard your first organization when you sign in for the first time
 - From the left navigation, choose **Onboarding** and click **+ AWS** at the top of the page.
2. Provide the following information about your AWS account:
 - Name for the root account

This name is what will appear in CloudSecure. The name should be descriptive so that you can easily identify it in CloudSecure.
 - The AWS ID of the root account you are onboarding into CloudSecure

NOTE:

The page contains a toggle below the Account ID field to specify the type of access CloudSecure will have to your AWS account. Choosing *Yes* grants the Illumio Cross Account Role permission to view your AWS account resources and to apply policy to them. Choosing *No* provides the Illumio Cross Account Role read-only access. To view the permissions you are granting CloudSecure to your AWS account, click **Download Permissions**.

When done completing these settings, click **Next**.

The wizard advances to step two: **Set up Access**.

3. Select or create a service account.

NOTE:

During onboarding, you configure a service account for CloudSecure. CloudSecure uses this digital identity to interact with your AWS account. The service account has read/write access, which you granted in the first step of the wizard.

If you haven't onboarded any accounts yet, click **Add a new Service Account** in the *Service Account* drop-down list and specify a name and description (optional) and click **Create**.

A pop-up dialog box appears displaying information about the credentials created for the service account.

You cannot copy information from the dialog box. Click **Download Credentials** to save this information locally, then click **Close**.

IMPORTANT:

- Make a note of the CloudSecure Tenant Id. It will be needed for running the template in AWS Console.
- Open the downloaded credentials file (*Service-Account-
<name>.txt*) for the service account and copy the value in the `serviceAccountKeyId` and `serviceAccountToken` fields. You will need these values when creating the CloudFormation stack or stack-set in AWS. CloudSecure provides these credentials for download only during this step of the onboarding wizard.

NOTE:

Alternatively, you can select an existing service account from a previous onboarding. When you use an existing service account, you must still have access to the downloaded credentials file and service account secret. If you do not have access to that file, you must create a new service account.

4. In step two (**Set up Access**) of the onboarding wizard, select **Download Cloud Formation Stack** and click **Download**.
CloudSecure downloads an AWS Integration YAML file to your local system.
5. Click **Next**. The final step of the wizard appears.
6. Review the account information and if everything looks correct, click **Save and Confirm**. If you see issues you need to correct, click **Back** and return to that wizard step.

Create Roles in the AWS Console by Running a Stack

NOTE:

Choose this option when you want to onboard accounts under root account, *and* the root account itself. Then follow the subsequent instructions in [Create Roles in the AWS Console for Accounts Under the Root Account](#).

In order to create the Assume role and provide CloudSecure with read/read-write permission to resources in your AWS root account, follow these steps to run the template as stack in the root account.

Create the Stack

1. Log into your AWS console with the required permissions (root account) to run a CloudFormation stack or provide the file to members of your organization who have the required AWS root account access.
2. Under *Services*, click **CloudFormation**.
3. Click on **Create Stack** and choose the **With new resources** option.
4. In the *Create stack* page, select the **Template is ready** and **Upload a template file** options, and click **Choose File**. (The CloudSecure YAML file provided by Illumio

is a valid stack template file.)

5. Upload the CloudSecure YAML file and click **Next**.

Specify the Stack Details

1. In the *Specify stack details* page, enter the Stack name. The stack name must be unique and not the same name used to create previous stacks.
2. In the *IllumioServiceAccountKey* and *IllumioServiceAccountSecret* text boxes, enter the *serviceAccountKeyId* and *serviceAccountToken*, respectively, from the downloaded *ServiceAccount* file.
3. Enter the *CloudSecureTenantId* in the form. The *IAMRoleName* field will auto-populate with a default, but you can modify the name if needed.
4. Click **Next** to continue.

Configure, Review, and Run the Stack

1. In the *Configure stack option* page, allow the default values and click **Next**.
2. In the *Review* page, select the acknowledgment check box and click **Submit**.

The stack will run, creating the resources needed to create the IAM Assume role and will make a callback to CloudSecure with the *RoleARN*, *ExternalId*, *OrgId*, and *MasterAccountId*. The *RoleARN* and *ExternalId* will be used by CloudSecure to connect with the account and sync resources. The *OrgId* and *MasterAccountId* will be used by CloudSecure to create a mapping between the root account and the accounts under it.

When the stack command finishes running in AWS and you've successfully created the stack, a CloudSecure script will notify CloudSecure that the stack was successfully created and CloudSecure will detect that the organization was onboarded and begin synchronizing the organization resources with CloudSecure. A new row for that organization will appear in the *Onboarding* page.

Create Roles in the AWS Console for Accounts Under the Root Account

NOTE:

Choose this option solely when you want to onboard only the accounts under root account, but *not* the root account itself.

If you want to onboard the accounts under the root account, *and* the root account itself, you must first perform the steps in [Create Roles in the AWS Console by Running a Stack](#) before performing these steps.

In order to create the Assume role and provide CloudSecure with read/read-write permission to resources in your AWS accounts under the root account, follow these steps to run the template as a stackset in the root account.

Create a Stackset

1. Log into your AWS console with the required permissions (root account) to run a CloudFormation stack or provide the file to members of your organization who have the required AWS root account access.
2. Under *Services*, click **CloudFormation**.
3. Click **Create StackSet**.

Choose a Template

1. In the *Choose a template* page, select the **Service-managed permissions, Template is ready**, and **Upload a template file** options, and click **Choose File**.
2. Upload the CloudSecure YAML file and click **Next**. (The CloudSecure YAML file provided by Illumio is a valid stack template file.)

Specify the Stackset Details

1. In the *Specify stackset details* page, enter the Stackset name. The stack name must be unique and not the same name used to create previous stacks.
2. Add a description in the *Stackset description* field.
3. In the *IllumioServiceAccountKey* and *IllumioServiceAccountSecret* text boxes, enter the *serviceAccountKeyId* and *serviceAccountToken*, respectively, from the downloaded *ServiceAccount* file.

4. Enter the *CloudSecureTenantId* in the form. The *IAMRoleName* field will auto-populate with a default, but you can modify the name if needed.
5. Click **Next** to continue.

Configure, Review, and Run the Stackset

1. In the *Specify regions options* page, choose the region under which the stacks are set to be deployed. This will allow CloudSecure to access resources in all regions, so selecting only one region is preferable.

In the *Set deployment options* page, assuming that only one region was chosen, allow the default values and click **Next**.

2. In the *Review* page, select the acknowledgment check box and click **Submit**.

The stackset will run, creating the resources in all accounts under the root account to create the IAM Assume role and will make a callback to CloudSecure with the *RoleARN*, *ExternalId*, *OrgId*, and *MasterAccountId*. The *RoleARN* and *ExternalId* will be used by CloudSecure to connect with the account and sync resources. The *OrgId* and *MasterAccountId* will be used by CloudSecure to create mapping between the root account and accounts under it.

When the stack command finishes running in AWS and you've successfully created the stack, a CloudSecure script will notify CloudSecure that the stack was successfully created and CloudSecure will detect that the organization was onboarded and begin synchronizing the organization resources with CloudSecure. Clicking the organization in the *Onboarding* page will let you see the accounts under it.

What's Next?

For the next steps after onboarding organization, see [After Onboarding Cloud Accounts](#) and [What to Do Next](#).

Edit the Accounts in the Organization

1. In the *Onboarding* page, click on the organization.
2. Click **Edit**.
3. You can change read/write access permissions if you like.
4. Select the individual account in question and click **Enable**, **Disable**, or **Remove** as needed.
5. In the dialog that appears, click to confirm.

Remove the Integration

You can delete the integration for a given organization by selecting the it in the *Onboarding* page and clicking **Remove > Remove**. However, you will need to then manually delete the CloudFormation stack and/or stackset in AWS.

Note: Once an AWS organization is deleted, the accounts under the account will also be removed.

Remove the Stack in AWS

1. Login to the AWS Console and choose **Services > CloudFormation**.
2. Select **Stacks**, and, in the list of stacks, choose the stack name you used while onboarding Cloudsecure and click **Delete**.

Initially the stack deletion will fail. The CloudFormation template provided by CloudSecure creates Lambda-backed custom resources, which AWS does not automatically clear.

3. If it fails, select the stack and click **Delete** again.

A pop-up window appears with the option to retain the resources that are failing to delete.

4. Choose that checkbox option and click **Delete**.

Note: Although you selected the option to retain resources, custom resources are specific to CloudFormation and they will be cleared upon the deletion of the stack. Ref: <https://repost.aws/knowledge-center/cloudformation-lambda-resource-delete>.

The Stack will be deleted, removing all the resources (Role, Lambda, Custom Resource) created when running the stack.

Remove the Stackset in AWS

1. Login to the AWS Console and choose **Services > CloudFormation**.
2. Select **StackSet**, and, in the list of stacksets, choose the stackset name you used while onboarding Cloudsecure.
3. From the *Actions* drop-down menu, select **Delete stacks from StackSet**. (This must be done before you can delete the stackset.)

4. In the *Set deployment options* page, *Organization units (OUs)* section, enter the AWS OU ID (the organization ID, which can be found in the organization service).
5. In the *Set deployment options* page, *Specify Regions* section, select the region. This will be the region you selected when you created the stackset.
6. Leave the rest of the options with their defaults and click **Next**.
7. In the *Review* page, click **Submit**.
This will remove all the stacks from the stackset. To monitor the status of the operation, select the stackset and click the **Operations** tab.
8. If the action fails, it means that the individual stacks in the accounts under the master account are failing to delete. If that happens, login to the specific accounts (not the root account) and follow the same steps seen in [Remove the Stack in AWS](#). Once that is done, repeat the instructions to [Remove the Stackset in AWS](#).
Once the stacks are completely removed, select the stackset again and choose **Delete StackSet** from the *Actions* drop-down menu.

Onboard an Azure Cloud Subscription

This topic explains how to onboard an Azure subscription.

Prerequisites

- In Azure, copy the subscription ID and its parent management group ID for the subscription you want to onboard. You must provide them in the first step of the onboarding wizard.
- The custom role must be set up properly before the running the onboarding PowerShell script mentioned in [Onboard a Subscription](#) below.
- For simplicity, the user who onboards Azure with CloudSecure must have owner permissions or user access administrator privileges in Azure for that subscription. Having these permissions is required for CloudSecure to create a custom role in Azure named "Illumio Network Security Administrator." CloudSecure needs this custom role so that it has read/write permissions to the subscription resources.
- If the user who onboards Azure with CloudSecure does *not* have the above permissions, submit a request to your group that has them so they can create a custom role using the recommended name of "Illumio Network Security

Administrator" and containing the NSG write permissions defined in [Azure Requirements](#). The custom role is created when you run the onboarding script.

Onboard a Subscription

1. If this is the first time you are logging in, click **+ Azure** on the Onboarding page to onboard your first account.

If you've already onboarded other accounts, choose **Onboarding** from the left navigation. The Onboarding page appears. Click **+Add Azure** at the top of the page.

The *Add Azure Cloud Account* wizard starts and displays the first step: **Connect to Azure**

2. Provide the following information about your Azure account:
 - **Name:** You specify a name for the account; this name is what will appear in CloudSecure. The name should be descriptive so that you can easily identify it in CloudSecure.
 - **Tenant ID:** Paste the parent management group ID that you copied from Azure.
 - **Subscription ID:** Paste the subscription ID that you copied from Azure.

NOTE:

The page contains a toggle below the Subscription ID field to specify the type of access CloudSecure will have to your Azure subscription. Choosing *Yes* grants the Illumio Cross Account Role permission to view your Azure subscription resources and to apply policy to them. Choosing *No* provides the Illumio Cross Account Role read-only access. To view the permissions you are granting CloudSecure to your Azure subscription, click **Download Permissions**.

3. When done completing these settings, click **Next**.
4. Select a service account that you want to use or create a new one. Make sure to download the credentials, as they will be needed for the PowerShell script to return the Azure AD app credentials back to CloudSecure.
5. Enter the ServiceAccountToken in the appropriate field.

The wizard advances to step two: **Set up Access**

1. The Set up Access step includes a field containing a PowerShell command to run the `illumio-init.ps1` script in Azure. Illumio securely hosts the script so that it can run during the onboarding process. The PowerShell command automatically appends the subscription ID you entered in the first step of the wizard.
2. To the left of the PowerShell command field, click the copy icon. The icon refreshes with a check mark on a green field indicating you successfully copied the command.
3. In a new browser window, open your Azure portal.
4. From the top taskbar, click the **Cloud Shell** icon to open a console; select the PowerShell option.
5. After Azure finishes building your Azure drive, paste the copied PowerShell command.

When you run the script in Azure, it creates an AD app registration named “Illumio-CloudSecure-Access.” The script also creates a custom role named “Illumio Network Security Administrator.” Additionally, the app registration includes Reader roles.

Creation of the AD app registration and the roles allows CloudSecure access to the subscription resources. CloudSecure will be able to discover subscription resources and write policies for them.

For the complete list of permissions granted to CloudSecure for your account, see [Azure Requirements](#).

The script sends the Client ID and Client Secret to CloudSecure. CloudSecure accesses your Azure subscription so that you don't have to repeatedly provide your Azure credentials.

CloudSecure can access your Azure subscription so that you don't have to repeatedly provide your Azure credentials.

6. Leave your Azure portal and return to CloudSecure. The **Set up Access** step in the onboarding wizard should still be displayed.
7. Select the check box indicating that the “deployment” script has finished running in Azure, and click **Next**.
8. The final step of the wizard appears. This step displays a summary of the subscription information you just specified for onboarding.

9. Review the subscription information and if everything looks correct, click **Save and Confirm**. If you see issues you need to correct, click **Back** and return to that wizard step.

NOTE: CloudSecure can read flow logs from several NSGs going to the same storage account. With Azure, you can configure NSG flow logs in the same region, despite being from multiple VNets residing in different subscriptions, to be sent to a single storage account in the same region residing in a single subscription. By providing access to that specific storage account, CloudSecure can obtain and analyze flow logs for all the NSGs residing in different subscriptions. For more information on flow logs, see [Grant Flow Log Access](#).

What's Next?

When finished, the **Onboarding** page opens and displays a new row for that account.

For the next steps after onboarding an account, see [After Onboarding Cloud Accounts](#) and [What to Do Next](#).

Onboard an Azure Cloud Tenant

This topic explains how to onboard an Azure tenant. Onboarding an Azure tenant allows you to connect all the subscriptions and resources under the tenant with CloudSecure. Running the PowerShell script for Azure Tenant onboarding will create a new AD application with the tenant scope. This will allow CloudSecure to retrieve subscriptions and resources under the given tenant. After the Azure AD application is created and required permission are set, the PowerShell script will automatically send the necessary credentials (Client Id and Client Secret). CloudSecure requires these credentials to communicate with the your Azure tenant.

Prerequisites

- In Azure, you must copy the parent management group id (tenant id). It can be found under the Management Groups. You must provide it in the first step of the onboarding wizard.
- The user onboarding the tenant must have Owner permissions or the User Access Administrator Role for running the onboarding PowerShell script mentioned in [Onboard a Tenant](#) below
- You must download the newly created CloudSecure Service Account or have access to the credentials of the existing CloudSecure Service Account

Onboard a Tenant

1. If this is the first time you are logging in, click **+ Azure** on the Onboarding page onboard your first account.
2. If you've already onboarded other accounts, choose **Onboarding** from the left navigation. The Onboarding page appears. Click **+Add Azure** at the top of the page.
3. The *Add Azure Cloud Tenant* wizard starts and displays the first step: **Connect to Azure**
4. Provide the following information about your Azure account:
 - **Name:** You specify a name for the account; this name is what will appear in CloudSecure. The name should be descriptive so that you can easily identify it in CloudSecure.
 - **Tenant ID:** Paste the parent management group ID that you copied from Azure.

NOTE:

The wizard contains the following toggles:

- A toggle to enable all member subscriptions along with the tenant. If you want to onboard only some subscriptions in the tenant, set this toggle to **No**. Then go to the Onboarding page to onboard those subscriptions individually.
- A toggle to specify the type of access CloudSecure will have to your Azure tenant and the subscriptions in it. Choose **Yes** to grant the Illumio Cross Account Role permission to view your Azure tenant resources and to apply policy to them. Choose **No** to provide the Illumio Cross Account Role read-only access. To view the permissions you are granting CloudSecure to your Azure tenant, click **Download Permissions**.

The wizard advances to step two: **Set up Access**

1. Select a service account that you want to use or create a new one. Make sure to download the credentials, as they will be needed for the PowerShell script to return the Azure AD app credentials back to CloudSecure.
2. Enter the ServiceAccountToken in the appropriate field.
3. The Set up Access step includes a field containing a PowerShell command to run the `illumio-init.ps1` script in Azure. Illumio securely hosts the script so that it

can run during the onboarding process. The PowerShell command automatically appends the subscription ID you entered in the first step of the wizard.

4. To the left of the PowerShell command field, click the copy icon. The icon refreshes with a check mark on a green field indicating you successfully copied the command.
5. In a new browser window, open your Azure portal.
6. From the top taskbar, click the **Cloud Shell** icon to open a console; select the PowerShell option.
7. After Azure finishes building your Azure drive, paste the copied PowerShell command.
8. The script runs, creating a Azure AD application with a tenant scope. It adds Reader permissions and an Illumio Network Security Administrator-<subscriptionId> custom role (if you chose the ReadWrite option with the toggle mentioned above in [Provide the following information about your Azure account](#): Name: You specify a name for the account; this name is what will appear in CloudSecure. The name should be descriptive so that you can easily identify it in CloudSecure. Tenant ID: Paste the parent management group ID that you copied from Azure.).
9. Creation of the AD app registration and the roles allows CloudSecure access to the tenant resources. CloudSecure will be able to discover tenant resources and write policies for them.
10. For the complete list of permissions granted to CloudSecure for your account, see [Azure Requirements](#).
11. The script sends the Client ID and Client Secret to CloudSecure. CloudSecure accesses your Azure tenant so that you don't have to repeatedly provide your Azure credentials.
12. Leave your Azure portal and return to CloudSecure. The **Set up Access** step in the onboarding wizard should still be displayed.
13. Select the check box indicating that the "deployment" script has finished running in Azure.

NOTE:

The wizard UI obscures the pasted Client Secret. To view the secret and confirm you pasted the correct value, click the icon to the right of the field. You can only click this icon once. The secret remains obscured in all CloudSecure pages.

14. The final step of the wizard appears. This step displays a summary of the subscription information you just specified for onboarding.
15. Review the subscription information and if everything looks correct, click **Save and Confirm**. If you see issues you need to correct, click **Back** and return to that wizard step.

NOTE: CloudSecure can read flow logs from several NSGs going to the same storage account. With Azure, you can configure NSG flow logs in the same region, despite being from multiple VNets residing in different subscriptions, to be sent to a single storage account in the same region residing in a single subscription. By providing access to that specific storage account, CloudSecure can obtain and analyze flow logs for all the NSGs residing in different subscriptions. For more information on flow logs, see [Grant Flow Log Access](#).

What's Next?

When finished, the **Onboarding** page opens and displays a new row for that tenant .

For the next steps after onboarding a tenant, see [After Onboarding Cloud Accounts](#) and [What to Do Next](#).

Caveats

After tenant onboarding is complete, it will show a list of subscriptions. If a subscription belonging to a tenant is onboarded before the tenant onboarding, it will not show in the tenant's list of subscriptions. If you wish to see a subscription that you onboarded prior to the tenant onboarding, you need to delete the onboarded subscription. Upon tenant onboarding, it will automatically sync and onboard the subscription.

Grant Flow Log Access

This topic provides an overview of allowing Illumio CloudSecure access to your cloud account flow logs.

CloudSecure uses flow logs to display the flows. Granting access to flow logs allows CloudSecure to use these flow logs. For AWS you can enable SG flow logs, and for Azure you can enable NSG and VNet flow logs. For instructions on how to grant flow log access to Illumio, see the in-application help. For instructions on how to set up flow logs, see [Set up Flow Logs in AWS and Azure](#).

Prerequisites

To use this feature of the *Onboarding* page, you need the following items, which you used when you onboarded your cloud accounts:

- AWS Flow Logs

To grant access, you will need:

- Your Account ID, which you can select from a list
- Your service account name, which you can select from a drop-down menu in the *Grant Access...* dialog box
- Your CloudFormation Stack, which you need to create or download, similar to how you created or downloaded it when you onboarded your AWS account. See [Onboard an AWS Cloud Account](#) for information.

- Azure Flow Logs

To grant access, you will need:

- Your Account ID, which you can select from a list
- Your service account name, which you can select from a drop-down menu in the *Grant Access...* dialog box
- Your service account token
- Your Azure portal open in a browser window, so that you can run the PowerShell script in the *Grant Access...* dialog box. See [Onboard an Azure Cloud Account](#) for information.

Review and Grant Flow Log Access

The main *Flow Log Access* page shows your account IDs, the type of access currently in effect (None, Partial, or Full), and the available flow logs. Click **Flow Log Access** on the *Onboarding* page to begin. This action opens the *Flow Log Access* page. You can:

- Mouse over the entries to see how many logs are accessible for a given account ID, view individual log basic details, and copy individual log destinations
- Click on the "+ n more" element to expand the full list of flow logs available for a given account ID
- Filter your results by Account ID, Access, and Cloud

Review Flow Log Access Details and Grant Access

Both before and after you grant access, you may want to review the flow log access in more detail.

1. Find the *Account ID* you want and click on that row.
2. Review the details per the below guidelines. You will see two tabs where you can filter results:
 - By Log Destination Account
This tab lists destinations belonging to the selected account, along with the list of log sources pushing flows to their respective destinations. You can filter by Destination, Source, Region, and Access.
 - By Log Source Account
This tab lists flow logs going to the selected account, along with the source and the destination to which they are sent and stored. You can filter by Source, Destination, Region, and Access.
3. If you wish to grant access, click **Grant Access** and use the above prerequisite information in the *Grant Access...* dialog box, as explained in the in-application help. After granting access, you can test that access as described below in [Test Flow Log Access](#).

Guidelines for Reviewing Flow Log Access Details

The following items are guidelines for reviewing access details:

- Before granting access, the access is not granted by default, but you will still see the Flow Log ID, VPC, S3 Bucket, Region, etc.
- Once you grant access, you will see either *Granted Access* or *Partially Granted* in the *Access Status* column
- If you see that access is partially granted for an account, you might want to review your cloud account settings for any child flow logs that are listed as not granted in the access details
- You might need to refresh the *Flow Log Access Details* page immediately after granting access to make sure that the updated status appears promptly

Test Flow Log Access

You grant Illumio CloudSecure permission to access flow logs using the Grant Access feature, for which you run a script (CloudFormation for AWS and PowerShell for Azure). To ensure that Illumio CloudSecure has relevant permissions to the flow logs for which you enabled Full or Partial access, click the **Test Access** button.

This feature is enabled at the account level, where it checks relevant permissions to all the destinations (S3/storage accounts) in that account. For each account, the response tells you if the permissions are compatible with granting access. If not, you

will receive an appropriate error message detailing the issue for each destination. You can then take actions to make sure that all relevant permissions are provided to Illumio CloudSecure.

Note: The *Test Access* button is enabled for accounts with access described as Full/Partial.

Caveats

To use the traffic analysis feature in CloudSecure, you need to provide access to flow logs. It is important to note that CloudSecure does not enable or configure flow logs in your accounts during onboarding. This is something that needs to be set up in the cloud, either before or after onboarding. Note that for AWS, CloudSecure can read flows from S3 buckets only, so it is important to configure these accordingly. Once the flow logs are configured in the cloud console, the flow details will be displayed in the flow log access page of CloudSecure. It's worth noting that this might take up to 5 minutes to appear on the page. By granting access to flow logs, you will allow CloudSecure to read the flows and provide details about network traffic in the traffic analysis page.

Note the following regarding the Test Flow Log Access feature:

- For AWS, the feature works only for S3 buckets as destinations
- If the flow logs in one account are configured to be sent to a destination in a different account, the feature will give an error saying that CloudSecure cannot access it

The following are known limitations of CloudSecure's flow log reading capability:

- In AWS, CloudSecure supports reading flow logs that are stored in S3 buckets only. Currently, other storage destinations are not supported.
- For S3 storage, CloudSecure does not support reading from custom paths or nested folders inside S3 buckets
- For both AWS and Azure, if the VPC/NSG flow logs from one account are configured to be stored in S3/storage accounts in another account, then the destination account should be onboarded into CloudSecure. If the account that owns the S3 bucket is not onboarded, CloudSecure will not be able to fetch the flow logs of that S3 bucket.

Every 10 minutes the map ingests traffic flows in 60-minute chunks. Flows are shown only for completed chunks. This means that if flow log access has just been enabled, you would need to wait at least an hour to see the flows in the Cloud Map, Traffic, and

Inventory pages. However, if you enabled flow log access some time ago and already have previous 60-minute flow chunks, you would see the updated flow within 10 minutes.

Set up Flow Logs in AWS and Azure

This topic provides an overview of setting up flow logs for use by Illumio CloudSecure.

CloudSecure uses flow logs to display the flows. Granting access to flow logs allows CloudSecure to use these flow logs. For instructions on how to grant flow log access to Illumio, see the in-application help. For instructions on how to enable flow logs see [Grant Flow Log Access](#). To manually configure CloudSecure to fetch flow logs stored in custom S3 bucket paths, see [Manually Configure CloudSecure to Fetch Flow Logs](#).

AWS

You can set up flow logs in AWS using the console, a CloudFormation template, or the command line.

Using the Console

To configure flow logs for a VPC in the AWS console:

1. Go to the VPC console at <https://console.aws.amazon.com/vpc/> and select the region to which the VPC belongs.
2. Select the VPC for which flow logs are to be enabled.
3. Under the *VPC details* page, select the *Flow logs* page and click the **Create flow log** button.
4. Provide the following details in the flow log configuration page:
 - Name for the flow log config
 - Type of traffic to be filtered. For more insights, select **All**.
 - The time interval can be set to 10 minutes
5. Select **Send to an Amazon S3 bucket** and paste the ARN of the S3 bucket. It also provides the option to create a new S3 bucket from there.
6. For log record format, select any value. For more details, select **Custom format** and select all attributes. Use defaults for all other values.
7. After entering the required information click the **Create flow log** button.

Using the CloudFormation Template

To enabled flowlogs for a VPC using the CloudFormation template:

1. Go to the VPC console page at <https://console.aws.amazon.com/vpc/>, select the VPC for which the flow logs are to be enabled, and copy the VPC ID.
2. Go to the S3 console page at <https://console.aws.amazon.com/s3/> and select the bucket in which the flow logs are to be stored. Under the *Properties* tab, copy the name.
3. Save the following CloudFormation Template to a file named `enabling-vpc-flow-logs.yaml`.

```
AWSTemplateFormatVersion: "2010-09-09"
Description: "Enable Flow logs for a vpc"
Parameters:
    VpcId:
        Type: String
        Description: VPC Id for which flow logs are to be enabled
    BucketName:
        Type: String
        Description: Name of the bucket in which flow logs are to be
stored.
Resources:
    FlowLog:
        Type: AWS::EC2::FlowLog
        Properties:
            ResourceId: !Ref VpcId
            ResourceType: "VPC"
            TrafficType: "ALL"
            LogDestination: !Join
                - ""
                - ["arn:aws:s3:::", !Ref BucketName]
            LogDestinationType: "s3"
            LogFormat: "${version} ${account-id} ${interface-id}
${srcaddr} ${dstaddr} ${srcport} ${dstport} ${protocol} ${packets}
${bytes} ${start} ${end} ${action} ${log-status}"
```

```
MaxAggregationInterval: 600

Tags:
  - Key: "Name"
    Value: "FlowLogsForIllumioCloudSecure"
  - Key: "Purpose"
    Value: "Alltrafficvizualizationmap"
```

Outputs:

```
FlowLogArn:
  Description: The ARN of the created flow log
  Value: !Ref FlowLog
```

For more information, see the vendor documentation:

<https://docs.aws.amazon.com/vpc/latest/userguide/working-with-flow-logs.html>

Running the CloudFormation Template

1. Go to AWS CloudFormation service and use the template file to create a new stack with new resources (standard).
2. Select **Template is Ready** and then **Upload a template file**. Upload the enabling-vpc-flowlogs.yaml file.
3. In the next page, enter a desired stack name followed by the bucket name and VPC ID you copied before.
4. Click **Next** and leave default values in the successive pages. In the final page click **Create stack**.

