TECHNICAL
DOCUMENTATION

Illumio Segmentation for the Cloud User Guide

Publication date November 4, 2025

This documentation portal contains information about the Illumio Segmentation for the Cloud release. It provides the complete set of content for onboarding your cloud accounts, defining your cloud environments in Illumio Segmentation for the Cloud, and creating security policy.

Table of Contents

Release Notes

The content in this category provides information about key release issues.

Current Illumio Cloud what's new and release notes

These release notes describe the new features, enhancements, resolved limitations, and known limitations for Illumio Segmentation for the Cloud.

What's new in this release - October 16, 2025

Issue Number	Feature	Description
CLOUD-13843	Onboarding	**Fixed issue with configuring flow access to AWS buckets with subfolder paths** An issue preventing flow access from being granted to AWS destinations using subfolder paths is resolved. Previously during onboarding, configuration with the CloudFormation template may have failed when granting flow access to an AWS bucket with VPC flow logs configured to use subfolder paths.

Previous Illumio Cloud what's new and release notes for 2025

These prior release notes describe the new features, enhancements, resolved limitations, and known limitations for Illumio Segmentation for the Cloud in previous 2025 releases.

What's new in this release - September 23, 2025

Issue Number	Feature	Description
CLOUD-12799	Onboarding	The issue causing errors when using the Remove and Disable buttons during onboarding is resolved. Users can now update folders and projects in bulk without errors, helping them manage cloud accounts more effectively.
CLOUD-12877	Policy	The policy impact feature now functions correctly without server errors, allowing users to view policy impacts as expected. This fix improves reliability when creating or editing policy rules in the production environment.
CLOUD-12789	Connector	Illumio Segmentation for the Cloud now correctly sends Resource Count data to Salesforce instead of the previously used converted Workload count. This update makes sure that reporting is accurate and improves data consistency for your security management processes
CLOUD-12589	Policy	The Enforcement Point API now correctly returns unique workload IDs with matching IP addresses, ensuring accurate policy data retrieval. This update prevents confusion caused by duplicate IDs and improves policy enforcement consistency.
CLOUD-12521	Policy	An issue prevented the policy impact screen from displaying data due to outdated enforcement point identifiers. The system environment was resynced to restore proper alignment, allowing policy impact to display correctly.
CLOUD-12435	User Interface	The Usage page now correctly displays billable workloads and resource data for authorized users. Users without proper access will no longer see inaccurate or zero values on the Usage page.
CLOUD-12234	Policy	Illumio resolved an issue where the impact display incorrectly showed unrelated network security rules for Azure policies. This fix makes sure that impact results correctly reflect actual changes, improving accuracy and confidence during policy management.
CLOUD-11413	Policy	An issue with IP addresses missing from certain inbound and outbound policy rules in segmented virtual environments is resolved. This correction makes sure that proper network access controls are applied, improving security policy enforcement.

What's new in this release - August 22, 2025

Feature	Description
Policy	Illumio Segmentation for the Cloud now supports security reviews, which ensure that users review policy enforcement on Azure subscriptions and tenants, and AWS accounts and organizations, reducing the risk of implementing ineffective rules. See Security reviews [267].
GCP Policy	Illumio Segmentation for the Cloud supports GCP policy, including the following: See Onboarding GCP [119]. See Writing application policy [255]. See Policy enforcement and resource types [262]. See Illumio visibility for resource types [193].
Policy	Illumio Segmentation for the Cloud now supports using the UI to easily enable read and write permissions for Illumio to enforce Azure subscription and AWS account policies after onboarding. See Enable read-write permissions [269]

What's new in this release - August 8, 2025

Feature	Description
OCI Policy	Illumio Segmentation for the Cloud supports OCI policy: See Onboarding OCI [127]. You must be a BETA participant to see these OCI features. Please reach out to your account team if you want to participate in the OCI BETA program. By participating in the BETA program for OCI features you agree that your company's use of the BETA version of OCI features will be governed by Illumio's Beta Terms and Conditions. Please refer to the feature documentation to understand the supported functionality and limitations related to OCI BETA.

What's new in this release - July 31, 2025

Feature	Description
GCP Visibility	Illumio Segmentation for the Cloud for GCP is generally available and supports visibility, including the following: • Onboarding. See Onboarding GCP [119]. • Flow Log Access using PubSub. See Grant flow log access to your CSPs [141]. • Inventory View. See Inventory [157]. • Traffic View. See Traffic [219]. • Map View. See Cloud Map navigation [171]. • Tag to Label Mapping. See Cloud Tag to Label Mapping [241]. • Application Discovery. See Define an application individually [235]. • Resource Visibility. See Illumio visibility for resource types [193]. Please refer to the feature documentation to understand the supported functionality and limitations.

What's new in this release - July 30, 2025

Feature	Description
Agentless Containers	Agentless Containers now supports: • OpenShift OVN (Open Virtual Networking)-Kubernetes • Azure AKS • Google GKE See Agentless Containers overview [148].

What's new in this release - July 14, 2025

Feature	Description
AI Labeling	Update: AI label-based recommendations draw from a pool of over 300 labels, including role and application labels. See Use AI Labeling [244].

What's new in this release - July 10, 2025

Feature	Description
GCP Visibility (BETA)	Illumio Segmentation for Cloud now supports GCP visibility under the Beta program. Visibility includes the following: • Onboarding. See Onboarding GCP [119]. • Flow Log Access using PubSub. See Grant flow log access to your CSPs [141]. • Inventory View. See Inventory [157]. • Traffic View. See Traffic [219]. • Map View. See Cloud Map navigation [171]. • Tag to Label Mapping. See Cloud Tag to Label Mapping [241]. • Application Discovery. See Define an application individually [235]. Please refer to the feature documentation to understand the supported functionality and limitations related to GCP Beta. You must be a Beta participant to see these GCP features. Please reach out to your account team if you want to participate in GCP Beta program. By participating in the BETA program for GCP features you agree that your company's use of the BETA version of GCP features will be governed by Illumio's Beta Terms and Conditions.
Onboarding	The flow log destination review process is streamlined with an improved interface. It provides more information and lets you sort destinations by different traffic directions with one click. You can assess your flow logs more efficiently before granting access. See Review destinations before granting flow log access [145].

What's new in this release - June 6, 2025

Feature	Description
AI Labeling	Update: AI label-based recommendations now draw from a pool of over 300 labels instead of 60. See Use AI Labeling [244].

What's new in this release - May 30, 2025

Feature	Description
Onboarding	The flow log destination review process is streamlined with an improved interface. It provides more information and lets you sort destinations by different traffic directions with one click. You can assess your flow logs more efficiently before granting access. Contact your Customer Success and Account teams to request access to this feature.

What's new in this release - May 13, 2025

Feature	Description
Azure Firewalls	This feature is no longer in beta and is now generally available. You can now get a clear view of your Azure Firewall inventory, with details about your firewalls and their current policies. Gain insights into your firewall network flows using the Traffic page, and see a visual representation of your firewall hub and spoke topology in the Map. Click on traffic paths in the Map to see traffic between firewalls and connections between VNets and clouds. In addition, you can now use Azure Firewalls to enforce policy on VNets. See Azure Firewalls Overview [167].

What's new in this release - May 8, 2025

Feature	Description
Resources	Illumio Segmentation for the Cloud now supports visibility for the following resources: • App Service (Web App, Function App) • SQL Managed Instance • Key Vault See Illumio visibility for resource types [193].

What's new in this release - May 7, 2025

Feature	Description
Policy	Illumio Segmentation for the Cloud now supports the following resources for enforcing policy on Azure Private Endpoints: • App Service (Web App, Function App) • SQL Managed Instance • Key Vault See Policy enforcement and resource types [262].

What's new in this release - April 24, 2025

No.	Feature Category	Feature List
1.	Agentless Containers	Illumio Segmentation for the Cloud now extends security to Kubernetes containers, delivering visibility and control for both cloud-managed and self-managed Kubernetes clusters. Use advanced filtering and dynamic mapping capabilities to navigate large-scale Kubernetes environments and pinpoint clusters, nodes, namespaces, and workloads. The onboarding process is streamlined to allow rapid integration of new Kubernetes clusters with Illumio Segmentation for the Cloud. Once you onboard your clusters, Illumio Segmentation for the Cloud provides insights into your containerized inventory, application dependencies, and network traffic patterns. These insights help you enforce security postures and micro-segmentation policies across dynamic workloads. See Agentless Containers overview [148]. See Onboard and Offboard Kubernetes Clusters [151]. See Navigating the Map Kubernetes View [184]. See Kubernetes Resources Inventory [165]. See Illumio IP addresses accessed by the Kubernetes Cloud Operator [344].

What's new in this release - March 31, 2025

No.	Feature Category	Feature List
1.	Policy	The Policy Preference page now lets you set your preferences for enforcement points to include Azure Firewalls, along with subnet and NIC NSGs. See Policy preferences [282]. Contact your Customer Success and Account teams to request access to this feature. **NOTE** Illumio Segmentation for the Cloud does not support Classic Azure Firewall.
2.	Resources	NOTICE: Illumio Segmentation for the Cloud no longer supports Azure network IP configurations as a visibility-supported resource type. However, the Inventory resource details tab for Microsoft network interfaces will still display network IP configurations as applicable. See Illumio visibility for resource types [193]. See Inventory Details [162].

What's new in this release - March 6, 2025

No.	Feature Category	Feature List
1.	Terraform for Applications	You can create a Terraform onboarding application and use it to onboard Azure subscriptions. Use Terraform to save time and effort instead of manually onboarding. See Create a Terraform Illumio Onboarding Application for Azure [75]. See Onboard an Azure Subscription using a Terraform Illumio Onboarding Application [78].
2.	Policy	You can write policy using Azure Firewalls. This allows you to define and apply network and application level security rules across multiple virtual networks and subscriptions. This also allows you to effectively manage all traffic filtering from a single point of control through its centralized policy management capabilities. See Writing Azure Firewall policy [259]. This is a beta feature. Review the beta terms and conditions on the Illumio website.

What's new in this release - February 20, 2025

No.	Feature Category	Feature List
1.	Terraform for Applications	You can now use Terraform to create, edit, and delete tag to label mappings, deployments, and applications. For applications, you can also use Terraform to auto-approve, allow cloud service source metadata, and populate a field with a list of deployments for the application definitions. Using Terraform saves you the time and effort of manually performing the work. See the Illumio Terraform website.
2.	Visibility	You can now get a clear view of your Azure Firewall inventory, with details about your firewalls and their current policies. Gain insights into your firewall network flows using the Traffic page, and see a visual representation of your firewall hub and spoke topology in the Map. Click on traffic paths in the Map to see traffic between firewalls and connections between VNets and clouds. In addition, you can now use Azure Firewalls to enforce policy on VNets. See Azure Firewalls Overview [167]. See Navigating Azure Firewalls [179]. See Inventory [157]. This is a beta feature. Review the beta terms and conditions on the Illumio website.
3.	Policy	You can now enforce policy using Azure Firewalls. See Illumio visibility for resource types [193]. See Policy enforcement and resource types [262]. This is a beta feature. Review the beta terms and conditions on the Illumio website.

What's new in this release - February 4, 2025

No.	Feature Category	Feature List
1.	Visualization	The Traffic page now has a time slider visualization that shows the total number of active flows for any time period you specify. You can zoom into a given time period and view detailed information about the flows such as flow status, the source and destination, and the times when the flows were detected. See Search traffic [220].

Resolved limitations in Illumio Segmentation for the Cloud

- **Update the Downloadable file "Download Permissions" file on CS UI Onboarding (C-8304)**

  Illumio Segmentation for the Cloud now provides both a read-only and a read and write permissions file that you can download during onboarding and from the documentation. Previously, customers who wanted to update their read-only permissions might have used the read and write permissions update script. You can now take the necessary steps to correct the permissions given by using the relevant scripts. See the Updating Permissions on the Assume Role section of Prerequisites for Onboarding AWS [94].

What's new in this release - January 23, 2025

No.	Feature Category	Feature List
1.	Visualization	Cloud now lets you drill down into EKS Clusters using the Resource Map tab. Use it to view your EKS Clusters' Node Groups and individual EC2 instances along with their traffic. See Navigating AKS, EKS, and GKE clusters [178]. Updated: You can now also enforce policy on EKS Clusters. See Policy enforcement and resource types [262].
2.	Documentation Updates	Download the combined User Guide and Release Notes PDF from the top navigation pane. 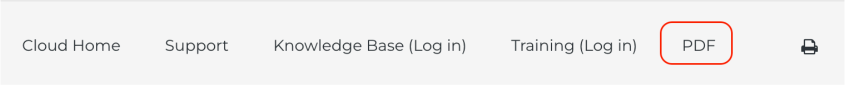 Cloud Home Support Knowledge Base (Log in) Training (Log in) PDF

Resolved limitations in Illumio Segmentation for the Cloud

- **Viewing Azure flow logs in storage accounts (C-8005)**
  Illumio Cloud now supports custom resource groups that generate VNET and NSG flow logs for Azure Network Watcher. This resolves a previous limitation where there was a mismatch in the folder structure generated by the flow service and the storage account.

What's new in this release - January 9, 2025

No.	Feature Category	Feature List
1.	Connector	Illumio Segmentation for the Cloud's Connector feature now lets you choose between CSV and JSON formats when exporting data to an onboarded S3 bucket. See Connector [274].

Previous Illumio Cloud what's new and release notes for 2024

These prior release notes describe the new features, enhancements, resolved limitations, and known limitations for Illumio Segmentation for the Cloud in previous 2024 releases.

Illumio Cloud is an agentless SaaS solution that provides visibility into your AWS and Azure network flows to define Zero Trust Segmentation policies in the public cloud, with the following features

- Multi-cloud coverage
- Fast breach containment
- Ease of use
- Low total cost of ownership

What's New in This Release - December 12th, 2024

No.	Feature Category	Feature List
1.	Policy	Illumio Segmentation for the Cloud supports the following Azure resource types for policy enforcement: • DBforPostgreSQL flexible server • DBforPostgreSQL server • Load balancer • Private endpoint • Private link service • SQL server See Policy enforcement and resource types [262].

What's New in This Release - November 25th, 2024

No.	Feature Category	Feature List
1.	Onboarding	Added new topic that describes how to update the service principal used by Illumio Cloud for accessing customer Azure Resources and how to rotate the secret for an existing client when the secret expires. See Update Service Principals for Onboarded Azure Subscriptions and Tenants [315].

What's New in This Release - November 21st, 2024

No.	Feature Category	Feature List
1.	Reports	Illumio Segmentation for the Cloud now lets you enable or disable report schedules. See Generated reports [212].
2.	Onboarding	The AWS account onboarding wizard now gives you the option to auto-discover your AWS CodeDeploy applications that already exist, and include them as Illumio Cloud applications. See Onboard an AWS Cloud account [113].

What's New in This Release - November 14th, 2024

No.	Feature Category	Feature List
1.	Policy	Illumio Segmentation for the Cloud now lets you optimize policy by setting your preferred enforcement points. See Policy preferences [282].
2.	Resources	Illumio Segmentation for the Cloud now supports Azure Agent Pool Machines and OCI Autonomous Databases. See Illumio visibility for resource types [193].

What's New in This Release - October 31st, 2024

No.	Feature Category	Feature List
1.	Administration	Illumio Segmentation for the Cloud now has a Terraform provider available, which provides first-class API support. See Illumio Terraform source [281].
2.	Inventory	The Inventory page now lets you search your resources by IP address. See Inventory [157].

What's New in This Release - October 28th, 2024

No.	Feature Category	Feature List
1.	Reports	The Reports page now lets users create ad hoc reports from existing schedules on demand. See Generated reports [212].
2.	Inventory	The Inventory page now displays policy sync status and policy last updated in the tool tip for SGs and NSGs in the Security Control column. This information also appears on the details for SGs and NSGs. See Inventory [157].

What's New in This Release - October 17th, 2024

No.	Feature Category	Feature List
1.	Reports	The Reports page now lets you export the following: • System Reports • Inventory Reports • Usage Reports • Application Reports • Application Definition Reports The page also now lets you filter Audit reports. See Generated reports [212].

What's New in This Release - October 10th, 2024

No.	Feature Category	Feature List
1.	Reports	The Reports page now allows users to edit existing scheduled reports. See Generated reports [212].
2.	Connector	The Connector Feature now lets you edit notification rules. See Connector [274].
3.	Policy	Two new features, Drift Detection and Tamper Protection, work in tandem to monitor and protect your cloud environment by checking for rules changes. Contact your Customer Success and Account teams to request access to this feature. See Drift Detection [271] and Tamper Protection [272].
4.	Policy	Cloud now checks for conflicting rules between organization policies and application policies. It also shows inherited aspects from other rules. Contact your Customer Success and Account teams to request access to this feature.

What's New in This Release - October 3rd, 2024

No.	Feature Category	Feature List
1.	Onboarding, Inventory, Traffic, Map, Application	Illumio Segmentation for the Cloud now provides visibility into Oracle Cloud Infrastructure (OCI) tenants. See Onboard an OCI tenant [129], Set up flow logs in your CSP environment [132], Grant flow log access to your CSPs [141], Prerequisites for onboarding OCI [127], OCI Flow Log Access IP Addresses [343], Inventory [157], Traffic [219], Map [168], and Deployments and Applications [223].
2.	Connector	The Connector feature now lets you connect Illumio Segmentation for the Cloud to your AWS S3 buckets so that you can export Illumio Segmentation for the Cloud traffic flows to them. See Connector [274].

What's New in This Release - September 26th, 2024

No.	Feature Category	Feature List
1.	Labeling	The AI labeling feature has been moved to the Labels page and updated for improved usability. The AI labeling feature is available for select US cloud resources only. Contact your Customer Success and Account teams to request access to this feature. See Use AI Labeling [244].

What's New in This Release - September 19th, 2024

No.	Feature Category	Feature List
1.	Flow Logs	Illumio Segmentation for the Cloud now includes new data planes with IP addresses for you to allow for flow log access, which may have an impact on customers with Azure tenants. See Prerequisites for Onboarding AWS [94] and Prerequisites for Onboarding Azure [60].

What's New in This Release - September 12th, 2024

No.	Feature Category	Feature List
1.	Visualization	The Usage page now lets you export reports. See Usage [211].
2.	Visualization	The Reports page now lets you configure exported traffic reports to match conditions. See Connector [274].
3.	Visualization	Illumio Segmentation for the Cloud now lets you schedule report generation in advance, either once or at regular intervals. See Generated reports [212].

What's New in This Release - September 5th, 2024

No.	Feature Category	Feature List
1.	Applications	The Applications page now lets you export reports. See View and approve an application [238].
2.	Administration	The Connector page automation feature for Slack now allows you to test the new trigger rule before saving. See Connector [274].

What's New in This Release - August 29th, 2024

No.	Feature Category	Feature List
1.	Onboarding	You no longer need to manually configure Illumio Segmentation for the Cloud to fetch flow logs from custom S3 paths, as Illumio now automatically accommodates fetching flow logs stored there. The manual configuration steps have been removed from the documentation portal.
2.	Applications	The Application Definitions tab now lets you export reports. See Define an application automatically [231].
3.	Visualization	Cloud now retains system events for 30 days. See Events [192].

Erratum

The previous release notes erroneously stated that report scheduling was available as a feature. The statement has been removed from the previous release notes.

What's New in This Release - August 21st, 2024

No.	Feature Category	Feature List
1.	Administration	The Connector page now lets you automate notifications based on triggers you select. See Connector [274].

What's New in This Release - August 15th, 2024

No.	Feature Category	Feature List
1.	Administration	The Illumio Virtual Advisor (IVA) is an AI chatbot that helps organizations understand and reduce their risk posture by using natural language questions to generate quick answers and actions.
2.	Visualization	These resources are now visible on the Inventory page: • AWS VPN Connection • Azure Virtual Hub IP Configuration

What's New in This Release - August 8th, 2024

No.	Feature Category	Feature List
1.	Visualization	The Traffic page now supports: • Linking directly to the Reports page to download a traffic list report as soon as you create it • Filtering by known networks See Reports [212], Traffic, [219] and Search Traffic [220].
2	Visualization	The Events page now lets you export system event reports. See Events [192].
3.	Visualization	These AWS resources are now visible on the Cloud Map page: • DocumentDB DB Cluster • DocumentDB Instance

What's New in This Release - August 2nd, 2024

No.	Feature Category	Feature List
1.	Visualization	The Insights page now includes 24 insight tiles, so you can use out-of-the-box queries to gain quick insights into resources within minutes of onboarding them: • Cross Cloud Traffic • Cross Region Traffic • Cross Tenant Traffic • Cross Tenant Traffic • Account to Malicious External IP Traffic • Account to Unknown External IP Traffic • Account to Known External IP Traffic • Tenant to Malicious External IP Traffic • Tenant to Unknown External IP Traffic • Tenant to Known External IP Traffic • Region to Malicious External IP Traffic • Region to Unknown External IP Traffic • Region to Known External IP Traffic • Azure Tenant to AWS IP Traffic • Account to External Cloud Traffic • Tenant to External Cloud Traffic • Region to External Cloud Traffic • Account to External Geo Traffic • Tenant to External Geo Traffic • Region to External Geo Traffic • Cross Talking Peering Connections • Internet Exposed EC2 Instances • Unprotected Resources • Traffic Blind Spots
2	Visualization	The Insights page now provides descriptions when you hover over the insight tiles.