After the stack creation is complete, go to the VPC console and verify the flow logs being created.

NOTE:

The template must be run in the same region in which the VPC belongs. Choose the appropriate region on top right before running CloudFormation template.

Using the Command Line

See the vendor documentation:

<https://docs.aws.amazon.com/cli/latest/reference/ec2/create-flow-logs.html>

Azure

Using the Console

See the vendor documentation:

<https://learn.microsoft.com/en-us/azure/network-watcher/nsg-flow-logs-tutorial>

Manually Configure CloudSecure to Fetch Flow Logs

This topic provides an overview of manually setting up Illumio CloudSecure to fetch flow logs stored in custom S3 bucket paths.

CloudSecure uses flow logs to display the flows. Granting access to flow logs allows CloudSecure to use these flow logs. For instructions on how to grant flow log access to Illumio, see the in-application help. For instructions on how to enable flow logs see [Grant Flow Log Access](#).

Overview

CloudSecure supports flow logs stored in a custom S3 bucket path when the permissions to read this path are provided to the CloudSecure role. You can manually add permissions to the Illumio CloudSecure role so it can fetch the flow logs present in these custom directories and provide traffic data.

When you grant permission to read flow logs from custom S3 buckets, the CloudFormation template creates two new policy documents:

- `IllumioCloudBucketListPolicy` - Grants Permission to List the items within the bucket and get the location in which bucket is located.
- `IllumioCloudBucketReadPolicy` - Grants permission to Get the objects from the list of provided buckets.

`IllumioBucketReadPolicy` does not require any change. It adds read permission to the bucket to that specific path. However, `IllumioBucketListPolicy` needs to be modified as described below.

Add a Destination with a Custom Path to the CloudFormation Template

By default, the VPC flow logs going to the S3 bucket are shown in the By Log Destination view. However, the flow log which is sent to an S3 destination within a custom

directory will be shown only in the By Log Source view. Use the following steps to add a destination:

1. Copy the name of destination, with the full path, to the custom directory e.g., `arn:aws:s3:::bucketname/customdirectory`.
2. From the By Log Destination view, click **Grant Access** to open the CloudFormation stack page in the AWS console.
3. Add the S3 bucket with the custom destination in the `S3Buckets` parameter, along with the values already populated by CloudSecure. (The parameter is a comma separated list, so add a comma before adding the S3 bucket with custom path to the existing value.)
4. Run the template.

Once the template runs successfully, the status will show as granted for the flow log configuration in CloudSecure under the By log Source page.

Update the IllumioBucketListPolicy Document

To update the document:

1. In the AWS Console, open **Services > IAM > Roles** and select the Role name. The default Role name is `IllumioCloudIntegrationRole`. Illumio created this role name when you onboarded Illumio and granted flow log access permissions.
2. Under the Permissions, select `IllumioCloudBucketListPolicy`.
3. The contents of `IllumioCloudBucketListPolicy` might look something like this:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::bucket1",
```



```

        "arn:aws:s3:::bucket2",
        "arn:aws:s3:::bucket3",
        "arn:aws:s3:::bucket4/custompath/second-dir",
    ],
    "Effect": "Allow",
    "Sid": "IllumioBucketListAccess"
}
]
}

```

4. Remove the S3 bucket with the custom prefix from the common statement, and replace it with `Sid: IllumioBucketListAccess2`. Create new statements under the policy document as seen below.

```

{
    "Action": [
        "s3:GetBucketLocation"
    ],
    "Resource": [
        "arn:aws:s3:::bucket4"
    ],
    "Effect": "Allow",
    "Sid": "IllumioBucketListAccess2"
},
{
    "Action": [
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::bucket4"
    ],

```

```

    "Condition": {
      "StringLike": {
        "s3:prefix": "custompath/second-dir/*"
      }
    },
    "Effect": "Allow",
    "Sid": "IllumioBucketListAccess3"
  }

```

5. Finally, ensure that the `IllumioCloudBucketListPolicy` file has the policies seen below.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::bucket1",
        "arn:aws:s3:::bucket2",
        "arn:aws:s3:::bucket3",
      ],
      "Effect": "Allow",
      "Sid": "IllumioBucketListAccess"
    },
    {
      "Action": [
        "s3:GetBucketLocation"

```

```

    ],
    "Resource": [
        "arn:aws:s3:::bucket4"
    ],
    "Effect": "Allow",
    "Sid": "IllumioBucketListAccess2"
},
{
    "Action": [
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::bucket4"
    ],
    "Condition": {
        "StringLike": {
            "s3:prefix": "custompath/second-dir/*"
        }
    },
    "Effect": "Allow",
    "Sid": "IllumioBucketListAccess3"
}
]
}

```

Adding this permission grants CloudSecure the required permissions to read flow logs from the custom directories under an S3 bucket.

NOTE: If you add a new flow log configuration with flows going to a new bucket (not to the old bucket for which the access is already granted) and run the CloudFormation template to grant flow access, these manual changes will be overwritten and you will need to make these changes again to resume fetching flow logs from the custom directory.

What to Do Next

This topic explains what to do next after onboarding your public cloud accounts and enabling flow log access for them.

1. Visualize Your Cloud Resources

After you've onboarded your public cloud accounts, CloudSecure begins the process of discovering and ingesting their resources.

Before defining your public cloud environment, Illumio strongly recommends that you review the resources ingested into CloudSecure. Reviewing your ingested resources helps you gain an understanding of how your cloud resources are utilized and how they are communicating.

If you are a visual person and want to analyze a graphical, hierarchical display, go to the Cloud Map and browse your resources. If you are most comfortable working with lists and sorted data, go to the Inventory page and browse your discovered resources.

- **About the Cloud Map**

In CloudSecure, the Cloud Map displays a view of your cloud inventory as a network topology map for the cloud infrastructure. The map displays the relationships between your resources by using cloud native constructs. Go to the map to view your entire state of cloud resources from the cloud accounts you have onboarded with CloudSecure.

Use the Cloud Map to view your cloud topology and analyze the traffic flow data CloudSecure captures. The map helps you visualize your cloud resources and provides an understanding of the traffic flows between them.

CloudSecure will synchronize the data in cloud accounts you have onboarded, and display the data in the Inventory, Traffic, and Cloud Map pages.

For more information, see [CloudSecure Map](#).

- **About the Inventory Page**

This page lets you view the cloud resources from accounts that CloudSecure has discovered in your environment.

The search function allows you to search and filter cloud resources on different parameters. You also view preset filters and set custom columns for viewing by selecting the options under the Cloud Details drop-down menu.

For more information, see [Inventory](#).

- **Review Your Traffic Flows**

After you onboard your cloud accounts and configure your flow log access, CloudSecure discovers all their resources and looks for traffic.

Before writing policy rules to either allow or block traffic, Illumio recommends you determine if there are any traffic flows between resources. The *Traffic* page lets you filter your resources by flow status, source labels and addresses, destination labels and addresses/ports, and so forth.

For more information see [Traffic](#).

2. Define Your Public Clouds in CloudSecure

Defining your public clouds in CloudSecure is a multi-step process:

- a. **Define your deployment stacks:**

In CloudSecure, you may decide to create deployment stacks as part of specifying which applications in your cloud account to protect with CloudSecure.

After onboarding your cloud accounts, you may begin by defining the environments you're using in the cloud. In CloudSecure, we refer to this as "adding deployment stacks." In the cloud, stacks provide a way to manage your resources as a single, atomic unit.

When you define a deployment, CloudSecure doesn't discover anything about your applications. You defined your deployment stacks separately.

For more information, see [Define a Deployment](#).

- b. **Define your applications:**

Defining an application follows a similar process. You begin by specifying an *Application* label. Then, you associate cloud resources to that label by selecting the appropriate cloud tags or cloud metadata associated with that application.

For more information, see [Define an Application](#).

c. **Approve your application definitions:**

CloudSecure separates the process of defining an application from the ability to create policy for it.

In this way, CloudSecure ensures other key stakeholders are in the loop to approve your application definitions.

For more information, see [View and Approve an Application](#).

3. Create Policy in CloudSecure

Now that you've reviewed your ingested cloud resources and defined your cloud environment in CloudSecure, you are ready to create security policy for your public clouds.

The *Policies* page lists all the different policies you have created in CloudSecure. The page contains two types of policies:

- Organization policies
- Application policies

a. **Create your organization policies:**

You can think of organization policies as guardrail policies that prevent application policies from allowing undesired traffic, or that are additive to application policies allowing desired traffic. An organization policy can exist all by itself, but these policies are also evaluated during policy computation for any application policy.

Organization policies are broader policies that you write that are independent of applications. They can override application policies, including any future application policies, that may have overly permissive allow rules.

For more information, see [Writing Organization Policy](#).

b. **Create your application policies:**

Illumio allows or denies traffic between applications using policies that you write. You can think of application policies as segmentation policies to control network traffic using Illumio labels, services, and IP/IP lists to define what can talk to applications.

For more information, see [Writing Application Policy](#).


About the Illumio Virtual Advisor

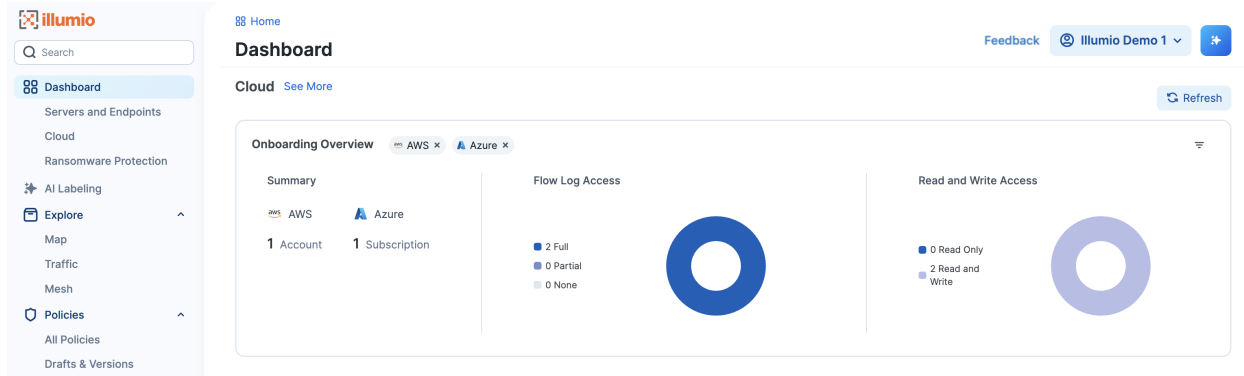
The Illumio Virtual Advisor AI chatbot helps you understand your risk exposure by using natural language questions to generate quick answers and actions to reduce your risk.

- Use IVA to visualize all traffic from a compromised server so you can take quick action to stop the risk. You can also use IVA to see all traffic from risky ports so you can preempt and prevent breaches.
- Illumio's AI-powered zero trust segmentation helps enhance workload visibility and ensure more proactive segmentation controls so you can stop attacks before they happen
 - Illumio does this with actionable guidance, automated labeling, and robust policy recommendations
 - By reducing the time and effort required to operate a Zero Trust Segmentation platform, Illumio empowers organizations to respond to the threat of AI-powered attacks by containing spread and getting insights faster
- Illumio's AI-powered segmentation strengthens cyber resilience by minimizing the attack surface
 - Complex tasks like labeling and policy creation are automated so that you can receive actionable guidance to ensure continuous protection. The internal attack surface is minimized with segments and policies that are identified automatically to protect against threats.

IVA can answer questions related to Illumio products and provide best practices recommendations to improve security based on feedback and questions generated from Illumio support.

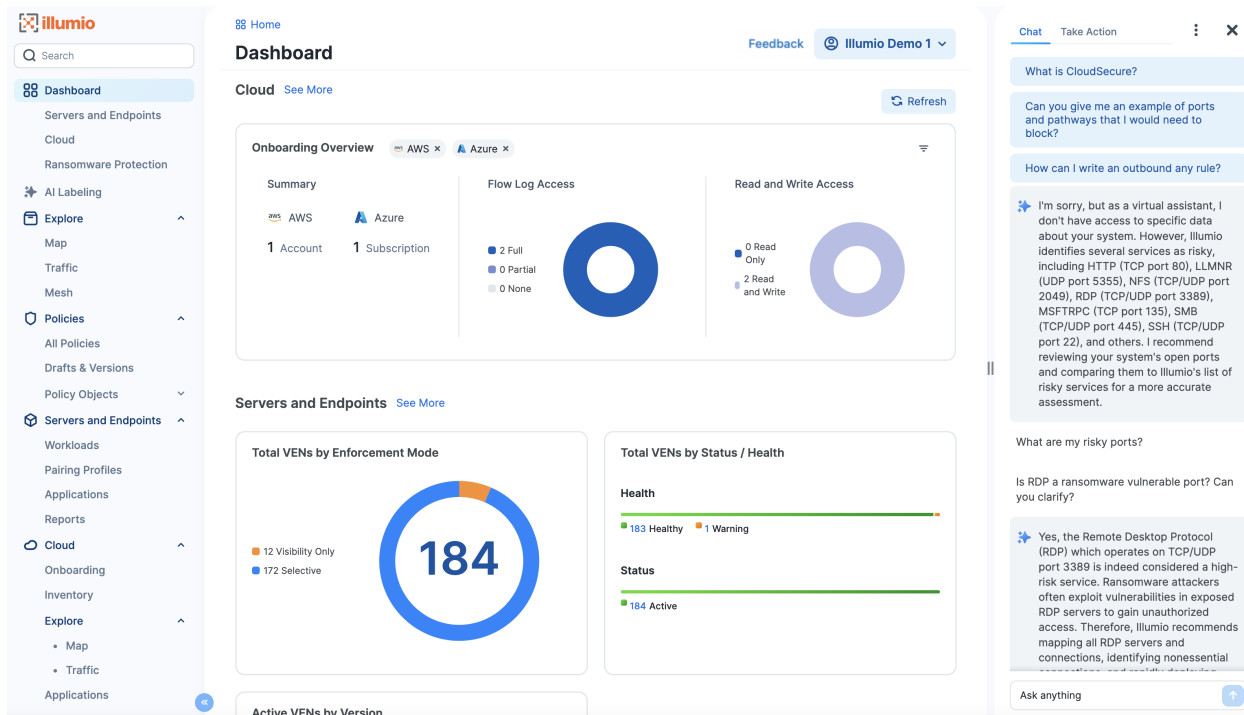
Use IVA

1. Click the IVA  located next to your name on the top, right side.
2. Ask general questions in the Chat tab, or ask specific questions about maps and web traffic in the Take Action tab.

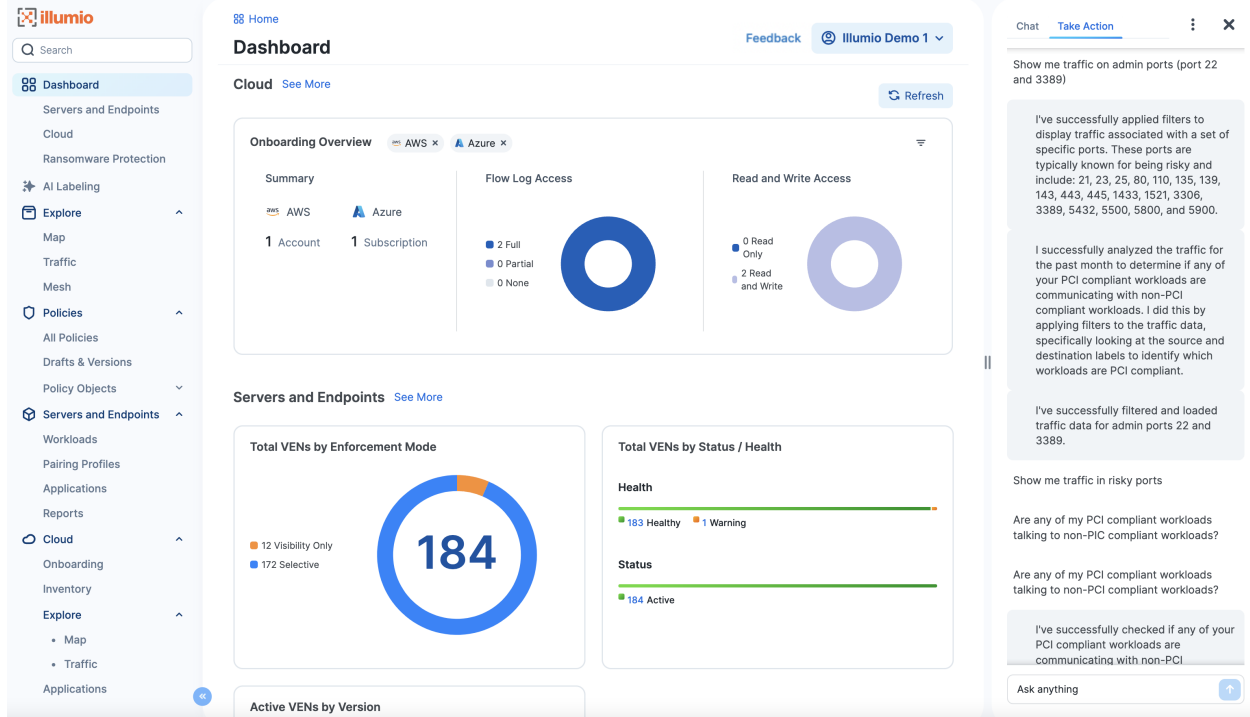


There are two tabs you can view:

- Chat view: Ask general questions and IVA provides an answer using natural language processing.



- Take Action view: This applies to questions related to Maps and Traffic only. For example, "Show me all web traffic in my production environment for the past week." Based on your input, IVA will set appropriate filters such as specific ports and generate web traffic data.



Chapter 2

Visualize Your Cloud in CloudSecure

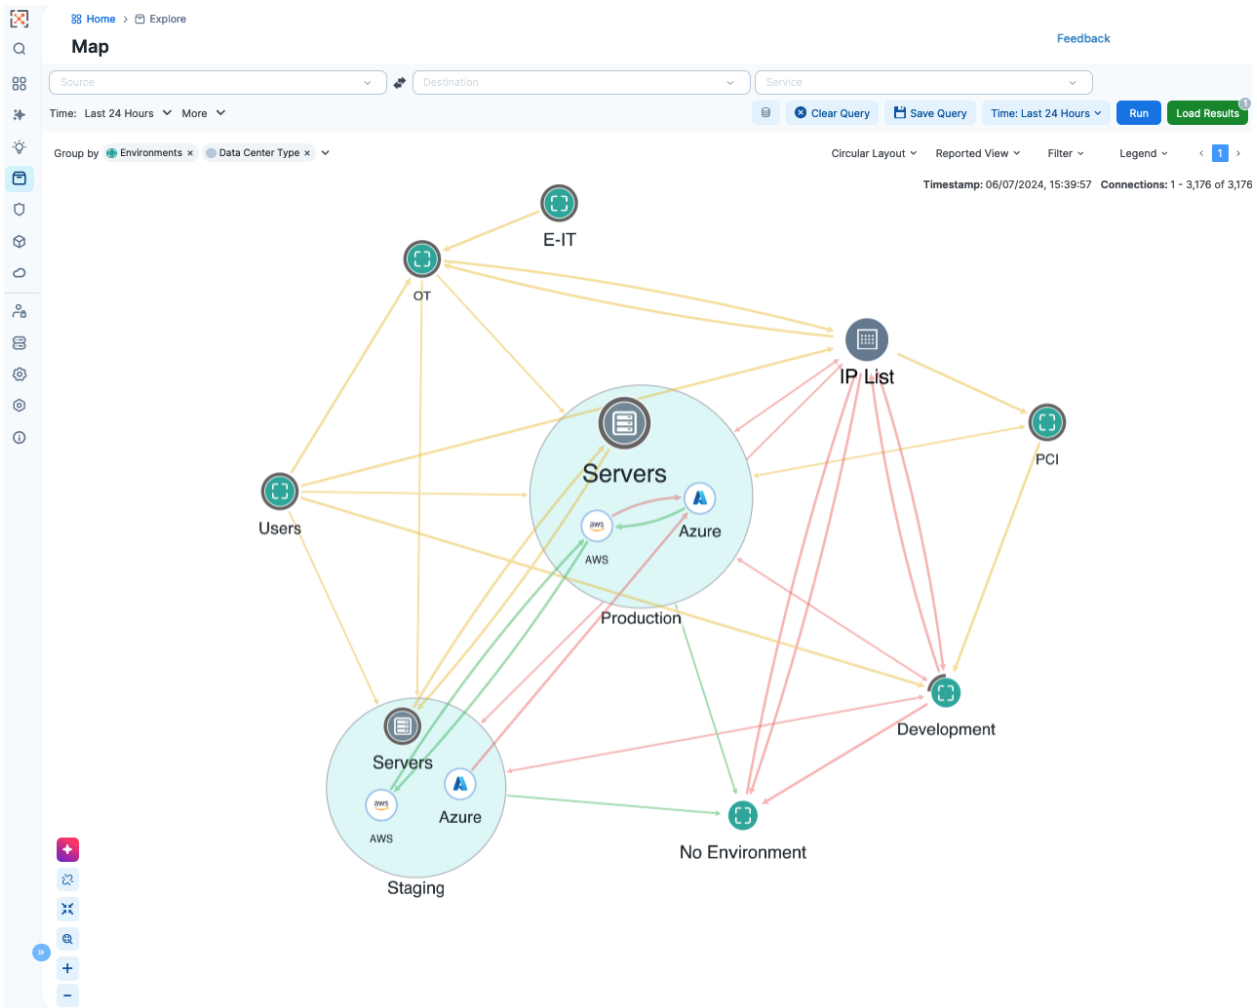
This chapter contains the following topics:

Map	50
Cloud Map	59
Inventory	75
Traffic	85
CloudSecure Dashboard	91
CloudSecure Search	93
Product Usage	95
Events	96
Reports	97

This section explains all the ways to analyze your cloud resources in Illumio CloudSecure. You work with these tools after onboarding your public cloud accounts.

Map

This topic describes the purpose of the Illumio *Map* page, found in the left navigation menu. Use the Map to visualize workloads that form logical groups (based on labels attached to workloads) and to better understand the traffic flows between workloads.



- You can hover your mouse over a cloud item, such as a region. Illumio will display information about it such as the number of resources and applications. Right-click items to see additional details.
- Left-click items to write policy for them. See [Writing Organization Policy](#).

Grouping in the Map

Groups in the Map represent a collection of workloads or services that communicate with each other and for which you can write rules. Groups are displayed in the Map after you pair workloads.

The Map displays three different types of groups: a group based on a single label, an app group, or a common set of labels.

Once you pair VENs to create workloads or connect to cloud accounts to get the cloud resources and traffic logs, PCE analyzes the workload data and the traffic data. Based on the traffic flows among your workloads, the Map organizes them into

groups. A group could represent an instance of an application running in your data center, such as an HRM application running in the Test environment in your North America data center, or a Web store in Production with its web workloads hosted in AWS and its databases hosted in your private data center.

The Map lets you group by labels, locations, app groups, etc. It also lets you split the view when in Map view mode by selecting items on the Map.

Configurable Grouping

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

NOTE:

For optimal scale and performance, if there are two connections with the same source workload, destination workload, destination port, and protocol but the process or service names are different, the two connections are combined in the Map. The process or service name that was part of the most recently reported connection is displayed.

Tips for Grouping in Your Map

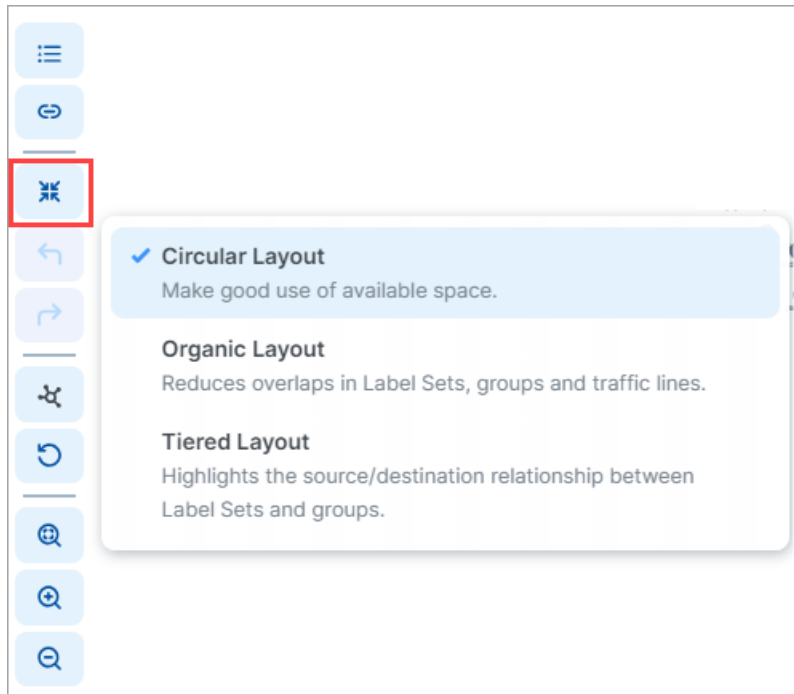
- Each group is a label set. Every workload which has the same set of labels is grouped into one of those label-sets.
- Mousing over a group in the Map displays a pop-up dialog box with the list of labels and the number of workloads using the labels.



- In the **Group by** drop-down list, you can drag and drop labels in the list to re-order how the Map displays groups. Labels at the top of the list control the prominence of those groups in the Map.
- The UI displays the groups in your Map using the colors you've selected for your labels. Use these colors to help orient yourself on the Map.

Map Layout Options

You can choose how the UI displays the Map:



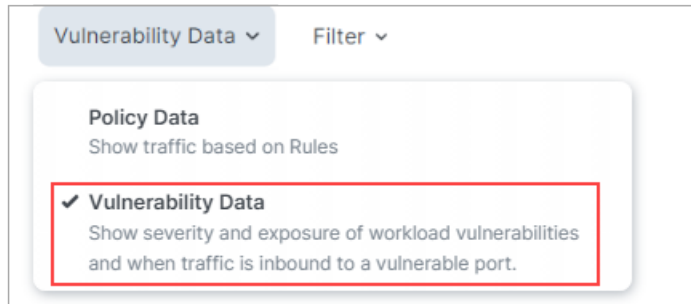
Not every layout choice is good for your Map data. See the descriptions of each layout in the Layout menu.

For example, the Organic Layout option attempts to organize groups so that the workloads that are connected are grouped together and displays less cross traffic. Workloads that are communicating are grouped together on one side of the Map and the traffic links aren't crossing as much.

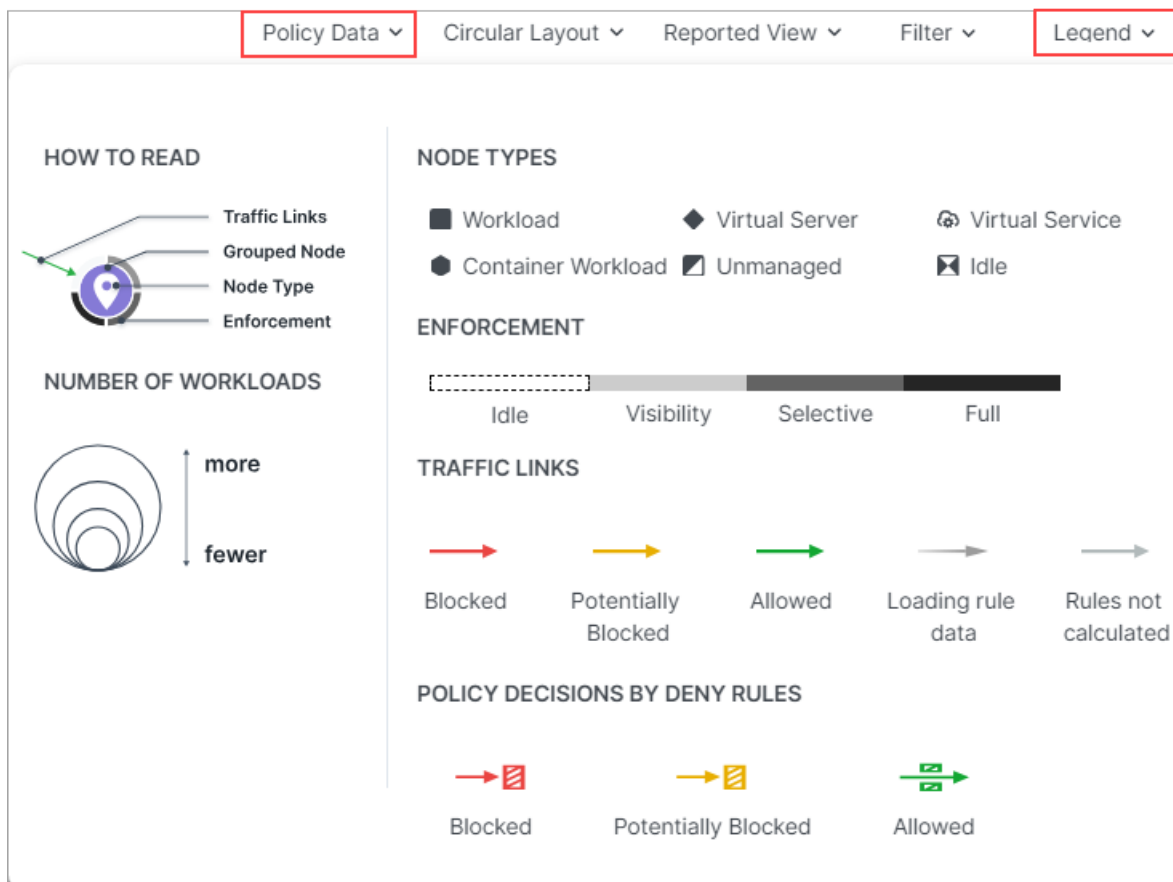
The Tiered Layout option provides a sense of traffic flow from top to bottom. The Tiered Layout option is better for smaller data sets than larger ones.

How to Read the Map Symbols

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



Legend - Policy Data



Map Symbols Explained

Number of Workloads (Policy Data and Vulnerability Data modes)

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

Enforcement (Policy Data mode)

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

- Workloads and groups inside fully dark lines are in *FullEnforcement* mode.
- Workloads and groups inside semi-dark lines are in *SelectiveEnforcement* mode.
- Workloads and groups inside light gray lines are in *Visibility only* mode.
- Workloads and groups not surrounded by any of the above-described lines are in *Idle* mode.
- The completeness of the ring around a group denotes the proportions of different enforcement states

As you navigate into the groups, you notice that the workloads also have borders indicating their enforcement modes.

Traffic Links (Policy Data mode)

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

- **Red:** Traffic is blocked
- **Yellow:** Traffic is potentially blocked
- **Green:** Traffic is allowed
- **Gradient arrows:** The light color is next to the source and dark next to the destination. Gradient arrows are used while the rule data is still loading from the traffic.
- **Grey:** Rules are not calculated

Map Reported View

The Illumio UI displays the traffic on the Map using red, orange, or green lines to indicate whether the workload had a rule that allows the traffic when the connection was attempted.

- A green line indicates that the workload had an explicit rule to allow the traffic when the connection was attempted
- A red line indicates that the workload did not have an explicit rule to allow the traffic when the connection was attempted
- An orange line indicates that no explicit rule exists, but because of the enforcement state of the workloads the traffic is not blocked when provisioned.

NOTE:

When a policy change occurs, only flows that are created after the policy change are displayed in red or green based on the new policy. Flows created before the policy change might continue to be displayed in red or green using the old policy.

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

Rules created for existing or live traffic don't change the color of the traffic lines in the Reported view, even when they are provisioned, until new traffic is detected.

Map Draft View

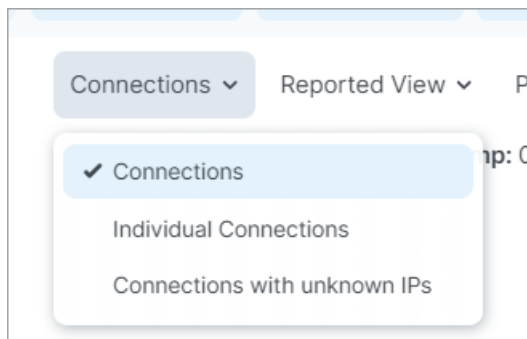
This view also displays the traffic using red, green, and orange lines to indicate whether Illumio has a rule to allow the connection that was reported by the workload. This way, you can add rules and see their anticipated effect in real-time before the rules are implemented. In the Draft view of the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

- A green line indicates that Illumio had an explicit rule (in either a draft or an active policy) to allow traffic when the connection was attempted.
- A red line indicates that Illumio did not have an explicit rule (in either a draft or an active policy) to allow traffic when the connection was attempted.
- An orange line indicates that no explicit rule exists, but because of the enforcement state of the workloads, the traffic will not be blocked when the rules are provisioned.

Filtering the Map

Connections Menu

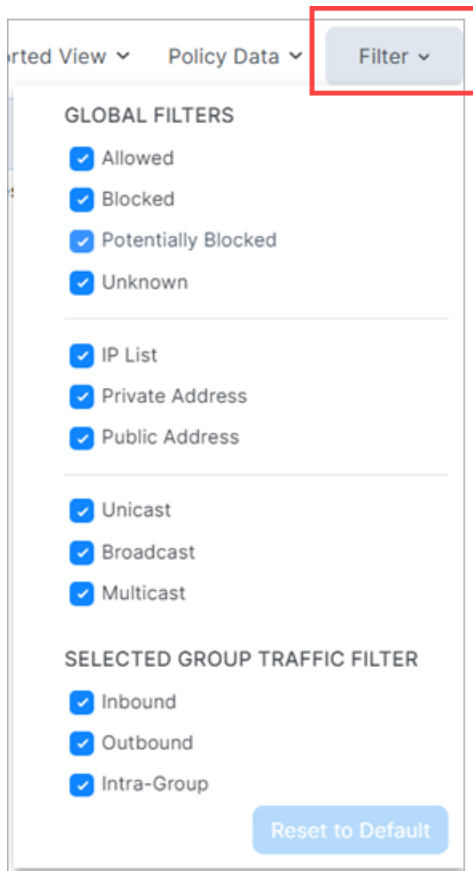
When viewing the Traffic tab in on the Connections Menu allow you to view aggregated or individual connections.



Filter drop-down

Options in the Filter drop-down allow you to control which traffic information is displayed on the Map. This is useful for controlling the overall complexity of the visual

information, making it easier to focus on the types of traffic you're interested in at any given time.



The Filter dropdown presents two types of filters:

Global Filters

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

Selected Group Filters

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

Panels in the Map

TIP:

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

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

Summary Tab

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

Traffic Tab

The Traffic tab is a summary version of the main Traffic table and filtered by what you've selected in the Map. The Traffic tab appears regardless of what you select in the Map: group types, workloads, IP lists, private addresses, public addresses, or links. By default, the Traffic tab displays the following columns.

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

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

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

Workloads Tab

The Workloads tab displays a list of all workloads in the selected group and the following information for each workload:

- Connectivity
- Enforcement
- Visibility
- Name

- Policy Sync status
- Ransomware Exposure
- Protection Coverage Score
- Labels
- When the policy was last applied

As you drill in and out of the groups in the Map, the Workloads tab adjusts to show the workloads in the super set group.

Virtual Services Tab

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

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

You can add or remove columns by using the *Customize columns* drop-down list.

Cloud Map

This topic explains how to work with the Cloud Map in CloudSecure, found in the *Cloud > Explore* menu.

What is the Cloud Map?

Organizations can find it difficult to understand their cloud topology. For example, understanding the relationships between the objects and related components such as security groups, tags, and other metadata in your cloud accounts is challenging. CloudSecure is designed to handle this challenge. CloudSecure analyzes these relationships to provide a view of assets with proper cloud hierarchy.

In CloudSecure, the Cloud Map displays a view of your cloud inventory as a network topology map for the cloud infrastructure. The map displays the relationships

between your resources by using cloud native constructs. Go to the map to view your entire state of cloud resources from the cloud accounts you have onboarded with CloudSecure.

Use the Cloud Map to view your cloud topology and analyze the traffic flow data CloudSecure captures. The map helps you visualize your cloud resources and provides an understanding of the traffic flows between them.

CloudSecure will synchronize the data in cloud accounts you have onboarded, and display the data in the Inventory, Traffic, and Cloud Map pages.

Supported Resource Types

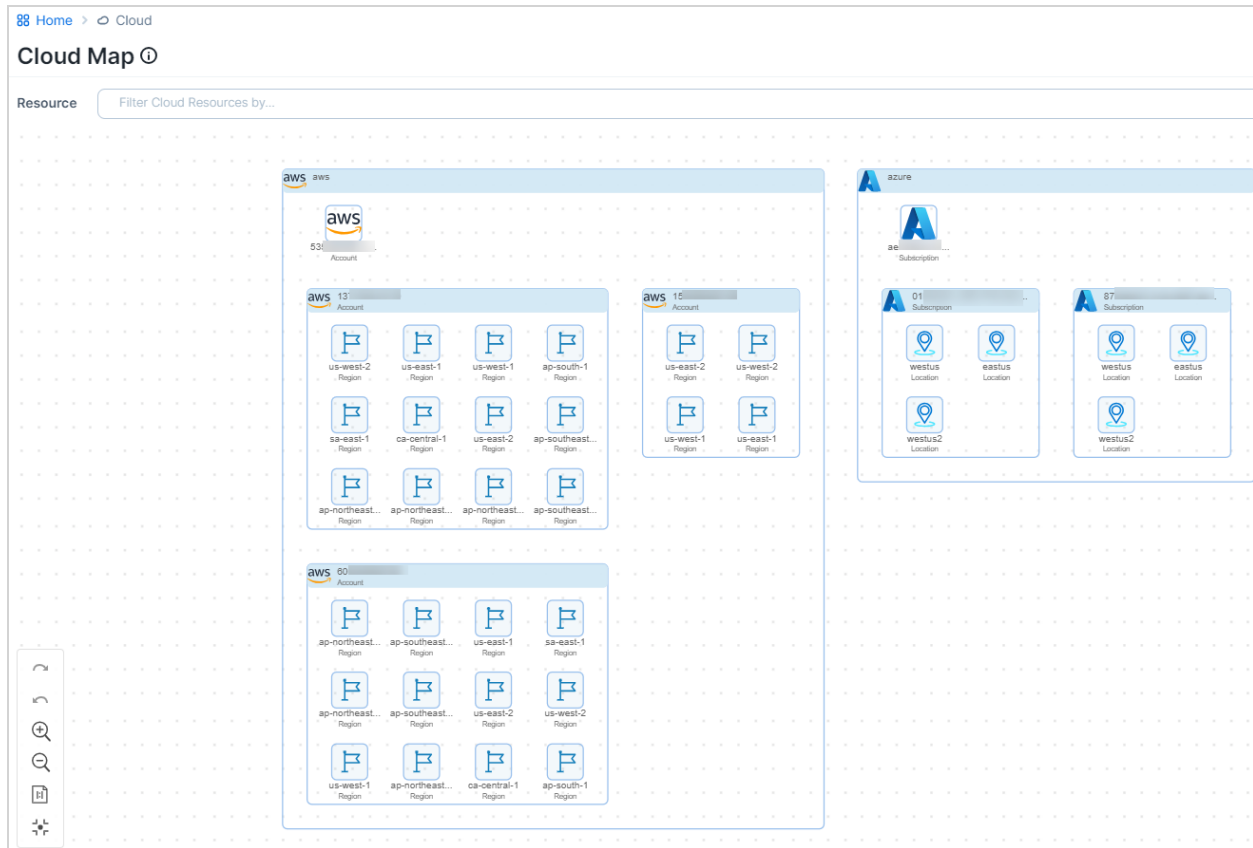
See [Cloud Map Supported Resources](#). For a list of resources against which you can write policy, see [Resources that Support Policy](#).

How the Cloud Map is Organized

CloudSecure organizes the map first by cloud — AWS versus Azure. Each public cloud has its own grouping in the map.

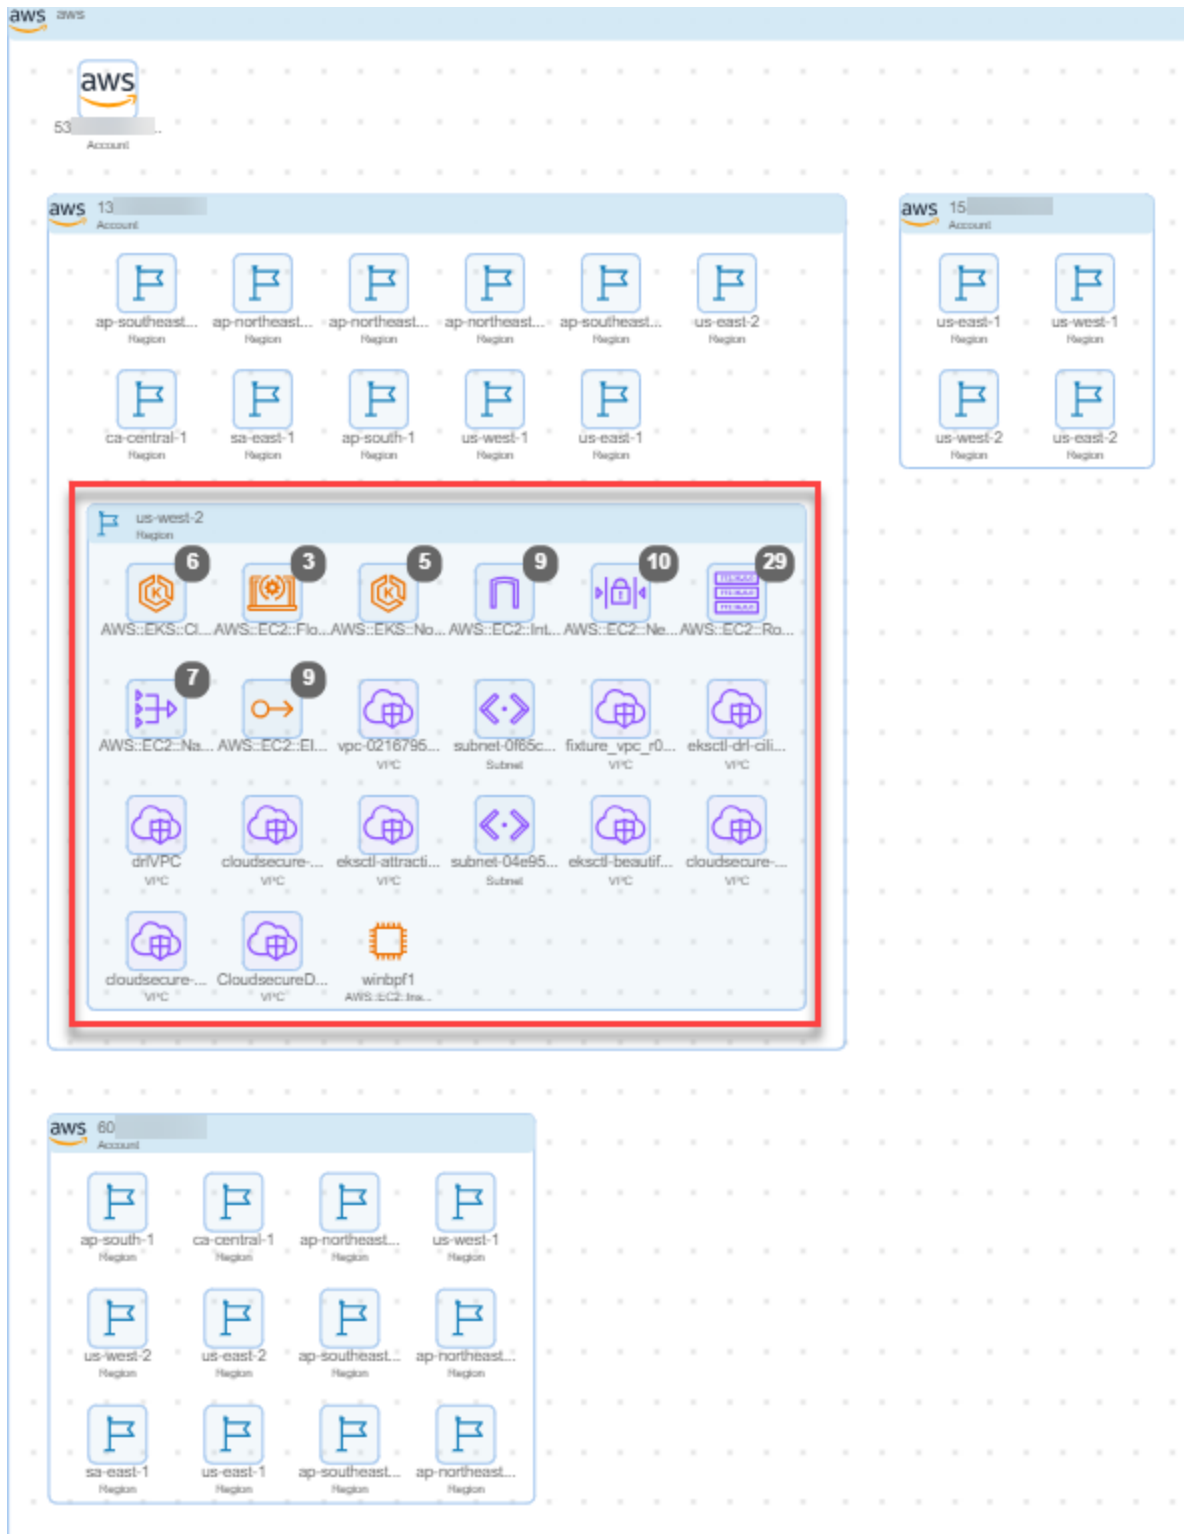
The map organization continues to get progressively more granular and displays resources in this hierarchy:

Region (Location) → VPC (VNet) → Subnet → Resources



The map displays your resources within the regions. This example shows us-west-2 region in your AWS 13##### account.

When you zoom in to view a region, you see the number of resources in that region. The map tells you the count of the resources.



Each region of the map contains the following types of objects:

- **Cloud Hierarchy Combo**

This can be a cloud, account, region, VPC, or subnet that contains other resources. For example, a VPC combo can contain a subnet, and a subnet combo can contain an EC2 instance.

- **Resource Combo**

This is a group of resources of the same type, indicated with a number.

- **Resource Node**

This is an individual resource.

How to Navigate the Cloud Map

Ways to Move Around Your Map

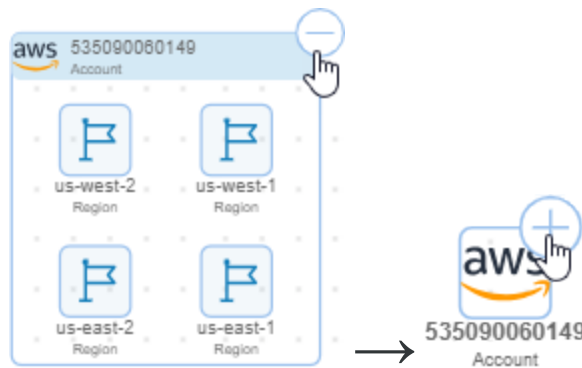
You have these ways to navigate the Cloud Map:

- Use the filters at the top of the page to locate and zoom in to specific areas or resources; see [Filtering the Map](#).
- Click anywhere in the map to refocus the view to that level. For example, you have zoomed in to an object. Click outside the cloud groups to refocus on the full map.
- Use the built-in map tools to zoom in:
 - Click the plus (+) for a group to expand it:




To collapse a group so that you it's not expanded and you see the

resources within it, click the minus (-) icon.



- Click the white space within a group to zoom in:



- Use the map tools in the bottom left-hand corner to:
 - Map Configurations
See [Configuring Your Map View](#).
 - Zoom map so everything fits on the page
 - Zoom in
 - Zoom out




Filtering Your Map Resources

At the top of the page, the map includes a **Resource** filter. You can set one of several filters to show or hide different elements of your data and focus your map on what is most important to you.

The **Resource** filter includes several options, including Cloud, Account ID, Region, Object Type, VPC/VNET ID, Subnet ID, Cloud Tags, and others.

By default, when you first open your Cloud Map, the **Resource** filter is empty. The map displays groups for each of the clouds you have onboarded — AWS and Azure. Next, it displays the accounts you've onboarded from each of those public clouds.

When you are filtering for resources that support displaying traffic flows, the map includes a traffic filter to help you narrow the traffic flows to display:

**IMPORTANT:**

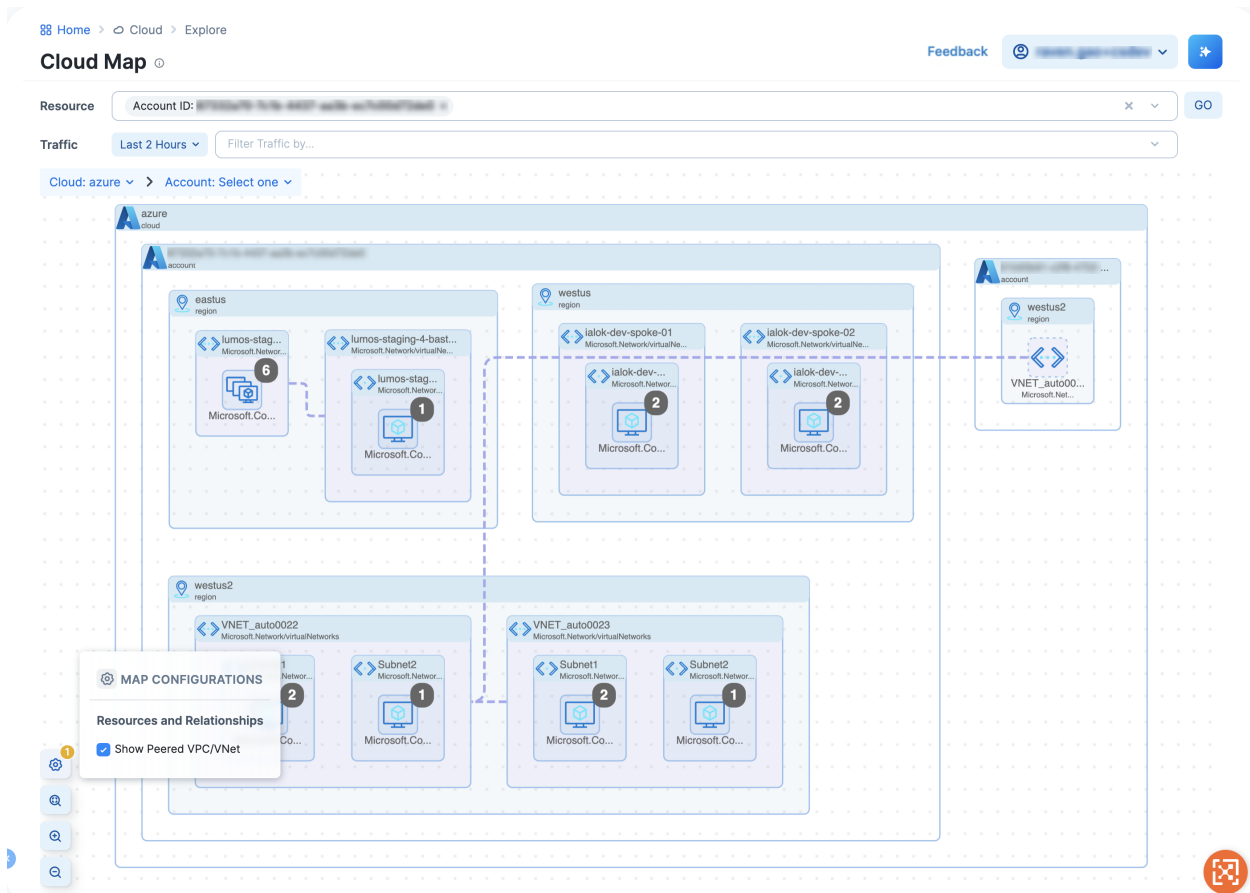
As you use the filters to manipulate the map display and display details about accounts and the resources in them, CloudSecure might display a message that it can't display all the results for your query because your filter results would display more than 2,000 resources or more than 10,000 traffic flows. When this happens, refine your query so that it is more focused and returns fewer results.

For information, see [Limitations for Using the Cloud Map](#).

Configuring Your Map View

Click the **Map Configurations** button, which has the gear icon, to open the *Map Configurations* panel. Under the *Resources & Relationships* portion of the panel you will see checkboxes for showing relationships between specific types of resources. These are unchecked by default. If you check one or more of them, the map will stop displaying anything from your filtered results that does not correspond to checked boxes.

For example, if you check the box for *Show peered VPC/VNet*, all resources *not* associated with a peered VPC/VNet will be hidden, as seen in the following figure.



Some things to remember:

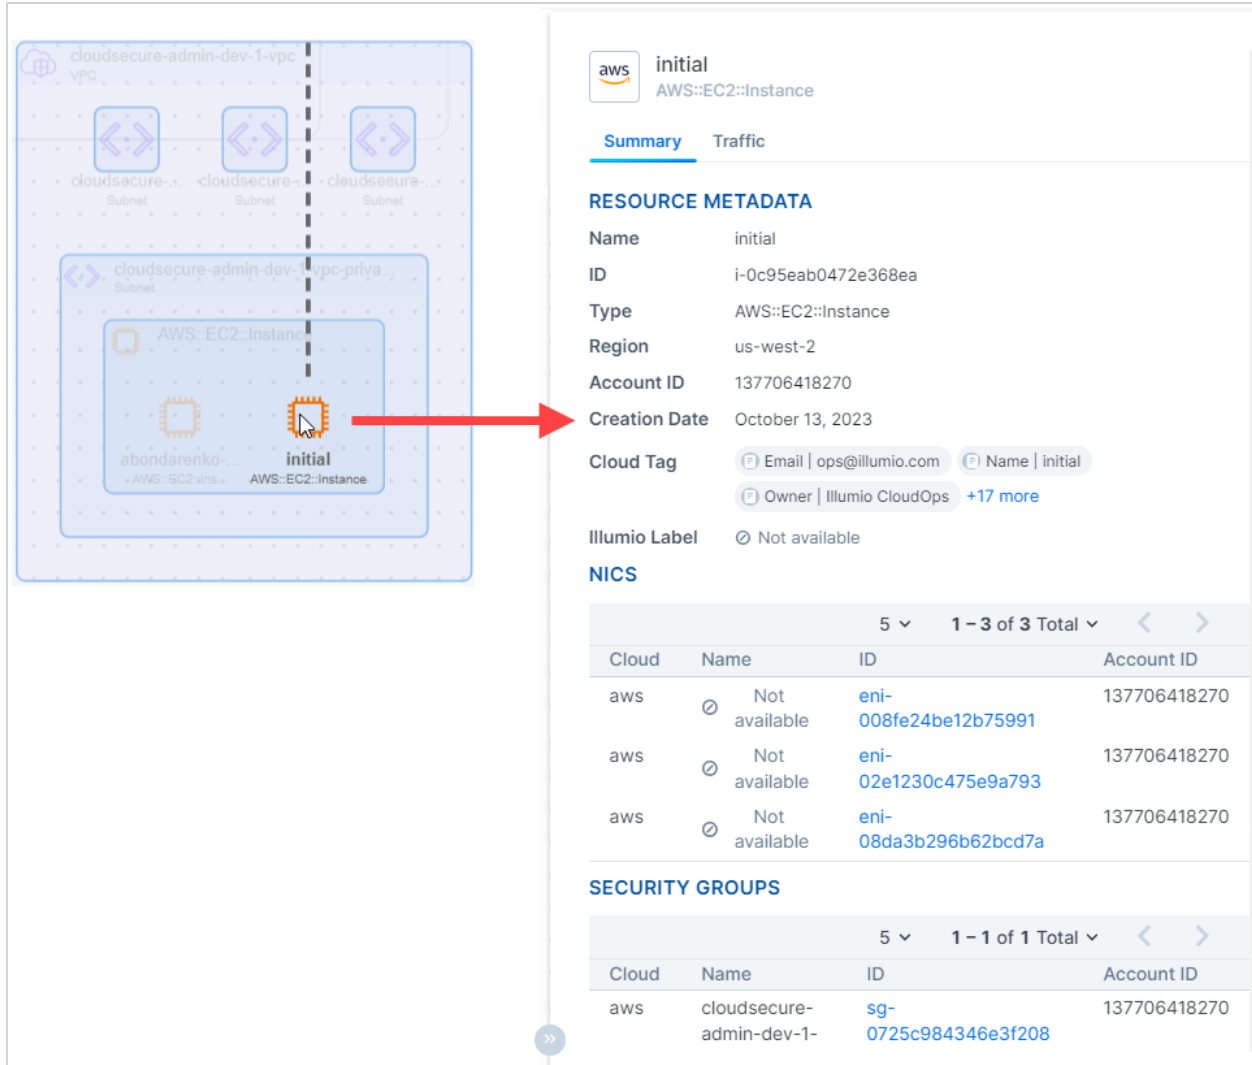
- The *Map Configurations* button will appear only when your filtered results contain items that have corresponding checkboxes in the *Map Configurations* panel
- If you check one or more boxes in the panel, a numeral appears in the upper right-hand corner of the *Map Configurations* button to remind you that you have non-default view configurations in place

The panel has checkboxes for the following resources:

- Peered VPCs/VNets
- VPC/Private Endpoints and associated resources

Display Resource Side Panel

When you click a resource in the map, CloudSecure opens a right-side panel that displays the resource metadata. For example, you can click an EC2 instance to see a summary information about the resource.