What's New in This Release - August 1st, 2024

No.	Feature Category	Feature List
1.	Visualization	These Azure resources are now visible on the Inventory page: • Virtual Hub • Virtual Hub Connection
2.	Visualization	The Inventory page now lets you export reports. See Inventory [157].
3.	Visualization	The Reports page now lets you edit reports. See Reports [212].
4.	Visualization	The Traffic and Cloud Map pages now lets you filter IP addresses by CIDR blocks. See Search Traffic [220] and Cloud Map [168].
5.	Visualization	Illumio Segmentation for the Cloud now removes flow information more than 90 days old. See Traffic [219].
6.	Visualization	The Traffic page now lets you refresh your filter results to clear stale data without refreshing the browser page. See Search Traffic [220].
7.	Visualization	The Usage page now has updated terminology for the data displayed See Product Usage [211].
8.	Applications	Illumio Segmentation for the Cloud now has an Application Summary tab in the application details panel. See View and Approve an Application [238].
9.	Onboarding	Illumio Segmentation for the Cloud now has a Service Accounts page for adding and deleting service accounts and their secrets. See Service Accounts [280].

What's New in This Release - July 25th, 2024

No.	Feature Category	Feature List
1.	Visualization	This Azure resource is now visible on the Map page: • Redis Cache
2.	Visualization	These Azure resources are now visible on the Inventory page: • Virtual Network Gateway • Virtual Network Gateway connection
3.	Visualization	The Events page now performs cleanup after events become seven days old. See Events [192].
4.	Visualization	The Reports page now lets you delete reports in bulk. See Reports [212].

What's New in This Release - July 18th, 2024

No.	Feature Category	Feature List
1.	Administration	The Show Impact filter now lets you filter by network access control lists. See Writing Application Policy [255].
2.	Visualization	These AWS resources are now visible on the Inventory page: • Transit Gateway • Transit Gateway Attachment • Transit Gateway Route Table • Transit Gateway Multicast Domain • DocumentDB Cluster • DocumentDB Instance • Document DB Elastic Cluster • VPC Endpoint Service

What's New in This Release - July 11th, 2024

No.	Feature Category	Feature List
1.	Administration	Role-based access control (RBAC) is now available. See Role-Based Access Control [279].
2.	Visualization	The Traffic and Application Traffic basic filter has been replaced by the advanced filter. See Search Traffic [220].
3.	Visualization	The Traffic page now has an export option. See Traffic [219].
4.	Visualization	These Azure resources are now visible on the Inventory page: • Firewall Policy • Rule Collection Group • Diagnostic Setting
5.	Visualization	The Reports page now lets you generate Risk reports. It also now lets you delete reports. See Reports [212].
6.	Visualization	The Events page now has a System Events Tab, which lets you view system-generated events. See Events [192].

What's New in This Release - June 28th, 2024

No.	Feature Category	Feature List
1.	Visualization	The Reports page gives you the ability to generate event audit reports asynchronously. You can then download and share the reports. See Reports [212].
2.	Visualization	Cloud's Inventory page provides the following new resource properties for the AWS resources listed: • Security Group Rule: The properties for his resource now also provide a table of security group rules. • ElasticLoadBalancingV2 Load Balancer: The properties for this resource now also provide schemes. • RAM Resource Share: The properties for this resource now also provide resource owner ID and resource type. • Route Table: The properties for this resource now also provide a list of associated rules. • VPC Peering: The name and ID for this resource now appear as a hoverable resource tile instead of static text. • Shared VPC and Shared Subnets: The properties for this resource now also provide owner IDs and shared status. • All resources: The properties for all resources now also provide resource group information if it exists.
3.	Visualization	Cloud's Inventory page provides the following new information or presentation: • Route table relationships • Category and region information in attached resource tables • Resource names and IDs now appear as a hoverable resource tile instead of static text • Account IDs now appear as a hoverable resource tile instead of static text
4.	Visualization	These Azure resources are now visible on the Inventory page: • Redis Cache • Private Links
5.	Visualization	The Inventory, Application Inventory, and Cloud Map filters now let you search for resource groups and resource names. See Cloud Search [190].
6.	Flow Log Access	You can now manually add permissions to the Cloud role so that it can fetch flow logs that you may have stored in custom S3 bucket directories. The topic link has been removed as Cloud has since begun automatically adding such permissions.

What's New in This Release - June 21st, 2024

No.	Feature Category	Feature List
1.	Visualization	The advanced traffic filter now also appears on the Applications page Traffic tab. Use it to search traffic to and from the selected application. See Search Traffic [220].

What's New in This Release - June 18th, 2024

Note: If you are on a single product offering and are interested in learning more about the new Platform offering, contact Illumio Customer Success.

No.	Feature Category	Feature List
1.	Visualization	With Illumio Segmentation for the Cloud Insights, you can use out-of-the-box queries to gain quick insights into resources within minutes of onboarding them, including: • Networks where flows logs are not enabled • EC2 instances that are directly reachable from the Internet • Cross-account, cross-region communications enabled with peered network connections

What's New in This Release - June 13th, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Flow Log Access	Illumio Segmentation for the Cloud now lets you filter Flow Log Access tables in the following places: • Flow Log Access page • Flow Log Access Details By Log Destination Account tab • Flow Log Access Details By Log Source Account tab See Grant Flow Log Access [141].
2.	Visualization	• Illumio Segmentation for the Cloud now lets you add additional search terms without having to delete existing terms in the Traffic page Advanced Filter. See Search Traffic [220]. • This resource now appears on the Inventory Details page as an attached resource for VPCs and subnets: • AWS: • RAM ResourceShare
3.	Policy	• Illumio Segmentation for the Cloud now lets you write single allow rules for multiple ports for AWS security groups so long as there is no deny rule prohibiting the allow rule. • Illumio Segmentation for the Cloud now lets you create a single rule for multiple IPs if they belong to a CIDR. • Policies now let you select 'All Resources' instead of 'All Workloads.' See Writing application policy [255]Writing Application Policy.

Note: If you are on a single product offering and are interested in learning more about the new Platform offering, contact Illumio Customer Success.

No.	Feature Category	Feature List
1.	Visualization	Illumio now lets you achieve unified visibility with the Map: • You can view the traffic between resources • You can right-click on a resource to write policy • You can distinguish between AWS, Azure, and server datacenter types • You can view the cloud metadata in search filters like Account ID, Region, Resource Type, and more See Map [203].
2.	Policy	Policy can now be authored and enforced for all datacenters and cloud workloads. 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. See Unified Policy [265].
3.	Administration	The Illumio Virtual Advisor (IVA) is an AI chatbot that helps organizations understand and reduce their risk posture by using natural language questions to generate quick answers and actions.
4.	Labeling	• Use AI labeling to label assets based on metadata and flow logs to make sure you have accurate and consistent labeling. This method speeds up deployments and ensures consistent enforcement. See Use AI Labeling [244]. • Rule-based labeling allows you to assign labels to one or more workloads when their attributes match conditions that you specify in easily-configurable rules. You can perform the following tasks with this feature: • You can create a basic rule to match workloads running on a specific operating system • You can create a rule with multiple values to match workloads with a hostname containing any of the entered values (up to 20) • You can create an IP Address rule to match workloads within an IP address range • You can create a CIDR block rule to match workloads within a CIDR block See Rule-Based Labeling.

What's New in This Release - June 6th, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Flow Log Access	Cloud now gives destination-based view to grant flow log access. The new view provides a list of flow log destinations that are used for storing flow logs on a per-account basis. You can also see a list of log sources sent from different accounts. For centralized flow logs, you can now grant access to the log archive account's destination so that Cloud can read and process the logs. See Grant Flow Log Access [141].
2.	Visualization	This resource is now visible on the Inventory page: AWS: • RAM ResourceShare

Resolved Limitations in Illumio Segmentation for the Cloud

- **[Policy Services UI] Do not highlight Delete button with a resource when you create a new Service** (C-3944)

  When provisioning new services, users saw the Remove button automatically gain focus with a numeral '1.' Clicking Remove deleted the new service.

What's New in This Release - May 30th, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Flow Log Access	Cloud now lets you access Azure VNet flow logs. See Grant Flow Log Access [141].
2.	Policy	Organization policies now let you select 'All Workloads' that allow you to write organization policies for all resources in onboarded cloud accounts. See Writing Application Policy [255].

Resolved Limitations in Illumio Segmentation for the Cloud

- **[Policy Services UI] Do not highlight Delete button with a resource when you create a new Service** (C-3944)

  When provisioning new services, users saw the Remove button automatically gain focus with a numeral '1.' Clicking Remove deleted the new service.
- **Error shown when users attempt to add an existing user to their account** (C-3083)

  When a user tried to add existing users to their existing Cloud account, Cloud correctly prevented the action, but did not issue an error message. For example, if a customer had one live Cloud account and also one trial account, trying to add an existing trial user to the live account silently failed.

What's New in This Release - May 23rd, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Visualization	The Dashboard now lets you ingested resources at a glance. See Cloud Dashboard [188].
2.	Labeling	The Label Mapping page now lets you view a list of the following system-generated labels at a glance: • ServiceCategory describes resources by their categories • ServiceRole describes resources according to their roles See View System Labels [243].

Resolved Limitations in Illumio Segmentation for the Cloud

- **Reselecting custom traffic filter will reset the time span** (C-1978)

  When users adjusted the time filter after searching for flows in a given time span, the filter reset to the previous day.

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)

  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)

  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Middle, right, or control click to open in new tab do not work** (C-2398)

  Middle click, right click, and control click sometimes do not open the specific desired Cloud tab.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)

  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)

  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - May 16th, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Visualization	The Traffic page now lets you view traffic flow source and destination details. See Traffic [219].

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)

  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)

  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.

- **Middle, right, or control click to open in new tab do not work** (C-2398)
  Middle click, right click, and control click sometimes do not open the specific desired Cloud tab.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - May 13th, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Visualization	The *Inventory* page now has a Service Category filter for searching your inventory of resources. See Cloud Search [190].

Resolved Limitations in Illumio Segmentation for the Cloud

- **Map not matching Azure VM topology** (C-2720)
  Sometimes the Cloud Map showed VMs as the child of a location instead of a subnet. The subnet was found, but the VM did not show up in the subnet.

What's New in This Release - May 2nd, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Labeling	You can now use the following system-generated labels: • ServiceCategory describes resources by their categories • ServiceRole describes resources according to their roles See Labels [250].
2.	Visualization	The Traffic page Beta Advanced Filter now lets you search by VPC, subnet, and resource type. See Search Traffic [220].
3.	Policy	The following resources now support policy: • AWS: • Redshift Clusters • RDS DB Instances • ElastiCache CacheClusters • Lambda Functions • Azure: • Virtual Machine ScaleSets

Resolved Limitations in Illumio Segmentation for the Cloud

- **Empty page should have string called "No integrations"** (C-983)
  When the Onboarding page was empty, there was no text string. If the page lacks data, it now says "No data to display."

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)
  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)
  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Middle, right, or control click to open in a new tab do not work** (C-2398)
  Middle click, right click, and control click sometimes do not open the specific desired Cloud tab.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - April 25th, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Policy	The following resources now support policy: • AWS: • ElasticLoadBalancingV2 Load Balancer • Azure: • SQL Server
2.	Flow Log Access	You can now test your accounts' flow log access. See Grant Flow Log Access [141].
3.	Visualization	• The Usage page now lets you select a custom time range, going back to day zero. See Product Usage [211]. • The Audit Events page now shows up to 10,000 results. See Audit Events [192].

Resolved Limitations in Illumio Segmentation for the Cloud

- **Missing Feature: Day 0 Map, Inventory, and Traffic Views** (C-2913)
  The Day 0 Cloud Map and Traffic pages did not show the Add Cloud Banner. It instead gave a "no resource/traffic available" message.

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)
  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)
  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Middle, right, or control click to open in a new tab do not work** (C-2398)
  Middle click, right click, and control click sometimes do not open the specific desired Cloud tab.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - April 18th, 2024

The following new features are available in this release:

No.	Feature Category	Feature List
1.	Visualization	• Cloud, Region, and Account are now available as filterable categories in the Traffic page Beta Advanced Filter mode. See Traffic [219]. • This resource is now visible on the Inventory page, and appears on the Cloud Map page as an attached resource for EC2 Instances and ElasticLoadBalancingV2 Load Balancers: • AWS: • ElasticLoadBalancingV2 Target Group
2.	Deployments	You can now edit your deployments. See Define a Deployment [228].

Resolved Limitations in Illumio Segmentation for the Cloud

- **Traffic flow filter by status not working as expected** (C-3566, C-3686)

  Users navigating the Cloud Map sometimes also saw denied traffic included on a node Details page despite filtering for allowed traffic.
- **Error onboarding Azure Flow Logs** (C-2890)

  Users would sometimes get an error when onboarding Azure flow logs due to Cloud not understanding that flow log destination access was already granted.
- **No description in Azure "Forbidden" onboarding message** (C-2023)

  When encountering an Azure onboarding error message, users did not get sufficient information to readily resolve the problem.

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS may not have ENI** (C-3265)

  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)

  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Middle, right, or control click to open in new tab do not work** (C-2398)

  Middle click, right click, and control click sometimes do not open the specific desired Cloud tab.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)

  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)

  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - April 11th, 2024

The following new features are available in the April 11th, 2024 release:

No.	Feature Category	Feature List
1.	Visualization	• The Audit Events page shows you a running list of different events in your environment such as onboarding, policy, labeling and user actions. See Audit Events [192] for information. • The Inventory page is no longer limited in the number of resources it can display. See Inventory [157] for information.
2.	Applications	You can now bulk-delete applications that were created using Application Discovery Rules. See Define an Application [231] for information.
3.	Labels	The Tag to Label Mapping page now shows both the Illumio label type and the labels to which you have mapped your CSP cloud tag keys. See Cloud Tag to Label Mapping [241] for information.

Resolved Limitations in Illumio Segmentation for the Cloud

- **Map is empty when no regions returned in top down view** (C-2982)

  When users filtered the Cloud Map in a way that excluded regions, it would appear empty. This limitation is resolved.

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)

  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)

  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)

  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)

  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - April 4th, 2024

The following new features are available in the April 4th, 2024 release:

No.	Feature Category	Feature List
1.	Visualization	• These resources are now visible on the Cloud Map page: • AWS: • VPC Endpoints • Azure: • Private Endpoints • You can now see colored traffic lines on the Cloud Map page indicating allowed (green), denied (red), and mixed (orange) See the Cloud Map [168] documentation on the portal. • The Inventory Details page now shows inbound and outbound rules for AWS Network ACLs • The resources documentation now contains a Category column. See the Inventory [157], Cloud Map [168], and Traffic [219] documentation on the portal. • You can now see the IP addresses of certain types of resources in the Inventory Details page and the Map page Details pane. Such resources include Redshift Clusters and Load Balancers. See the Inventory [157] and Cloud Map [168] documentation on the portal. • You can now filter your Traffic page searches with labels. See the Traffic [219] documentation on the portal.

Resolved Limitations in Illumio Segmentation for the Cloud

- **Cloud Map is only showing some VNET peering links** (C-3428)

  Sometimes the Inventory page showed additional peers that did not show up on the Map page. This limitation is resolved.
- **Security group names not showing up in console** (C-1875)

  Discovered EC-2 instances did not show security group names. This limitation is resolved.
- **AWS Security Group Rules not rendered on UI** (C-3466)

  The Inventory detail page displayed security group details, but the rules were missing. This limitation is resolved.

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)

  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)

  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)

  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)

  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - March 28th, 2024

The following new features are available in the March 28th, 2024 release:

1.	Visualization	• The Inventory page Resource Graph tab now lets you view a graphical representation of resource that you select. The graph contains the following: • Your selected resource • An inner ring around your selected resource, depicting each of its attached resources • An outer ring, depicting the individual instances of the attached resources shown in the inner ring • A series of incoming flow lines from the left, depicting sources for which your selected resource is the destination • A series of outgoing flow lines to the right, depicting destinations for which your selected resource is the source See the Inventory [157] documentation on the portal.
2.	Applications	Cloud now lets you bulk delete application definitions. See the Define an Application [231] documentation on the portal.

Resolved Limitations in Illumio Segmentation for the Cloud

- **Azure NAT Gateway not showing up in Cloud Map** (C-3427)

  Azure NAT gateways appeared on the Inventory page but did not show up on the Cloud Map page. This limitation is resolved.
- **Allow multiple rules with empty prefix** (C-3339)

  There was previously a constraint enforced where two rules could not have the same prefix, even if the prefix were left blank. This limitation is resolved.

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)

  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)

  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)

  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)

  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - March 21st, 2024

The following new features are available in the March 21st, 2024 release:

No.	Feature Category	Feature List
1.	Visualization	- The Usage page now lets you choose the graph style and includes the following additional workload hour categories: - Container Hosts - Serverless Containers - Serverless Functions See the Product Usage [211] documentation on the portal. - You can now filter your searches with operators (AND, OR, =, !=, etc.) - See the Cloud Search [190] documentation on the portal. - These resources are now visible on the Inventory page: - Azure: - Azure NAT Gateway - Azure publicIPAddress - Network Security Groups Default Security Rule For a full list of all supported resources visible on the Inventory page, see the Inventory documentation on the portal. - These resources are now visible on the Cloud Map page: - Azure: - Azure NAT Gateway (Azure public IP prefixes will appear on the Details panel) - For a full list of all supported resources visible on the Cloud Map page, and the VPC/VNet peering described below, see the Cloud Map [168] documentation on the portal - You can now view VPC/VNet peering in detail on the Cloud Map page.
2.	Flows	The Risk Report feature on the Traffic page now lets you toggle which details you wish to include. See the Traffic [219] documentation on the portal.
3.	Traffic	The filter now lets you use the Beta Advanced Filter mode, which lets you use joiners and operators while searching for sources, destinations, categories, etc. See the Traffic [219] documentation on the portal.
4.	Onboarding	- The AWS onboarding process now lets you download a text file containing the permissions indicated by the read/write toggle - The AWS account onboarding process now lets you see the Illumio Segmentation for the Cloud ID you will need if you share CloudFormation stacks. See the Onboard an AWS Cloud Account [113] documentation on the portal. - The Azure onboarding process is now more streamlined, so that you no longer need to manually enter client IDs and secrets - For Azure, Illumio Segmentation for the Cloud can read now flow logs from several NSGs going to the same storage account. See the Onboard an Azure Cloud Tenant [79] and Onboard an Azure Cloud Subscription [65] documentation on the portal.
5.	Applications	Illumio Segmentation for the Cloud now lets you automatically approve application definitions in two places. The Application Definition page lets you toggle whether you want Illumio Segmentation for the Cloud to automatically approve all discovered applicable deployments and resources. Similarly, the Application Discovery Rule page lets you toggle whether you want Illumio Segmentation for the Cloud to automatically approve all discovered application definitions, as well as any updates made to their deployments and resources. See the Define an Application [231] documentation on the portal. Either of these methods will skip the manual approval process for those applications.

Resolved Limitations in Illumio Segmentation for the Cloud

- **Editing discovery rules inserts extra dash (-) automatically** (C-3337)
  When modifying discovery rules, an extra dash was added automatically to the prefix. This limitation is resolved.
- **Deleting T2L mapping does not delete label dimension** (C-2646)

When users deleted a tag to label mapping, any labels that were assigned to resources using that mapping were not removed. Deleting the mapping kept those mapped labels on the resources, resulting in the label never being deleted. This limitation is resolved.

Known Limitations in Illumio Segmentation for the Cloud

- **AWS PaaS resources may not have ENI** (C-3265)
  Illumio Segmentation for the Cloud uses DNS lookup on the fully qualified domain name to get the elastic network interface relationships, which is not guaranteed to get a match. The potentially affected AWS resources are RDS DBInstances, RDS DBClusters, ElasticLoadBalancingV2 load balancers, MemoryDB clusters, ElastiCache for Redis clusters, and Redshift clusters.
- **Error shown when users attempt to add an existing user to their account** (C-3083)
  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - March 11th, 2024

The following new features are available in the March 11th, 2024 release:

No.	Feature Category	Feature List
1.	Onboarding	You can now onboard Azure tenants in addition to individual subscriptions as before. See the Onboard an Azure Cloud tenant - default setup [79] documentation on the portal.
2.	Visualization	These resources are now visible on the Inventory page: AWS: • VPC Azure: • VNet For a full list of all supported resources visible on the Inventory page, see the Inventory [157] documentation on the portal.
3.	Policy	• You can now preview a policy's impact before provisioning it • This resource is now available for policy enforcement: • AWS RDS DB clusters See the Writing application policy [255] documentation on the portal.
4.	Applications	You can now approve application deployments and resources in bulk on the application definition page. See the View and approve an application [238] documentation on the portal.

Resolved Limitations in Illumio Segmentation for the Cloud

- **Slice bug on Flow Log Access page** (C-3080)
  A conditional check was missing for sliced items. Therefore, users might have gotten a blank screen. This limitation is resolved.
- **406 errors should be displayed when deleting tag to label mappings** (C-3217)
  When users deleted a tag to label mapping, any errors returned by the delete response were not shown in the UI. This limitation is resolved.
- **Application has 0 resources, but the map is rendering resources** (C-3041)
  When users selected an application on the Cloud Map, the map would sometimes indicate resources despite there not being any. This limitation is resolved.
- **Go button does not refresh data unless filters change** (C-2296)
  When users executed a query on the Traffic, Inventory, or Cloud Map pages, the Go button did not re-run the same query on fresh data. To re-run the same query, users had to change the filter and change it back again before re-running the query. This limitation is resolved.
- **Avoid label create/delete race conditions** (C-2957)
  When users deleted and re-created an application or deployment in quick succession, Cloud sometimes deleted the label that was re-used by the re-created app/deployment. Users ended up with an application or deployment linked to a deleted label. This limitation is resolved.
- **Events in Cloud UI should show the latest events at the top** (C-2946)
  The Events page would show the oldest events at the top rather than at the bottom. This limitation is resolved.
- **Editing Azure subscription integrations showed child account list** (C-2920)
  When users edited their Azure subscriptions, the user's child accounts were mistakenly listed. This limitation is resolved.

Known Limitations in Illumio Segmentation for the Cloud

- **Error shown when users attempt to add an existing user to their account** (C-3083)
  When a user tries to add existing users to their existing Cloud account, Cloud correctly prevents the action, but does not issue an error message. For example, if a customer has one live Cloud account and also one trial account, trying to add an existing trial user to the live account will silently fail.
- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - February 29th, 2024

The following new features are available in the February 29th, 2024 release:

No.	Feature Category	Feature List
1.	Visualization	• The new Usage feature displays workload hours and flow log storage usage. • These resources are now visible on the Cloud Map page: • AWS: • DynamoDB tables • Lambda For a full list of all supported resources visible on the Cloud Map page, see the Cloud Map documentation on the Illumio documentation portal. • This resource is now visible on the Inventory page: • AWS: • Lambda For a full list of all supported resources visible on the Inventory page, see the Inventory documentation on the Illumio documentation portal.

Resolved Limitations in Illumio Segmentation for the Cloud

- **App approval status filters do not show correct results** (C-2945)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.

Known Limitations in Illumio Segmentation for the Cloud

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.

- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - February 22nd, 2024

The following new features are available in the February 22nd, 2024 release:

No.	Feature Category	Feature List
1.	Application Definition	Application Discovery Rules now allow full editing.
2.	Policy	Allow rules are now available for organization policies.
3.	Visualization	• The new Usage feature displays workload hours and flow log storage usage. • These types of resources are now visible on the Cloud Map page: • AWS: • Redshift clusters • Azure: • Microsoft.Web/sites • Microsoft.Web/sites/functions For a full list of all supported resources visible on the Cloud Map page, see the Cloud Map documentation on the Illumio documentation portal. • These types of resources are now visible on the Inventory page: • AWS: • Redshift clusters • DynamoDB tables • Azure • Microsoft.Web/sites • Microsoft.Web/sites/functions For a full list of all supported resources visible on the Inventory page, see the Inventory documentation on the Illumio documentation portal.

Resolved Limitations in Illumio Segmentation for the Cloud

- **Tag to label mapping must be defined before an app is defined** (C-2997)
  User did not have the ability to write policies on labels created using tag to label mapping if those labels were not associated with any application. This limitation is resolved.
- **Editing proxy username is not supported** (E-113332)
  Cloud did not support updating the username. Due to this limitation, name editing was disabled in existing tenants and all the new users added to existing tenants. The edit user function in the User detail page and the My Profile page were disabled. For new tenants and users in new tenants, editing the user is now supported. This limitation is resolved.
- **Traffic doesn't show labeled workloads** (C-2559)
  When users went to the Traffic tab, flows sometimes erroneously lacked labels. When users searched for labeled traffic flows, sometimes no results were returned. This limitation is resolved.

Known Limitations in Illumio Segmentation for the Cloud

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.

- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - February 15th, 2024

The following new features are available in the February 15th, 2024 release:

No.	Feature Category	Feature List
1.	Visualization	Public IPs are supported for Azure VM flows but not policies.

Resolved Limitations in Illumio Segmentation for the Cloud

- **The username is incorrectly displayed on the main page and within the user grid** (C-2897)
  User's names displayed incorrectly after being added. This limitation is resolved.
- **Resources not shown for pending approval apps** (C-2887)
  When creating applications either individually or using a discovery rule, resources were not visible on the Application Definition page resources link while the applications were pending. This limitation is resolved.
- **UI must validate application deployment inputs** (C-2797)
  Users were allowed to add deployment types without any values. If a user did not enter any values, a UI page crash occurred and/or the backend rejected the request. The UI now disables the Add button when no values are selected. This limitation is resolved.
- **Tried to onboard an AWS account previously onboarded and offboarded, getting errors in cloudformation template creation** (C-2715)
  Offboarding AWS accounts did not completely remove the stack. Workaround: Follow the Remove the Integration instructions on the Illumio documentation portal.

Known Limitations in Illumio Segmentation for the Cloud

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.

- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - February 8th, 2024

The following new features are available in the February 8th, 2024 release:

No.	Feature Category	Feature List
1.	Visualization	• These types of resources are now visible on the Cloud Map page: • AWS: • S3 bucket VPC endpoints in will appear in the Detail panel, but multiple VPC endpoints for a single S3 bucket are not supported • ElasticLoadBalancingV2 Load balancer • Azure: • Azure spot VMs • VM scale sets • Storage account private endpoints will appear in the Detail panel as attachments • Azure postgreSQL • Microsoft.DBforPostgreSQL/serverGroupsv2 • Microsoft.DBforPostgreSQL/flexibleServers (databases will appear in the Detail panel as attachments) • Microsoft.DBforPostgreSQL/servers (databases will appear in the Detail panel as attachments) • Microsoft.DocumentDB/cassandraClusters • Microsoft.DocumentDB/mongoClusters • (databases will appear in the Detail panel as Azure SQL servers attachments) For a full list of all supported resources visible on the Cloud Map page, see the Cloud Map documentation on the Illumio documentation portal. • These types of resources are now visible on the Inventory page: • Azure • Microsoft.DBforPostgreSQL/flexibleServers/databases • Microsoft.DBforPostgreSQL/servers/databases For a full list of all supported resources visible on the Inventory page, see the Inventory documentation on the Illumio documentation portal. • In the Inventory page you will see two additional tabs: Inbound Rules and Outbound Rules. These tabs appear in your AWS Security Groups' and Azure Network Security Group's Detail panels as attachments.
2.	Onboarding	You can now onboard AWS organizations in addition to individual accounts as before.
3.	Applications	Although Cloud 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 will automatically define applications when new resources are added that meet the rule parameters. You can also now use accounts, in addition to cloud tags or virtual networks and subnets, to define your applications.

Resolved Limitations in Illumio Segmentation for the Cloud

- **NSGs attached to subnet is not included in vm > nsg relationship** (C-2594)

  Cloud was programming only network security groups associated with a NIC. This limitation is resolved. Now Cloud will program both network security groups associated with a subnet and those associated with a NIC.
- **Label search within an application shows resources that do not belong to the application** (C-2568)

  A label search within an application showed all resources instead of showing the resources for only the selected application. This limitation is resolved.
- **Dashboard Traffic Summary tile forgets user's previous filter selection** (C-2387)

  When users filtered by a specific CSP and a specific timeframe, and went away from the Dashboard page, the Traffic Summary tile would reset to the 24-hour default, with all CSPs selected. This limitation is resolved.

Known Limitations in Illumio Segmentation for the Cloud

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - February 1st, 2024

The following new features are available in the February 1st, 2024 release:

No.	Feature Category	Feature List
1.	Visualization	• These types of resources are now visible on the Cloud Map and Inventory pages: • AWS: • RDS DB clusters • RDS DB instances • EC2 VPCs, subnets, NAT gateways, Internet gateways, spot fleet requests and spot instance requests • ECS clusters • ECS container instances • Glacier vaults • ElastiCache clusters • MemoryDB clusters • Azure: • Virtual networks and their subnets • Storage accounts • Application gateways • Load balancers • Azure firewalls • Virtual network gateways • VPN gateways • NAT gateways • DocumentDB database accounts • Additional types of resources are visible on the Traffic page: • AWS • ENIs • Azure • Network interfaces
2.	Flows	The Risk Report feature on the Traffic page lets you generate 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
3.	Onboarding	When onboarding CSP accounts or subscriptions, you can now select read-only access.

Known Limitations in Illumio Segmentation for the Cloud

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.

- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

Previous Illumio Cloud what's new and release notes for 2023

These prior release notes describe the new features, enhancements, resolved limitations, and known limitations for Illumio Segmentation for the Cloud in previous 2023 releases.

Illumio Cloud is an agentless SaaS solution that provides visibility into your AWS and Azure network flows to define Zero Trust Segmentation policies in the public cloud, with the following features:

- Multi-cloud coverage
- Fast breach containment
- Ease of use
- Low total cost of ownership

What's New in This Release - December 14th, 2023

The following new features are available in the December 14th, 2023 release:

No.	Feature Category	Feature List
1.	Application	• Cloud automatically resynchronizes pending applications with any resource changes. This means you can add or drop a cloud tag in such a way that it applies to an additional resource, and Cloud will automatically re-synchronize the application to include the new resource. • Cloud lets you edit application definitions
2.	Policy	• Cloud allows users to specify all services for the destination service when writing rules • Cloud allows users to specify 'Any' (0.0.0.0/0 and ::/0) for the source or destination IP address when writing rules
3.	Visualization	Cloud now features a tiled dashboard, which displays Traffic statistics
4.	Inventory	Cloud now has context-based search for inventory-based filters, such as the ones on these pages: • Inventory list page • Traffic list page • Map page • Application inventory page • Application traffic page

Resolved Issues in Illumio Segmentation for the Cloud

- **Inventory returning deleted resources when filtering by tags** (C-2250)
  When users queried inventory by tags, the query returned deleted resources that had one of the tags assigned to it. This meant that it would appear in the Inventory list. When

querying inventory by that deleted resource's `resource_id,` the response was empty as expected. This issue is resolved.
- **sgpolicyenforcement svc fails to enforce all-svc rules in AWS SG** (C-2232)
  An Illumio AWS security group service encountered errors when attempting to enforce rules covering all services. This issue is resolved.
- **Getting logged out after login** (C-2198)
  Illumio was not properly pruning sessions, which occasionally resulted in users getting logged out involuntarily. This issue is resolved.
- **Running the CloudFormation Template (CFT) in AWS does not work** (C-2255, C-2079)
  Illumio created the flow stack when users ran the template, but the link did not open in a new window. If the user clicked "Download," nothing would happen. This issue is resolved.
- **Application used in Ruleset was allowed to be deleted** (C-1607)
  Users were allowed to delete applications even if they were part of a ruleset. This issue is resolved. A message now displays telling the user that the application is currently used by ruleset, and blocks application removal.
- **Empty destinations shown in flow log page** (C-1275)
  In the flow log list page, empty destinations were shown for AWS accounts. This affected giving permission for S3 buckets. This issue is resolved.

Known Issues in Illumio Segmentation for the Cloud

UI Components

- **Azure data is truncated and unable to create a deployment** (C-1405)
  Because Azure has long cpid's, UI is truncating the cipd so a user can't find the correct object when creating a deployment stack.

Functionality

- **Issue with relationships building in sync with cloudsync new resource events** (C-2168)
  The resource resync for applications created using VPC and subnet may not work.
- **Policy not removed after removing rules** (C-1878)
  When the user removed rules from a policy, the security group rules were not updated after exceeding the limit on security groups.
- **All previous deleted apps are showing up in policy creation** (C-1487)
  Applications that were previously created and removed are showing up in application policy creation.
- **Azure VM instances does not contain public IPs required for policy** (C-1219)
  Currently, policy generated for Azure resources only contains private IP addresses. Azure VM instances do not contain public IPs required for policy.

Known Limitations in Illumio Segmentation for the Cloud

Functionality

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **The application definition isn't showing the deployment if the deployment is added afterward** (C-1118)
  Deployment stacks need to be created before Application Definition.
- **Policy not applied to resources aside from Azure VMs/ AWS EC2s** (C-1114)
  Cloud discovers many resources in inventory, but policy can only be written on Azure Virtual Machines/AWS EC2 instances.

- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - December 11th, 2023

The following new features are available in the December 11th, 2023:

No.	Feature Category	Feature List
1.	Application	Cloud automatically resynchronizes pending applications with any resource changes. This means you can change a tag in such a way that it applies to an additional resource, and Cloud will automatically re-synchronize the application to include the new resource.
2.	Policy	• Cloud allows users to specify all services for the destination service when writing rules. • Cloud allows users to specify 'Any' (0.0.0.0/0 and ::/0) for the source or destination IP address when writing rules.

Resolved Issues in Illumio Segmentation for the Cloud

- **sgpolicyenforcement svc fails to enforce all-svc rules in AWS SG** (C-2232)
  An Illumio AWS security group service encountered errors when attempting to enforce rules covering all services. This issue is resolved.
- **Getting logged out after login** (C-2198)
  Illumio was not properly pruning sessions, which occasionally resulted in users getting logged out involuntarily. This issue is resolved.
- **Running the CloudFormation Template (CFT) in AWS does not work** (C-2255, C-2079)
  Illumio created the flow stack when users ran the template, but the link did not open in a new window. If the user clicked "Download," nothing would happen. This issue is resolved.
- **Application used in Ruleset was allowed to be deleted** (C-1607)
  Users were allowed to delete applications even if they were part of a ruleset. This issue is resolved. A message now displays telling the user that the application is currently used by ruleset, and blocks application removal.
- **Empty destinations shown in flow log page** (C-1275)
  In the flow log list page, empty destinations were shown for AWS accounts. This affected giving permission for S3 buckets. This issue is resolved.

Known Issues in Illumio Segmentation for the Cloud

UI Components

- **Azure data is truncated and unable to create a deployment** (C-1405)
  Because Azure has long cpid's, UI is truncating the cipd so a user can't find the correct object when creating a deployment stack.
- **Application policies last modified by should show username, not email** (C-1060)
  Application policies 'last modified by' should show username, not email.

Functionality

- **Policy not removed after removing rules** (C-1878)
  When the user removed rules from a policy, the security group rules were not updated after exceeding the limit on security groups.

- **All previous deleted apps are showing up in policy creation** (C-1487)

  Applications that were previously created and removed are showing up in application policy creation.
- **Azure VM instances does not contain public IPs required for policy** (C-1219)

  Currently, policy generated for Azure resources only contains private IP addresses. Azure VM instances do not contain public IPs required for policy.

Known Limitations in Illumio Segmentation for the Cloud

Functionality

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)

  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **The application definition isn't showing the deployment if the deployment is added afterward** (C-1118)

  Deployment stacks need to be created before Application Definition.
- **Policy not applied to resources aside from Azure VMs/ AWS EC2s** (C-1114)

  Cloud discovers many resources in inventory, but policy can only be written on Azure Virtual Machines/AWS EC2 instances.
- **Competing application definition (multiple app-def using same tags)** (C-1095)

  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

What's New in This Release - December 4th, 2023

The following new features are available in the December 4th, 2023 release:

No.	Feature Category	Feature List
1.	Cloud Map View	Illumio remembers your map browse sequence for your browser session. (Note that refreshing the page may cause Illumio to forget your map browse sequence.)
2.	Label Creation	Easier navigation for cloud tag to label mapping.

Resolved Issues in Illumio Segmentation for the Cloud

- **Unable to approve applications with multiple deployments** (C-2077)

  If an application discovered two or more deployments, and Illumio approved the first one (success), then approving the second one failed. This issue is resolved.
- **Empty applications match to resources** (C-2046)

  If an application contained no tags, VNets, or subnets, the application would match to all resources instead of none. This issue is resolved.
- **Resource sync doesn't sync associated labels** (C-2045)

  Associated labels were not captured when synchronizing an application's resources during cloudsync events. This issue is resolved.
- **All protocols not handled properly in some AWS SG cases** (C-2011)

  When computing the intersection of deny rules with allow rules in AWS SG, Illumio did not correctly handle the case where the allow rule has all protocols and the deny rule has a specific protocol. This issue is resolved.

Known Issues in Illumio Segmentation for the Cloud

UI Components

- **Azure data is truncated and unable to create a deployment** (C-1405)
  Because Azure has long cpid's, UI is truncating the cipd so a user can't find the correct object when creating a deployment stack.
- **Application policies last modified by should show username, not email** (C-1060)
  Application policies 'last modified by' should show username, not email.

Functionality

- **All previous deleted apps are showing up in policy creation** (C-1487)
  Applications that were previously created and removed are showing up in application policy creation.
- **Azure VM instances does not contain public IPs required for policy** (C-1219)
  Currently, policy generated for Azure resources only contains private IP addresses. Azure VM instances do not contain public IPs required for policy.

Known Limitations in Illumio Segmentation for the Cloud

- **Application sometimes gets mapped to the wrong deployment's env label** (C-1257)
  The resources have multiple cloud tags, the tag in the application definition label doesn't align with the one used in the environment label.
- **The application definition isn't showing the deployment if the deployment is added afterward** (C-1118)
  Deployment stacks need to be created before Application Definition.
- **Policy not applied to resources aside from Azure VMs/ AWS EC2s** (C-1114)
  Cloud discovers many resources in inventory, but policy can only be written on Azure Virtual Machines/AWS EC2 instances.

- **Competing application definition (multiple app-def using same tags)** (C-1095)
  Cloud allows users to create multiple application definition with the same rules, i.e., same set of tags can be shared for two applications.

Get Started

Congratulations on choosing Illumio Segmentation for the Cloud to protect the applications in your environment. The content in this category introduces you to Illumio Segmentation for the Cloud and explains many of the key concepts. In particular, this category explains how to onboard your public cloud accounts into Illumio Segmentation for the Cloud and provides an overview of the task flow for successfully working with Illumio Segmentation for the Cloud for your public cloud security needs.

Welcome to Illumio Cloud

You can prevent breaches and ransomware from becoming cyber disasters in the public cloud by getting visibility and insights into your cloud workloads and traffic activities. Use these insights to author effective policies. Note that the product name has changed to Illumio Segmentation for the Cloud.

Use the following features to gain visibility into your traffic patterns:

- **Multi-Cloud:** Single-pane-of-glass view into multi-cloud environments.
- **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. Use the Cloud Map to visualize resource hierarchy and traffic flows, and locate resources and security controls for network security reviews.
- **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.

Typical Cloud Workflow

1. Onboard a cloud account.
   - See Onboarding AWS Cloud [94].
   - See Onboarding Azure [59].
   - See Onboarding OCI [127].
2. Onboarding begins the process of discovering and ingesting resources from your environment.

   See Cloud Discovers Your Application Environments [225].
3. Go to the Cloud Map to see your discovered resources. Go to the Inventory page to see a list of your discovered resources.

   See Cloud Map [168] and Inventory [157].
4. Set up your deployments (optional) and define your applications.

   See Define a Deployment [228] and Define an Application [231].
5. Review the application definition before you apply security policies to your application definition.

See Cloud Policy Model [246] and Organization Policy versus Application Policy [252].

Content enhancements

Learn about documentation updates and enhancements we've made to enhance your experience with the documentation library.

New and updated documents

- October 2025: The documentation has been updated with these changes:
  - More GCP Dataplane Region IP addresses: GCP Flow Log Access IP Addresses [319]
  - Additional supported resource types for OCI: Illumio visibility for resource types [193]
  - Added how to upgrade the Kubernetes Operator in agentless containers: Upgrade the Kubernetes Operator in Agentless Containers [150]
- September 2025: Added support for two new OCI resources.
  - Illumio visibility for resource types [193]
- August 2025: The documentation has been updated to show how to set up an Azure Firewall flow log.
  - Set up flow logs in Azure [136]
- August 2025: The documentation has been updated to show how to verify an Azure Firewall Policy is attached to your Azure Firewall.
  - Verify Azure Firewall Policy resource [261]
- August 2025: GCP now supports policy.