The screenshot shows the Cloud Map interface. On the left, a VPC diagram titled 'cloudsecure-admin-dev-1-vpc' contains several subnets: 'cloudsecure-admin-dev-1-vpc-private', 'cloudsecure-admin-dev-1-vpc-public', and 'cloudsecure-admin-dev-1-vpc-private'. Inside the private subnet, there are two EC2 instances: 'abondarenko-...' and 'initial'. A red arrow points from the 'initial' instance to the right-hand panel.

The right-hand panel displays the details for the 'initial' resource, which is an AWS::EC2::Instance. The panel has two tabs: 'Summary' (selected) and 'Traffic'. The 'Summary' tab shows the following information:

- RESOURCE METADATA**
 - Name: initial
 - ID: i-0c95eab0472e368ea
 - Type: AWS::EC2::Instance
 - Region: us-west-2
 - Account ID: 137706418270
 - Creation Date: October 13, 2023
 - Cloud Tag: Email | ops@illumio.com, Name | initial, Owner | Illumio CloudOps, +17 more
 - Illumio Label: Not available
- NICS**

Cloud	Name	ID	Account ID
aws	Not available	eni-008fe24be12b75991	137706418270
aws	Not available	eni-02e1230c475e9a793	137706418270
aws	Not available	eni-08da3b296b62bcd7a	137706418270
- SECURITY GROUPS**

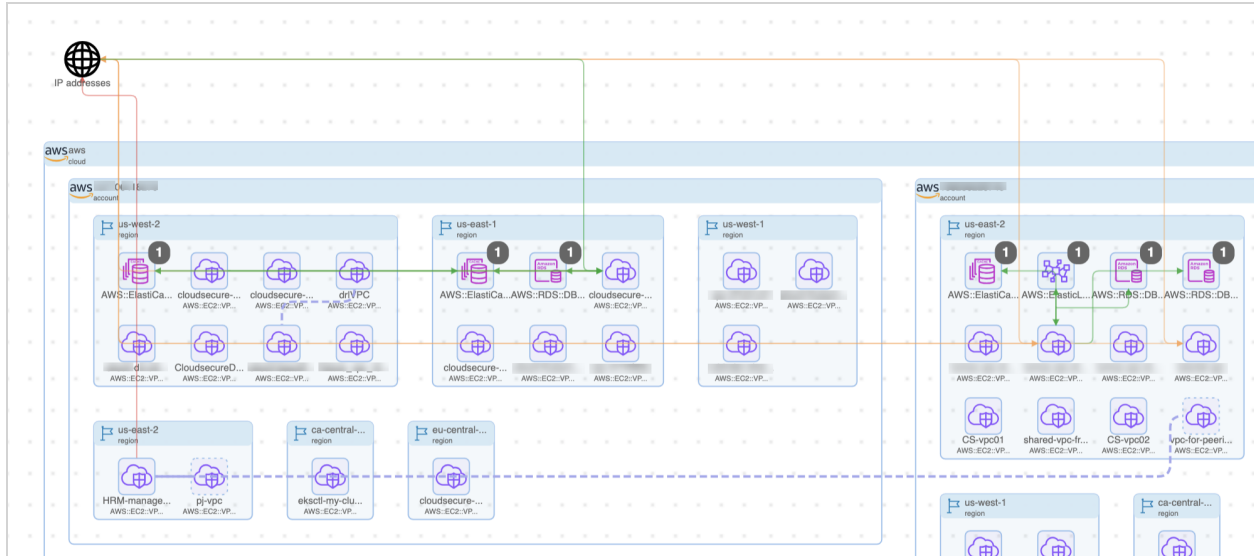
Cloud	Name	ID	Account ID
aws	cloudsecure-admin-dev-1-	sg-0725c984346e3f208	137706418270

When you open a VM (Azure) or an EC2 instance (AWS) the right panel will include a **Traffic** tab. The Traffic tab displays when that resource is sending or receiving traffic. In the tab, you can view information for the flows, such as source and destination, label sets, port/protocol, associated security groups, packet counts, etc.

At this time, the Cloud Map only supports displaying traffic data for VM (Azure) and EC2 instance (AWS) resources. For resources that don't support displaying traffic flows, the panel includes a **Summary** tab only.

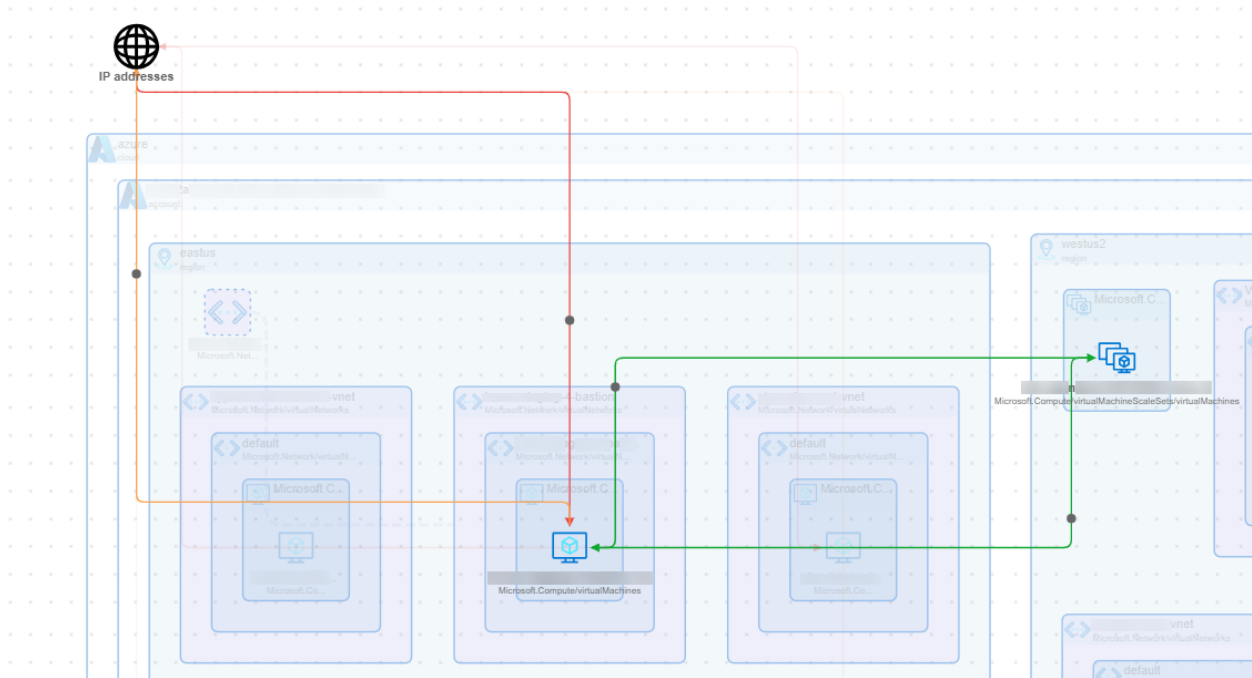
Cloud Map Traffic Lines

The Cloud Map includes solid traffic lines for resources that are sending and/or receiving traffic. Flows that are one direction are displayed with a single arrow line. Bidirectional flows have dual arrows.



Orange lines indicate mixed state (both denied and allowed) traffic. Green lines indicate allowed traffic. Red lines indicate denied traffic. These traffic lines are displayed from the lowest level node selected. For example, you may have green lines between two regions, indicating that strictly regional traffic is enabled. However, if you drill down, you might see a pair of resources, one in each region, with mixed state traffic between them. Dotted lines indicate relationships rather than flows.

When you select a traffic line, a *Traffic Details* panel will open, showing flow status, source, destination, and the like. When you hover over a traffic line, the map shows an animation of the traffic flow for just that traffic line. Similarly, when you hover over a resource displaying a traffic line, the map refreshes with an animation of the traffic flow for just that resource. This animation isolates the traffic flow for only the resource of interest. Using hover is a good way to isolate a resource and see at a glance all the flows from that point of view coming from and going to that resource. To stop the animation, simply move your cursor to another part of the map.



Limitations for Using the Cloud Map

After onboarding an account, the resources within the map begin to display within the cloud map within five to ten minutes. During this time, your map displays the message “No resources available yet.”

When the map loads, CloudSecure limits on the number of objects that the map will display.

- **Resources:** 2,000 objects
- **Traffic:** 10,000 flows

These display limitations are not configurable. After you onboard your cloud accounts, CloudSecure discovers all their resources. To provide optimal map display performance, Illumio sets these display limitations. These limitations are a UI limitation only. You can filter your map to retrieve data about resources that aren’t initially displayed when you elect to view your full map. See [Filtering Your Map Resources](#) for information.

When you encounter this display limitation, the map includes a information message informing you to filter your map to see more resources. For example, the following message indicates the current map view is not displaying all traffic flows.

Note The traffic results are partial due to the current limitation of 10000 results. Please refine the filters.

Caveats

Every 10 minutes the map ingests traffic flows in 60-minute chunks. Flows are shown only for completed chunks. This means that if flow log access has just been enabled, you would need to wait at least an hour to see the flows in the Cloud Map, Traffic, and Inventory pages. However, if you enabled flow log access some time ago and already have previous 60-minute flow chunks, you would see the updated flow within 10 minutes.

Cloud Map Supported Resources

The *Cloud Map* displays the following resources. For information on the resources for which Illumio CloudSecure supports policies, see [Resources that Support Policy](#)

AWS

Resource	Category	Attached Resources on Details Panel	Map Display Notes	Flow Support
VPC	Network Management	Elastic Network Interface (ENI), Subnet, Virtual Private Cloud (VPC) peering, peered VPC	Displays at Region level	No
Subnet	Network Management	Subnet	Displays at VPC level	No
EC2 Instance	Compute	ENI, Subnet, VPC, Security Group (SG), Elastic IP, Elastic Block Storage (EBS) Volume, Target Group	Displays at Subnet level	Yes
EKS Cluster	Containers	ENI, Subnet, VPC, SG, Node Group	Displays at Region level	Yes
EKS Nodegroup	Containers	Cluster	Displays at Region level	No
VPC Endpoint	Network Routing	S3, S3 Bucket Policy, VPC, ENI, SG, Subnet	Displays at VPC level	Yes
RDS DB Cluster	Databases	ENI, Subnet, VPC, SG	Displays at VPC level	Yes
RDS DB Instance	Databases	ENI, Subnet, VPC, SG	Displays at	Yes

			VPC level	
ECS Cluster	Containers	N/A	Displays at Region level	No
ECS Container Instance	Containers	N/A	Displays at Region level	No
Glacier Vault	Storage	N/A	Displays at Region level	Yes
ElastiCache CacheCluster	Databases	N/A	Displays at Region level	Yes
MemoryDB Cluster	Databases	N/A	Displays at Region level	Yes
Spot Fleet Request	Compute	EC2 Instance, SpotInstanceRequest	Displays at VPC level	No
Spot Fleet Instance Request	Compute	EC2 Instance	Displays at VPC level	No
S3 Bucket	Storage	Bucket Policy, VPC Endpoint, VPC, ENI, SG, Subnet	Displays at Region level	No
ElasticLoadBalancingV2 Load Balancer	Network Routing	ENI, Subnet, VPC, SG, Target Group	Displays at Region level	No
VPC Peering	Network Management	VPC	Does not display as a resource node, but does display as a relationship	No
DynamoDB Table	Databases	N/A	Displays at Region level	No
Redshift Cluster	Data Warehouse	ENI, Subnet, VPC, VPC Endpoint, SG, Network Interface Controller (NIC)	Displays at VPC level	No
Lambda Function	Serverless	Subnet, VPC, SG, Key Management Services (KMS) key	Displays at VPC level	No

Azure

Resource	Category	Attached Resources on Details Panel	Map Display Notes	Flow Support
Virtual Network	Network Management	NIC, IP Config, Subnet	Displays at Region level	No
Subnet	Network Management	VNet, IP Config, NIC, Network Security Group (NSG)	Displays at VNet level	No
Storage Account	Storage	Private Endpoint, NIC, Subnet, VNet, NSG	Displays at Region level	No
Virtual Machine	Compute	NIC, NSG , IP Config, Subnet, VNet, VM ScaleSet	Displays at Subnet level	Yes
SQL Server	Databases	Private Endpoint, NIC, Subnet, VNet, NSG, SQL Server Databas	Displays at Region level	No
Virtual Machine ScaleSet	Compute	VM Scaleset	Displays at Region level	No
Load Balancer	Compute	N/A	Displays at Region level	No
Private Endpoint	Network Management	NIC, Subnet, VNet, Azure PaaS resources	Displays at Subnet level	No
DocumentDB Database Account	Databases	Private Endpoint, NIC, Subnet, VNet, NSG, SQL Database, DocumentDB Table, Document DB Gremlin Database, DocumentDB Cassandra Keyspace, DocumentDB Mongo Database	Displays at Region level	No
DocumentDB Cas-	Databases	N/A	Displays at	No

sandra Cluster			Region level	
DocumentDB Mongo Cluster	Databases	Private Endpoint, NIC, Subnet, VNet, NSG	Displays at Region level	No
DBforPostgreSQL ServerGroup V2	Databases	Private Endpoint, NIC, Subnet, VNet, NSG	Displays at Region level	No
DBforPostgreSQL Flexible Server	Databases	Private Endpoint, NIC, Subnet, VNet, NSG, DBforPostgreSQL Flexible Server Database	Displays at Region level	No
DBforPostgreSQL Server	Databases	Private Endpoint, NIC, Subnet, VNet, NSG, DBforPostgreSQL Server Database	Displays at Region level	No
Web Site	Serverless	Subnet, Web Site Function	Displays at Region level	No
Web Site Function	Serverless	Web Site	Displays at Region level	No
Virtual Network Peer-Networking	Network Management	VNet	Does not display as a resource node, but does display as a relationship	No
NAT Gateway	Network Management	Subnet, Public IP, Public IP Prefix	Displays at Region level	No
Redis Cache	Databases	VNet, Subnet, Private Endpoint	Displays at Region level	No

NOTE:

Because CloudSecure may not always discover elastic network interfaces (ENIs), a flow search based on resource IDs will not work for the following supported resources if their *Details* page does not display the ENI. The workaround is to search using the IP address of the associated ENI, if known:

- AWS RDS DBInstances
- AWS RDS DBClusters
- ElasticLoadBalancingV2 load balancers
- AWS MemoryDB clusters
- AWS ElastiCache for Redis clusters
- AWS Redshift clusters

Inventory

This topic describes the purpose of the Illumio CloudSecure Inventory feature, and provides a general example of how you would use it. For instructions on how to use the search function in the *Inventory* page, see the pop-ups in the CloudSecure GUI.

Supported Resource Types

See [Inventory Supported Resources](#). For a list of resources against which you can write policy, see [Resources that Support Policy](#).

Use Case and Example

Illumio CloudSecure discovers your resources when cloud onboarding is done. This feature lets you search through a table of your discovered resources. You might want to confirm general expectations of what resources you have, want to know what is in a given region, or be interested in a specific type of resource.

For example, suppose you are interested in reviewing a particular virtual machine, like an AWS EC2 instance. The following steps illustrate how you would do that.

1. The first part of the sequence might be to filter by *Object Type* and select **AWS::EC2::Instance**:

This filter would return a list of EC2 instances. Depending on how you customize your columns, you might see:

- Cloud type
- Name and ID

- Resource State
- Account ID

And many other characteristics. You can also choose one of the preset column customizations, including *Cloud Details*, *Labels and Cloud Tags*, and *Security Controls*.

2. The next step in the sequence would be to click one of the entries in the *Name and ID* column. In the case of an EC2 instance or VM, you will see additional information, beyond the general information, listed in the *Attached Resources* tab. That tab displays the following information:
 - NICs
 - Security Groups
 - Subnets
 - Traffic

Selecting an *ID* column entry in any of those headings will show details for that entry such as its state or creation date.

For more information on *Inventory* page search, see [CloudSecure Search](#).

VPC/VNet Peering Details

VPC and VNet peering connection details are provided in the *Details* pages of VPC and VNET resources in the inventory list.

VPC/VNet Peering Guidelines

- You can click on any of the peered VPCs or VNets to see further details
- The requester/acceptor is defined by the peering connection, so the current VPC or VNet can either be a requester or an acceptor
- VPCs and VNets can be peered across accounts. For example, this means you could have two VPC connections, with one VPC in each of the two accounts, but only one peering relationship. Note that to see the full details, you must have *both* accounts onboarded. For cross-account VPC/VNet connections, if you do not have both accounts onboarded, you will still see the peering connection, but the details of the non-onboarded peer (attached resource) will display only its CSP ID rather than a link to an inventory resource.
- If you do not have both accounts onboarded, you will still see the peering connection, but the details of the non-onboarded peer (attached resource) will display only its ID rather than a link

- Cross-account peering connections for AWS VPCs have the same CSP ID, but cross-account peering connections for Azure VNets will have a different CSP ID for each VNet because Azure CSP IDs include account information within the CSP ID

Security Control Resource Details

Inbound/Outbound rules are featured for security control resources, including:

- AWS Security Groups
- Azure Network Security Groups
- AWS Network ACLs

On the *Details* page of any security control resource, you will see two additional tabs: *Inbound Rules* and *Outbound Rules*.

- Inbound rules: these control the incoming traffic that's allowed to reach the instances associated with the security group
- Outbound rules: these control the outgoing traffic from your instances

Each of these rules will contain information such as source/destination, port/port range, protocol, etc.

NOTE:

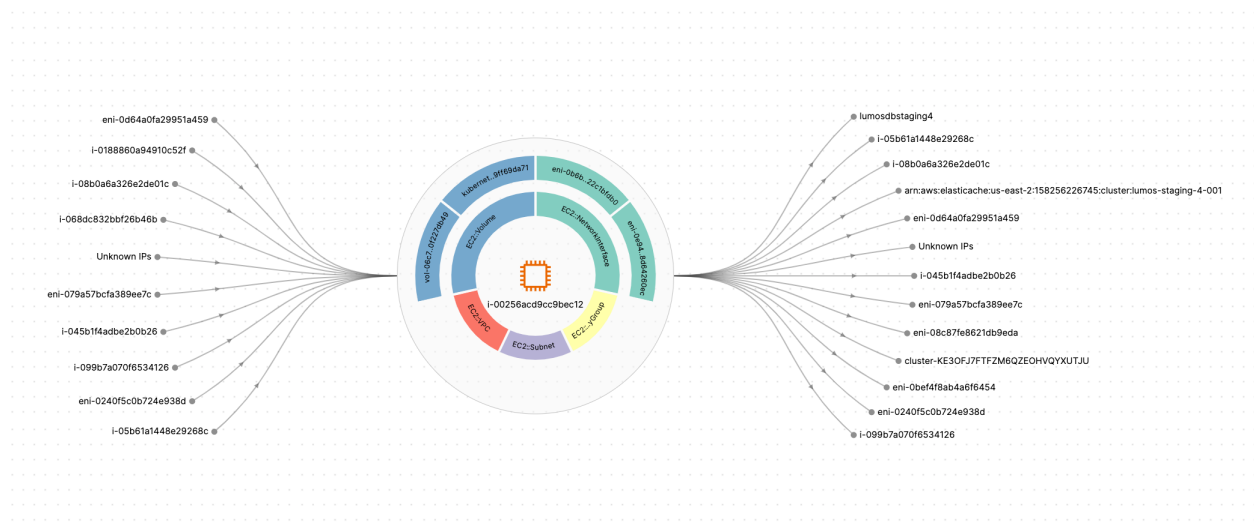
Although AWS security group rules and Azure network security rules are visible on the *Details* page for AWS security groups and Azure network security groups, Azure network security group rules created before July, 2021 will not appear in the *Details* page. This is because CloudSecure does not ingest rules created without resource IDs. If any of your rules do not appear due to this issue, recreating the rule will allow it to display.

Details Resource Graph

When you click on the details for a given resource, you can go to the *Resource Graph* tab for a visual representation of that resource's relationships to sources, destinations, and attached resources. For example, if you selected the graph for an EC2 instance you could see:

- The EC2 instance depicted in the center of a series of concentric rings
- An inner ring, depicting each of the attached resources such as subnets, VPCs, security groups, and network interfaces
- An outer ring, depicting the individual instances of the attached resources shown in the inner ring. For example, you might see an outer ring listing one or more individual network interfaces and their ID numbers.
- A series of incoming flow lines from the left, depicting sources such as other EC2 instances, ENIs, IPs, and so forth, for which the EC2 instance in the center is the destination
- A series of outgoing flow lines to the right, depicting destinations such as RDS DB clusters, ENIs, IPs and so forth, for which the EC2 instance in the center is the source

The following figure provides an example.



Inventory Supported Resources

The *Inventory* page displays the following resources. For information on the resources for which Illumio CloudSecure supports policies, see [Resources that Support Policy](#)

AWS

Resource	Category	Attached Resources on Details Page	Flow Support
----------	----------	------------------------------------	--------------

VPC	Network Management	Elastic Network Interface (ENI), Subnet, Virtual Private Cloud (VPC) peering, peered VPC, RAM ResourceShare	No
Subnet	Network Management	Subnet, RAM ResourceShare	No
Network Interface	Network Management	All AWS workload types, Subnet, VPC, SG	Yes, on <i>Traffic</i> page
Security Group	Network Security	All AWS workload types, ENI	No
Security Group Rules	Network Security	N/A	No
Route Table	Network Routing	NAT Gateway, VPN Gateway, VPC, Subnet, Internet Gateway, VPC peering, EC2 Instance, ENI	No
Network ACL	Network Security	N/A	No
NAT Gateway	Network Routing	N/A	No
Internet Gateway	Network Routing	N/A	No
Elastic IP	Network Management	ENI, EC2 Instance	Yes, on <i>Traffic</i> page
EC2 Instance	Compute	ENI, Subnet, VPC, Security Group (SG), Elastic IP, Elastic Block Storage (EBS) Volume, Target Group	Yes
Flow Log	Network Monitoring	N/A	No
EKS Cluster	Containers	Subnet, VPC, SG, EKS Node Group	Yes
EKS Nodegroup	Containers	EKS Cluster	No
Carrier Gateway	Network Routing	Gateway	No

Customer Gateway	Network Routing	Gateway	No
EC2 Instance Connect Endpoint	Network Routing	N/A	No
VPC Endpoint	Network Routing	S3, S3 Bucket Policy, VPC, ENI, SG, Subnet	Yes
VPN Gateway	Network Routing	Gateway	No
Egress Only Internet Gateway	Network Routing	Gateway	No
EBS Volume	Storage	N/A	No
RDS DB Cluster	Databases	ENI, Subnet, VPC, SG, KMS Key	Yes
RDS DB Instance	Databases	ENI, Subnet, VPC, SG, KMS Key	Yes
RDS DB Security Group	Network Security	N/A	No
ECS Cluster	Containers	N/A	No
ECS Container Instance	Containers	N/A	No
Glacier Vault	Storage	N/A	Yes
ElastiCache CacheCluster	Databases	N/A	Yes
MemoryDB Cluster	Databases	N/A	Yes
Spot Fleet Request	Compute	EC2 Instance, Spot Instance Request	No
Spot Fleet Instance Request	Compute	EC2 Instance	No
S3 Bucket	Storage	Bucket Policy, VPC Endpoint, VPC, ENI, SG, Subnet	No
ElasticLoadBalancingV2 Load Balancer	Network Routing	ENI, Subnet, VPC, SG, Target Group	No
ElasticLoadBalancingV2 Target Group	Network Routing	Load Balancer, VPC, EC2 Instance	No
VPC Peering	Network Management	VPC	No
DynamoDB Table	Databases	N/A	No
Redshift Cluster	Data Warehouse	ENI, Subnet, VPC, VPC Endpoint, SG, Network Interface Controller (NIC)	No

Lambda Function	Serverless	Subnet, VPC, SG, Key Management Services (KMS) key	No
KMS Key	Security Infrastructure	DB Cluster, DB Instance, EBS Volume, Redshift Cluster, Lambda Functions	No
IAM Account	Account Management	N/A	No
IAM User	Account Management	N/A	No
RAM Resource Share	Resource Management	Subnet	No
Document DB Elastic Cluster	Databases	Subnet, Security Group	No
Transit Gateway	Network Routing	Transit Gateway Attachment, Transit Gateway Route Table, Transit Gateway Multicast Domain, RAM Resource Share	No
Transit Gateway Attachment	Network Routing	Transit Gateway Attachment, Transit Gateway Multicast Domain, Transit Gateway, Subnet, VPC	No
Transit Gateway Route Table	Network Routing	Transit Gateway Attachment, Transit Gateway, Subnet, VPC	No
Transit Gateway Multicast Domain	Network Routing	RAM Resource Share, Transit Gateway Attachment, Transit Gateway, ENI, Subnet, VPC	No
VPC Endpoint Service	Network Routing	ElasticLoadBalancingV2 Load Balancer	No

Azure

Resource	Category	Attached Resources on Details Page	Flow Support
IP Configuration	Network Management	N/A	Yes, on <i>Traffic</i> page
Network Interface	Network Man-	Virtual Machine (VM), Net-	No

	agement	work Security Group (NSG), IP Config, Subnet, VNet, Public IPs, VM ScaleSet	
Network Security Group	Network Security	Network Interface Con- nection (NIC) Subnet, VM, VM ScaleSetVM, NSG Rules	No
Virtual Network	Network Man- agement	NIC, IP Config, Subnet, Vir- tual Network Gateway	No
Subnet	Network Man- agement	VNet, IP Config, NIC, NSG	No
Network Watcher	Network Mon- itoring	N/A	No
Flow Logs	Network Mon- itoring	N/A	No
Resource Group	Infrastructure Management	N/A	No
Storage Account	Storage	Private Endpoint, NIC, Sub- net, VNet, NSG	No
Virtual Machine	Compute	NIC, NSG , IP Config, Subnet, VNet, VM ScaleSet	Yes
SQL Server	Databases	Private Endpoint, NIC, Sub- net, VNet, NSG, SQL Server Database	No
Virtual Machine ScaleSet	Compute	VM ScaleSet	No
VirtualMachineScaleSet Vir- tual Machine	Compute	VM, VM ScaleSetVM	No
Network Security Groups Rule	Network Security	N/A	No
Application Gateway	Network Routing	N/A	No
Load Balancer	Network Routing	N/A	No
Route Table	Network Routing	N/A	No
Application Security Group	Network Security	N/A	No
Azure Firewall	Network Security	Subnet, Firewall Policy	No
Virtual Network Gateway	Network Routing	Connection, IP Address, Sub-	No

		net, Virtual Network	
Connections	Network Routing	Virtual Network Gateway	No
VPN Gateway	Network Routing	N/A	No
NAT Gateway	Network Routing	N/A	No
Private Endpoint	Network Man- agement	NIC, Subnet, VNet, Azure PaaS resources	No
Managed Cluster	Containers	N/A	No
Container Group	Containers	N/A	No
DocumentDB Database Account	Databases	Private Endpoint, NIC, Sub- net, VNet, NSG, SQL Data- base, DocumentDB Table, Document DB Gremlin Data- base, DocumentDB Cas- sandra Keyspace, DocumentDB Mongo Data- base	No
DocumentDB Cassandra Cluster	Databases	N/A	No
DocumentDB Mongo Cluster	Databases	Private Endpoint, NIC, Sub- net, VNet, NSG	No
DBforPostgreSQL Server- Group V2	Databases	Private Endpoint, NIC, Sub- net, VNet, NSG	No
DBforPostgreSQL Flexible Server	Databases	Private Endpoint, NIC, Sub- net, VNet, NSG, DBforPostgreSQL Flexible Server Database	No
DBforPostgreSQL Flexible Server Database	Databases	DBforPostgreSQL Server	No
DBforPostgreSQL Server	Databases	Private Endpoint, NIC, Sub- net, VNet, NSG, DBforPost- greSQL Server Database	No
DBforPostgreSQL Server Database	Databases	DBforPostgreSQL Server	No
SQL Server Database	Databases	SQL Server	No
Web Site	Serverless	Subnet, Web Site Function	No
Web Site Function	Serverless	Web Site	No

DocumentDB Database Account	Databases	DocumentDB Database Account	No
DocumentDB Table	Databases	DocumentDB Database Account	No
DocumentDB Gremlin Database	Databases	DocumentDB Database Account	No
DocumentDB Cassandra Keyspace	Databases	DocumentDB Database Account	No
DocumentDB Mongo Database	Databases	DocumentDB Database Account	No
Virtual Network Peering	Network Management	VNet	No
DBforPostgreSQL Server-Group V2 Server	Databases	DBforPostgreSQL Server-Group V2	No
Subscription	Account Management	N/A	No
Network Security Groups Default Security Rule	Network Security	N/A	No
Public IP Address	Network Management	N/A	No
Public IP Prefix	Network Management	Public IP Address	No
Redis Cache	Databases	VNet, Subnet, Private Endpoint	No
Private Link Service	Network Management	Private Endpoints, NIC, NSG, IP Config, Subnet, VNet, Azure PaaS Resource	No
Firewall Policy	Network Management	Firewall Policy, Rule Collection Group	No
Rule Collection Group	Network Management	Firewall Policy	No
Diagnostic Setting	Network Management	Firewall, Storage Account	No

NOTE:

Because CloudSecure may not always discover elastic network interfaces (ENIs), a flow search based on resource IDs will not work for the following supported resources if their *Details* page does not display the ENI. The workaround is to search using the IP address of the associated ENI, if known:

- AWS RDS DBInstances
- AWS RDS DBClusters
- ElasticLoadBalancingV2 load balancers
- AWS MemoryDB clusters
- AWS ElastiCache for Redis clusters
- AWS Redshift clusters

NOTE:

Although they will appear, EKS Clusters/Nodegroups and S3 buckets will not have flows. Only AWS EC2 instances, AWS RDS DBClusters, AWS RDS DBInstances, and Azure VMs will have flows.

Traffic

This topic describes the purpose of the Illumio CloudSecure traffic feature, found in the *Cloud > Explore* menu, and provides a general example of how you would use it. For instructions on how to use the search function in the *Traffic* page, see [Search Traffic](#).

The traffic page lets you view denied and allowed flows in a table. Click on a table row to see more details about the source and destination of the flow, such as IP Addresses, account IDs, labels, categories, resource types, etc.

Supported Resource Types

See [Traffic Supported Resources](#). For a list of resources against which you can write policy, see [Resources that Support Policy](#).

Exporting Traffic Lists

Click **Export** to export the filtered traffic data to one of the following formats:

- CSV
- JSON

You can enter a name and select a time range.

Go to the [Reports](#) page to download the exported traffic list (report).

Limitations for Displaying Traffic

In the *Traffic* page, the list displays only 10,000 results. This display limit is not configurable. This may cause you to see only the most recent 10,000 flows, irrespective of the earliest time you select, because collection of flows starts from the current day. For example, if the current day already has 10,000 flows, then irrespective of your time selection (such as the last 7 or 14 days), it will show only the first 10,000 flows from the current day. Illumio set this display limitation to provide optimal page display performance. You can filter your traffic list to retrieve data about traffic that isn't initially displayed when you elect to display everything. CloudSecure does not display your traffic in any specific order. When you don't filter your traffic, the page will typically display the most recent 10,000 results.

Generating Risk Reports

This is an overview of the Risk Report feature. For instructions on generating a Risk Report, see the in-application help on the *Traffic* page. For a list of services that Illumio considers to be at risk, see [Risky Services](#). The *Risk Report* tab lets you download a .PDF report summarizing the following at the account/subscription level:

- Total count of ransomware-susceptible traffic flows
- Total count of resources in your cloud environment affected by such flows

Before you click **Download**, you can toggle to include or exclude the following details from the report:

- Top Sources/Destinations
- Top Conversations

You can also select the time frame and whether to sort by byte count or flow count.

When generating the report, CloudSecure reviews your traffic against a list of services that are susceptible to ransomware attacks. It provides an executive summary. If it finds any susceptible services, it displays the following details:

An *Onboarded Account Summary* table, containing the following columns:

- Cloud
- Number of Accounts with Risk
- Number of Accounts
- An *Observed Risky Activities Summary* table, containing the following columns:
 - Service
 - Port
 - Protocol
 - Severity
 - Active Accounts
- A *Ransomware Risky Services Detected* table for each at-risk service, with the following columns:
 - Account, tallying all accounts identified as affected by the risk
 - Flow Count, tallying all traffic flows identified as affected by the risk
 - Byte Count, tallying the volume identified as affected by the risk
 - Resource Count, tallying all resources identified as affected by the risk
- If enabled, a *Top Sources By Flow/Byte Count* table for each service, with the following columns:
 - Top Sources By Flow/Byte count, ordering all sources identified as affected by the risk
 - CSP Resource ID
 - Account
 - Flow Count, tallying all traffic flows identified as affected by the risk
 - Byte Count, tallying the volume identified as affected by the risk
 - Origin, indicating if the risk is external or internal
- If enabled, a *Top Destinations By Flow/Byte Count* table for each account, with essentially the same columns as the top sources tables
- If enabled, a *Top Conversation Flow/Byte Count* table for each account, with essentially the same columns as the top sources/top destinations tables

If CloudSecure does not find any of your traffic in the list of services it considers risky, it displays a *Ransomware Risky Services Not Detected* section, containing a table the following details:

- Heading row, containing the following columns:
 - Severity
 - Service
 - Port
 - Protocol

Search Traffic

This topic describes the steps for searching the Illumio CloudSecure traffic feature, found in the *Cloud > Explore* menu, and provides a general example of how you would use it. For an overview of the *Traffic* page, including Risk Report generation, see [Traffic](#). For instructions on how to use the search function in the *Traffic* page, see the in-application pop-ups in the CloudSecure GUI.

Supported Resource Types

See [Traffic Supported Resources](#). For a list of resources against which you can write policy, see [Resources that Support Policy](#).

Searching Traffic Guidelines

Using the Filter

The following are guidelines for using the Filter, which is also available on the Applications page Traffic tab:

- You have the option of using operators such as != and =, but note that ! does not work with labels
- If you want to switch the automatically inserted joiners from OR to AND, or the reverse, select **Match All Conditions (AND)** or **Match Any Conditions (OR)** as appropriate. You can add additional search terms without having to delete existing terms.
- You can filter by:
 - Source/Destination (this menu can change depending on Category selection)
 - Category (these include Cloud, Account, Region, Label, Flow Status, IP Address, Port, Subnet, VPC, and Resource Type)
 - Operator (this menu can change depending on Category selection)

- Value (these include label name, port, and IP address). Note that if you type an IP address, the numerals appear in the search bar before they appear in the value field in search menu.

Risky Services

This topic lists services that Illumio considers to be at risk. For information on the *Traffic* page and Risk Reports, see [Traffic](#). For instructions on how to use the search function in the *Traffic* page, see [Search Traffic](#).

Ransomware Risk Services

The following is a list of services that Illumio considers to be at risk for ransomware penetration and lateral movement.

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

WSD S-WSD TCP/UDP 3702 Medium

Traffic Supported Resource Types

See [Traffic Supported Resources](#). For a list of resources against which you can write policy, see [Resources that Support Policy](#).

Traffic Supported Resources

The *Traffic* page displays the following resources. For information on the resources for which Illumio CloudSecure supports policies, see [Resources that Support Policy](#).

AWS

Resource	Category	Attached Resources on Details Panel	Flow Support
Network Interface	Network Management	All AWS workload types, Subnet, VPC, SG	No
Elastic IP	Network Management	ENI, EC2 Instance	No
EC2 Instance	Compute	ENI, Subnet, VPC, Security Group (SG), Elastic IP, Elastic Block Storage (EBS) Volume	Yes
RDS DB Cluster	Databases	ENI, Subnet, VPC, SG, KMS Keys	Yes
RDS DB Instance	Databases	ENI, Subnet, VPC, SG, KMS Keys	Yes
ElastiCache CacheCluster	Databases	N/A	Yes
MemoryDB Cluster	Databases	N/A	Yes

Azure

Resource	Category	Attached Resources on Details Panel	Flow Support
Network Interface	Network Management	Virtual Machine (VM), Network Security Group (NSG), IP Config, Subnet, VNet, Public IPs, VM ScaleSet	No
Virtual Machine	Compute	NIC, NSG , IP Config, Subnet, VNet, VM ScaleSet	Yes

NOTE:

Because CloudSecure may not always discover elastic network interfaces (ENIs), a flow search based on resource IDs will not work for the following supported resources if their *Details* page does not display the ENI. The workaround is to search using the IP address of the associated ENI, if known:

- AWS RDS DBInstances
- AWS RDS DBClusters
- ElasticLoadBalancingV2 load balancers
- AWS MemoryDB clusters
- AWS ElastiCache for Redis clusters
- AWS Redshift clusters

CloudSecure Dashboard

This topic explains how to work with the Dashboard in CloudSecure.

What is the Dashboard?

The Dashboard provides a quick way to understand your cloud presence at a glance. The Dashboard includes the following tiles:

Onboarding Overview

This tile has three sections: *Summary*, *Flow Log Access*, and *Read and Write Access*.

- Click the *Summary* section to go to the *Onboarding* page and see the list of providers types along with additional details. See [About Onboarding Cloud Accounts](#).
- Click the *Flow Log Access* section to go to the *Flow Log Access* page. See [Grant Flow Log Access](#).
- Click the *Read and Write Access* section to go to the *Onboarding* page and see the list of permissions
- Click the filter icon in the upper right-hand corner to change the tile's filters

Traffic Flow Summary

This tile has two sections, *Allowed Traffic* and *Denied Traffic*.

- Click the count of *Allowed Traffic* to go to the [Traffic](#) page and see only flows with *ALLOWED* status within the selected time-frame

- Click the count of *Denied Traffic* to go to the *Traffic* page and see only flows with *DENIED* status within the selected time-frame
- Click the rest of the body of the tile to go to the *Traffic* page and see every flow within the selected time-frame
- Click the filter icon in the upper right-hand corner to change the tile's filters. Note that when one single Cloud Service Provider (CSP) is specified, the CSP is also filtered in the *Traffic* page.

Summary of Ingested Resources

This tile has two main sections, *Resource Category* and *Count*.

- Click the filter icon in the upper right-hand corner of the tile to change the tile's filters, including CSPs and column sorting preferences. Note that when one Cloud Service Provider (CSP) is specified, the CSP is also filtered in pages reached by clicking in this tile.
- Click any of the following in the *Resource Category* column to see the applicable resources in the *Inventory* page:
 - Account Management
 - Compute
 - Containers
 - Databases
 - Infrastructure Management
 - Network Monitoring
 - Network Management
 - Network Routing
 - Network Security
 - Security Infrastructure
 - Serverless
 - Storage

The categories displayed may vary depending on the cloud environment and may change over time.

- The *Count* column for resources is static. It indicates the raw counts for the listed resource categories. These counts reflect the totals across all onboarded CSPs, even if you filter by CSP in the tile. Hover over the numbers to see the percentage of the total that each represents.

The Illumio CloudSecure dashboard gets updated over time, so check here as new tiles are added.

CloudSecure Search

This topic explains how to work with the following search features in CloudSecure:

- Context-based search
- Filter-based search

What is Context-based Search?

This feature provides a quick way to view *Inventory*-based filtered information on a number of pages in CloudSecure. Context-base search shows resources in the context of what is being searched. For example, if you were to select AWS instead of Azure in a filter on a given page, a subsequent search would give results for only your AWS cloud resources but not your Azure cloud resources.

Pages with the Feature

- *Inventory*
- *Traffic*
- *Cloud Map*
- *Application > Inventory* tab
- *Application > Traffic* tab

Using Context-based Search

- Use the dropdown search field dropdown to select filters, such as Cloud: AWS, Cloud: Azure, Region: ap-east-1, or Status: Allowed in the case of the *Traffic* or *Cloud Map* pages. Other pages will offer filters relevant to the context of those pages. They are too numerous to list here.
- You may also have the option of using a time-frame dropdown, depending on which page you are viewing
- When you are ready to activate the search with your selected parameters, click **GO**

Context-based Search Caveats

- The *Inventory* page search feature accepts limited types of filters (in other words, these search values narrow down all inventory supplied metadata values bar tags):
 - Cloud
 - Account ID
 - Region
 - Resource Type
 - VPC/VNET ID
 - Subnet ID
 - Cloud Tags
 - Labels
 - Categories
 - Resource Group
 - Resource Name
- CSP id, VPC id, VNET id, subnet id, cloud tag, Resource Group, and Resource Name metadata types are supplied by *Inventory*, but are not accepted as metadata list filters, so the selection of these values has no impact on context-based search
- Account IDs are supplied by integrations, and support only cloud context
- Labels are supplied by the labeling service and support no context

What is Filter-based Search?

This feature provides a quick way to view filtered policy information on the *Policies*, *Onboarding*, *Application Discovery*, *Tag to Label Mapping* pages and their child tabs.

Policies Page Tabs with the Feature

The following example shows what is available on the *Policies* page. Other pages will offer filters relevant to the context of those pages. They are too numerous to list here.

- *Organization Policies*, which has the following filter terms:
 - Name, Status, Provision Status
- *Application Policies*, which has the following filter terms:
 - Name, Environment, Status

- *Services*, which has the following filter terms:
 - Name, Ports, Protocol
- *IP Lists*, which has the following filter terms:
 - Name, Provision Status

Using Filter-based Search

Use the dropdown search field dropdown to select filters, such as *Provision Status*: Added in the case of the *Organization Policies* tab. Other tabs offer filters relevant to the context of those tabs, as shown in the preceding list.

Filter-based Search Caveats

- This search makes no API calls
- This search is strictly a client side filter

Product Usage

This topic describes the purpose of the Illumio CloudSecure product usage feature and provides a general example of how you would use it.

Displaying Product Usage

In the *Product Usage* page, the graphs display the following:

- Drop-down menu for a defined time window (30 days or 90 days) or a custom time range going back to day zero
- Drop-down menu for the presentation style (line chart or area chart)
- Daily Illumio Workload Hours by date
 - Total Illumio Workload Hours
 - Compute Workload Hours
 - Database Workload Hours
 - Container Hosts
 - Serverless Containers
 - Serverless Functions
- Daily Log Storage by date (volume in GB)

NOTE:

A workload represents an Illumio-managed resource in your environment.
A workload hour represents the number of hours for which a workload was managed.

The workloads hours and log storage display according to your time selection (such as the last 30 or 90 days). You can see mouse-over text for the data by moving your cursor over the dots on the graph lines. Click **Export** to export the workload and log storage data to a .csv file, which will contain data for the time selection.

Events

This topic describes the purpose of the Illumio CloudSecure product event log feature and provides a general example of how you would use it. Note that the user interface displays up to 10,000 events, and removes events older than seven days.

Displaying Events

Audit Events Tab

This tab shows user initiated-events such as logging in or creating a policy. You can use filtering parameters to choose what the report includes.

In the *Events* page, the Audit Events tab displays the following information about user-initiated events:

- Filter drop-down menu so you can search for different properties
- Success/Fail (green check mark/red X; column and filter property)
- Timestamp (column and filter property)
- Event Type (column and filter property)
 - User Login/Logout
 - User Added/Removed
 - Authentication Failure
 - Access changes (various)
 - Account changes (various)
 - Application changes (various)
- Category (column and filter property)
 - Onboarding
 - Policy

- Labeling
- User (column and filter property)
- Details (column and filter property)
- CSP (filter property)
- Account ID (filter property)

Click a row to see the event properties, including CSP, Account ID, and Tenant ID, listed in a panel.

System Events Tab

The System Events tab displays system-initiated events such as synchronizing resources or auto-creating an application based on a discovery rule. The information listed is similar to that seen in the Audit Events tab.

Exporting an Events Report

Click **Export** to export the event data to one of the following formats:

- CSV
- JSON

Go to the [Reports](#) page to download the exported report.

Reports

This topic describes the purpose of the Illumio CloudSecure reports feature and provides a general example of how you would use it. You will be able to generate both Audit Event and Risk reports.

Reports

Reports Table

The Reports page displays a list of your previously generated reports with the following columns:

- Name
- Report Type
- Generated At
- Generated By
- Status

- Expiration Date
- Action

Adding Reports

Use the **Add Report** button to generate a new report. The menu lets you choose between Risk or Audit reports. In the dialog that appears in a side panel, you can choose a name, and, if applicable for that type of report, parameters like the following:

- Time range
- Details to include
- Export format
- Filters
- Retention duration
- Etc.

Using Reports

Use the following features to interact with your generated reports:

- Click **Download** in the Action column to access your reports. The reports are generated asynchronously. You are prompted when the report creation has completed.
- Select reports and click **Remove** to delete them. A confirmation dialog displays, asking if you wish to proceed.

Chapter 3

Define Your Cloud Resources in CloudSecure

This chapter contains the following topics:

Deployments and Applications	99
Cloud Tag to Label Mapping	117
Rule-Based Labeling	119
Use AI Labeling	132

This section explains the requirements and steps for defining your public cloud accounts in Illumio CloudSecure. You perform these tasks after onboarding your cloud accounts and gaining an understanding of resources, topology, and communications.

Deployments and Applications

This topic explains defining deployments and defining applications in CloudSecure. Additionally, it explains why you may want to define deployments to segment your applications. Note that although deployments are recommended, they are optional.

To further explain these concepts, the topic includes an example of how to define an application and the three deployments hosting it.

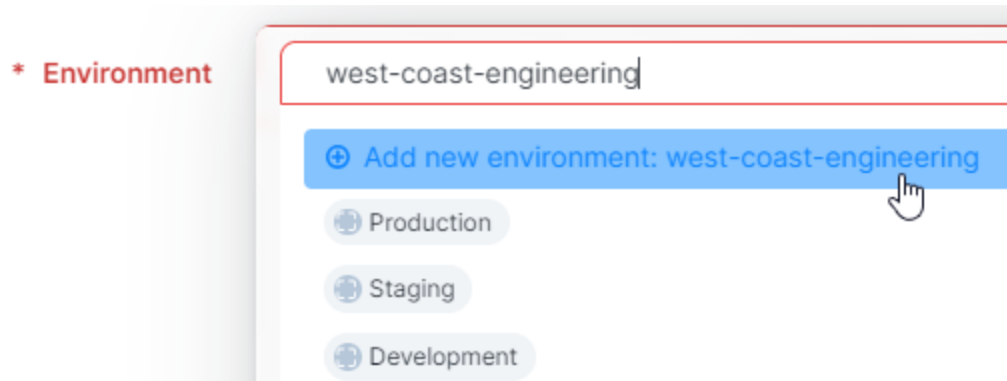
What is a Deployment in CloudSecure?

In CloudSecure, you may decide to create deployment stacks as part of specifying which applications in your cloud account to protect with CloudSecure.

After onboarding your cloud accounts, you may begin by defining the environments you're using in the cloud. In CloudSecure, we refer to this as "adding deployment

stacks.” In the cloud, stacks provide a way to manage your resources as a single, atomic unit.

In CloudSecure, a deployment stack correlates with the stages that organizations use to manage their product development lifecycle and defines the boundaries of application deployment. To define these boundaries in CloudSecure, you create your deployment stacks by selecting an Environment label. Then, associate that Environment label with cloud metadata (such as tags), and resources to define the boundaries of that environment.



For example, you might realize that you want a deployment stack in CloudSecure equal to your development environment that exists for your AWS us-west-1 region. Perhaps, this environment is constrained to operate only on a specific subnet. Typical environments include development, staging, and production.

Relationship between Deployments and Applications

To more fully use CloudSecure, you need to understand the relationship between applications and their deployments.

In CloudSecure, you will typically work with two special label types to manage security for your cloud resources. These label types are also related to deployments as defined above. As explained above, deployment stacks use an *Environment* label. You associate attributes to that label to set the boundaries of the stack. Recall that deployment stacks are optional but often helpful.

Defining an application follows a similar process. You begin by specifying an *Application* label. Then, you associate cloud resources to that label by selecting the appropriate cloud tags or cloud metadata associated with that application.

Ultimately, the process of getting the most out of your application definition involves:

- (Optional) First, defining your deployment stacks

You only need to define deployment stacks first if:

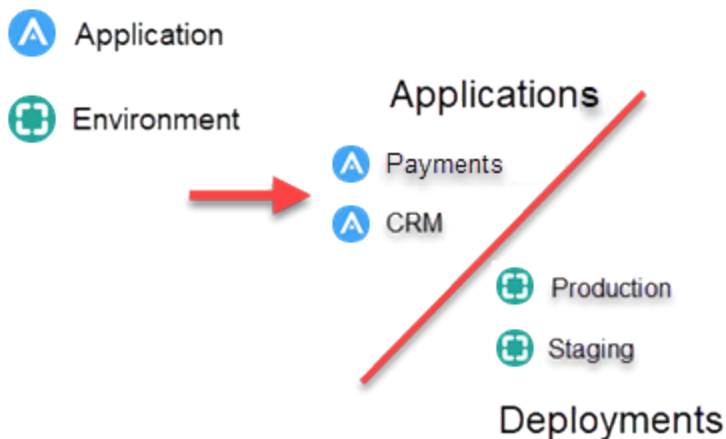
- You wish to define your application with deployment stacks (that is to say define them in part by environment and environment-specific resources)
 - You haven't previously defined any. If you've already defined your deployment stacks, simply select them when defining your applications.
- (Optional) Second, creating tag to label mappings. See [Cloud Tag to Label Mapping](#) for information.

CloudSecure analyzes each of these types of definitions and sees the unions between them. For example, it's able to detect that the CRM application you defined is hosted in the Staging and Production deployments running in your Azure and AWS clouds, respectively.

CloudSecure gets you started by including *Environment* labels for Production, Staging, and Development in its deployment definitions page.

Organizations also create their own environment-specific definitions around how they have deployed applications; for example, they might have an environment for Eastern European Engineering.

Label Types



To break these concepts down further, see how CloudSecure utilizes Application and Environment labels.

Application Labels

You define an application using cloud tags and/or cloud metadata so that CloudSecure can discover the deployments and resources for that application.

For example, you have two applications — a Payment application and a CRM application. In CloudSecure, you define each application and assign them an Application label.

Environment Labels

Your company has different deployments of your applications. You can think of these environments as different instances of each application based on where they reside. For example, your company has a staging environment and a production environment.

In the illustration above, your Payments and CRM applications reside in two environments — production and staging. These two applications are “deployed” in both production and staging. In this way, you assign these applications to the correct deployments.

CloudSecure Discovers Your Application Environments

When you define a deployment, CloudSecure doesn’t discover anything about your applications. You defined your deployment stacks separately. Then, after you defined each application, CloudSecure analyzes them by reviewing the associated cloud metadata, such as account, VPC, subnet, tags, etc. CloudSecure recognizes the union of those separate definitions and determines the environments where your applications are running, if you have defined deployments. This union defines each application’s environment boundary.

Say you create an application (in this example, test2) and save it to CloudSecure. CloudSecure begins the process of discovering the environments that it’s running in. The *Application Definitions* page refreshes and includes the new application. The *Deployments* column indicates that CloudSecure is discovering all the defined deployments that host this application.

Application Label	Deployments	Cloud Tags	Resources	Approval Status
123	test002	Name anand-vpc	1 Resources	Approved
test2	Discovering...	alpha.eksctl.io/eksctl-version 0.126.0		

CloudSecure has discovered all the environments in which the test2 application is running. When the discovery process finishes, the list includes all the deployments where CloudSecure discovered matching cloud metadata.

Application Label	Deployments	Cloud Tags	Resources	Approval Status
123	test002	Name anand-vpc	1 Resources	Approved
test2	Payments	alpha.eksctl.io/eksctl-version 0.126.0	16 Resources	Pending Approval
test2	west-coast-engineering	alpha.eksctl.io/eksctl-version 0.126.0	2 Resources	Pending Approval

CloudSecure will not populate the Deployments column if you choose not to define any for that application.

Why is Environment Discovery Important?

CloudSecure treats each environment where an application runs as a separate application instance. This functionality allows you to define policy tailored for the environment.