  Current Illumio Cloud what's new and release notes [10]
- August 2025: Security review and enable read-write are now available.
  - Security reviews [267]
  - Enable read-write permissions [269]
- August 2025: OCI now supports policy under the Beta program.

  Current Illumio Cloud what's new and release notes [10]
- July 2025: GCP is now GA.
- Current Illumio Cloud what's new and release notes [10]
- July 2025: Agentless Containers now supports OpenShift OVN (Open Virtual Networking)-Kubernetes, as well as visibility for AKS and GKE.

  Agentless Containers overview [148]
- July 2025: The documentation has been updated to reflect the new product name, Illumio Segmentation for the Cloud.
- Onboard an OCI tenant [129]

  July 2025: The onboarding wizard now allows write permissions.
- Use AI Labeling [244]

  July 2025: The use of application labels and the size of the pool of recommendations have been added.
- Onboarding GCP [119]

  July 2025: Illumio Segmentation for the Cloud now supports visibility under the Beta program.
  - Prerequisites for onboarding GCP [119]
  - Permissions for onboarding GCP [320]
  - Onboard a GCP project [120]
  - Onboard a GCP organization [122]
  - Onboard a GCP folder [125]

  - Prerequisites for granting flow log access to your CSPs [140]
  - Set up flow logs in your CSP environment [132]
  - Grant flow log access to your CSPs [141]
  - GCP Flow Log Access IP Addresses [319]
  - Cloud Tag to Label Mapping [241]
  - Define an application individually [235]
  - Illumio visibility for resource types [193]
  - Cloud Map navigation [171]
  - Traffic [219]
- Review destinations before granting flow log access [145]

  July 2025: The flow log destination review process is now streamlined
- Drift Detection [271]

  June 2025: Content was added to distinguish between Drift Detection and Tamper Protection
- Tamper Protection [272]

  June 2025: Content was added to distinguish between Tamper Protection and Drift Detection, and tamper protection messages were updated
- Grant flow log access to your CSPs [141]

  June 2025: Content was updated to reflect increased traffic flow ingestion frequency

"Show me how" videos and animations

For best results, Illumio recommends viewing videos in Chrome.

Onboarding		
Onboarding an AWS Cloud Organization	Onboard an Azure Cloud Tenant	Onboard an Azure Cloud Subscription
Grant flow log access to your CSPs	Review destinations before granting flow log access	Onboard a GCP Project
Onboard a GCP Organization	Onboard a GCP Folder	
Kubernetes and Agentless Containers		
Agentless Containers overview	Onboard and Offboard Kubernetes Clusters	
Deployments and applications		
Define an application manually	Define an application automatically	Define a deployment
View and approve an application		
Reports and filtering		
Generate reports	Cloud search and filtering	Export a usage report
Traffic	Search traffic	Inventory
Events	Export an application report	Define an application automatically
Administration		
User management	Role-based access control	Service accounts

Activating your Account and Signing in

Illumio is set up for Okta multi-factor authentication (MFA) for sign-in. You do not need to acquire Okta separately.

After you sign up for a free trial, Illumio sends you an onboarding email with a link to activate your account. The invitation lasts for seven days. If you do not accept it by then, Illumio can send another invitation, which lasts 24 hours.

Note that the first user (with the same domain as the organization) to onboard an account becomes the owner for your organization.

Activate your Account and Sign In

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

1. Click the link in your onboarding email to begin.
2. In the Okta page, activate your account by setting a password.

After setting your password in Okta, Illumio sends you an email with this URL:

`https://console.illum.io`

3. Click the console URL.
4. Enter your email address and click **Continue**.
5. If you want to stay signed in, select that option. Click **Next**. When signing in to your console session, please note:
   - Illumio keeps your Okta one-time-password for 15 to 30 seconds
   - Console sessions last two hours by default
6. Verify your identity. If prompted, re-enter your password and click **Verify**.

   The **Connecting to** page refreshes with a message that you are being signed in.
7. Next, view your profile, roles, and set up users. See My Profile [278], Role-based access control [279], My Roles [278], and User management [276].

After you complete these steps, onboard your cloud service provider (CSP) accounts.

- See Onboarding AWS Cloud [94].
- See Onboarding Azure [59].
- See Onboarding OCI [127].

Onboarding Azure

Review the onboarding workflow for your cloud environment:

Azure onboarding workflow - default setup

If you have the permissions to run the onboarding PowerShell script, follow the default setup for your Azure subscriptions and tenants.

1. Before you begin, review the prerequisites and permissions
   - See Prerequisites for Onboarding Azure [60].
   - See Permissions for Onboarding Azure [61].
2. Use the wizard to onboard subscriptions and tenants.
   - See Onboard an Azure Cloud subscription - default setup [65].
   - See Onboard an Azure Cloud tenant - default setup [79].
3. Set up flow logs and grant access.
   - See Set up Flow Logs [132].
   - See Grant flow log access to your CSPs [141].

After you onboard your Azure subscriptions or tenants, you can visualize your resources, define your public cloud environments, and create policies. See After onboarding your accounts [146].

Azure onboarding workflow - guided setup

If you cannot onboard an Azure tenant or subscription using the default method due to permissions limitations, such as being unable to run the onboarding PowerShell script, follow this guided setup.

1. Before you begin, review the prerequisites and permissions.
   - See Prerequisites for Onboarding Azure [60].
   - See Permissions for Onboarding Azure [61].
2. Use the guided setup to onboard subscriptions and tenants.
   - See Onboard an Azure Cloud Tenant - Guided Setup [84].
   - See Onboard an Azure Cloud Subscription - Guided Setup [67].
3. Set up flow logs and grant access.
   - See Set up Flow Logs [132].
   - See the Grant Flow Log Access - Guided Setup section in Onboard an Azure Cloud Tenant - Guided Setup [90]. Do not refer to the default method for granting flow log access.

After you onboard Azure tenants using the guided setup, you can visualize your resources, define your public cloud environments, and create policies. See After onboarding your accounts [146].

Azure tenant onboarding workflow - Terraform Application setup

If you wish to use the Illumio Terraform module to automate and simplify Azure tenant onboarding, follow this Terraform setup.

1. Before you begin, review the prerequisites and permissions.
   - See Prerequisites for Onboarding Azure [60].
   - See Permissions for Onboarding Azure [61].
2. Create a Terraform application.
   - See Create a Terraform Illumio Onboarding Application for Azure [75].
3. Use the Terraform application to onboard tenants.
   - See Onboard an Azure Subscription using a Terraform Illumio Onboarding Application [78].
4. Set up flow logs and grant access.
   - See Set up Flow Logs [132].
   - See Grant flow log access to your CSPs [141].

After you onboard Azure tenants using the Terraform Application setup, you can visualize your resources, define your public cloud environments, and create policies. See After onboarding your accounts [146].

Prerequisites for Onboarding Azure

Review these prerequisites before you begin onboarding your Azure tenants or subscriptions.

Before you begin onboarding Azure

Once you review these prerequisites, return to Onboarding Azure [59] for next steps.

□ Log into an Azure account. The onboarding wizard flow assumes that you are already logged into an Azure account.

□ The default installation assumes that you have Owner access or the User Access Administrator Role for assigning the Reader Role at the Tenant scope to complete the Azure Entra ID App registration. Check your permissions for the subscription or tenant you are onboarding, using the Azure portal Access Control (IAM) page. If you do not have access, see Onboard an Azure Cloud Tenant - Guided Setup [84] and Onboard an Azure Cloud Subscription - Guided Setup [67].

- □ Know your parent management group (tenant and/or subscription) IDs that you want to onboard. They can be found under the Management Groups in the Azure portal. The tenant ID is also known as the parent management group ID.
- □ If you are restricting public access to flow logs, make certain ports and IP addresses available to Illumio Segmentation for the Cloud. See Azure Flow Log Access IP Addresses [308].
- □ If you will be onboarding Azure with PowerShell and the illumio-init.ps1 script, you must be running PowerShell version 7.4 with the Az module.

Required Azure Permissions

If you use the guided method described in Onboard an Azure Cloud Tenant - Guided Setup [84] and Onboard an Azure Cloud Subscription - Guided Setup [67], set required permissions using the Azure console.

See Permissions for Onboarding Azure [61].

Permissions for Onboarding Azure

This section describes the set of permissions that you grant to the Illumio Segmentation for the Cloud App that is registered in Azure Active Directory.

These permissions are required, irrespective of whether you use the default method provided by the wizard or the guided method.

- If you are onboarding using the default method described in Onboard an Azure Cloud tenant - default setup [79] and Onboard an Azure Cloud subscription - default setup [65], it automatically provisions the permissions described here.
- If you are onboarding using the guided method, which does not involve running the PowerShell script provided in the wizard, or if you lack Owner access, described in Onboard an Azure Cloud Tenant - Guided Setup [84]) and Onboard an Azure Cloud Subscription - Guided Setup [67], you need to set these permissions via the Azure Console.

Azure permission descriptions

Permission Type	Permission Name	Notes
Read	Reader - role	This role gives Illumio Segmentation for the Cloud the permissions to read data or resources from your subscription or tenant. This role allows the viewing of all resources, but does not allow modification.
Write	Writer - role	This role gives Illumio Segmentation for the Cloud the permissions to modify data or resources in your subscription or tenant. This role allows the modification of resources.
NSG, Azure Firewall	Multiple, see below.	Use these permissions to create custom roles. Define any custom roles with elevated permissions, as part of the PowerShell script that is run when you onboard an Azure subscription. If the user onboarding Azure has Owner permissions, these permissions are automatically assigned to the "Illumio Network Security Administrator" custom role that is created when the onboarding PowerShell script is run. However, if the user onboarding Azure does *not* have Owner permissions, you must create the"Illumio Network Security Administrator" custom role with these NSG and Azure Firewall permissions *before* the onboarding PowerShell script is run.
Flow	Storage Blob Data Reader – role	

Azure read and write policy

When you grant read and write permissions to Illumio Segmentation for the Cloud, the following roles are created in the Azure tenant.

```
Reader Role - Built In Role
{
  "assignableScopes": [
    "/"
  ],
  "description": "View all resources, but does not allow you to make any
changes.",
  "id": "/providers/Microsoft.Authorization/roleDefinitions/
acdd72a7-3385-48ef-bd42-f606fba81ae7",
  "name": "acdd72a7-3385-48ef-bd42-f606fba81ae7",
  "permissions": [
    {
      "actions": [
        "*/read"
      ],
      "notActions": [],
      "dataActions": [],
      "notDataActions": []
    }
  ],
  "roleName": "Reader",
  "roleType": "BuiltInRole",
  "type": "Microsoft.Authorization/roleDefinitions"
}
Illumio Network Security Administrator Role - Custom Role
{
    "properties": {
        "roleName": "Illumio Network Security Administrator",
        "description": "Illumio Network Administration Role",
        "assignableScopes": [
            "/"
        ],
        "permissions": [
            {
                "actions": [
                    "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.Network/networksecuritygroups/providers/
Microsoft.Insights/diagnosticSettings/read",
                    "Microsoft.Network/networksecuritygroups/providers/
Microsoft.Insights/diagnosticSettings/write",
                    "Microsoft.Network/networksecuritygroups/providers/
Microsoft.Insights/logDefinitions/read",
                    "Microsoft.Network/networkWatchers/securityGroupView/
action",
                    "Microsoft.Network/networkSecurityGroups/*",
                    "Microsoft.Network/networkInterfaces/read",
                    "Microsoft.Network/networkInterfaces/write",
                    "Microsoft.Network/virtualNetworks/read",
                    "Microsoft.Network/virtualNetworks/subnets/write",
                    "Microsoft.Authorization/locks/*",
                    "Microsoft.Compute/virtualMachines/read"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
}
Illumio Firewall Administrator Role - Custom Role
{
    "properties": {
        "roleName": "Illumio Firewall Administrator",
        "description": "Illumio Firewall Administrator role",
        "assignableScopes": [
            "/"
        ],
        "permissions": [
            {
                "actions": [
                    "Microsoft.Network/azurefirewalls/read",
                    "Microsoft.Network/azurefirewalls/learnedIPPrefixes/
action",
                    "Microsoft.Network/azureFirewalls/
applicationRuleCollections/write",
                    "Microsoft.Network/azureFirewalls/
applicationRuleCollections/delete",
                    "Microsoft.Network/azureFirewalls/
applicationRuleCollections/read",
                    "Microsoft.Network/azurefirewalls/providers/
Microsoft.Insights/logDefinitions/read",
                    "Microsoft.Network/azureFirewalls/natRuleCollections/
write",
                    "Microsoft.Network/azureFirewalls/natRuleCollections/
read",
                    "Microsoft.Network/azureFirewalls/natRuleCollections/
delete",
                    "Microsoft.Network/azureFirewalls/
networkRuleCollections/read",
                    "Microsoft.Network/azureFirewalls/
networkRuleCollections/write",
                    "Microsoft.Network/azureFirewalls/
networkRuleCollections/delete",
                    "Microsoft.Network/azureFirewallFqdnTags/read",
                    "Microsoft.Network/azurefirewalls/providers/
Microsoft.Insights/metricDefinitions/read",
                    "Microsoft.Network/firewallPolicies/read",
                    "Microsoft.Network/firewallPolicies/write",
                    "Microsoft.Network/firewallPolicies/join/action",
                    "Microsoft.Network/firewallPolicies/certificates/
action",
                    "Microsoft.Network/firewallPolicies/delete",
                    "Microsoft.Network/firewallPolicies/
ruleCollectionGroups/read",
                    "Microsoft.Network/firewallPolicies/
ruleCollectionGroups/write",
                    "Microsoft.Network/firewallPolicies/
ruleCollectionGroups/delete",
                    "Microsoft.Network/firewallPolicies/ruleGroups/read",
                    "Microsoft.Network/firewallPolicies/ruleGroups/write",
                    "Microsoft.Network/firewallPolicies/ruleGroups/delete",
                    "Microsoft.Network/ipGroups/read",
                    "Microsoft.Network/ipGroups/write",
                    "Microsoft.Network/ipGroups/validate/action",
                    "Microsoft.Network/ipGroups/updateReferences/action",
                    "Microsoft.Network/ipGroups/join/action",
                    "Microsoft.Network/ipGroups/delete"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
}
```

Azure flow log support

Illumio Segmentation for the Cloud supports NSG Flow logs version 2 (includes flow state and byte counts), but does not support version 1. It also supports VNet flow logs and Azure Firewall flow logs.

See Set up flow logs in your CSP environment.

Onboard an Azure Cloud subscription - default setup

Learn how to onboard an Azure subscription. Before you begin:

- Review the prerequisites. See Prerequisites for Onboarding Azure [60].
- See Permissions for Onboarding Azure [61] for the list of required permissions. Onboarding a subscription with the wizard automatically provides the required permissions.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Azure_Onboarding.mp4

1. If this is the first time logging in to Illumio Segmentation for the Cloud, 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 Illumio Segmentation for the Cloud. The name should be descriptive so that you can easily identify it.
   - **Tenant ID:** Paste this ID that you copied from Azure. The tenant ID is also known as the parent management group ID.
   - **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 Illumio Segmentation for the Cloud 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 Illumio Segmentation for the Cloud 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 Illumio Segmentation for the Cloud.
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-Cloud-Secure-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 Illumio Segmentation for the Cloud access to the subscription resources. Illumio Segmentation for the Cloud is able to discover subscription resources and write policies for them.

   For the complete list of permissions granted to Illumio Segmentation for the Cloud for your account, see Permissions for Onboarding Azure [61].

   The script sends the Client ID and Client Secret to Illumio Segmentation for the Cloud. It accesses your Azure subscription so that you don't have to repeatedly provide your Azure credentials.
6. Leave your Azure portal and return to Illumio Segmentation for the Cloud. 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.

Onboard a subscription with Illumio

> **NOTE**
>
> Illumio Segmentation for the Cloud 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, Illumio Segmentation for the Cloud 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 [141].

What's next after onboarding your subscription?

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

For the next steps after onboarding a subscription, set up and enable flow logs. See Onboarding Azure [59]. Once you set up and enable flow logs, see After onboarding your accounts [146].

If you originally set the permissions to read only, and wish to change them to read and write, see Change Azure permissions from read to read and write [93].

Updating your Service Account Principals

If you need to update your service account principals when they expire, see Update Service Principals for Onboarded Azure Subscriptions and Tenants [315].

Onboard an Azure Cloud Subscription - Guided Setup

Learn how to onboard an Azure subscription using the guided setup if you cannot onboard it using the default setup described in Onboard an Azure Cloud subscription - default setup [65].

> **NOTE**
>
> If you don't have permissions to run the PowerShell script for Azure subscription onboarding, you will not be able to use the default setup.

1. Review the prerequisites. See Prerequisites for Onboarding Azure [60].
2. Use your Azure console to provide the required permissions. See Permissions for Onboarding Azure [61] for a list of required permissions.

To provide permissions in your Azure console, see the Microsoft website.

3. Onboard your subscription using the guided setup. See Using the Guided Setup [68].

When you onboard an Azure subscription, the service principal allows Illumio Segmentation for the Cloud to retrieve subscriptions and resources. After you create the Azure Entra ID application, set the required reader permissions at the subscription scope, and provide the client ID and client secret credentials, invoke the API using a PowerShell script. The credentials are required to communicate with your Azure subscription.

Using the Guided Setup

Connect to Azure with the Wizard

The steps you take in this first part of the Illumio onboarding wizard are mostly the same as described in Onboard an Azure Cloud subscription - default setup [65]. The exception is that that if you choose to give Illumio Segmentation for the Cloud more than read-only access, you *must* download the permissions as described in Step 4.

1. If you are logging in for the first time, click **+ Azure** on the Onboarding page to onboard your first account.
2. If you've already onboarded other accounts, choose **Onboarding** from the left navigation. Click **+Add Azure** at the top of the page.
3. The **Add Azure Cloud Subscription** wizard starts and displays the first step: **Connect to Azure**

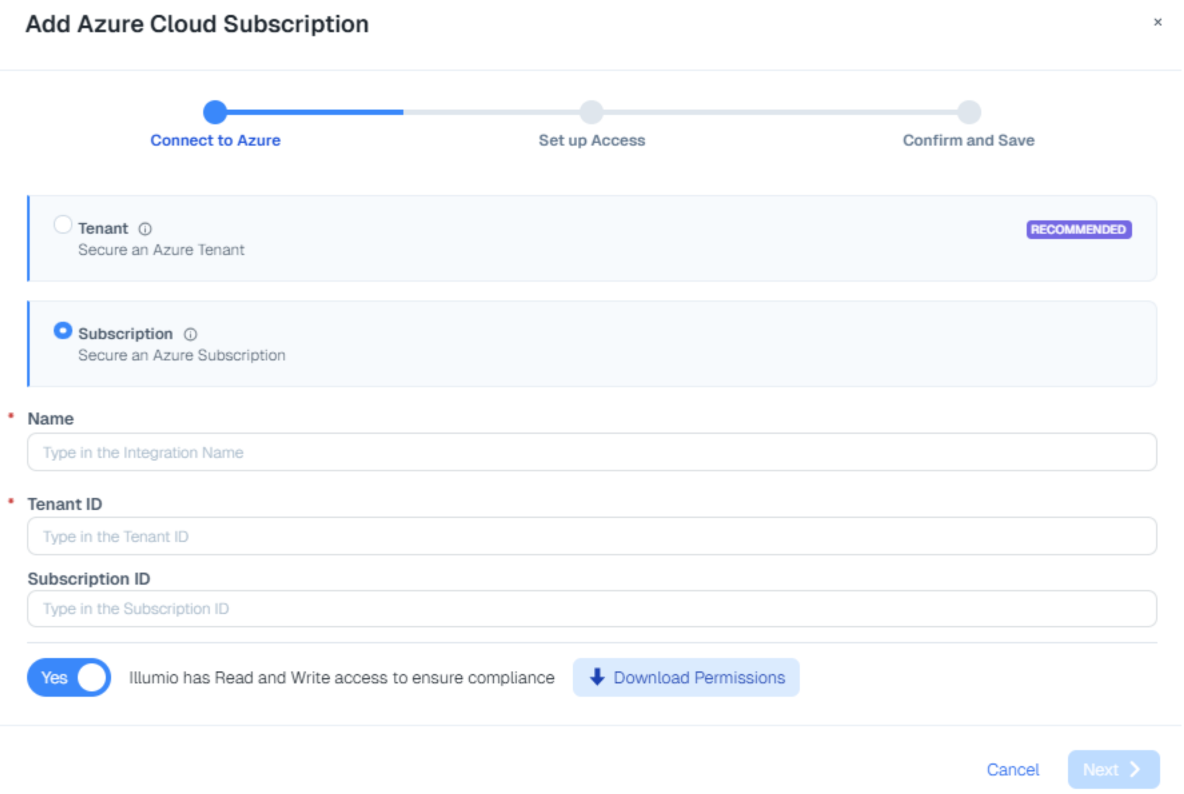

Add Azure Cloud Subscription ×

Connect to Azure — Set up Access — Confirm and Save

○ Tenant ⓘ RECOMMENDED
Secure an Azure Tenant

● Subscription ⓘ
Secure an Azure Subscription

* Name: Type in the Integration Name

* Tenant ID: Type in the Tenant ID

Subscription ID: Type in the Subscription ID

Yes — Illumio has Read and Write access to ensure compliance — Download Permissions

4. Provide the following information about your Azure account:
   - **Name:** Specify a descriptive name for the account. This name appears in Cloud.
   - **Tenant ID:** Paste the parent management group ID that you copied from Azure.
   - **Subscription ID:** Paste the subscription ID that you copied from Azure.
   - **Onboarding toggle options:**

Onboarding Toggle Option	Action
Illumio has Read and Write access to ensure compliance: Yes	To grant the Illumio Cross Account Role permission to view your Azure subscription resources and to apply policy to them.
Illumio has Read and Write access to ensure compliance: No	To provide the Illumio Cross Account Role read-only access.