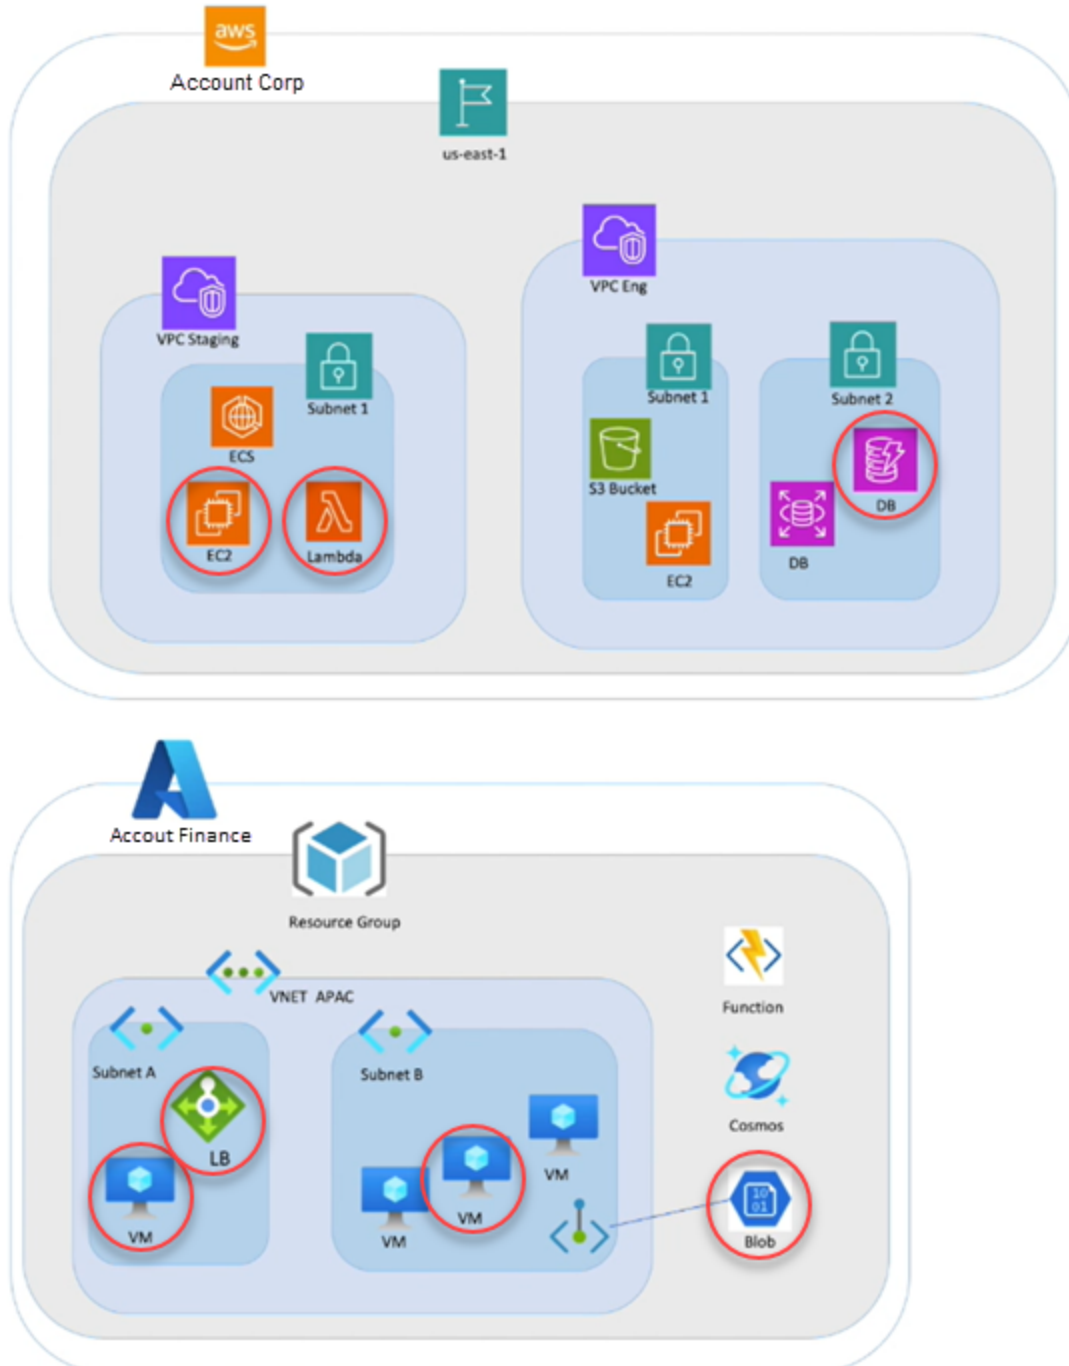
For example, you might want very flexible and open security policy for applications running in your Development environment. However, when those applications move to production, you may require very controlled policy to eliminate risk.

Example Deployment and Application Definitions

In this example, a company has the standard development, staging, and production environments. It manages a travel ticketing application on its corporate website. The company uses both Amazon AWS and Microsoft Azure to host this application and its environments.

In AWS, the company's "Corp" account has a VPC that they use as their staging environment (the "VPC Staging") and a VPC they use for their development environment (the "VPC Eng"). Their "Finance" account in Azure has a resource group that they use for their production environment.

The travel ticketing application has resources in both the AWS and Azure accounts and in all three environments. It uses two resources in the AWS Staging VPC and a database in the AWS VPC Eng. They host their production environment in their Azure Resource Group, and use these resources: a load balancer, two VMs, and a storage account.



In CloudSecure, the company defines the resources that are part of this application. They begin by defining all three deployments in CloudSecure — development, staging, and production. See [Define a Deployment](#). Then, they are ready to define the application in CloudSecure. When they define the application, they specify the scope for what comprises the application. See [Define an Application](#).

In the application definition, they specify cloud tags and metadata as follows:

- The AWS Corporate account → US East 1 region → 2 VPC s - Staging and Engineering → Subnets 1 and 2
- The Azure Finance account → APAC VNet → Subnets A and B and a storage blob

They have already defined their environments: the development and staging environments in AWS and their production environment in Azure. CloudSecure can now determine that the application has resources in the AWS development and staging environments and resources in the Azure production environment.

In CloudSecure, the travel ticketing application appears as three separate instances and each instance can have its own security policy.

Define a Deployment

This topic explains how to define a deployment in CloudSecure.

For an explanation of how CloudSecure uses deployments and why they are helpful despite not being required, see [Deployments and Applications](#).

Prerequisites

Before you define a deployment, you must have onboarded one or more public cloud accounts and given CloudSecure time to synchronize with them so that you can configure the correct boundaries for the deployment. See [About Onboarding Cloud Accounts](#) for information.

Before defining a deployment, Illumio recommends that you review your Inventory and Cloud Map topology to gain an understanding of how your cloud resources are utilized and how they are communicating. See [CloudSecure Map](#) and [Inventory](#) for information about using these features.

Define a Deployment (Optional)

1. From the left navigation, choose **Application Discovery > Application Definitions**.

If necessary, select the **Deployments** tab.


If you haven't defined any cloud deployments, the page contains a button to add your first deployment.

2. Click **Add**. The Deployment page appears.
3. From the *Environment* field, select an existing label or create a new one.

By default, CloudSecure includes Environment labels for “Production,” “Staging,” and “Development.” If you select a label that already has a deployment defined for it, CloudSecure displays a message that the selected label is already assigned to a deployment. Click the red **X** at the end of the field to clear the value.

To create a new Environment label, simply type the name in the field and select it when it appears in the drop-down list.

4. (Optional) Provide a description so that other members of your organization understand how you are defining the boundaries of the deployment.
5. Click **Add** to open the drop-down menu of the resource types to use to define the deployment scope.



When you select an item from the resource drop-down menu, the **Add Deployment Stacks** dialog box opens.

6. In the **Includes** field, select the resource to use from the pre-populated drop-down list. The list includes resource that CloudSecure discovered after you onboarded your cloud accounts.

You can select multiple resources of that type. When finished including resources, click **Add**. The dialog box closes and those resources are added to the list under the *Deployment Stack*.

You can continue defining the deployment stack with other types of resources by selecting from the **Add** drop-down list, then selecting the specific resources. The types of resources you’ve already included are unavailable in the list. For example, if you already created a row for subnets and specified several, the subnets type is grayed out in the list.

7. When done fully setting the boundaries for this deployment, click **Save**.

The new deployment appears in your list of deployments.

Edit a Deployment

1. From the left navigation, choose **Application Discovery > Application Definitions**.

If necessary, select the **Deployments** tab.

If you have defined any cloud deployments and wish to modify one, select it.


2. Click **Edit**. The *Deployment Edit* page appears.

3. From the *Environment* field, select an existing label or create a new one.

By default, CloudSecure includes Environment labels for “Production,” “Staging,” and “Development.” If you select a label that already has a deployment defined for it, CloudSecure displays a message that the selected label is already assigned to a deployment. Click the red **X** at the end of the field to clear the value.

To create a new Environment label, simply type the name in the field and select it when it appears in the drop-down list.

4. If desired, edit the description.
5. To change the resource types to use to define the deployment scope, click **Add** to open the *Deployment Stack* resource drop-down menu.



When you select an item from the resource drop-down menu, the **Add Deployment Stacks** dialog box opens.

6. To change the resource to use, select the **Includes** field to open the pre-populated drop-down list. The list includes resource that CloudSecure discovered after you onboarded your cloud accounts.

You can select multiple resources of that type. When finished including resources, click **Add**. The dialog box closes and those resources are added to the list under the *Deployment Stack*.

You can continue editing the deployment stack with other types of resources by selecting from the **Add** drop-down list, then selecting the specific resources. The types of resources you've already included are unavailable in the list. For example, if you already created a row for subnets and specified several, the subnets type is grayed out in the list.

7. When done editing the boundaries for this deployment, click **Save**.

Delete a Deployment

IMPORTANT:

Before you can delete a deployment, you must ensure that none of your application definitions are running in that deployment. CloudSecure won't let you delete a deployment that is in use by any ruleset. (Each deployed application features a ruleset).

1. From the left navigation, choose **Application Discovery > Application Definitions**. The *Applications Definitions* page appears and the *Application Definitions* tab is selected.
2. Select the **Deployments** tab.
3. Select the deployment you want to delete and click **Remove**. (You can remove multiple deployments if necessary.)

A confirmation dialog box appears displaying the deployment you are deleting.

4. Verify that you are deleting the correct deployment and click **Remove**.

What's Next?

Now that you've defined your deployments, which are optional, you can begin creating any application definitions that rely on them. You do *not* need to have a deployment defined in order to define an application. See [Define an Application](#).

Define an Application

This topic explains how to define an application in CloudSecure.

For an explanation of CloudSecure application definitions and how they relate to deployments, see [Deployments and Applications](#).

Prerequisites

Before you define an application, you must have onboarded at least one cloud account. Defining a deployment is optional. For information about defining a deployment, see [Define a Deployment](#).

Define Applications Automatically

Although CloudSecure has always allowed you to define applications individually, you can now automatically create multiple applications by defining an Application Discovery Rule. This feature runs in the background, so the rule you create automatically defines applications when new resources are added that meet the rule parameters.

Create an Application Discovery Rule

The user interface presents the following ways to begin creating such rules:

- If you have not defined your first application, either individually or with an Application Discovery Rule, the *Application Definitions* page displays a banner inviting you to add your first application definition using either method. Click **Add Application Definition**, select **Application Discovery Rule**, and click **Confirm** to begin.
- If you have already added an application individually, but not yet with the Application Discovery Rule method, a banner invites you to do so. Click **Create** to begin.
- If you have already created an Application Discovery Rule, navigate to the *Application Discovery* menu item

The in-application pop-up guide instructs you on how to proceed.

Application Discovery Rule Guidelines

- Choose your rule name carefully, to make it clear what sort of applications you are automatically defining
- You can add a prefix to the name of all applications discovered with the rule
- The prefix and name may be changed when editing the rule. Other parts of the rule are not editable.

When editing the rule, if the edit does not affect an existing application definition, you do not need to modify or re-approve the application. If the edit affects an existing application definition, then the following apply:

- If the change is to only the prefix in the rule, rename the existing application label to reflect the new prefix. You do not need to re-approve application.
- If any of the change is to metadata (the type of rule, such as account/subscription, virtual network, etc.), you may need to review and approve new or existing application deployments

When you save your rule edits, the application approval or re-approval workflows begin. A prompt tells you to review and remove any policies associated with application labels that were previously associated with the rule.

- The rule's exact behavior may vary depending on the rule type you select.
 - Cloud Tags: If you choose this rule type, a *Cloud Tag Keys* dropdown menu appears
 - Cloud Accounts: If you choose this rule type, it applies to *all* the available account/subscription across all accounts and ties an application to each account/subscription with the relevant resources. You have the option to specify the CSPs to which the rule applies. You can have only one rule of this type.
 - Virtual Networks: If you choose this rule type, it applies to *all* the available virtual networks across all accounts and ties an application to each virtual network with the relevant resources. You have the option to specify the CSPs to which the rule applies. You can have only one rule of this type.
 - Subnet: If you choose this rule type, it applies to *all* the available subnets across all of your accounts, and therefore ties any application to each subnet with the relevant resources. You have the option to specify the CSPs to which the rule applies. You can have only one rule of this type.
- Application Discovery Rules cannot be disabled or paused once added. There are two workarounds:
 - You can delete the rule, which will also delete all application definitions created with the rule
 - You can modify individual application definitions for those created with the rule, which decouples the application definition with the rule

- Once you create an Application Discovery Rule, you can browse to *Discovery Rules > View* details to edit it.
- Application definitions have contexts for how they were created, viewable on their respective detail pages, either individually or using an Application Discovery Rule

NOTE:For any application definitions created with an Application Discovery Rule, the approval process begins as described in [View and Approve an Application](#), *unless* you click the *Auto Approve Setting* toggle to **ON**. Do this if you want CloudSecure to automatically approve all discovered application definitions, as well as any updates made to their deployments and resources. This skips the manual approval process for automatically defined applications. If you click the toggle to **OFF**, you must approve the discovered application definitions manually. See [View and Approve an Application](#) for information.

Application Label Conventions

- Tag-based application labels are generated in the format Prefix-<TagValue> e.g., infosec-payment
- Account/Subscription-based application labels are generated in the format Prefix-<unique account/sub identifier> e.g., InfoSec-Act123
- VPC/VNet based-metadata application labels are generated in the format Prefix-<unique virtual network identifier> InfoSec-VirtualNetwork123
- Subnet based metadata application labels are generated in the format Prefix-<unique subnet identifier> InfoSec-Subnet123

Define Applications Individually

1. From the left navigation, choose **Application Discovery > Application Definitions**.
2. Click **Add**. A page with the fields to define the application appears.
3. Enter a name and description (optional) for the application.

This name is what appears in CloudSecure. The name should be descriptive so that you can easily identify it in CloudSecure.

Though optional, providing a description helps other members of your organization understand the purpose of this application.

4. Click **Add Resources Using Cloud Metadata**.

Cloud metadata contains information about the instances of your running cloud resources and can include subnets and virtual networks. CloudSecure obtains your cloud tags directly from your cloud accounts. This data is the label that you assigned to a cloud resource along with an optional tag value.

You do not define your application instances using Illumio CloudSecure labels. Your applications are defined for CloudSecure purely based on cloud properties.

The *Application Definition* dialog box appears.

5. In the top-most drop-down list, choose whether to use cloud tags, virtual networks and subnets, or accounts to define the application.
6. In the **Filter By Cloud Accounts** field, select the accounts that are hosting the application resources. Continue selecting accounts until you've specified them all. To clear an account from the field, click backspace or click the **X** to clear them all.
7. In the **Select** field, select the specific tags or metadata (depending on the type you chose) that defines the application.

TIP:

The list is pre-populated with values that CloudSecure discovered after you onboarded your cloud accounts. Depending on the size of your cloud environments, the list can get quite long. You can scroll the list to locate the values you want or type a value in the *Select* field to filter the list. The list refreshes with values matching your search criteria.

When done adding data, click **Add to Selection**. The tags or metadata move to the selected section.

You can continue this process to add as many tags or metadata as required to define this application.

8. When done, click **Confirm Selection**. The dialog box closes, and your selected tags or metadata appears in the *Selected* section.

If necessary, repeat the process using the other type of data until you've fully defined all resources for the application. For example, you chose to locate all the relevant clouds tags first and then repeated the process adding the relevant metadata.

9. Click the *Auto Approve Setting* toggle to **ON** if you want CloudSecure to automatically approve all discovered deployments and resources for this application. This skips the manual approval process for applications.
If you click the toggle to **OFF**, you must approve the application definition manually. See [View and Approve an Application](#) for information.
10. When you have defined the application with enough specificity, click **Save**.

The *Application Definitions* page refreshes and includes the new application: The *Deployments* column indicates that CloudSecure is discovering any defined deployments that host this application.

When the discovery process finishes, the list includes any deployments where CloudSecure discovered matching cloud tags or metadata.

CloudSecure does not populate the *Deployments* column if you choose not to define any for that application.

When CloudSecure finishes discovering your saved application definition, and your application is listed as pending approval, you can still modify the resources defined for the application. For instance, you can add or drop cloud tags in the application definition in such a way that it applies to an additional resource, and CloudSecure automatically re-synchronizes the application to include the new resource. Once an application is approved i.e., no longer pending, any subsequent resource modifications could trigger a new pending approval state for the application deployment.

Edit an Application Definition

You may wish to update or otherwise edit an application you have already defined. Use the following steps to do so.

1. From the *Application Discovery > Application Definitions* tab, find the application label for which you want to edit the definition.
2. Click **View Details** for the application of interest.
3. Click **Edit**. The in-application pop-up guide instructs you on how to proceed. Note that if during editing you change the *Auto Approve Setting* toggle, you must confirm and save to retain the toggle change.

Delete Application Definitions

When you delete applications that are pending approval, CloudSecure simply deletes the application definitions.

When you delete approved applications, CloudSecure deletes the application definitions and the rulesets (policies) associated with the application definitions and the application instances. CloudSecure also disassociates any related resources from the application definitions being removed.

Delete Individually Created Application Definitions

1. From the left navigation, choose **Application Discovery > Application Definitions**. The *Applications Definitions* page appears and the *Application Definitions* tab is selected.
2. Select all the application definitions that you want to delete and click **Remove**. A confirmation dialog box appears displaying the applications you are deleting.
3. Verify that you are deleting the correct applications and click **Remove** in the dialog box.

Delete Application Discovery Rule-Created Application Definitions

Note that deleting a discovery rule automatically deletes all application definitions associated with the rule. You may also choose to manually delete associated application definitions, as follows:

1. From the left navigation, choose **Application Discovery > Discovery Rules**. The *Application Discovery* page appears and the *Discovery Rules* tab is selected.
2. For the Application Discovery Rule in question, select the **View Details** link in its table row. The *Details* page for that rule appears.
3. In the *Discovered Application Definitions* section of the *Details* page, select all the application definitions that you want to delete and click the **Remove** button in the upper right of the *Discovered Application Definitions* section. This is different than the *Remove* button at the very top of the page, which is grayed-out when you select an application definition.
A confirmation dialog box appears displaying the applications you are deleting.
4. Verify that you are deleting the correct applications and click **Remove** in the dialog box.

What's Next

Approve your application. (Each instance of the application in different deployments requires approval.) See [View and Approve an Application](#) for information.

Begin creating policy for your application. See [Writing Application Policy](#) for information.

View and Approve an Application

This topic explains how to approve an application definition after you've created it. See [Define an Application](#) for information.

Prerequisites

This topic assumes that you've already onboarded your cloud accounts and have created an application definition.

Why is Approval Required?

CloudSecure separates the process of defining an application from the ability to create policy for it.

After you define an application, it appears in the *Application Definitions* list. First, if you have defined a deployment, CloudSecure discovers any environments where the application is running. See [CloudSecure Discovers Your Application Environments](#) for information.

When the discovery process finishes, the list will include any deployments where CloudSecure discovered matching cloud tags or metadata.

NOTE: The *Application Definition* page lets you toggle whether you want CloudSecure to automatically approve all discovered applicable deployments and resources. Similarly, the *Application Discovery Rule* page lets you toggle whether you want CloudSecure to automatically approve all discovered application definitions, as well as any updates made to their deployments and resources. See the [Define an Application](#) documentation on the portal.

Either of these methods will skip the manual approval process for applications as described here.

For applications definitions that are not automatically approved, you can see that each of the application instances needs to be approved; meaning, you've defined an application but the status is still Pending Approval." In this way, CloudSecure ensures other key stakeholders are in the loop to approve your application definitions.

CloudSecure will not populate the *Deployments* column if you choose not to define any deployments for that application.

Approve a Given Application Definition

1. From the left navigation, choose **Application Discovery > Application Definitions**.

The list of defined applications appears.

2. Select the application that you want to review and/or approve. Note that if you select just one application definition, it will allow you to approve it if it is pending approval. However, if you select more than one application, the *Approve* button will become grayed-out because bulk application approval is not supported at this time.

The **Approve** button becomes enabled.

3. Click **Approve**. A confirmation dialog box appears displaying the application you are approving.
4. Verify that you are approving the correct application and click **Confirm**.

The dialog box closes and the *Approval Status* column updates and shows that the application definition is approved.

The application becomes part of the applications displayed in the *Applications* page, meaning you can now create policy for that application.

Approve Application Deployments and Resources in Bulk

You can have a single application that has multiple resources or deployments, such as staging and production. For example, you could have two application definitions associated with that application, one for each deployment. CloudSecure lets you approve two or more such application deployments in bulk.

1. From the left navigation, choose **Application Discovery > Application Definitions**.

The list of defined applications appears.

2. Select the application that you want to review and/or approve.
3. Click **Approve**. A confirmation dialog box displays the application's associated deployments and/or resources..
4. Select the checkboxes for the deployments and resources you wish to approve. For example, you may wish to choose an AWS us-west -1 resource on staging and production, but not development.
5. Verify your selections and click **Confirm**. Illumio will then create the approved definitions for that application based on the deployments and resources you

selected. Using the above example, you would have two approved definitions for the application, one using the staging deployment and the other using the production deployment.

CloudSecure does not let you bulk approve application definitions associated with different applications as their basis.

What's Next

Once you have approved an application, you can map your cloud tags to Illumio labels and write policy rules for it. Although mapping cloud tags to Illumio labels is not strictly required for creating policies, it will assist you in making your policies specific.

See [Cloud Tag to Label Mapping](#).

See [Writing Application Policy](#).

Cloud Tag to Label Mapping

This section describes the purpose of the Illumio CloudSecure cloud tag to label mapping feature, and provides a general example of how you would use it.

IMPORTANT:

Cloud tags are required to use this feature. For instructions on how to use the cloud tag to label mapping interface, see the pop-up notes in the CloudSecure UI.


Use Case and Example

If you have a tagging strategy in your cloud environment, this feature lets you associate more than application and environment labels with your resources. You can use this feature to associate additional labels with your resources too, allowing for more granularity when writing policies. You can create up to 20 such mappings.

Note that tag to label mapping can map labels to resources that are not part of an application. In this way, application approval is not required to complete the tag to label mapping process. Unlike the application approval process, the tag to label mapping process occurs immediately, without the need for approval.

For example, if you have cloud tags such as “Risk,” “Cost Center,” “Compliance,” etc., you can map those cloud tags to Illumio labels. Once you map these additional tags to Illumio labels, you will be able to associate these labels with resources in Illumio CloudSecure. Considering this example, if you have resources that have the cloud tag

“Risk,” those resources will associate them with the Illumio “Risk” label. The following diagram illustrates how you could use this feature:



In the diagram, cloud tag keys (“Risk,” “Threat,” and “RiskStatus”) are mapped to the Illumio label type “Risk.” This mapping enables different values of cloud tag keys to automatically map to the value of the Illumio label key. The following instructions simplify the process steps by focusing on mapping the cloud tag key “Risk” to the Illumio label “Risk.”

1. The first part of the sequence is to create one or more tag to label mappings, such as the following mappings:

- Cloud tag key Risk mapped to Illumio Label Risk

You can also map multiple cloud tag keys to one Illumio label type, i.e., mapping cloud tag keys “Compliance,” “Regulations,” or “Guidelines” to the Illumio label type “Compliance.” Note that the relationship between cloud tags to label types is that you can have multiple mappings using the same cloud tag keys, but there can be only one mapping for each label type. Defining the mapping from a cloud tag key to an Illumio label type automatically assigns the corresponding cloud tag values to Illumio label values. These Illumio labels can then be associated with resources in CloudSecure. Although you do not need to have resources associated with an application when mapping cloud tags to Illumio labels, you may choose to do so. The following example supposes that you do have an application that you wish to define using resources that you have associated with tag to label mappings. For example:

- Cloud tag 1 on Resource 1: Risk:Critical
- Label: tag key: Risk, value: Critical

2. Any cloud tags that were mapped to Illumio labels for the desired resources will then be *notionally* associated with any applications or deployments using those resources. Note that although the labels are notionally associated with an application possessing those resources in order to provide context, such labels are not in fact functionally associated with the application. These mapped labels are functionally associated with the resources only.

Assume the label `Application: Payment` has the following deployments:
`env:dev/staging/prod`.

If any resources within the `Payment` application are mapped to the label `Risk:Critical`, the Illumio “Risk” label will be notionally associated to the application. The *Tag to Label Mapping* page will show the Illumio label type and the labels to which you have mapped your CSP cloud tag keys.

3. Then, you could write granular policies using specific labels, such as the Illumio “Risk” label. Note that those policies will reference only the resources in question, and not the notionally associated application itself.

Cloud tags are required for this degree of granularity. Without cloud tag to label mapping, you can still write policies, but those policies would be coarser with broad Illumio labels such as `app` or `environment`.

View System Labels

This section describes the purpose of the system labels feature. View system labels on the *Label Mapping* page *System Generated Labels* tab. Use the filter to search for labels by their properties. See the in-application help for instructions.

For more information on system labels, see [Labels](#). For information on tag to label mapping, see [Cloud Tag to Label Mapping](#).

Category Labels

You can view cloud service categories mapped to Illumio labels. CloudSecure creates these system labels automatically, based on your cloud environment. You cannot edit or delete these system labels.

Service Role Labels

You can view cloud service roles mapped to Illumio labels. CloudSecure creates these system labels automatically, based on your cloud environment. You cannot edit or delete these system labels.

Rule-Based Labeling

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

Before you begin

- Label assignment:
 - You can assign system default and user-defined labels to matching workloads.
 - You can assign only one label of a given type to a workload.
 - Rule-Based Labeling assigns labels to workloads but doesn't replace existing labels already assigned to workloads. For example, if a matching workload has an existing Location label of New York and your labeling rule specifies a Location label of London, the existing New York Location label is preserved and the London Location label is bypassed.
- Depending on how many workloads match labeling rules, it may take a few minutes for the labels to be assigned to all of them. To verify that labels have been assigned to the matching workloads, check the **Workloads** page (**Servers & Workloads**).
- An event is created when a rule-based label is assigned to a workload. The name format of the event differs depending on how the label is assigned:
 - When assigned from the PCE UI: `label_mapping_rules_run.assign_labels`
 - When assigned from a system job: `system_task.automatic_label_application_for_new_vens`

Typical Labeling Rule Workflow

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

Step 1: Add a Labeling Rule

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

Step 2: Find and review matching workloads

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

Step 3: Assign labels to matching workloads


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

Add and Manage Labeling Rules

This section includes procedures for adding and managing labels, finding and matching workloads, and exporting a list of labeling rules to a CSV file.

Add a Labeling Rule


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

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

Find and Review Matching Workloads

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

1. Click **Apply Rules** and then choose **Review and Assign Labels**.



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

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

Assign labels to matching workloads

This procedure describes the different ways to assign labels to the workloads that match your labeling rules.

1. Make sure the **Workloads that match criteria** side panel is open (see [Find and Review Matching Workloads](#)).
2. Choose how you want to assign labels to matching workloads.
 - **Immediately:** From the **Workloads that match criteria** side panel, click **Assign** if you want to assign labels immediately. The message **Labels have been assigned to _ workloads** appears.
 - **At specific times and intervals.** See [Schedule Label Assignment](#).
 - **When VENs are activated.** See [Schedule Label Assignment](#).

Schedule Label Assignment

You can assign labels on a recurring schedule and/or when VENs are activated.

- a. From the **Apply Rules** drop-down list, click **Schedule Label Assignment**.

- b. In the **Recurring Rule Application** dialog box, select one or both of the following options:
- **Apply rules when triggered.** Select if you want to automatically assign the specified label(s) to the matching workload(s) whenever a VEN is activated. (For details, see [Pairing and Activating the VEN](#)).
 - **Apply rules regularly.** Select if you want the specified label(s) to be applied automatically according to a schedule that you configure.

Recurring Rule Application ×

APPLY RULES WHEN TRIGGERED

☒ On Apply rules whenever a VEN is activated

APPLY RULES REGULARLY

☒ On Apply rules at a specified time at specified intervals

Date and Time

10:38 PM ⌚

Daily ▾

- ✓ Daily
- Weekly
- Monthly

Cancel Done

3. Click **Done** when you are finished.


Remove a Labeling Rule

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

Edit a Labeling Rule


You can edit a rule's condition and label(s). To learn more about rule components, see [Terminology](#).