NOTE

To view the permissions you are granting Illumio Segmentation for the Cloud to your Azure subscription, click **Download Permissions**. This is *required* for you to onboard the subscription if you are unable to run the script, because you must manually provide the listed permissions.

Setup Access using the Wizard without the Script

1. Select an existing service account or create a new one by clicking **Add a new Service Account**. You can use spaces, underscores, numbers, or other characters (such as !,@,#, and so on) for the account name. Make it something you can remember.

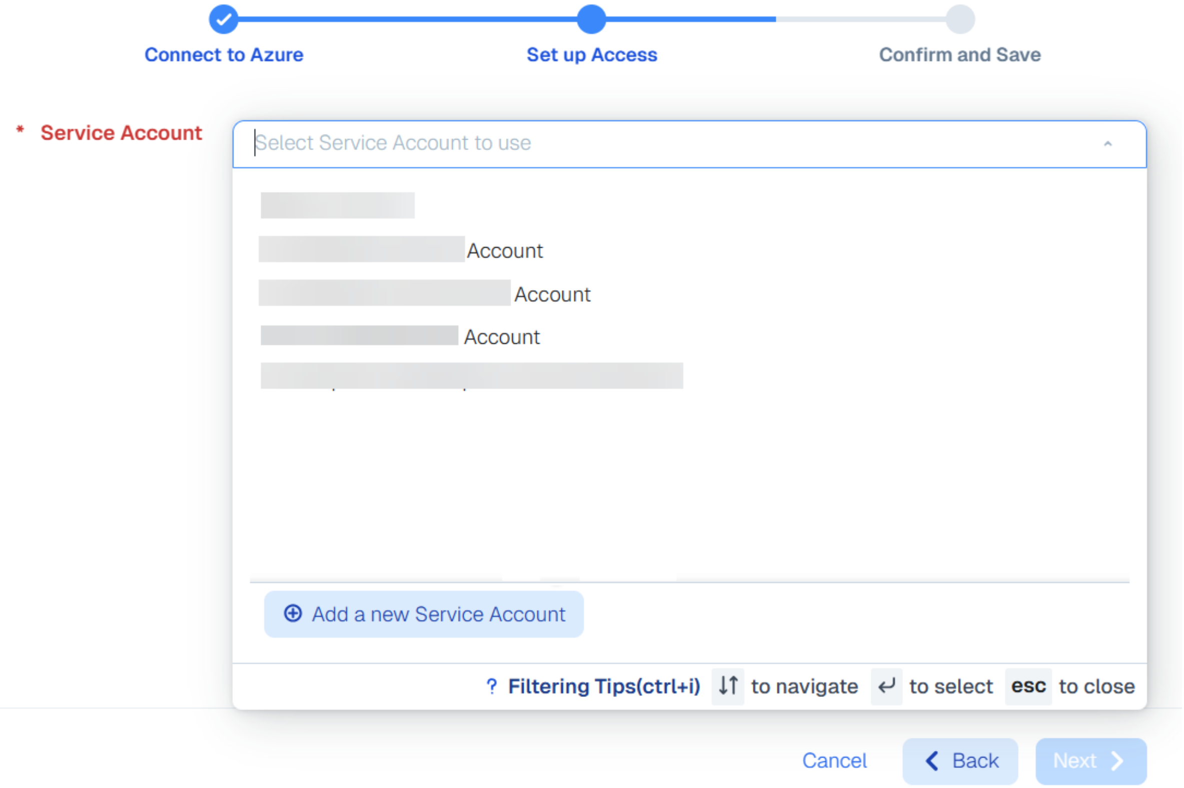

2. If you created a new service account, click **Download Credentials**. They are needed for the callback API to return the Azure Entra ID app credentials back to Illumio Segmentation for the Cloud.

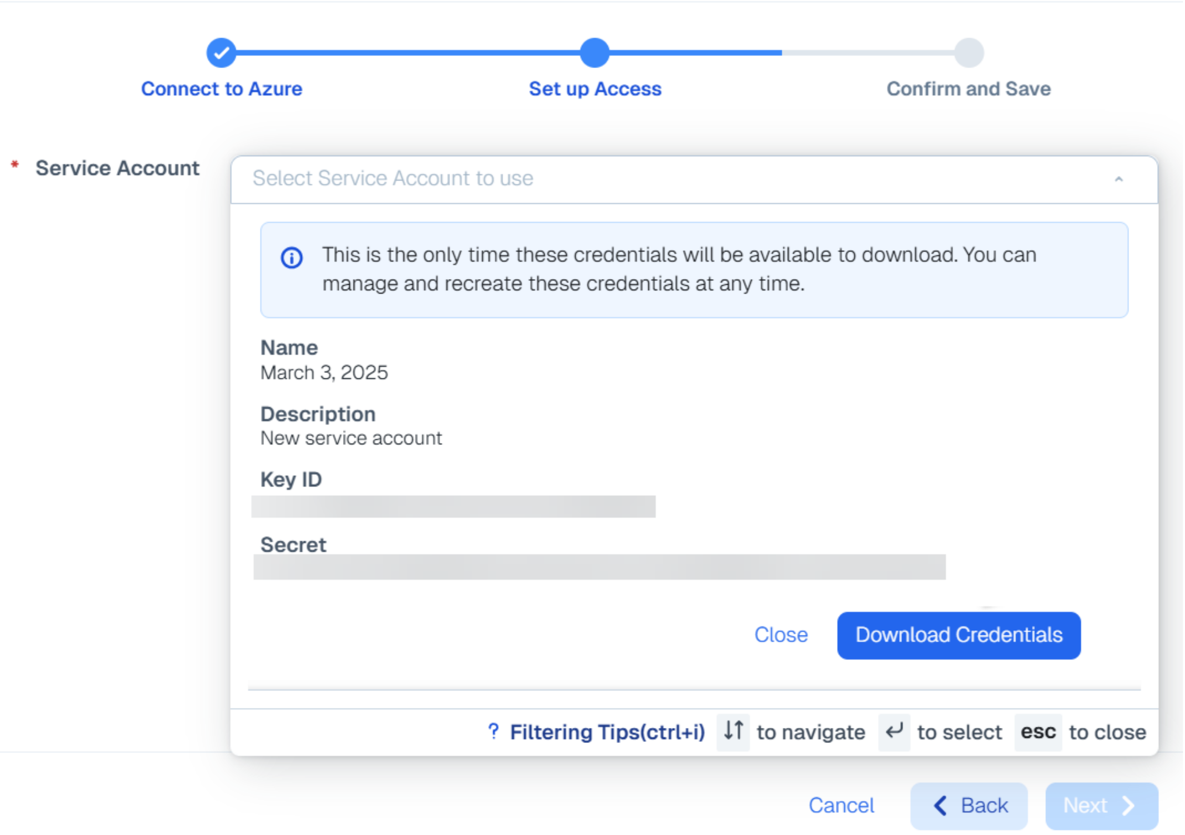

This is the only time these credentials will be available to download. You can manage and recreate these credentials at any time.

Name
March 3, 2025

Description
New service account

Key ID
[redacted]

Secret
[redacted]

3. Enter the ServiceAccountToken.
4. Do not use the PowerShell script offered in the wizard because this onboarding method assumes you lack the permissions to run it. Instead, follow the steps below in Manually Create an Azure Entra ID App and Assign the Reader RBAC Role [71] to create and register your Azure application.

   This gives Illumio Segmentation for the Cloud your newly created Azure application registrations, client, client id, and client secret.
5. Complete the fields.

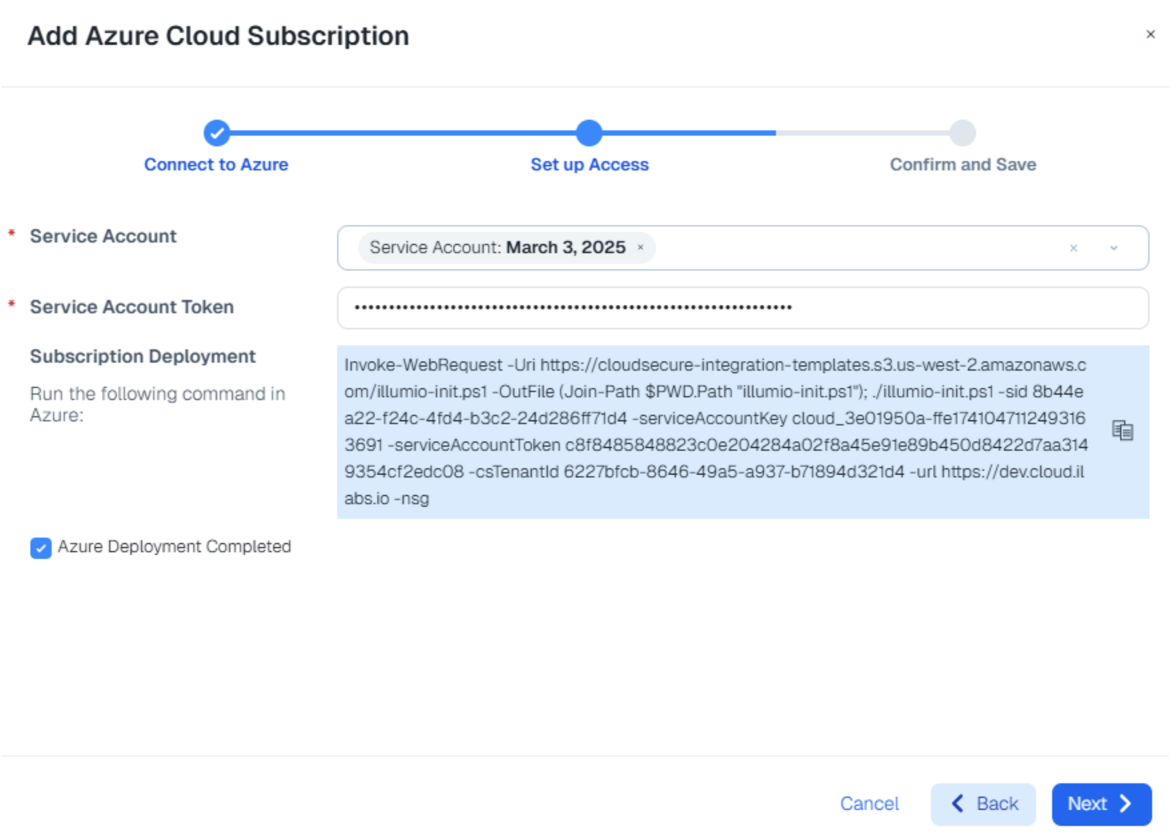

Add Azure Cloud Subscription

* Service Account: Service Account: **March 3, 2025**

* Service Account Token: ••

Subscription Deployment

Run the following command in Azure:

```
Invoke-WebRequest -Uri https://cloudsecure-integration-templates.s3.us-west-2.amazonaws.c
om/illumio-init.ps1 -OutFile (Join-Path $PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -sid 8b44e
a22-f24c-4fd4-b3c2-24d286ff71d4 -serviceAccountKey cloud_3e01950a-ffe17410471124931 6
3691 -serviceAccountToken c8f8485848823c0e204284a02f8a45e91e89b450d8422d7aa314
9354cf2edc08 -csTenantId 6227bfcb-8646-49a5-a937-b71894d321d4 -url https://dev.cloud.il
abs.io -nsg
```

☑ Azure Deployment Completed

6. Review the details and click **Save and Confirm**.

Manually Create an Azure Entra ID App and Assign the Reader RBAC Role

Use these steps to create an Azure Entra ID App manually.

1. In a new browser window, open your Azure portal.
2. Create the Azure Entra ID App registration as described on the Microsoft website. Skip the redirect Uniform Resource Identifier (URI) in this step.
3. Once the Entra ID App registration is created, create a new client secret as documented on the Microsoft website.
4. Once the Entra ID App and secret are created, assign the Reader RBAC role to the App at the subscription scope, as documented on the Microsoft website.
5. Ensure that the role is assigned at the subscription scope and not the tenant scope.
6. After you create the reader RBAC role, assign the API permissions for Illumio Segmentation for the Cloud. You can see the permissions in the .txt file you downloaded from the wizard.

Run the Callback API to Illumio Segmentation for the Cloud

After you connect to Azure and set up access, run the following PowerShell callback to the Illumio Segmentation for the Cloud API in your Azure console to complete the subscription onboarding. If the callback is successful, no output is printed.

NOTE

The following code is just a reference PowerShell script (Web_request.ps1), so update it according to your environment, but ensure that the subscription ID is empty.

```
# Set your service account key ID, token, and client secret
$serviceAccountKeyId = "<YourServiceAccountKeyId>"
$serviceAccountToken = "<YourServiceAccountToken>"
$clientSecret = "<YourClientSecret>" # The actual client secret to be
encoded

# Combine the key ID and token with a colon and base64 encode for the
Authorization header
$authString = "$($serviceAccountKeyId):$($serviceAccountToken)"
$encodedAuthString =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($authString))

# Base64 encode the client secret separately
$encodedClientSecret =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($clientSecret))

# Construct the headers with the encoded Authorization header
$headers = @{
  "X-Tenant-Id"   = "<CloudsecureTenantId>"
  "Content-Type"  = "application/json"
  "Authorization" = "Basic $encodedAuthString"
}

# Construct the request body with the encoded client secret
$body = @{
  "type"             = "AzureRole"
  "client_id"        = "<ClientId>"
  "client_secret"    = $encodedClientSecret  # Use the base64 encoded client
secret
  "subscription_id" = "<SubscriptionId>" # remove this and use
azure_tenant_id if onboarding the entire tenant.
  "azure_tenant_id" = "<AzureTenantId>" # both azure tenant id and
subscription_id should be present for subscription onboarding.
} | ConvertTo-Json -Depth 10

# Send the POST request
$response = Invoke-WebRequest -Uri 'https://cloud.illum.io/api/v1/
integrations/cloud_credentials' -Method Post -Headers $headers -Body $body


# Output the response
Write-Host $response
```

Set up and Enable Azure Flow Logs after Onboarding

The **Onboarding** page opens and displays a new row for that subscription.

Set up and enable flow logs.

- To set up flow logs before enabling them, see Set up Flow Logs [132].
- To enable flow logs using the guided setup, see Grant Azure Flow Log Access - Guided Setup [73].

Grant Azure Flow Log Access - Guided Setup

Run the following script in your Azure console to manually grant flow log access for your subscription. Note that "csTenantId" is the tenant ID and that "subscriptionId" is the subscription ID for Illumio Segmentation for the Cloud.

```
#params for running the script
# to run the script use the command
# ./grant_flow.ps1 -serviceAccountKey <key> -serviceAccountToken
<token> -csTenantId <cs_tenant_id> -subscriptionId <subscription_id>
-storageAccount <storage_account1,storage_account2,...> -url <url>
# default url is https://cloud.illum.io
param(
    $serviceAccountKey = "",
    $serviceAccountToken = "",
    $csTenantId = "",
    $subscriptionId = "",
    [String[]]$storageAccounts,
    $url = "")

#check if the given inputs are valid
if ([string]::IsNullOrEmpty($serviceAccountKey) -or
[string]::IsNullOrEmpty($serviceAccountToken)) {
    Write-Host "Service Account Token or Key cannot be empty"
-ForegroundColor Red
    exit
}

if ([string]::IsNullOrEmpty($csTenantId)) {
    Write-Host "csTenantId cannot be empty" -ForegroundColor Red
    exit
}

if ([string]::IsNullOrEmpty($subscriptionId)) {
    Write-Host "subscriptionId cannot be empty" -ForegroundColor Red
    exit
}


if ($storageAccounts.Count -eq 0) {
    Write-Host "storage accounts list cannot be empty" -ForegroundColor Red
    exit
}

# Combine the key ID and token with a colon and base64 encode for the
Authorization header
$authString = "$($serviceAccountKey):$($serviceAccountToken)"
$encodedAuthString =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($authString))


# Construct the headers with the encoded Authorization header
$headers = @{
    "X-Tenant-Id"   = $csTenantId;
    "Content-Type"  = "application/json";
    "Authorization" = "Basic $encodedAuthString";
}

# Construct the request body with the storage accounts
# list of storage accounts which has flow logs and for which access has
been granted.
```

```
# the storage account names must be the entire namespace as shown in
cloudsecure flow logs page.
# Eg: /subscriptions/<subscription_id>/resourceGroups/<resource_group_name>/
providers/Microsoft.Storage/storageAccounts/<storage_account_name>

$body = @{
    "type"            = "AzureFlow";
    "subscription_id" = $subscriptionId; # subscription in which the
storage accounts are present
    "destinations"     = $storageAccounts;
} | ConvertTo-Json -Depth 10

# Send the POST request
$endPoint = "/api/v1/integrations/cloud_credentials"

if ([string]::IsNullOrEmpty($url) ) {
    Write-Host "url is empty hence using default url https://
cloud.illum.io" -ForegroundColor Yellow
    $url = "https://cloud.illum.io"
}

$url = $url + $endPoint

Write-Host "making API call to $($url)"

$response = Invoke-WebRequest -Uri $url -Method Post -Headers $headers
-Body $body

# Output the response
if ($($response.StatusCode -eq 200)) {
    Write-Host "API call to cloudsecure successful"
} else {
    Write-Host "Error making API call to cloudsecure. Status code: $
(response.StatusCode)"
}
```

What's Next after Onboarding your Subscription?

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

For the next steps after onboarding a subscription, After onboarding your accounts [146].

If you originally set the permissions to read only, and wish to change them to read and write, see Change Azure permissions from read to read and write [93].

Updating your Service Account Principals

If you need to update your service account principals when they expire, see Update Service Principals for Onboarded Azure Subscriptions and Tenants [315].

Create a Terraform Illumio Onboarding Application for Azure

Learn to create an onboarding application for Azure subscriptions using Terraform. Illumio provides a Terraform module to automate Azure application creation and grant the application the necessary permissions for integration with Illumio Segmentation for the Cloud. Create the Terraform Azure application before you onboard your Azure subscription using Terraform. See Onboard an Azure Subscription using a Terraform Illumio Onboarding Application [78].

Here's an overview of the workflow to onboard Azure subscriptions using Terraform.

1. Create and Register a Terraform Illumio Onboarding Application [76].
2. Get your Terraform Illumio Onboarding Application Client Secret [77].
3. Set Permissions for your Terraform Illumio Onboarding Application [77].
4. Assign Roles and Access for your Terraform Illumio Onboarding Application [77]

Create and Register a Terraform Illumio Onboarding Application

1. Launch the Microsoft Azure Portal and sign in.
2. Browse to **Microsoft Entra ID (formerly Azure Active Directory)** > **Properties**.
3. Copy the tenant ID and save it in a text file. You'll need it later when you modify your Terraform script for onboarding Azure.