To add a statement to an existing rule

1. Click the **Edit** icon for the rule you want to edit.
2. Click the down arrow to activate the Condition selectors.
3. Specify the statement you want to add.
4. If needed, add or remove label(s) in the **Label** field.
5. Click **Save** .

To delete a value from an existing rule

1. Click the **Edit** icon for the rule you want to edit.
2. On the condition you want to delete, click the **X** to delete it.



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

To edit a value in an existing condition

NOTE:

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

For example, if you want to change the IP range


10.13.0.26-10.13.8.26

to . . .

10.13.0.26-10.92.8.26

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


1. Click the **Edit** icon for the rule you want to edit.
2. Click the down arrow to activate the Condition selectors.

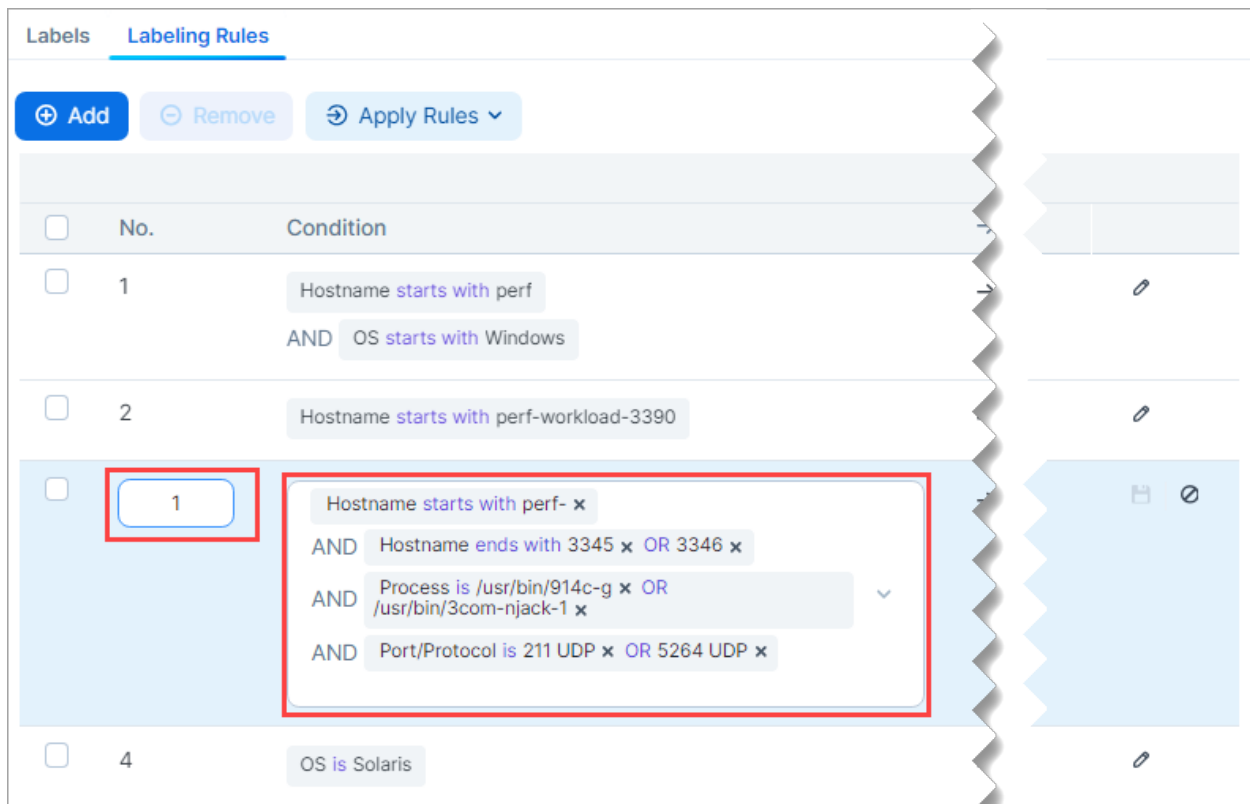
3. Add the new statement.
4. Delete the original value.
5. If needed, edit label(s) in the **Label** field.
6. Click **Save** .

Reorder Labeling Rules

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

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

1. Click the **Edit** icon for the rule you want to move. The rule number becomes an editable field.
2. Enter the new rule number in the field.
3. Click **Save** .



The screenshot shows the 'Labeling Rules' interface. At the top, there are buttons for 'Add', 'Remove', and 'Apply Rules'. Below this is a table of rules. The first rule is selected, and its condition is being edited. The rule number '1' is highlighted in a red box, and the condition is being modified to 'Hostname starts with perf- x AND Hostname ends with 3345 x OR 3346 x AND Process is /usr/bin/914c-g x OR /usr/bin/3com-njack-1 x AND Port/Protocol is 211 UDP x OR 5264 UDP x'.

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

Note that reordering rules changes the precedence of other rules.

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

Labels		Labeling Rules	
		⊕ Add ⊖ Remove ↻ Apply Rules ▼	
<input type="checkbox"/>	No.	Condition	
<input type="checkbox"/>	1	Hostname starts with perf- AND Hostname ends with 3345 OR 3346 AND Process is /usr/bin/914c-g OR /usr/bin/3com-njack-1 AND Port/Protocol is 211 UDP OR 5264 UDP	
<input type="checkbox"/>	2	Hostname starts with perf AND OS starts with Windows	
<input type="checkbox"/>	3	Hostname starts with perf-workload-3390	
<input type="checkbox"/>	4	OS is Solaris	

Export a Workload-Label-Review List

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

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

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

5. Open and review the CSV file.

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

Learn More:

- [How Label Matching Works](#)
- [Labeling Rule Examples](#)

How Label Matching Works

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

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

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

Terminology

- **Rule:** Rules consist of a condition and one or more label(s). If a workload matches the rule's condition, it is assigned the corresponding label(s), provided the workload has not already been assigned a label of the same type.
- **Condition:** Conditions are the user-defined criteria that workloads must match to be eligible for label assignment. A condition consists of one or more statements connected by AND, ensuring that workloads must satisfy all statements of the condition to match the rule.
- **Statement:** Statements define the specific workload attributes, operators, and values that are evaluated. Multiple values within a statement are considered using OR, allowing you to specify match criteria flexibly.
- **Precedence:** Rules are numbered, with Rule 1 having the highest precedence. A workload is evaluated against the rules in order, ensuring that rules with the labeling criteria most important to you are considered first.

Matching Logic

Example: Workload and Rule Evaluation

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

Learn More:

[Add and Manage Labeling Rules](#)

[Labeling Rule Examples](#)


Labeling Rule Examples

This section provides several detailed examples of adding labeling rules.

Keep in mind the following as you add labels:

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

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

1. Select **Hostname** in the **Attribute** field.
2. Select **contains** in the **Operator** field.
3. Enter **AWS** in the **Values** field.
4. Click **Close**.
5. Select one or more labels in the **Label** field.
6. Click **Save** .


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

TIP:

Match on OS version or release


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

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


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

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

1. Select **IP Address** in the **Attribute** field.
2. Select **is in** in the **Operator** field.


3. In the **Value** field, enter a narrow range such as `10.2.0.0 - 10.2.200.0`
4. Click **Close**.
5. Select one or more labels in the **Label** field.
6. Click **Save** .

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

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

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

1. Specify a hostname:
 - Select **Hostname** in the **Attribute** field.
 - Select **contains** in the **Operator** field.
 - Enter details in the **Values** field.
2. Specify an operating system:
 - Select **OS** in the **Attribute** field.
 - Select **contains** in the **Operator** field.
 - Select an operating system in the **Values** field.
3. Specify an IP address:
 - Select **IP Address** in the **Attribute** field.
 - Select **is in** in the **Operator** field.
 - In the **Values**, field enter an IP range or CIDR block.
4. Specify a listening port and/or protocol:

- Select **Port/Protocol** in the **Attribute** field.
 - In the **Operator** field, select **is** for a specific port/protocol; select **is in** to specify a range.
 - In the **Values** field, enter either a specific port/protocol or a range as appropriate.
5. Specify a process path:
 - Select **Process** in the **Attribute** field.
 - In the **Operator** field, select an appropriate operator.
 - In the **Values** field, enter all or part of a process path according to your selected operator.
 6. Click **Close**.
 7. Select one or more labels in the **Label** field.
 8. Click **Save** .

Learn More:

[How Label Matching Works](#)

[Add and Manage Labeling Rules](#)

Use AI Labeling

Overview

Through machine learning, Illumio recommends day one security policies for critical workloads.

- Use AI Labeling to apply the label to the CloudSecure resource
- Optionally, you can ignore the label recommendation

88 Home

AI Labeling

Feedback Illumio Demo

Apply Label Ignore Label

50 per page		1 – 38 of 38 Total			
<input type="checkbox"/>	Resource	Label Type	Label Value		
<input type="checkbox"/>	Name: Mongo1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Mongo1	Role	MongoDB-ai	Apply	Ignore
<input type="checkbox"/>	Name: NetBack1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/NetBack1	Role	Backup-ai	Apply	Ignore
<input type="checkbox"/>	Name: Ticketing-Prod-DB5 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Prod-DB5	Role	MySQL-ai	Apply	Ignore
<input type="checkbox"/>	Name: Zabbix1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Zabbix1	Role	Monitoring-ai	Apply	Ignore
<input type="checkbox"/>	Name: DNS1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/DNS1	Role	DNS-ai	Apply	Ignore
<input type="checkbox"/>	Name: DC1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/DC1	Role	LDAP-ai	Apply	Ignore
<input type="checkbox"/>	Name: Ticketing-Prod-DB6 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Prod-DB6	Role	MySQL-ai	Apply	Ignore
<input type="checkbox"/>	Name: Oracle1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Oracle1	Role	Oracle-ai	Apply	Ignore
<input type="checkbox"/>	Name: Ticketing-Staging-DB5 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Staging-DB5	Role	MySQL-ai	Apply	Ignore
<input type="checkbox"/>	Name: ticketing-dev-web8 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/ticketing-dev-web8	Role	Jumphost-ai	Apply	Ignore
<input type="checkbox"/>	Name: Ticketing-Staging-DB5 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Staging-DB5	Role	MySQL-ai	Apply	Ignore

Use AI Labeling

1. Select the resource for which you want to view the label details. Note that you can select one or more resources. Also note that a listed resource may have more than one entry, because each separate entry for that resource will have a different set of recommended labels.

illumio

Search

Dashboard

Servers and Endpoints

Cloud

Ransomware Protection

AI Labeling

Explore

Map

Traffic

Mesh

Policies

All Policies

Drafts & Versions

88 Home

AI Labeling

Apply Label Reject Label

4 Selected

☐

Resource

☒

Name: Ticketing-Prod-DB6
ID: /subscriptions/cf849aac-1813-4fc7-bc15-7fa92a69a77f/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Prod-DB6

☒

Name: Ticketing-Prod-DB5
ID: /subscriptions/cf849aac-1813-4fc7-bc15-7fa92a69a77f/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Prod-DB5

☒

Name: DNS1
ID: /subscriptions/cf849aac-1813-4fc7-bc15-7fa92a69a77f/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/DNS1

☒

Name: DC1
ID: /subscriptions/cf849aac-1813-4fc7-bc15-7fa92a69a77f/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/DC1

Ticketing-Prod-DB6

13e40907-5368-513b-97f1-...

NameTicketing-Prod-DB6

ID/subscriptions/cf849aac-1813-4fc7-bc15-7fa92a69a77f/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Prod-DB6

CloudAzure

Regioneast

CategoryCompute

TypeMicrosoft.Compute/virtualMachines

Resource StateSucceeded

Account IDcf849aac

Last Updated2024-05-20T04:19:45.776617Z

TagsApp | Ticketing Biz | 1234 Compliance | PCI Env | Prod Role | db

LabelsTicketing Production db Compute ComputeVirtualMachines

2. Select the resource listing with the list of labels that you want to apply.
3. Click **Apply Label**.

The screenshot shows the Illumio AI Labeling interface. On the left is a navigation sidebar with options: Dashboard, Servers and Endpoints, Cloud, Ransomware Protection, **AI Labeling** (selected), Explore (with sub-options Map, Traffic, Mesh), and Policies (with sub-options All Policies, Drafts & Versions). The main content area is titled 'AI Labeling' and has tabs for 'Apply Label' (active) and 'Ignore Label'. Below the tabs, there's a table of resources with 3 selected. The table has columns for Resource, Label Type, and Label Value. Each resource row has checkboxes for 'Apply' and 'Ignore'.

Resource	Label Type	Label Value	Apply	Ignore
Name: Mongo1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Mongo1	Role	MongoDB-ai	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Name: Ticketing-Prod-DB5 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Ticketing-Prod-DB5	Role	MySQL-ai	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Name: Zabbix1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/Zabbix1	Role	Monitoring-ai	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Name: DNS1 ID: /subscriptions/92e7ed99-799a-4807-b9a4-36b86b4b8119/resourceGroups/rsa/providers/Microsoft.Compute/virtualMachines/DNS1	Role	DNS-ai	<input type="checkbox"/>	<input type="checkbox"/>

Chapter 4

Create Policy in CloudSecure

This chapter contains the following topics:

CloudSecure Policy Model	135
CloudSecure Policy Attributes Overview	137
Organization Policy versus Application Policy	141
Resources that Support Policy	150
Unified Policy	151

This section explains how to create security policy Illumio CloudSecure. Reviewing this content helps you understand the different types of policy and how to create and use them successfully.

CloudSecure 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 CloudSecure Policy Model

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

In contrast, Illumio security policy uses a multidimensional label system to sort and describe the function of resources. By describing resources functionally, policy statements are clear and unambiguous. Illumio users assign labels to their resources to identify their applications, environments, and regions. Additionally, users specify

labels with cloud tag to label mapping. See [Cloud Tag to Label Mapping](#) for information.

Together, labeling resources and creating the corresponding rules define the security policies for resources. Illumio converts these label-based security policies into the appropriate protection for the resources.

Security Policy Guidelines

The following guidelines are recommendations on how to create your security policy in Illumio CloudSecure. 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:

- Refine your initial policy to strengthen it by narrowing overly broad access
- Use provisioning to enact your policy

Understanding Rules

Rules are an integral component of Illumio security policy. Create the rules using labels, IP lists, and applications that identify aspects of your cloud environment. See [Overview of Policy Attributes](#) in this topic for more information.

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

Types of CloudSecure Policy

CloudSecure provides two types of policies — Organization and Application. For instructions on creating rules for policies, see the pop-ups in the CloudSecure GUI. For guidelines specific to each type, see the following topics:

- [Writing Application Policy](#)
- [Writing Organization Policy](#)

Overview of Policy Attributes

Illumio CloudSecure uses the following policy attributes that help you write your security policy:

- [Labels](#)
- [Services](#)
- [IP Lists](#)

CloudSecure Policy Attributes Overview

This chapter contains the following topics:

This section describes the policy attributes that you can use to write security policies. In addition to selecting applications when you create policy rules, you also select attributes like labels, IP lists, and services to identify aspects of your cloud environment. Note that you can create new IP lists and services as needed. The values you select for these attributes will specify how your rules deny or allow access to your resources.

IP Lists

IP lists allow you to create allow and deny rules using IP addresses, IP address ranges, or CIDR blocks. These values in your rules will deny or allow access to your resources. For instructions on selecting or creating IP lists, see the in-application help pop-ups.

Overview of IP Lists

After you define an IP list, you can use it in rulesets to create rules for traffic flows. When you provision the rulesets, the rules allow or deny traffic.

Rules that use IP lists are programmed on one side of the connection only. IP lists can be used as a destination and a source.

Examples of Different IP List Entries

Single IP

You can use IPv4 or IPv6.

Examples:

- 127.0.0.1
- 2001:0db8:0a0b:12f0:0000:0000:0000:0001

CIDR Block

Use a slash to indicate a CIDR Block.

Examples:

- 192.168.100.0/24
- 2620:0:860:2::/64

IP Ranges

Use a hyphen to indicate an IP range.

Example:

- 10.0.0.0-10.255.255.255

Comments

Use a hash symbol to indicate a line comment.

Example:

- 23.4.55.6 #Comment Text

Exclusions

Use an exclamation point to exclude an IP address, CIDR block or IP range.

The excluded IP addresses must be within the included IP range.

Examples:

- !192.168.100.0/30
- !3ffe:1900:4545:3:200:f8ff:fe21:67cf

More Information on 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 resources.

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.

IP List Exclusions Caveat

To add an IP address or subnet exclusion, use an exclamation point followed by the IP address, CIDR block, or IP range as shown above. However, the following caveat applies when using the exclamation point:

- For example, to add 192.16.0.0/12 as an allowed IP address but exclude an IP address from this CIDR block, enter the following value, without the exclamation point:
 - 192.31.43.0-192.31.43.100
- For example, 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.

Labels

The Illumio CloudSecure 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 primarily by using labels. This functionality helps you categorize your resources more quickly and makes it easier to set up your policy.

Label Types

Label	Description
Environment	This label type allows you to describe a deployment based upon its stage in the product development lifecycle, such as QA, staging and production.
Application	When you define your application, this label type is created, allowing you to describe the application composed of your resources. This functionality in turn allows Illumio CloudSecure to discover applicable deployments for applications.
ServiceCategory	This label type allows you to describe key resources by their cat-

	<p>egories, such as Databases, Data Warehouse, Storage, Network Management, Network Security, Network Routing, Security Infrastructure, Account Management, Compute, Serverless, and Containers. For a list of supported resources and their categories, see Inventory Supported Resources. To view your ServiceCategory labels, see View System Labels. Note that this label type is system-generated and cannot be edited or removed. Also note that CloudSecure does not apply this label type to AWS EC2 Snapshots, AWS ElasticLoadBalancingV2 Load Balancer Target Groups, or Azure Private Endpoints. For information on which ServiceCategory labels support policy authoring, see Writing Application Policy.</p>
ServiceRole	<p>This label type allows you to describe resources according to their roles. Examples include ServiceRole:S3 and ServiceRole:RDS. Note that this label type is system-generated and cannot be edited or removed. To view your ServiceRole labels, see View System Labels. Also note that CloudSecure does not apply this label type to AWS EC2 Snapshots, AWS ElasticLoadBalancingV2 Load Balancer Target Groups, or Azure Private Endpoints. For information on which ServiceRole labels support policy authoring, see Writing Application Policy.</p>
Other labels	<p>You can use cloud tag to label mapping to create any label that meets your organization's business needs. For example, you might want to label applications according to their function.</p>

Label Resources Using Cloud Tag to Label Mapping

If you have a tagging strategy in your cloud environment, this feature lets you associate labels other than application and environment resources with your application. This functionality allows for more granularity when writing policies.

For example, if you map a “risk” cloud tag key to the Illumio label type “Risk,” you could then create an application with a tag called “risk:Critical,” which would assign the Illumio “Risk” label to the application.

Illumio recommends that you use the cloud tag to label mapping feature before creating an application definition. This workflow is recommended but not mandatory. You can create your application definition independent of any associated labels.

See [Cloud Tag to Label Mapping](#) for information.

Create an Application Definition (Label and Auto-Discovery)

Once you have added at least one deployment, you can define your applications, which will create a label for that application (which is defined using cloud tags and metadata).

In effect, the application definition comprises an application label and auto discovery of application deployments. So, by defining your application, you are labeling it and allowing Illumio CloudSecure to discover applicable deployments that you previously added.

Once you define an application, the name you gave it will appear in the *Application Label* column on the *Application Definitions* page.

At this time, application definitions and their labels are not editable once created. You can delete application definitions, however, which will delete the associated application label.

See [Define an Application](#) for information.

Services

This is an overview of services. For instructions on creating or editing services, see the in-application pop-ups. All running processes and services are available for use when writing rules.

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

Services in a Rule

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

Organization Policy versus Application Policy

This topic explains the difference between organization and application policies.

For information about creating these types of policies, see [Writing Organization Policy](#) and [Writing Application Policy](#).

About CloudSecure Policies

The *Policies* page lists all the different policies you have created in CloudSecure. The page contains two types of policies:

- Organization policies
- Application policies

What Are Organization Policies?

Codify Organizational Network Security Policies as Guardrails

You can think of organization policies as guardrail policies that prevent application policies from allowing undesired traffic, or that are additive to application policies allowing desired traffic. An organization policy can exist all by itself, but these policies are also evaluated during policy computation for any application policy.

Organization policies are broader policies that you write that are independent of applications. They can override application policies, including any future application policies, that may have overly permissive allow rules.

Although you're not constrained by an application, you could still create an organization policy for an application if you wanted to. Conversely, you might want to create a broader policy such that applications in the development environment cannot talk to anything in the production environment, or block an entire set of IP ranges, or block all Telnet traffic. You could also write an organization policy using more fine-grained labels.

Define Organization Policies

Once you onboard your cloud accounts, you can define your organization policies. To write organization policies, go to **Policies > Organization Policies** tab. See [Writing Organization Policy](#).

What are Application Policies?

Security teams can drive segmentation policies to control network traffic using Illumio labels, services, and IP/IP lists to define what can talk to applications, what data can

be transferred from an organization's network, etc. Creating application policies is critical to minimizing an attacker's lateral movement.

Define Application Policies

If a policy addresses anything within an application, because you've now defined what an application is, it's an application policy and appears on the *Application Policies* tab.

Before you write application policies, you will want to first define services and IP lists by going to the *Policies* menu and selecting the **Services** and **IP List** tabs. See [Services](#) and [IP Lists](#) for information.

First, you will want to use the *Tag to Label Mapping* menu available in the left navigation under *Application Discovery*. Once you use the tag to label mapping feature, you can select the labels that you create when writing policy for your applications. See [Cloud Tag to Label Mapping](#) for information.

To write application policies, go to **Applications > your application > Policy** tab. See [Writing Application Policy](#) for information.

What Happens When Org and App Policies Conflict?

Consider the following example. If you write an application policy with an allow rule permitting all Telnet connections, but you have an organization policy with an override deny rule that blocks them all, the override deny rule will negate the allow rule.

Writing Organization Policy

This topic provides an overview of using rules to write organization Illumio CloudSecure policies. Organization policies are guardrail policies that prevent application policies from allowing undesired traffic, or that are additive to application policies allowing desired traffic. An organization policy can exist all by itself, but these policies are also evaluated during policy computation for any application policy.

For an overview of the Illumio CloudSecure policy model, see [CloudSecure Policy Model](#). For a list of resources against which you can write policy, see [Resources that Support Policy](#).

In order to write policy, you must create rules for the policy. Illumio CloudSecure has the following rule types for organization policies:

- Allow Rules

You can write rules that allow communication between sources and destinations. For example, if you have Allow Rule A in an organization policy and Allow Rule B in application policy, they will be combined and become Rule A and B for the application rule. Use cases examples include instances where:

- You want to allow SNMP traffic between two applications even if there are no such specific application policies with that allow rule
- You want to have an organization-wide allow rule that is more inclusive than present application policy allow rules dictate

- Override Deny Rules

This rule type is typically used to deny communication between sources and destinations that might inadvertently be given with allow rules created by another CloudSecure administrator. Override deny rules take precedence over all other types of rules, including organization policy allow rules. Use cases include instances where you do *not* want organization or application policies:

- Allowing development to talk to production
- Allowing public access to a database
- Allowing SSH anywhere
- Allowing Telnet anywhere

Differences between Organization and Application Policies

You can think of organization policies as guardrail policies that might need to be applied across your infrastructure. See [Organization Policy versus Application Policy](#) for information.

Unlike application rules, you do not start writing organization policy from an application seen in the *Applications* left navigation menu. Instead, go to the **Policies > Organization Policies** tab and click **Add** to begin. For instructions on creating rules for organization policies, see the pop-ups in the CloudSecure GUI.

Once you have saved your rule for the organization policy, the rule automatically enables, and the *Provision Status* column will have a green *Pending* icon. The *Policies > Organization Policies* tab will also show a green *Pending* icon in the *Provision Status* column. Depending on what you are doing to a given policy the icon may be red, green, or blue. See [Pending Icon Color Codes](#).

Pending Icon Color Codes

Color	Meaning
Red	Deletion pending
Blue	Update pending
Green	Addition pending

Guidelines, Permitted Combinations, Provisioning, and Caveats

These concepts for writing organization policy override deny rules are virtually the same as for application policies. See [Writing Application Policy](#) for information.

NOTE: Organization policies let you select All Applications for Allow Rule destinations.

Writing Application Policy

Illumio allows or denies traffic between applications using policies that you write. For an overview of the Illumio CloudSecure policy model, see [CloudSecure Policy Model](#). For a list of resources against which you can write policy, see [Resources that Support Policy](#).

In order to write application policies, you must create rules for the policy. Illumio CloudSecure has the following types of rules for application policies:

- **Override Deny Rules**

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

- **Allow Rules**

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

- **Deny Rules**

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

This topic provides an overview of using rules to write Illumio CloudSecure policies. For instructions on creating rules for policies, see the in-application help.

Differences between Application and Organization Policies

You can think of application policies as segmentation policies to control network traffic using Illumio labels, services, and IP/IP lists to define what can talk to applications. The guidelines below are generally applicable to writing both organization and application policies. For differences, see [Organization Policy versus Application Policy](#).

Guidelines

Use the following guidelines when creating rules for your policies:

- From the *Source* and *Destination* drop-down lists, you can select a combination of applications, labels, and IP lists. Note that when programming security groups, CloudSecure optimizes the rules by grouping a set of IPs into a CIDR block if possible.
- From the *Destination Services* drop-down list, you can select a combination of services and ports. Note that when there are adjacent rules i.e., adjacent ports, with all other parameters same, CloudSecure merges those rules. For example if you have Rule1 (ports 87100-87104), Rule2 (ports 87105), Rule3 (ports 87106-87110), then CS combines those rules to program a single rule with the port range 87100-87110.
- In the source or destination fields, select **All Resources** to include all resources at once instead of selecting them individually. By using All Resources in your source or destination, you can write organization policies for all resources in onboarded cloud accounts.
- The CloudSecure UI will prevent you from selecting disallowed source and destination combinations. For a full list of permitted source and destination combinations in a rule, see [Permitted Rule Writing Combinations](#).
- After completing your selections, click the **Save** icon at the end of the row for that rule
- To edit a rule, click the **Edit** icon at the end of the row
- After adding a rule, the *Status* column displays a green *Enabled* icon and the *Provision Status* column displays a green *Pending* icon
- Rules can be disabled or removed individually or in groups by selecting the check box next to a rule
- To enforce a rule, you must provision the policy. For more information about provisioning, see [Provisioning](#)

- Reverting a policy from the *Applications > your policy name > Policy* tab will cancel pending changes to the policy, including rules with a green *Pending* icon in the *Provision Status* column, and revert to the previously provisioned policy
- From the left navigation *Policies* menu item, reverting a policy that still has its provisioning pending will cancel that provisioning but leave previously saved policy rules intact

Permitted Rule Writing Combinations

Inter-Application and Inter-Deployment Policy

Illumio allows you to write rules between your applications and between your deployments. However, in order to write these rules, rules must be written in the context of the application on the inbound side of the rules. In other words, you can only write inbound inter-application and inter-deployment policy rules. However, when you do so, CloudSecure implicitly writes outbound rules for the security group containing the source application. This is to avoid the need for you or the security group owner to explicitly write a corresponding outbound rule.

Under a given application, if you want to specify an application in the destination of the rule, the application must match the application in the context. So, the destination application must be the application in which you are writing the rule. The source application can be a different application (or the same) than the application context in which you are writing the policy.

If a deployment is specified, the same principal applies to the destination deployment. You can write the rule for only the deployment context in which you are writing the policy. The source deployment can be a different deployment (or the same) than the application context in which you are writing the policy.

NOTE:

If the source does not match the application or deployment context, Illumio CloudSecure will take the meaning of the labels literally. For example, if the context is `app:CRM`, `deployment:PROD`, a rule with the source as `app:FINANCE` will represent all resources under `app:FINANCE`, regardless of deployment. This is true of all rule types (allow, deny, etc.).

	And
If Source is	Service Destination can be
	is

Application and/or deployment (any)	Any service	The application and/or deployment (if applicable), so long as it is the same as the one providing the context
IP List	Any service	Application or label

Intra-Application and Intra-Deployment Policy

In order to write rules within your application context, you can specify labels or IP lists on either side of the rule.

NOTE:

These labels must not be of the application or deployment types in order for the following to apply.

If a label is used on either side of the rule, Illumio CloudSecure will calculate which resources match both the context (application and/or deployment) and the label used, and create the rule accordingly.

If Source is	And Service is	Destination can be
Any label or IP list	Any service	Any label or IP list

Provisioning

When you provision updates, Illumio CloudSecure recalculates any changes made to rules, and then transmits those changes to all affected enforcement points. All the changes you make to those rules are considered to be in a “draft” state until you provision them.

Previewing the Impact of Provisioning a Policy

This section provides an overview of the *Show Impact* feature. For instructions on previewing policy impact, see the in-application help.

Before you provision a policy, you may wish to gauge what its impact will be. It can be difficult to see how it may map to destinations, various rules in security groups, enforcement points, and so forth.

To see such mappings on a policy that has not yet been provisioned, click **Show Impact**. You can then choose one of the following from the drop-down menu:

- All security controls
- Azure NSGs
- AWS Security Gateways
- Network Access Control Lists

Each of these will show you the following:

- CSP
- Resource
- Account ID
- Number of Protected Resources

If you select any particular affected AWS SG or Azure NSG , you may see rules that come from other applications and/or policies. Note that only Illumio-written rules will display.) The SG or NSG draft change summary will include the account name, as well as inbound and outbound rules with the following details :

- Provision Status (this can tell you whether a rule is being added, removed, or is already in place)
- Source
- Destination
- Port
- Protocol
- Action (deny, override deny, or allow)

Confirming

Once you have previewed the anticipated impact, you are ready to decide whether to proceed with provisioning.

You are given a description field for adding any comments when provisioning a policy. After you provision your changes, those changes become “active,” which is to say it is in enforcement mode. When you confirm by clicking **Confirm & Provision**, the *Policies* page *Provision Status* column displays the applications with policies, including those that are pending.

To see if CloudSecure experienced errors when provisioning your policies, click the **Provisioning errors** button in the upper right-hand corner of the page. The *Provisioning errors* page will display the cloud, name and ID, status, and modification date for both application and organization policies that experienced errors during provisioning.

Caveats

- Only rules that use the following attributes are supported: applications, labels, IP lists, and services
- As AWS does not have a deny rule concept for Security Groups, a CloudSecure override deny rule will only be implemented if there is a matching allow rule that is overlapping in scope. In effect, the override deny rule will constrict where the allow rule is implemented.
- You cannot write policy rules using metadata, but you can map cloud tags to Illumio labels and then write policy rules using that label
- Only the following ServiceCategory labels can be used when authoring policy: Compute, Serverless, and Network Management. ServiceRole labels can also be used when authoring policy, but the service roles must have resources that support policy. See [Resources that Support Policy](#).

NOTE:

Because CloudSecure may not always discover elastic network interfaces (ENIs), a flow search based on resource IDs will not work for the following supported resources if their *Details* page does not display the ENI. The workaround is to search using the IP address of the associated ENI, if known:

- AWS RDS DBInstances
- AWS RDS DBClusters
- ElasticLoadBalancingV2 load balancers
- AWS MemoryDB clusters
- AWS ElastiCache for Redis clusters
- AWS Redshift clusters

Resources that Support Policy

Illumio CloudSecure supports writing policy for the following types of resources. Note that policy enforcement is done through Security Groups on AWS and through Network Security Groups on Azure. For a list of all resources that appear in the *Inventory* page, and additional details such as flow support and attached resources, see [Inventory Supported Resources](#). For a list of all resources that appear in the *Cloud Map* and *Traffic* pages, see [Cloud Map Supported Resources](#) and [Traffic Supported Resources](#).

AWS

Resource

Category

EC2 Instances	Compute
RDS DB Clusters	Databases
RDS DB Instances	Databases
ElastiCache CacheClusters	Databases
MemoryDB Clusters	Databases
ElasticLoadBalancingV2 Load Balancers	Network Routing
Redshift Clusters	Data Warehouse
Lambda Functions	Serverless
Azure	
Resource	Category
Virtual Machines (inclusive of "spot" VMs)	Compute
Virtual Machine ScaleSets	Compute

NOTE:

Because CloudSecure may not always discover elastic network interfaces (ENIs), a flow search based on resource IDs will not work for the following supported resources if their *Details* page does not display the ENI. The workaround is to search using the IP address of the associated ENI, if known:

- AWS RDS DBInstances
- AWS RDS DBClusters
- ElasticLoadBalancingV2 load balancers
- AWS MemoryDB clusters
- AWS ElastiCache for Redis clusters
- AWS Redshift clusters

Unified Policy

This topic explains the unified policy capability of the Illumio Zero Trust Platform.

Overview

Three policy tabs are available in the Policies menu:

- All Policies: This tab includes all policy types, described below
- Organization Policies: Considered guardrail policies, they prevent application policies from allowing undesired traffic. These policies apply to all scopes. See [Writing Organization Policy](#).

- **Application Policies:** Security teams can drive segmentation policies to control network traffic using Illumio labels, services, and IP/IP lists to define what can talk to applications, what data can be transferred from an organization's network, and so forth. Creating application policies is critical to minimizing an attacker's lateral movement. See [Writing Application Policy](#).

The screenshot shows the 'Policies' page in the Illumio CloudSecure interface. The 'Application Policies' tab is selected. The table below lists the policies:

Provision Status	Status	Name	Scopes	Last Modified On
<input type="checkbox"/>	Enabled	Admin Policy	All	06/04/2024, 07:58:09
<input type="checkbox"/>	Pending	Demo	All	06/03/2024, 10:23:41
<input type="checkbox"/>	Enabled	Dev/Prod separation	All	06/04/2024, 07:54:02
<input type="checkbox"/>	Pending	Ecomm policy	ecomm	06/07/2024, 15:50:17
<input type="checkbox"/>	Disabled	pos	pos	06/04/2024, 08:12:07
<input type="checkbox"/>	Disabled	ringfence pos	All	06/04/2024, 08:12:07
<input type="checkbox"/>	Disabled	Staging	Staging	06/04/2024, 08:12:07
<input type="checkbox"/>	Disabled	test: block staging to production	All	06/04/2024, 07:52:43

For distinctions between organization and application policy, see [Organization Policy versus Application Policy](#).

From each of the above tabs, you can write policies for policy objects that span both the cloud and the datacenter:

- Services
- IP Lists
- Labels
- User Groups
- Label Groups
- Virtual Services
- Virtual Servers

Illumio allows or denies traffic between applications using policies that you write. In order to write application policies, you must create rules for the policy. Illumio has the following types of rules for application policies:

- Override Deny Rules

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

- Allow Rules

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

- Deny Rules

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

- Custom IPtables Rules

You can write rules for Linux workloads.

Notices

- Scopes and role-based access control (RBAC) remain the same as in previous releases
- Allow rules function the same as in previous releases
- Override Deny rules are now supported
- Override Deny rules take precedence over Allow Rules. They block traffic with no exceptions.
- Deny rules can be scoped and support RBAC
- Deny rules are introduced in policies to support scope and RBAC
- “Global” Deny rules (also known as enforcement boundaries) will be deprecated. Illumio recommends that you move legacy deny rules into policies. See the *Guidelines* section of [Writing Organization Policy](#).

Chapter 5

CloudSecure Administration

This chapter contains the following topics:

Connector	154
Events	155
Role-Based Access Control	156
User Management	157
My Profile	159
My Roles	159

This section describes the tasks you need to perform to administer CloudSecure for your organization; such as adding and managing CloudSecure users, reviewing CloudSecure events, and setting up a connector to receive notifications about CloudSecure activity.

Connector

This topic describes the purpose of the Illumio CloudSecure Connector feature, and provides a general example of how you would use it. For instructions on how to connect a specific workflow and incident management tool, such as Slack, using the *Connector* page, see the applicable pop-up in the CloudSecure GUI.

Use Case and Example

This feature lets you connect workflow and incident management tools, such as messaging applications or others, to CloudSecure. For example, you might want to receive a notification in your messaging application when a policy changes or when a deploy-

ment is removed. (For such notifications, a message banner displays the time and frequency of the aggregated alerts that CloudSecure delivers to the application.)

The following steps illustrate how you might set up a connection to such an application.

1. The first part of the sequence would be to browse to the **Settings > Connector > Apps** tab. From there, the pop-ups will give you instructions.

Depending on your application, you may need to provide the following:

- Channel Name (CloudSecure does not verify the name, so make sure it is correct.)
- Webhook URL (This would be how CloudSecure knows where to deliver the message.)

The dialog may have fields for other characteristics, depending on the application.

As soon as a channel is configured, any subsequent alerts would also be scheduled for your newly added channel.

Alerts are sent to all configured channels. In other words, the same alert message is sent to all of them if all the channels were added before the first alert of the day got triggered.

2. The next step would be to edit or delete your created channels if needed. Click the application tile to see a list of channels.

Different kinds of workflow and incident management tools will vary widely, so see the pop-up in the CloudSecure GUI that is specific to that particular one.

Events

As a CloudSecure customer, you can view a list of CloudSecure activities that occurred. CloudSecure displays events for two CloudSecure actions:

- When users onboarded cloud accounts for your organization
- When users deleted cloud accounts onboarded for your organization

If you are a customer for other Illumio products, such as Illumio Core, you view events for that Illumio product separately in the Illumio console. For more information, see the documentation for your Illumio product.

CloudSecure organizes events by account name and the dates and times that the events occurred. By default, CloudSecure sorts the list by displaying the oldest event

first. Each event includes a message indicating the user in your organization who performed the action.

View your Events

1. From the left navigation, choose **Settings > Events**.
2. (Optional) Customize how you want to display the events:
 - To sort events by account, by the date/time the event occurred, or by the type of event click the column heading.
 - To add or remove columns from the table, select the columns in the *Customize columns* menu.
 - To change the number of event displayed per page, select the number from the *Per page* menu.

Role-Based Access Control

This section describes how to manage user roles for your organization. For user management topics, see [User Management](#). For more information on roles, see [My Roles](#).

Add Roles

To add or remove user scopes and roles, navigate to **Access > Users** and click the user entry in question. A user detail panel opens.

1. Click **Add Role> Add Unscoped Role**. (Scoped roles are not applicable to CloudSecure.)
2. Select one or more of the roles listed below and click **Grant Access**.

The following roles are available for CloudSecure:

- Multi-product Roles

The user has these roles for CloudSecure and other Illumio products on the Illumio Console:

- Owner
 - Viewer
- CloudSecure-specific Roles

The user has these roles for CloudSecure only:

- Cloud Security Onboarding Administrator
- Cloud Security Policy Author
- Cloud Security Label Administrator
- Cloud Security Auditor
- Cloud Security Incident Responder
- etc.

Once a role is assigned to a user, you can click on the Role entry and see the detail page for that role. It lists all users with that role. You can then add or remove users to and from that role.

To view all available roles, browse to **Access > Roles**. This lists all the roles. Click on one of the roles to see all users assigned to it.

Remove Roles

To add or remove user scopes and roles, navigate to **Access > Users** and click the user entry in question. A user detail panel opens.

1. Select a user and click **Remove**.
2. In the dialog that appears, click **Remove**.

If a user has only one role, and you remove their access to that sole role, this removes the user account entirely. If the user has more than one role, removing a given role will not remove the user account.

User Management

This section describes how to manage users for your organization.

About Users

You add local users so that other members of your organization can use Illumio CloudSecure capabilities for their zero-trust segmentation programs. Some users have owner privileges, meaning that those users can perform the same tasks. All Illumio CloudSecure users are assigned a Cloud Security administrator role that provides them access to all the capabilities in the product, including the ability to invite additional users. Core administrative users are able to see the *Access Restriction* menu item.

You create local users in Illumio. You do not manage them outside the product using an IdP.

When you become a customer or trial user, you must sign in to add or remove users. For more information, see [Signing In](#).

This first user then sets up additional users. All users can add local users to their organization. Once a user is added to Illumio, they will need to complete their setup through the Okta activation process.

Add Users

Only users with an owner role can add other users.

1. From the left navigation, choose **Access > Users**. The list of users added to your organization appears.
2. Click **Add**. The *Add User* dialog box appears.
3. Enter the user's name. Only users see their name displayed in the UI when they sign in.
4. In the *Add User* dialog box, enter the user's email address and click **Add**.

The email address domain must match the domain used by your organization.

The new user enters this email address when they sign in.

Illumio uses this email address in the UI. It displays the user's email address to track user actions in the *Events* page.

The *Add User* dialog box closes and the list of users refreshes with the new user.

What Happens Next?

After you add a user, they receive an email from the Okta service with the subject "Welcome to Okta!" This email provides information about how the user can activate their Okta account. Illumio utilizes Okta to provide multi-factor authentication.

In addition to the local user account created in Illumio, users have access to an Okta dashboard where they can manage the security that Okta provides for sign-in. The Okta email includes a link to the user's Okta dashboard.

To access your user Okta dashboard, go to your Okta email and locate the URL in the line beginning with "Your organization's sign-in page is...."

For the next steps, see [Signing In](#).

Delete Users

Only users with an owner role can delete other users. However, they cannot delete their own user accounts. Not all users have administrative privileges, as role-based

access is possible for Core users who are not owners.

When a customer or trial user is provisioned access, there is a primary security administrator email that is associated with that account. This user cannot be deleted.

To delete users:

1. From the left navigation, choose **Access > Users**. The list of users added to your organization appears.
2. Select the users you want to delete.
3. Click **Remove**. The *Remove User* confirmation dialog box appears.
4. Confirm that you are removing the correct users and click **Remove**.

Add or Remove Roles

To add or remove roles, see [Role-Based Access Control](#).

My Profile

In addition to resetting your password, you can view the following in your profile:

- Email Address/Username
- First Name
- Last Name

My Roles

The upper right-hand corner has a dropdown menu where you see your login name. The *My Roles* menu item, CloudSecure customers view current roles by the following columns:

- Scopes - CloudSecure values may include: All
- Roles - CloudSecure values may include: Owner, Viewer, and other CloudSecure roles. A user may have a combination of any of these roles.

The following should be noted:

- Only unscoped roles are applicable to CloudSecure
- Global roles other than Owner are not applicable to CloudSecure
- Users with both CloudSecure and Core subscriptions may see roles and scopes that apply to only one of those products.

Chapter 6

CloudSecure Reference

This chapter contains the following topics:

CloudSecure Requirements	160
--------------------------------	-----

This section itemizes the permissions that your CloudSecure service accounts require to access specific entities in your public cloud environments.

These permissions are unique based on the type public cloud you onboard.

CloudSecure Requirements

This chapter contains the following topics:

This section itemizes the permissions required by CloudSecure so that the service accounts you create have the necessary permissions for Illumio CloudSecure to access specific entities in your public cloud environments. These necessary permissions are unique to each public cloud supported by CloudSecure.

AWS Requirements

Overview

The following information is important to understanding how Illumio interacts with AWS.

Service Accounts in the CloudSecure Context

Within the CloudSecure platform, a "service account" refers to an account used by CloudSecure to interact with its own services (CloudSecure services) rather than

directly with your AWS services. This account is primarily used for internal operations within CloudSecure, such as making API calls to the CloudSecure platform, and is separate from AWS IAM roles and permissions.

The IAM Role for AWS

For reading the current state of AWS resources, and writing security groups to the customer's AWS accounts, CloudSecure requires the creation of an identification and access management (IAM) role within the customer's AWS account. CloudSecure assumes this IAM role to perform actions in AWS, such as reading resources and managing policies. This is consistent with Amazon's recommended practice of using cross-account roles for granting external services access to AWS resources. The IAM role ensures secure and scoped access in accordance with the principle of least privilege.

AWS IAM Permissions

To onboard your AWS account into CloudSecure, you will need use the CloudFormation Stack to create an IAM role within your AWS account, which CloudSecure will assume to make API calls. This role must be granted permissions to specific AWS resources for CloudSecure to provide visibility and manage policies for those resources. It is important to note that CloudSecure relies on the cross-account role assumption methodology. Ensure that you regularly check this page for updates, as new policies may be required in the future.

CloudSecure IAM Role Configuration

To facilitate CloudSecure's access to your AWS environment, you must create an IAM role specifically for CloudSecure within your AWS account. This role must be assigned the following policies:

- **SecurityAudit (managed by AWS):** Permissions in this policy are required for CloudSecure to read the resources in your AWS account.
- **IllumioCloudAWSIntegrationPolicy:** Permissions in this policy are required for CloudSecure to read the resources in your AWS account.
- **IllumioCloudAWSProtectionPolicy:** Permissions in this policy are required for CloudSecure to write policies for your AWS account.

Read Only Policy

The following items are AWS IAM read permissions that you will need to grant to the Illumio AssumeRole:

ManagedPolicyArns: ["arn:aws:iam::aws:policy/SecurityAudit"]

Policies

- PolicyName: IllumioCloudAWSIntegrationPolicy
 - PolicyDocument:
 - Version: 2012-10-17
 - Statement:
 - Effect: Allow
 - Resource: '*'
 - Action:
 - 'apigateway:GET'
 - 'autoscaling:Describe*'
 - 'cloudtrail:DescribeTrails'
 - 'cloudtrail:GetTrailStatus'
 - 'cloudtrail:LookupEvents'
 - 'cloudwatch:Describe*'
 - 'cloudwatch:Get*'
 - 'cloudwatch:List*'
 - 'codedeploy:List*'
 - 'codedeploy:BatchGet*'
 - 'directconnect:Describe*'
 - 'dynamodb:List*'
 - 'dynamodb:Describe*'
 - 'ec2:Describe*'
 - 'ecs:Describe*'
 - 'ecs:List*'
 - 'elasticache:Describe*'
 - 'elasticache:List*'
 - 'elasticfilesystem:DescribeAccessPoints'
 - 'elasticfilesystem:DescribeFileSystems'

- 'elasticfilesystem:DescribeTags'
- 'elasticloadbalancing:Describe*'
- 'elasticmapreduce:List*'
- 'elasticmapreduce:Describe*'
- 'es:ListTags'
- 'es:ListDomainNames'
- 'es:DescribeElasticsearchDomains'
- 'fsx:DescribeFileSystems'
- 'fsx:ListTagsForResource'
- 'health:DescribeEvents'
- 'health:DescribeEventDetails'
- 'health:DescribeAffectedEntities'
- 'kinesis:List*'
- 'kinesis:Describe*'
- 'lambda:GetPolicy'
- 'lambda:List*'
- 'logs:TestMetricFilter'
- 'logs:DescribeSubscriptionFilters'
- 'organizations:Describe*'
- 'organizations:List*'
- 'rds:Describe*'
- 'rds:List*'
- 'redshift:DescribeClusters'
- 'redshift:DescribeLoggingStatus'
- 'route53:List*'
- 's3:GetBucketLogging'
- 's3:GetBucketLocation'
- 's3:GetBucketNotification'
- 's3:GetBucketTagging'
- 's3:ListAllMyBuckets'
- 'sns:List*'
- 'sqs:ListQueues'
- 'states:ListStateMachines'

- 'states:DescribeStateMachine'
- 'support:DescribeTrustedAdvisor*'
- 'support:RefreshTrustedAdvisorCheck'
- 'tag:GetResources'
- 'tag:GetTagKeys'
- 'tag:GetTagValues'
- 'xray:BatchGetTraces'
- 'xray:GetTraceSummaries'

Write Policy

The following items are AWS IAM write permissions that you will need to grant to the Illumio AssumeRole.

- PolicyName: IllumioCloudAWSProtectionPolicy
 - PolicyDocument:
 - Version: 2012-10-17
 - Statement:
 - Effect: Allow
 - Resource:
 - 'arn:aws:ec2:*:*:security-group-rule/*'
 - 'arn:aws:ec2:*:*:security-group/*'
 - Action:
 - 'ec2:AuthorizeSecurityGroupIngress'
 - 'ec2:RevokeSecurityGroupIngress'
 - 'ec2:UpdateSecurityGroupRuleDescriptionsIngress'
 - 'ec2:AuthorizeSecurityGroupEgress'
 - 'ec2:RevokeSecurityGroupEgress'
 - 'ec2:UpdateSecurityGroupRuleDescriptionsEgress'
 - 'ec2:ModifySecurityGroupRules'
 - 'ec2:DescribeTags'
 - 'ec2:CreateTags'
 - 'ec2:DeleteTags'

FLOW READ Policy

- 's3:ListBucket'
- 's3:ListBucketVersion'
- 's3:GetBucketLocation'
- 's3:GetObject'

Flow Logs

Supported Flow Log Fields

Illumio CloudSecure uses the following fields in the logs: srcaddr, srcport, dstaddr, dstport, protocol, action, bytes, start, action, log-status, packets, tcp-flags*, interface-id*, flow-direction*, pkt-srcaddr*, pkt-dstaddr*

Fields marked by * are optional, but their absence will lead to limited functionality. It is strongly recommended that the log to contain all used fields. This requires selecting **Custom format** for the *Log record format* option.

For example, you would choose the following from the list in AWS:

```
${action} ${bytes} ${dstaddr} ${dstport} ${end} ${flow-direction} ${interface-id}
${log-status} ${packets} ${pkt-dstaddr} ${pkt-srcaddr} ${protocol} ${srcaddr} ${s-
rcport} ${start} ${tcp-flags}
```

All the required (i.e., not marked by *) fields are in Version 2 (the default AWS set)

Flow Log Support Notes

For instructions on setting up flow logs, see *Set up Flow Logs* in [Grant Flow Log Access](#).

- Only the default "text" format is supported for S3 storage of flow logs
- There is no support for the "Hive-compatible S3 prefix"
- There is currently no support for the "optional prefix" (customer path prefix inside the S3 bucket) for flow log destinations
- How CloudSecure fetches the flow logs depends on your configuration (e.g., a central account or multiple accounts)
- Every 10 minutes the map ingests traffic flows in 60-minute chunks. Flows are shown only for completed chunks. This means that if flow log access has just

been enabled, you would need to wait at least an hour to see the flows in the Cloud Map, Traffic, and Inventory pages. However, if you enabled flow log access some time ago and already have previous 60-minute flow chunks, you would see the updated flow within 10 minutes.

CloudSecure IP Addresses for Flow Log Access

The CloudSecure control and data plane uses the following public IP addresses to reach customer networks, so add them to your firewall allowed list:

- 35.163.224.94
- 54.190.103.0
- 44.226.137.227
- 35.167.22.3
- 52.88.124.247
- 52.88.88.252

The CloudSecure UK data plane uses the following public IP addresses to reach customer networks, so add them to your firewall allowed list:

- 18.169.5.9
- 13.41.233.77
- 18.169.6.17

The CloudSecure APAC data plane uses the following public IP addresses to reach customer networks, so add them to your firewall allowed list:

- 13.54.140.138/32
- 52.63.108.169/32
- 52.64.120.98/32

Background

When you start the onboarding process and begin creating IAM roles from the CloudSecure user interface, the restricted area console lets you run the stack. The following operations will occur at that time:

- Creation of a role for Lambda execution function with new permissions
- Creation of a role for Illumio to talk to AWS
- Creation of a Lambda function

- Creation of a custom resource for Lambda invocation
- Return of the Amazon Resource Name (ARN) and external ID via the Lambda function role back to CloudSecure

Note that the Lambda role cannot be deleted after onboarding. If it is removed, then the roles will be deleted along with it, which prevents CloudSecure from synchronizing resources from the cloud.

Updating Permissions on Assume Role

CloudSecure has the ability to update and modify EC2 security groups on a continuous basis. Use these steps provide CloudSecure with permissions for the newly added resources.

1. Download the CFT template to update permissions from this link. <https://cloud-secure-onboarding-templates.s3.us-west-2.amazonaws.com/cloudsecure/aws-policy-update.yaml>
2. Login to the AWS console of account to which you need to update the permissions to run the cloudformation stack.
3. Under services click **CloudFormation**.
4. Click **Create StackSet**.
5. In the *Choose template* page select, template ready and upload a template file option, and upload the downloaded template and click **Next**.
6. In the *Specify stackset details* page, enter the stack name. The stack name must be unique and not the same name used to create previous stacks.
7. In the *IAMRoleName* box, enter the name of the assume role created in AWS when onboarding with CloudSecure. By default, the name is *Illu-mioCloudIntegrationRole*. Click **Next**.
8. If you had given a different name during onboarding, make sure to give the same name. (The name can be verified by going to Service->IAM→ roles and finding the role name.)
9. Click continue and in the *Review* page, select the acknowledgment check box and click **Submit**.

The stack will run and add the newly required permissions to the role.

Handling Failures or Other Errors

CloudFormation Template Failures

In the event of a CFT failure, perform the following steps:

1. Completely delete the previous deployment stack.
2. Ensure that the stack name and resources being created are not already present.

If these steps are not done, the CFT will continue to fail.

Azure Requirements

Azure Permissions

The following items are Azure permissions that will be need to be granted to the Illumio App that is registered in Azure Active Directory. Check this page for updates, as new permissions may be included in the future.

READ Permission

Reader - role

NSG Write Permission

Use these permissions to create custom roles. Any custom roles with elevated permissions need to be defined as part of the PowerShell script that is run when you onboard an Azure subscription. See [Onboard an Azure Cloud Subscription](#) for information.

For example, if the user onboarding Azure does *not* have owner permissions, the "Illumio Network Security Administrator" custom rule needs to be created with these NSG write permissions *before* the onboarding PowerShell script is run.

If the user onboarding Azure does have owner permissions, these permissions will be automatically assigned to the "Illumio Network Security Administrator" custom role that is created when the onboarding PowerShell script is run.

- "Microsoft.Network/networkInterfaces/effectiveNetworkSecurityGroups/action"
- "Microsoft.Network/networkSecurityGroups/read"
- "Microsoft.Network/networkSecurityGroups/write"
- "Microsoft.Network/networkSecurityGroups/delete"

- "Microsoft.Network/networkSecurityGroups/join/action"
- "Microsoft.Network/networkSecurityGroups/defaultSecurityRules/read"
- "Microsoft.Network/networkSecurityGroups/securityRules/write"
- "Microsoft.Network/networkSecurityGroups/securityRules/delete"
- "Microsoft.Net-work/networksecuritygroups/providers/Microsoft.Insights/diagnosticSettings/read"
- "Microsoft.Net-work/net-worksecuritygroups/providers/Microsoft.Insights/diagnosticSettings/write"
- "Microsoft.Net-work/networksecuritygroups/providers/Microsoft.Insights/logDefinitions/read"
- "Microsoft.Network/networkWatchers/securityGroupView/action"

FLOW

Storage Blob Data Reader – role

Flow Log Support

Illumio CloudSecure supports NSG Flow **logs** version 2 (includes flow state and byte counts), but does not support version 1. CloudSecure also supports VNet flow logs.

There is no support for other "flow" logs, e.g., Firewall logs.

For instructions on setting up flow logs, see [Set up Flow Logs in AWS and Azure](#).

CloudSecure IP Addresses for Flow Log Access

The CloudSecure control and data plane uses the following public IP addresses to reach customer networks, so add them to your firewall allowed list:

- 35.163.224.94
- 54.190.103.0
- 44.226.137.227
- 35.167.22.3
- 52.88.124.247
- 52.88.88.252

The CloudSecure UK data plane uses the following public IP addresses to reach customer networks, so add them to your firewall allowed list:

- 18.169.5.9
- 13.41.233.77
- 18.169.6.17

The CloudSecure APAC (Sydney, Australia) data plane uses the following public IP addresses to reach customer networks, so add them to your firewall allowed list:

- 52.64.120.98
- 52.63.108.169
- 13.54.140.138

Background

The Reader Role

This role gives CloudSecure the permissions to read data or resources from your subscription. According to Microsoft, the role is defined as follows: "View all resources, but does not allow you to make any changes."