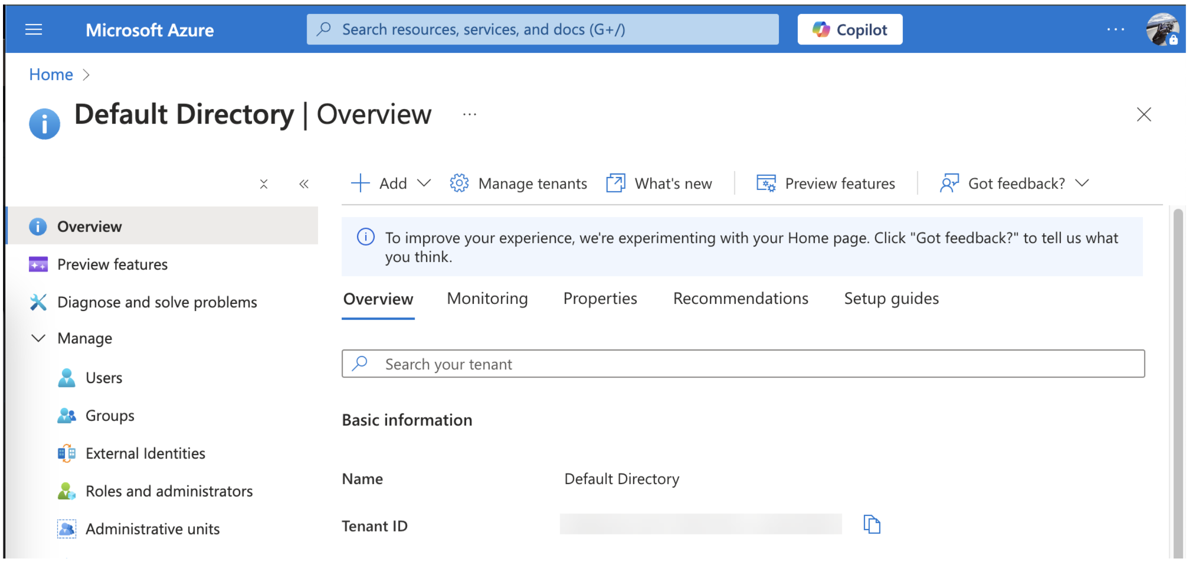

To improve your experience, we're experimenting with your Home page. Click "Got feedback?" to tell us what you think.

Basic information

Name Default Directory

Tenant ID [redacted]

4. Browse to **App registrations** > **New registrations**.
5. Enter the name and click **Register**. Terraform uses this application only to create another application, which provides Illumio Segmentation for the Cloud access to your tenant and its subscriptions.
6. Copy the Application (client) ID and save it in a text file. You'll need it later when you modify your Terraform script for onboarding Azure.

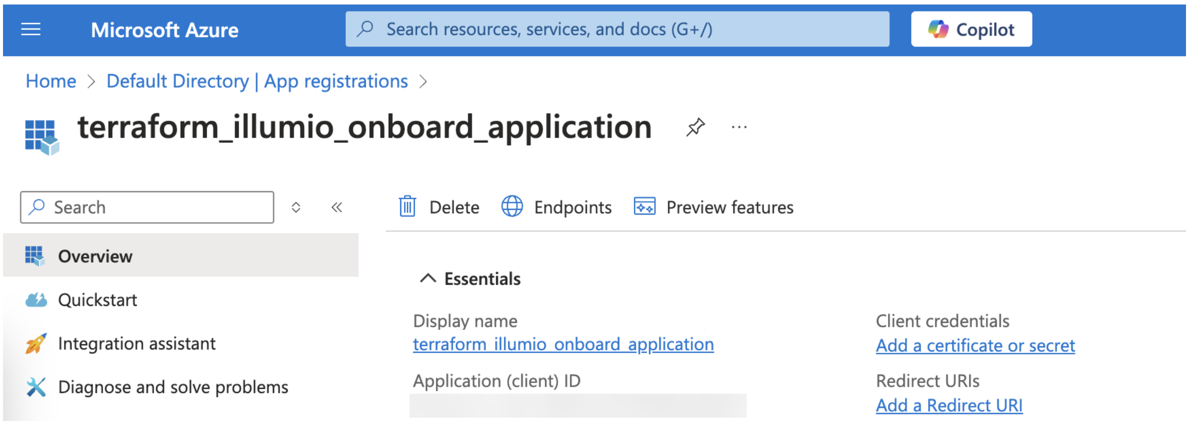
Microsoft Azure
Search resources, services, and docs (G+/)
Copilot
Home > Default Directory | App registrations >
terraform_illumio_onboard_application
Search
Overview
Quickstart
Integration assistant
Diagnose and solve problems
Delete
Endpoints
Preview features
Essentials

Display name
terraform_illumio_onboard_application

Application (client) ID
[redacted]

Client credentials
Add a certificate or secret

Redirect URIs
Add a Redirect URI

Get your Terraform Illumio Onboarding Application Client Secret

1. Click **Certificate & secrets**.
2. Under Client secrets, click **New client secret**.
3. Enter a description, select the recommended expiration, and click **Add**.
4. Copy the client secret value and save it in a text file. You'll need it later when you modify your Terraform script for onboarding Azure.

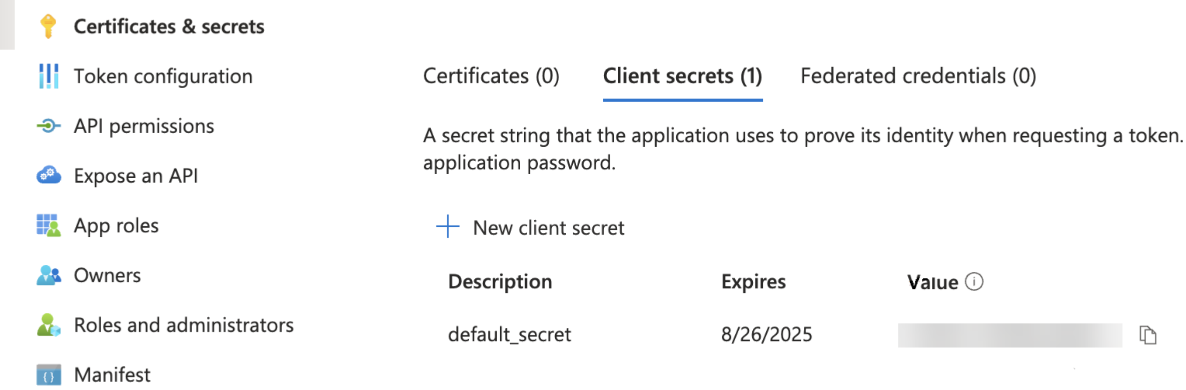

A secret string that the application uses to prove its identity when requesting a token. application password.

Description	Expires	Value ⓘ
default_secret	8/26/2025	[redacted]

Set Permissions for your Terraform Illumio Onboarding Application

1. Click **API permissions**.
2. Browse to **Configured permissions** > **Add a permission**.
3. Under Commonly used Microsoft APIs, browse to **Microsoft Graph** > **Delegated permissions**.
4. Expand the following:
   - Application: check the box for **Application.ReadWrite.All**
   - Directory: check the box for **Directory.ReadWrite.All**
5. Click **API permissions**.
6. Click **Grant admin consent for Default Directory** for Azure to grant the permission.

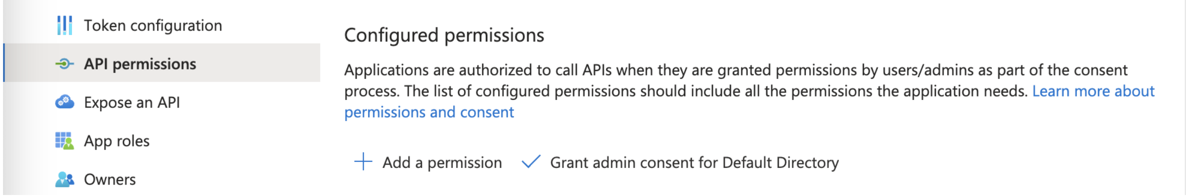

Applications are authorized to call APIs when they are granted permissions by users/admins as part of the consent process. The list of configured permissions should include all the permissions the application needs. Learn more about permissions and consent

Assign Roles and Access for your Terraform Illumio Onboarding Application

1. Navigate to the Subscription and copy the Subscription ID to save it in a text file. You'll need it later when you modify your Terraform script for onboarding Azure.

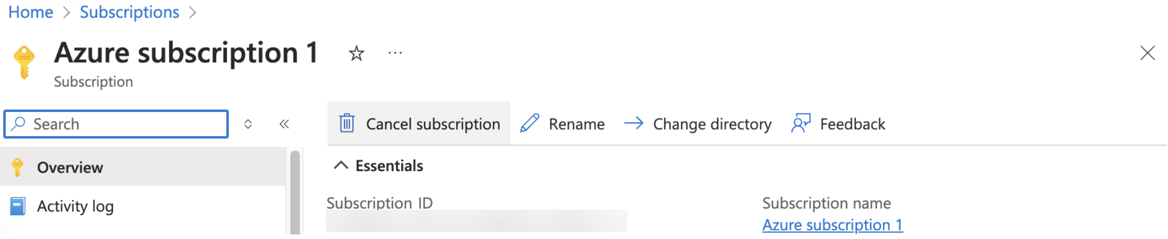

Subscription ID: [redacted]

Subscription name: Azure subscription 1

2. Click **Access Control (IAM)**.
3. Browse to Add > Add a role assignment.
4. Choose the following field values:

- Role tab: Privileged administrator roles: Owner
- Members tab: Assign access to: 'User, group, or service principal'

5. In the Members tab, click **Select Members**.
6. Enter the application name or ID and click **Select**.
7. Click **Save**.

   The application now includes the correct permissions with the correct identifiers and credentials. You will need these identifiers and credentials when you modify your Azure onboarding Terraform script. See Onboard an Azure Subscription using a Terraform Illumio Onboarding Application [78].

Onboard an Azure Subscription using a Terraform Illumio Onboarding Application

Learn how to onboard an Azure subscription using Terraform.

You must first create a Terraform Illumio Segmentation for the Cloud onboarding application. See Create a Terraform Illumio Onboarding Application for Azure [75].

Use your newly created application to onboard your Azure subscription with Illumio Segmentation for the Cloud. You need to onboard each subscription separately to onboard the entire tenant. To onboard your your Azure subscription with your newly created application, modify your Terraform script to resemble following script on GitHub.

Modify and run your Terraform script to onboard your subscription

Enter Your Saved Variable Values

The azure_subscription_dev module uses the following variables that you have saved to a text file as described in Create a Terraform Illumio Onboarding Application for Azure [75]. Modify your script so that it contains your saved variable values.

- azure_subscription_id
- azure_client_id
- azure_client_secret
- azure_tenant_id

Enter Your Variable Values Generated during Service Account Creation

The module also uses the following service account variables that you need to get from Illumio Segmentation for the Cloud when you create a service account:

- illumio_cloudsecure_client_id
- illumio_cloudsecure_client_secret

Create a service account in Illumio Segmentation for the Cloud and get the variable values:

1. Navigate to https://console.illum.io/ and sign in. See Activating your Account and Signing in [58].
2. Browse to to **Cloud** > **Settings** > **Service Accounts**.

3. Click **Add**, then enter a name and description, and click **Save**..

   The service account appears in your list with the name and client ID displayed. Copy the client ID.
4. Click on the service account to open its details page.
5. Under Secrets, click **Add**, then enter a name and click **Save**.

   A dialog appears with the client secret obscured. Copy the client secret.

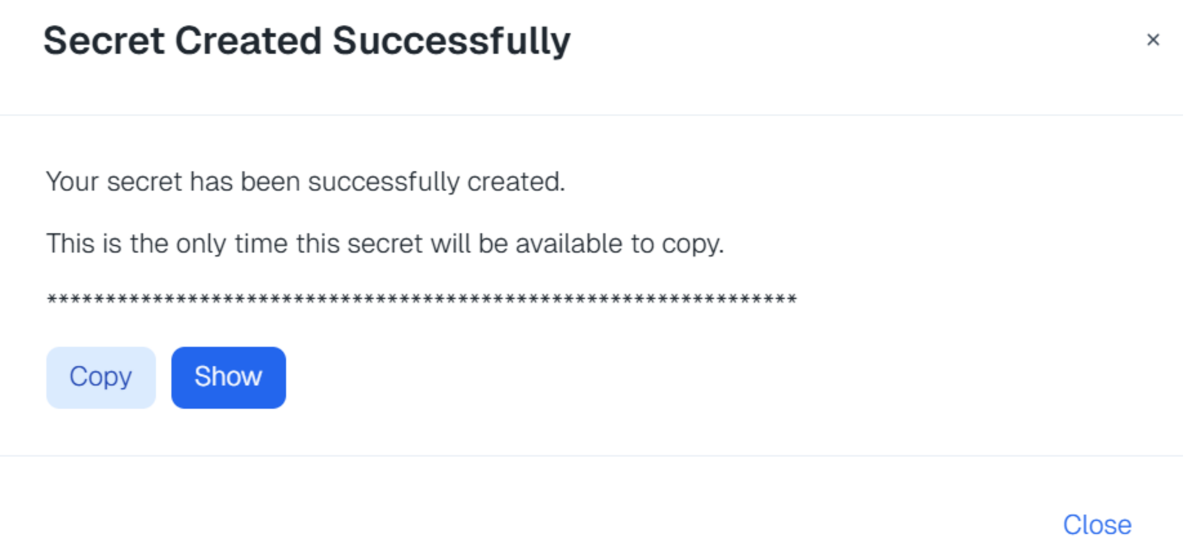

Secret Created Successfully

Your secret has been successfully created.

This is the only time this secret will be available to copy.

****

6. Modify your script so that it contains your copied variable values.

Run Your Terraform Script

1. Modify your Terraform script to specify read or read and write permissions.
2. Run your Terraform script.

   The Terraform module uses these credentials to create a new application in Azure. This application in turn gives your recently created application the permissions required for Illumio Segmentation for the Cloud to onboard your Azure subscription. The module checks for the permissions mode (read/read andwrite) to provision appropriate permissions.

What to do after onboarding Azure with Terraform

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

For the next steps after onboarding an account, see Onboarding Azure [59] and After onboarding your accounts [146].

Onboard an Azure Cloud tenant - default setup

Learn how to onboard an Azure tenant. Onboarding an Azure tenant allows you to connect all the subscriptions and resources under the tenant with Illumio Segmentation for the Cloud. Running the PowerShell script for Azure Tenant onboarding creates a new Entra ID application with the tenant scope, although this can be done manually on the Azure Portal. This service principal allows Illumio Segmentation for the Cloud to retrieve subscriptions and resources under the given tenant. After the Azure Entra ID application is created and the required Reader permission are set at the tenant scope, the PowerShell script automatically sends the necessary credentials (Client Id and Client Secret), although you can invoke the API via a PowerShell script to send them back to Illumio Segmentation for the Cloud. These credentials are required to communicate with your Azure tenant.

1. Review the prerequisites. See Prerequisites for Onboarding Azure [60].
2. Onboarding a tenant with the wizard automatically provides the required permissions. See Permissions for Onboarding Azure [61] for the list of required permissions.
3. Onboard your tenant using the default setup. See Onboard a Tenant with Illumio [80].

Onboard a tenant with Illumio

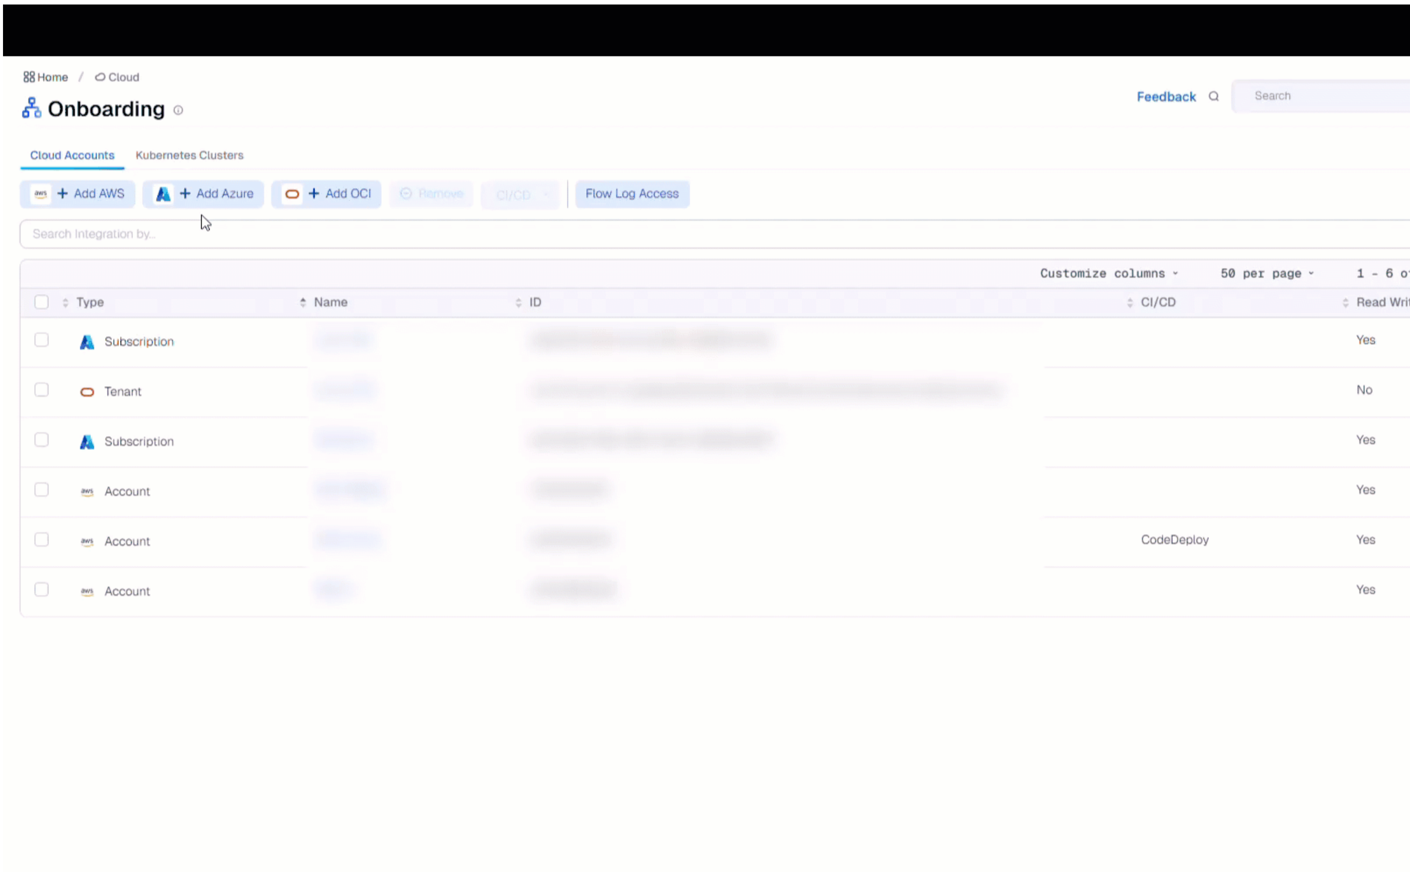

Type	Name	ID	CI/CD	Read Wri
Subscription	[illegible]	[illegible]		Yes
Tenant	[illegible]	[illegible]		No
Subscription	[illegible]	[illegible]		Yes
Account	[illegible]	[illegible]		Yes
Account	[illegible]	[illegible]	CodeDeploy	Yes
Account	[illegible]	[illegible]		Yes

1. If this is the first time logging in to Illumio Segmentation for the Cloud, click **+ Azure** on the Onboarding page to 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:** Specify a name for the account; this name is what appears in Illumio Segmentation for the Cloud upon onboarding. Use a descriptive name so that you can easily identify it.
   - **Tenant ID:** Paste this ID that you copied from Azure. The tenant ID is also known as the parent management group ID.
   - **Onboarding toggle options:**

Onboarding Toggle Option	Action
Onboard all subscriptions in the tenant: Yes	To onboard all member subscriptions along with the tenant, select this option.
Onboard all subscriptions in the tenant: No	To onboard only some subscriptions in the tenant, choose this option. Select those necessary for Illumio visibility and protection. Then go to the Onboarding page to onboard those subscriptions individually.
Illumio has Read and Write access to ensure compliance: Yes	To grant the Illumio Cross Account Role permission to view your Azure tenant resources and to apply policy to them, choose this option. To view the permissions you are granting Illumio Segmentation for the Cloud to your Azure subscription, click **Download Permissions**.
Illumio has Read and Write access to ensure compliance: No	To provide the Illumio Cross Account Role read-only access, choose this option. To view the permissions you are granting Illumio Segmentation for the Cloud to your Azure subscription, 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 callback API to return the Azure Entra ID app credentials back to Illumio Segmentation for the Cloud.
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. Run the PowerShell command in the Azure portal. This creates a new Entra ID application with the tenant scope, using the tenant ID you entered in the first step of the wizard.

   In summary, the PowerShell command does the following:
   - Creates an application registration called "Illumio-CloudSecure-Access"
   - Creates a custom role called "Illumio Network Security Administrator"
   - Assigns the application registration the role of "Storage Blob Data Reader"
   - Gives Illumio Segmentation for the Cloud your newly created application registrations, client, client id, and client secret
4. Ensure that all fields are completed like that shown in the following screen capture.

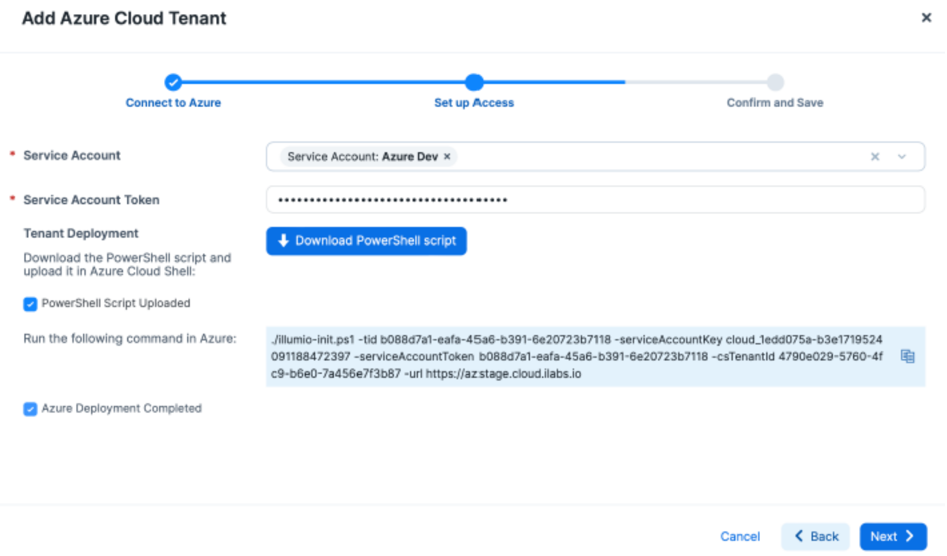

Add Azure Cloud Tenant

* Service Account: Service Account: **Azure Dev**

* Service Account Token: •••••••••••••••••••••••••••••••••••

Tenant Deployment

Download the PowerShell script and upload it in Azure Cloud Shell:

☑ PowerShell Script Uploaded

Run the following command in Azure:

```
./illumio-init.ps1 -tid b088d7a1-eafa-45a6-b391-6e20723b7118 -serviceAccountKey cloud_1edd075a-b3e1719524
091188472397 -serviceAccountToken b088d7a1-eafa-45a6-b391-6e20723b7118 -csTenantId 4790e029-5760-4f
c9-b6e0-7a456e7f3b87 -url https://az.stage.cloud.ilabs.io
```

☑ Azure Deployment Completed

5. In the Confirm and Save part of the wizard, review the details and click **Save and Confirm**.

Manually Create an Azure Entra ID App and Assign the Reader RBAC Role

If you did not use the PowerShell command as described above, use the following steps to perform the necessary work manually.

1. In a new browser window, open your Azure portal.
2. Create the Azure Entra ID App registration as mentioned in the step here: https://learn.microsoft.com/en-us/entra/identity-platform/howto-create-service-principal-portal#register-an-application-with-microsoft-entra-id-and-create-a-service-principal. Skip the redirect Uniform Resource Identifier (URI) in this step.
3. Once the Entra ID App registration is created, create a new client secret as documented in these steps: https://learn.microsoft.com/en-us/entra/identity-platform/howto-create-service-principal-portal#option-3-create-a-new-client-secret
4. Once the Entra ID App and secret are created, assign the Reader RBAC role to the App at the tenant scope: https://learn.microsoft.com/en-us/entra/identity-platform/howto-create-service-principal-portal#assign-a-role-to-the-application
5. Ensure that the role is assigned at the tenant scope and not at the subscription scope, as documented in the above steps.

Run the Callback API Call to Illumio Segmentation for the Cloud

After you connect to Azure and set up access, run the PowerShell callback to the Illumio Segmentation for the Cloud API to complete the tenant onboarding. If the callback is successful, no output is printed.

> **NOTE**
> The following code is just a reference PowerShell script (Web_request.ps1), so update it according to your environment, but ensure that the subscription ID is empty.

```
# Set your service account key ID, token, and client secret
$serviceAccountKeyId = "<YourServiceAccountKeyId>"
$serviceAccountToken = "<YourServiceAccountToken>"
$clientSecret = "<YourClientSecret>" # The actual client secret to be
encoded

# Combine the key ID and token with a colon and base64 encode for the
Authorization header
$authString = "$($serviceAccountKeyId):$($serviceAccountToken)"
$encodedAuthString =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($authString))

# Base64 encode the client secret separately
$encodedClientSecret =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($clientSecret))

# Construct the headers with the encoded Authorization header
$headers = @{
  "X-Tenant-Id"  = "<CloudsecureTenantId>"
  "Content-Type"  = "application/json"
  "Authorization" = "Basic $encodedAuthString"
}

# Construct the request body with the encoded client secret
$body = @{
  "type"            = "AzureRole"
  "client_id"       = "<ClientId>"
  "client_secret"   = $encodedClientSecret  # Use the base64 encoded client
secret
  "subscription_id" = "<SubscriptionId>" # remove this and use
azure_tenant_id if onboarding the entire tenant.
  "azure_tenant_id" = "<AzureTenantId>" # both azure tenant id and
subscription_id should be present for subscription onboarding.
} | ConvertTo-Json -Depth 10

# Send the POST request
$response = Invoke-RestMethod -Uri 'https://cloud.illum.io/api/v1/
integrations/cloud_credentials' -Method Post -Headers $headers -Body $body


# Output the response
Write-Host $response
```

Set up and enable Azure flow logs after onboarding

The **Onboarding** page opens and displays a new row for that tenant.

Now you set up and enable flow logs.

- To set up flow logs before enabling them, see Set up flow logs in your CSP environment [132].
- To enable flow logs, see Grant flow log access to your CSPs [141].

What's next after you onboard your tenant?

The **Onboarding** page opens and displays a new row for that tenant.

For the next steps after onboarding a tenant, set up and enable flow logs. See Onboarding Azure [59]. Once you set up and enable flow logs, see After onboarding your accounts [146].

If you originally set the permissions to read only, and wish to change them to read and write, see Change Azure permissions from read to read and write [93].

Updating your tenant

- After tenant onboarding is complete, it shows a list of subscriptions. If a subscription belonging to a tenant is onboarded before the tenant onboarding, it does not show in the tenant's list of subscriptions. To see a subscription that you onboarded prior to the tenant onboarding, you must delete the onboarded subscription. Upon tenant onboarding, it automatically syncs and onboards the subscription.
- You may someday need to update your service account principals due to expiry or other issues. If you must update your service account principals after onboarding, see Update Service Principals for Onboarded Azure Subscriptions and Tenants [315].

Onboard an Azure Cloud Tenant - Guided Setup

Learn how to onboard an Azure tenant if you cannot onboard the tenant using the default setup described in Onboard an Azure Cloud tenant - default setup [79].

NOTE

If you don't have permissions to run the PowerShell script for Azure tenant onboarding, you will not be able to use the default setup.

1. Review the prerequisites. See Prerequisites for Onboarding Azure [60].
2. Use your Azure console to provide the required permissions. See Permissions for Onboarding Azure [61] for a list of required permissions.

   To provide permissions in your Azure console, see Microsoft website.
3. Onboard your tenant using the guided setup. See Using the Guided Setup [84].

The method described here lets you manually create an application registration, assign permissions, and apply credentials.

When you onboard an Azure tenant, the service principal allows Illumio Segmentation for the Cloud to retrieve resources. After you create the Azure Entra ID application, set the required reader permissions at the tenant scope, and provide the client ID and client secret credentials, invoke the API using a PowerShell script. The credentials are required to communicate with your Azure tenant.

Using the Guided Setup

Connect to Azure with the Wizard

The steps you take in this first part of the Illumio onboarding wizard are mostly the same as described in Onboard an Azure Cloud tenant - default setup [79]. The exception is that that if you give Illumio Segmentation for the Cloud more than read-only access, you *must* download the permissions as described in Step 4.

1. If you are logging in for the first time, click **+ Azure** on the Onboarding page to onboard your first account.
2. If you've already onboarded other accounts, choose **Onboarding** from the left navigation. 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**

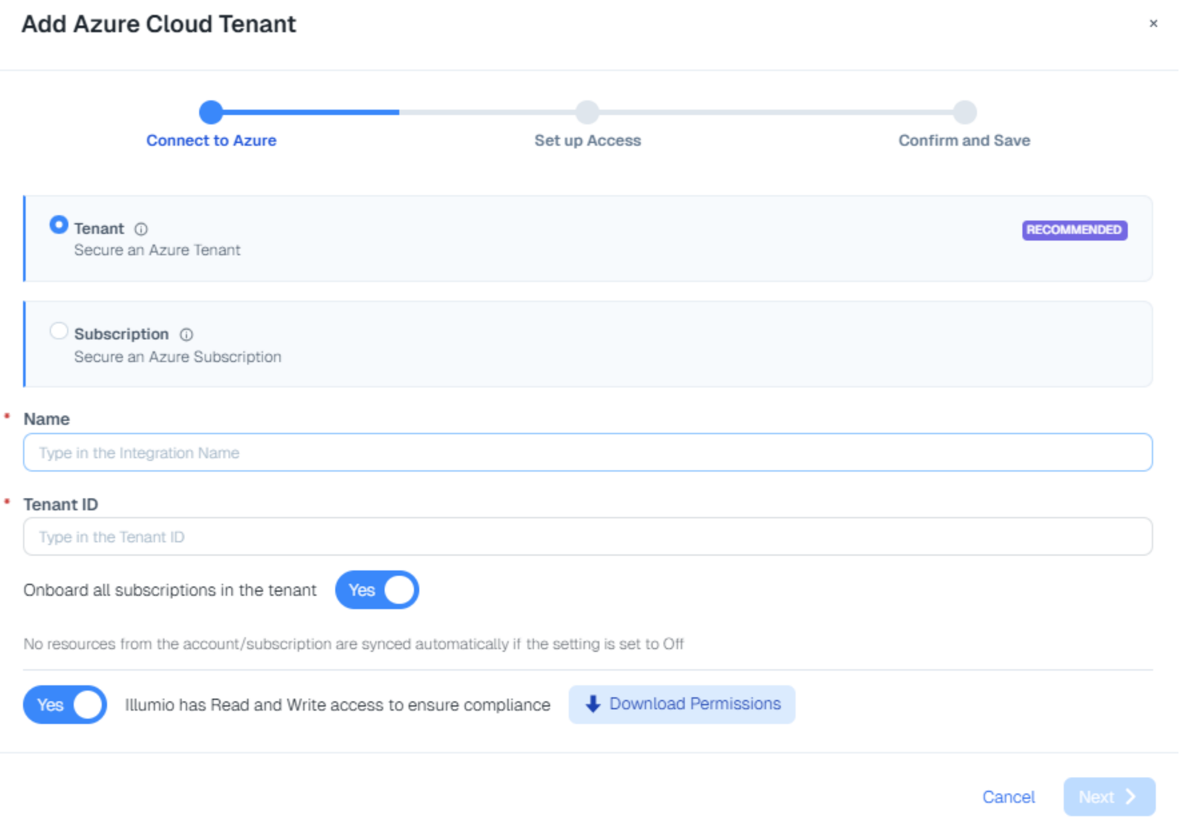

Add Azure Cloud Tenant

● Tenant — Secure an Azure Tenant — RECOMMENDED

○ Subscription — Secure an Azure Subscription

* Name: Type in the Integration Name

* Tenant ID: Type in the Tenant ID

Onboard all subscriptions in the tenant: Yes

No resources from the account/subscription are synced automatically if the setting is set to Off

Yes — Illumio has Read and Write access to ensure compliance

4. Provide the following information about your Azure account:
   - **Name:** Specify a name for the account; this name is what appears in Cloud upon onboarding. Use a descriptive name so that you can easily identify it in Illumio Segmentation for the Cloud.
   - **Tenant ID:** Paste the parent management group ID that you copied from Azure.
   - **Onboarding toggle options:**

Onboarding Toggle Option	Action
Onboard all subscriptions in the tenant: Yes	To onboard all member subscriptions along with the tenan.
Onboard all subscriptions in the tenant: No	To onboard only some subscriptions in the tenant, choose this option. Select those necessary for Illumio visibility and protection. Then go to the Onboarding page to onboard those subscriptions individually.
Illumio has Read and Write access to ensure compliance: Yes	To grant the Illumio Cross Account Role permission to view your Azure tenant resources and to apply policy to them.
Illumio has Read and Write access to ensure compliance: No	To provide the Illumio Cross Account Role read-only access, choose this option.

NOTE

To view the permissions you are granting Illumio Segmentation for the Cloud to your Azure tenant, click **Download Permissions**. This is *necessary* for you to onboard the tenant if you are unable to run the script, because you must manually provide the listed permissions. Save this list of permissions somewhere readily available.

Setup Access using the Wizard without the Script

1. Select an existing service account or create a new one by clicking **Add a new Service Account**. You can use spaces, underscores, numbers, or other characters (such as !,@,#, and so on) for the account name. Make it something you can remember.

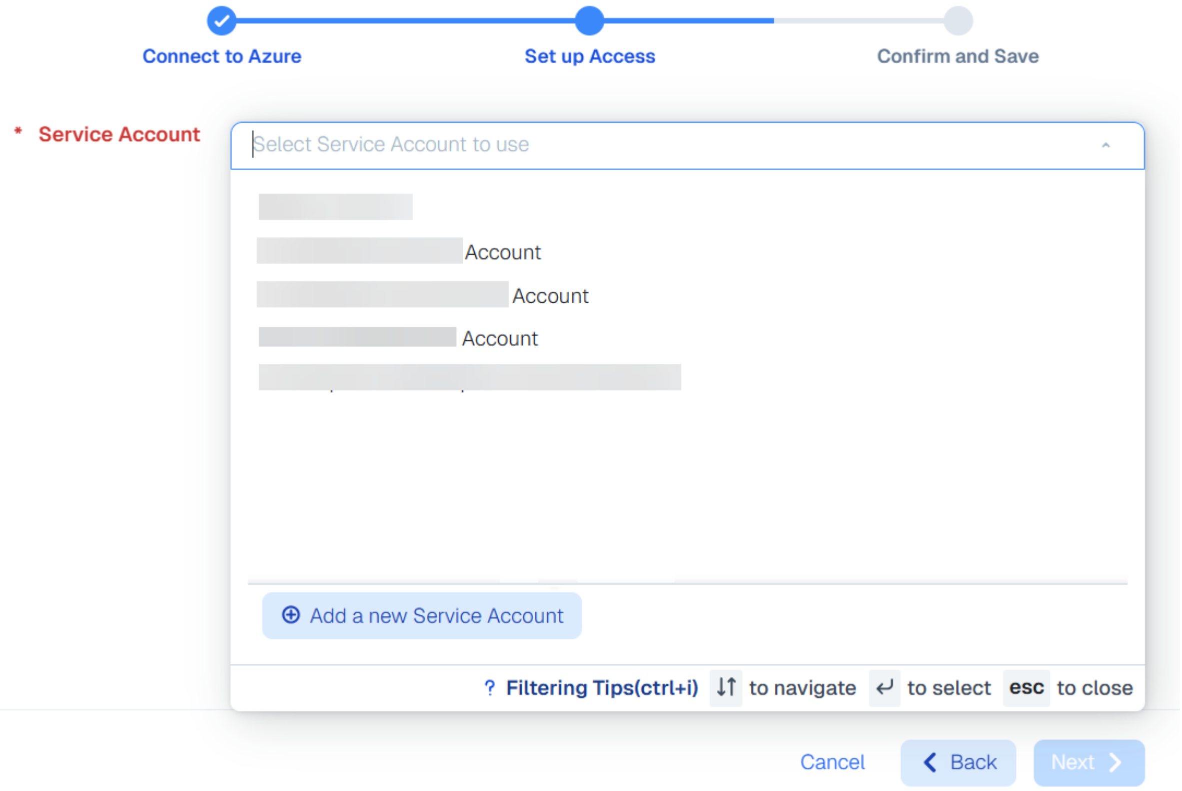

2. If you created a new service account, click **Download Credentials**. They are needed for the callback API to return the Azure Entra ID app credentials back to Illumio Segmentation for the Cloud.

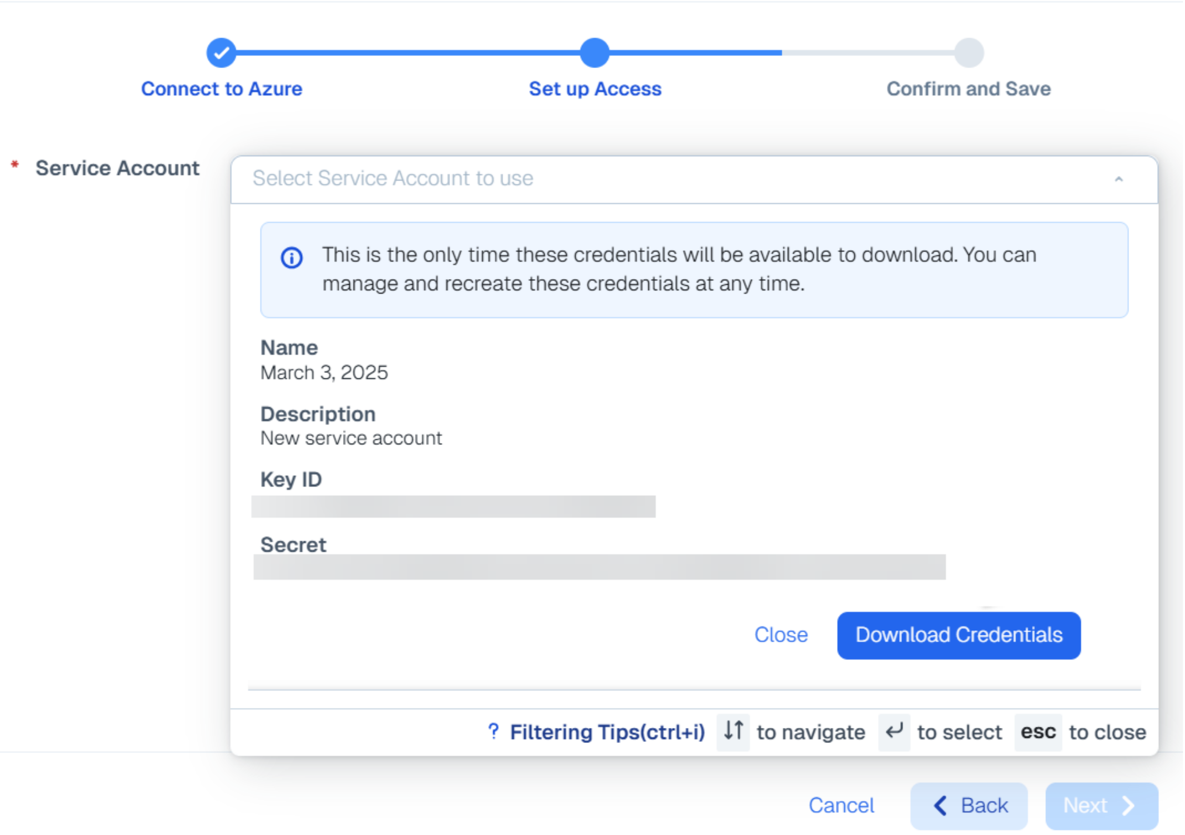

* Service Account: Select Service Account to use

This is the only time these credentials will be available to download. You can manage and recreate these credentials at any time.

Name
March 3, 2025

Description
New service account

Key ID
[redacted]

Secret
[redacted]

Close
Download Credentials
? Filtering Tips(ctrl+i) ↓↑ to navigate ↵ to select esc to close
Cancel
Back
Next

3. Enter the ServiceAccountToken.
4. Do not use the PowerShell script offered in the wizard because this onboarding method assumes you lack the permissions to run it. Instead, follow the steps below in Manually Create an Azure Entra ID App and Assign the Reader RBAC Role [88] to create and register your Azure application.

   This gives Illumio Segmentation for the Cloud your newly created Azure application registrations, client, client id, and client secret.
5. Complete the fields.

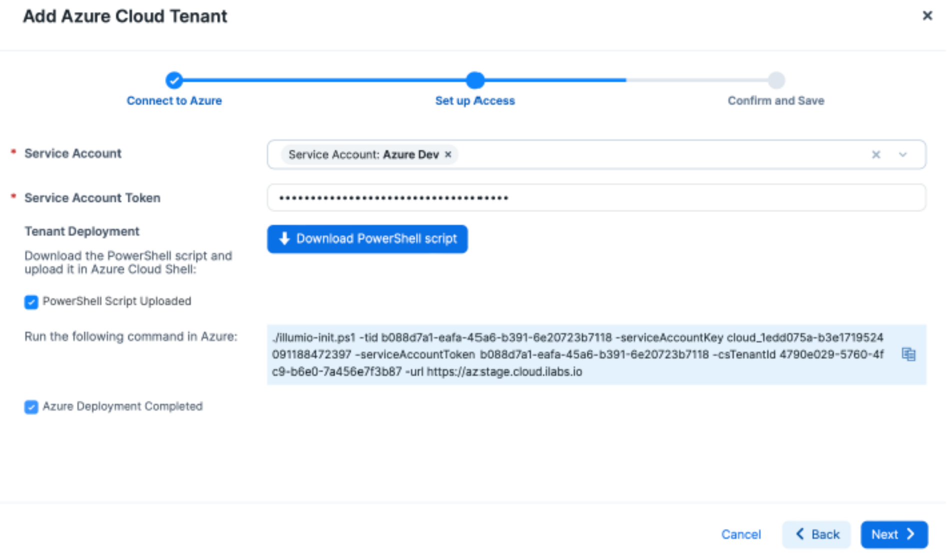

* Service Account: Service Account: **Azure Dev** ×

* Service Account Token: ••••••••••••••••••••••••••••••••••••

Tenant Deployment

Download the PowerShell script and upload it in Azure Cloud Shell:

☑ PowerShell Script Uploaded

Run the following command in Azure:

```
./illumio-init.ps1 -tid b088d7a1-eafa-45a6-b391-6e20723b7118 -serviceAccountKey cloud_1edd075a-b3e1719524
091188472397 -serviceAccountToken b088d7a1-eafa-45a6-b391-6e20723b7118 -csTenantId 4790e029-5760-4f
c9-b6e0-7a456e7f3b87 -url https://az.stage.cloud.ilabs.io
```

☑ Azure Deployment Completed

6. Review the details and click **Save and Confirm**.

Manually Create an Azure Entra ID App and Assign the Reader RBAC Role

Use these steps to create an Azure Entra ID App manually.

1. In a new browser window, open your Azure portal.
2. Create the Azure Entra ID App registration on the Microsoft website. Skip the redirect Uniform Resource Identifier (URI) in this step.
3. Once the Entra ID App registration is created, create a new client secret as documented on the Microsoft website.
4. Once the Entra ID App and secret are created, assign the Reader RBAC role to the App at the tenant scope, as documented on the Microsoft website.
5. Ensure that the role is assigned at the tenant scope and not at the subscription scope.
6. After you create the reader RBAC role, assign the API permissions for Illumio Segmentation for the Cloud. You can see the permissions in the .txt file you downloaded from the wizard.

Run the Callback API to Illumio Segmentation for the Cloud

After you connect to Azure and set up access, run the following PowerShell callback to the Illumio Segmentation for the Cloud API in your Azure console to complete the tenant onboarding. If the callback is successful, no output is printed.

NOTE

The following code is just a reference PowerShell script (Web_request.ps1), so update it according to your environment, but ensure that the subscription ID is empty.

```
# Set your service account key ID, token, and client secret
$serviceAccountKeyId = "<YourServiceAccountKeyId>"
$serviceAccountToken = "<YourServiceAccountToken>"
$clientSecret = "<YourClientSecret>" # The actual client secret to be
encoded

# Combine the key ID and token with a colon and base64 encode for the
Authorization header
$authString = "$($serviceAccountKeyId):$($serviceAccountToken)"
$encodedAuthString =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($authString))

# Base64 encode the client secret separately
$encodedClientSecret =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($clientSecret))

# Construct the headers with the encoded Authorization header
$headers = @{
  "X-Tenant-Id"  = "<CloudsecureTenantId>"
  "Content-Type"  = "application/json"
  "Authorization" = "Basic $encodedAuthString"
}

# Construct the request body with the encoded client secret
$body = @{
  "type"             = "AzureRole"
  "client_id"        = "<ClientId>"
  "client_secret"    = $encodedClientSecret  # Use the base64 encoded client
secret
  "subscription_id" = "<SubscriptionId>" # remove this and use
azure_tenant_id if onboarding the entire tenant.
  "azure_tenant_id" = "<AzureTenantId>" # both azure tenant id and
subscription_id should be present for subscription onboarding.
} | ConvertTo-Json -Depth 10

# Send the POST request
$response = Invoke-WebRequest -Uri 'https://cloud.illum.io/api/v1/
integrations/cloud_credentials' -Method Post -Headers $headers -Body $body


# Output the response
Write-Host $response
```

Set up and Enable Flow Logs after Onboarding

The **Onboarding** page opens and displays a new row for that tenant.

Set up and enable flow logs.

- To set up flow logs before enabling them, see Set up Flow Logs [132].
- To enable flow logs using the guided setup, see Grant Flow Log Access - Guided Set-up [90].

Grant Flow Log Access - Guided Setup

Run the following script in your Azure console to manually grant flow log access for your tenant. Note that "csTenantId" is the tenant ID and that "subscriptionId" is the subscription ID for Illumio Segmentation for the Cloud.

In summary, the script does the following:

- Creates an application registration called "Illumio-CloudSecure-Access"
- Creates a custom role called "Illumio Network Security Administrator"
- Assigns the application registration the role of "Storage Blob Data Reader"
- Gives Illumio Segmentation for the Cloud your newly created application registrations, client, client id, and client secret

```
#params for running the script
# to run the script use the command
# ./grant_flow.ps1 -serviceAccountKey <key> -serviceAccountToken
<token> -csTenantId <cs_tenant_id> -subscriptionId <subscription_id>
-storageAccount <storage_account1,storage_account2,...> -url <url>
# default url is https://cloud.illum.io
param(
    $serviceAccountKey = "",
    $serviceAccountToken = "",
    $csTenantId = "",
    $subscriptionId = "",
    [String[]]$storageAccounts,
    $url = "")

#check if the given inputs are valid
if ([string]::IsNullOrEmpty($serviceAccountKey) -or
[string]::IsNullOrEmpty($serviceAccountToken)) {
    Write-Host "Service Account Token or Key cannot be empty"
-ForegroundColor Red
    exit
}

if ([string]::IsNullOrEmpty($csTenantId)) {
    Write-Host "csTenantId cannot be empty" -ForegroundColor Red
    exit
}

if ([string]::IsNullOrEmpty($subscriptionId)) {
    Write-Host "subscriptionId cannot be empty" -ForegroundColor Red
    exit
}


if ($storageAccounts.Count -eq 0) {
    Write-Host "storage accounts list cannot be empty" -ForegroundColor Red
    exit
}

# Combine the key ID and token with a colon and base64 encode for the
Authorization header
$authString = "$($serviceAccountKey):$($serviceAccountToken)"
$encodedAuthString =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($authString))


# Construct the headers with the encoded Authorization header
$headers = @{
    "X-Tenant-Id"   = $csTenantId;
    "Content-Type"  = "application/json";
    "Authorization" = "Basic $encodedAuthString";
}

# Construct the request body with the storage accounts
# list of storage accounts which has flow logs and for which access has
been granted.
```

```
# the storage account names must be the entire namespace as shown in
cloudsecure flow logs page.
# Eg: /subscriptions/<subscription_id>/resourceGroups/<resource_group_name>/
providers/Microsoft.Storage/storageAccounts/<storage_account_name>

$body = @{
    "type"            = "AzureFlow";
    "subscription_id" = $subscriptionId; # subscription in which the
storage accounts are present
    "destinations"     = $storageAccounts;
} | ConvertTo-Json -Depth 10

# Send the POST request
$endPoint = "/api/v1/integrations/cloud_credentials"

if ([string]::IsNullOrEmpty($url) ) {
    Write-Host "url is empty hence using default url https://
cloud.illum.io" -ForegroundColor Yellow
    $url = "https://cloud.illum.io"
}

$url = $url + $endPoint

Write-Host "making API call to $($url)"

$response = Invoke-WebRequest -Uri $url -Method Post -Headers $headers
-Body $body

# Output the response
if ($($response.StatusCode -eq 200)) {
    Write-Host "API call to cloudsecure successful"
} else {
    Write-Host "Error making API call to cloudsecure. Status code: $
(response.StatusCode)"
}
```

What's next after you onboard your tenant?

The **Onboarding** page opens and displays a new row for that tenant.

If you originally set the permissions to read only, and wish to change them to read and write, see Change Azure permissions from read to read and write [93].

Updating your tenant

- After tenant onboarding is complete, it shows a list of subscriptions. If a subscription belonging to a tenant is onboarded before the tenant onboarding, it does not show in the tenant's list of subscriptions. To see a subscription that you onboarded prior to the tenant onboarding, you must delete the onboarded subscription. Upon tenant onboarding, it automatically syncs and onboards the subscription.
- You may someday need to update your service account principals due to expiry or other issues. If you must update your service account principals after onboarding, see Update Service Principals for Onboarded Azure Subscriptions and Tenants [315].

Change Azure permissions from read to read and write

Learn how to change your Azure permissions without re-onboarding your tenant or subscription.

> **NOTE**
> Use Case: You onboarded an Azure tenant or subscription with read-only permissions. You have now decided that you want Illumio Segmentation for the Cloud to have both read and write permissions so that it can apply policies to your tenant or subscription.

Prerequisites

You will need the following information before proceeding:

- clientId: To find and copy it, open Illumio Segmentation for the Cloud and browse to **Cloud > Onboarding > <the onboarding that you want to upgrade from read to read and write>**.
- Tenant ID: Copy this from Azure if you are onboarding either a subscription or a tenant. The tenant ID is also known as the parent management group ID.
- Subscription ID: Copy this from Azure if you are onboarding a subscription.
- Service Account Key: Copy this from wizard in the following steps.
- Service Account Secret (token): Copy this from wizard in the following steps.

Change to read and write - default method

Use this method if you have permissions to run the PowerShell script used to onboard Azure tenants and subscriptions.

1. Use the Azure tenant and/or subscription ID you copied to proceed through the Illumio Segmentation for the Cloud default Azure onboarding until you get to the wizard step that presents you with the onboarding PowerShell script. See Onboarding Azure [59].
2. If you do not have your original service account secret (token) readily available, create a new service account to get a secret. Enter the secret (token) to display the PowerShell script in the wizard.
3. Copy and save the PowerShell script so that you can modify it. Do not save the new integration (in other words, do not proceed further through the onboarding wizard).
4. Modify the PowerShell script by adding the following two parameters to the script:
   - `clientId` (you copied this value as described in Prerequisites [93])
   - `-nsg` (this flag adds the read and write permission)

   The PowerShell script to change a *tenant* to read and write should look something like this:

```
 Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-
west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -tid <azure-tenant-id>
-serviceAccountKey <illumio-service-account-key> -serviceAccountToken
<illumio-service-account-secret> -csTenantId <CloudSecure-tenant-id>
-url https://cloud.illum.io -nsg -clientId <client-id>
```

The PowerShell script to change a *subscription* to read and write should look something like this:

```
 Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-
west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -sid <subscription-id>
-serviceAccountKey <illumii-service-account-key> -serviceAccountToken
<illumio-service-account-token> -csTenantId <cloudsecure-tenant-id>
-url https://cloud.illum.io -nsg -clientId <client-id>
```

5. Save your changes to the PowerShell script and run it in the Azure portal.
6. Create a support ticket for Illumio to enable the read and write mode from the Illumio Segmentation for the Cloud end.

Change to read and write - guided method

Use this method if you don't have permissions to run the PowerShell script used to onboard Azure tenants and subscriptions.

1. Manually enable the read and write permissions for the Azure Active Directory service principal created during onboarding. See Permissions for Onboarding Azure [61] and Update Service Principals for Onboarded Azure Subscriptions and Tenants [315]. Use the clientId you copied in the prerequisites section to search for the Azure application you need to update.
2. Create a support ticket for Illumio to enable the read and write mode from the Illumio end.

Onboarding AWS Cloud

Review the onboarding workflow for your cloud environment:

1. Before you begin, review the prerequisites and permissions.
   - See Prerequisites for Onboarding AWS [94].
   - See Permissions for Onboarding AWS [95].
2. Use the wizard to onboard accounts and organizations.
   - See Onboard an AWS Cloud account [113].
   - See Onboard an AWS Cloud organization [107].
3. Set up flow logs and grant access.
   - See Set up flow logs in your CSP environment [132].
   - See Grant flow log access to your CSPs [141].

After you onboard your AWS organizations or accounts, you can visualize your resources, define your public cloud environments, and create policies. See After onboarding your accounts [146].

Prerequisites for Onboarding AWS

This is a list of things to have ready before you begin onboarding your AWS accounts or organizations.

Before you begin onboarding AWS

- □ You need the ability to log into an AWS account. The onboarding wizard flow assumes that you are already logged into an AWS account.
- □ You need the ability to create an IAM role in your AWS account and assign it permissions
- □ You need to know your account IDs that you want to onboard. If you are onboarding an organization, this will include the root account ID. You will need to specify your account IDs in the wizard.
- □ If you are restricting public access to flow logs, you need to make certain ports and IP addresses available to Illumio Segmentation for the Cloud. See AWS flow log access IP addresses [292].
- □ If onboarding an account, but not an organization, determine the method you want to use for onboarding the account, whether by using Illumio to launch the CloudFormation Stack or by using an Illumio-provided YAML file as a template to manually create the stack

Required AWS permissions

Onboarding requires certain permissions. Use the steps described in Onboard an AWS Cloud account [113] and Onboard an AWS Cloud organization [107] to automatically provision the permissions.

See Permissions for Onboarding AWS [95].

Permissions for Onboarding AWS

This page describes the set of required permissions that are created when onboarding AWS as described in Onboard an AWS Cloud account [113] and Onboard an AWS Cloud organization [107].

AWS IAM Permissions

To onboard your AWS account, you will need to use the CloudFormation Stack to create an IAM role within your AWS account, which Illumio assumes to make API calls. This role must be granted permissions to specific AWS resources for Cloud to provide visibility and manage policies for those resources. It is important to note that this 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.

Read and Write Permissions by Service and Category

Service	Category	Resource Types
Read and Write (IllumioCloudAWSProtectionPolicy)		
EC2	Network Security	DBSecurity Group, Network ACL, Security Group, Security Group Rule
RDS	Network Security	DB Security Group
Read (IllumioCloudAWSIntegrationPolicy)		
CodeDeploy	Infrastructure	Application, Deployment Group
DirectConnect	Network Routing	Connection, Gateway, Lag, Virtual Interface
DocumentDB	Database	Cluster
DynamoDB	Database	Table
EC2	Compute	Instance, Spot Fleet Request, Spot Instance Request
EC2	Network Management	EIP, Network Interface, Subnet, VPC, VPC Peering
EC2	Network Monitoring	Flow Log
EC2	Network Routing	Carrier Gateway, Customer Gateway, Egress Only Internet Gateway, Instance Connect Endpoint, Internet Gateway, Nat Gateway, Route Table, Transit Gateway, Transit Gateway Attachment, Transit Gateway Route Table, Transit Gateway Multicast Domain, VPC Endpoint, VPC Endpoint Service, VPN, VPN Connection, VPN Gateway
EC2	Network Security	Security Group
EC2	Storage	Volume
ECS	Containers	Cluster, Container Instance
EKS	Containers	Addon, Cluster, Fargate Profile, Node Group,
Elasticache	Database	Cache Cluster
ElasticLoadBalancingV2	Network Routing	Load Balancer
Glacier	Storage	Vault
Lambda Function	Serverless	Function
IAM	Account Management	Account, User
KMS	Security Infrastructure	Key
MemoryDB	Database	Cluster
Network Manager	Network Routing	Global Network, Core Network, Connect Attachment, VPC Attachment, Site To Site VPN Attachment, Transit Gateway Route Table Attachment, Transit Gateway Peering, Transit Gateway Registration

Service	Category	Resource Types
RAM	Resource Management	Resource Share
RDS	Database	DB Cluster, DB Instance, DBSecurityGroup
Redshift	Data warehouse	Cluster
S3	Storage	Bucket, Bucket Policy
Target Groups	Network Routing	Target Group

IAM Role Configuration

To facilitate access to your AWS environment, you must create an IAM role within your AWS account. This role must be assigned the following policies:

- **SecurityAudit (managed by AWS) and IllumioCloudAWSIntegrationPolicy:** Permissions in these policies are required to read the resources in your AWS account.
- **IllumioCloudAWSProtectionPolicy:** Permissions in this policy are required 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:

```
READ ONLY Policy

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*'
            - 'docdb-elastic:GetCluster'
            - 'docdb-elastic:ListTagsForResource'
            - 'dynamodb:List*'
            - 'dynamodb:Describe*'
            - 'ec2:Describe*'
            - 'ec2:SearchTransitGatewayMulticastGroups'
            - 'ecs:Describe*'
            - 'ecs:List*'
        - 'eks:DescribeAddon'
            - 'eks':ListAddons'
            - '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'
            - 'networkmanager:ListCoreNetworks'
            - 'networkmanager:GetCoreNetwork'
            - 'networkmanager:ListAttachments'
            - 'networkmanager:GetVpcAttachment'
            - 'networkmanager:GetSiteToSiteVpnAttachment'
            - 'networkmanager:GetConnectAttachment'
            - 'networkmanager:GetTransitGatewayRouteTableAttachment'
            - 'networkmanager:ListPeerings'
            - 'networkmanager:GetTransitGatewayPeering'
            - 'networkmanager:GetTransitGatewayRegistrations'
```

Write Policy

The following items are AWS IAM write permissions that you will need to grant to the Illumio AssumeRole.

```
READ ONLY Policy

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*'
            - 'docdb-elastic:GetCluster'
            - 'docdb-elastic:ListTagsForResource'
            - 'dynamodb:List*'
            - 'dynamodb:Describe*'
            - 'ec2:Describe*'
            - 'ec2:SearchTransitGatewayMulticastGroups'
            - 'ecs:Describe*'
            - 'ecs:List*'
            - 'eks:DescribeAddon'
            - 'eks':ListAddons'
            - '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'
            - 'networkmanager:ListCoreNetworks'
            - 'networkmanager:GetCoreNetwork'
            - 'networkmanager:ListAttachments'
            - 'networkmanager:GetVpcAttachment'
            - 'networkmanager:GetSiteToSiteVpnAttachment'
            - 'networkmanager:GetConnectAttachment'
            - 'networkmanager:GetTransitGatewayRouteTableAttachment'
            - 'networkmanager:ListPeerings'
            - 'networkmanager:GetTransitGatewayPeering'
            - 'networkmanager:GetTransitGatewayRegistrations'

WRITE Policy
- PolicyName: IllumioCloudAWSProtectionPolicy
  PolicyDocument:
    Version: 2012-10-17
    Statement:
      - Effect: Allow
        Resource:
          - 'arn:aws:ec2:*:*:security-group-rule/*'
          - 'arn:aws:ec2:*:*:security-group/*'
          - 'arn:aws:ec2:*:*:network-acl/*'
        Action:
          - 'ec2:AuthorizeSecurityGroupIngress'
          - 'ec2:RevokeSecurityGroupIngress'
          - 'ec2:UpdateSecurityGroupRuleDescriptionsIngress'
          - 'ec2:AuthorizeSecurityGroupEgress'
          - 'ec2:RevokeSecurityGroupEgress'
          - 'ec2:UpdateSecurityGroupRuleDescriptionsEgress'
          - 'ec2:ModifySecurityGroupRules'
          - 'ec2:DescribeTags'
          - 'ec2:CreateTags'
          - 'ec2:DeleteTags'
          - 'ec2:DescribeNetworkAcls'
          - 'ec2:CreateNetworkAclEntry'
          - 'ec2:ReplaceNetworkAclEntry'
          - 'ec2:DeleteNetworkAclEntry'
          - 'ec2:ModifyNetworkInterfaceAttribute'
          - 'ec2:CreateSecurityGroup'
          - 'ec2:DeleteSecurityGroup'
          - 'ec2:DescribeSecurityGroups'
```

AWS Resource Type Permissions

When you add permissions to the Illumio AssumeRole as described in Permissions for Onboarding AWS, you will need to ensure that your permissions for resource types match those seen in the .json file example. Make sure to frequently check with your Support or Customer Success teams to make sure you have updated your .json file. Each time Illumio Segmentation for the Cloud supports a new resource type, you will need to update your resource permissions.

FLOW READ Policy

```
's3:ListBucket'

's3:ListBucketVersion'

's3:GetBucketLocation'

's3:GetObject'
```

Service Accounts and IAM Roles for AWS

The following information is important to understanding how Illumio interacts with AWS.

Service Accounts in the Illumio Segmentation for the Cloud Context

Within the Illumio Segmentation for the Cloud platform, a "service account" refers to an account used by Illumio Segmentation for the Cloud to interact with its own services (Illumio Segmentation for the Cloud services) rather than directly with your AWS services. This account is primarily used for internal operations within Illumio Segmentation for the Cloud, such as making API calls to the Illumio Segmentation for the Cloud 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, Illumio Segmentation for the Cloud requires the creation of an identification and access management (IAM) role within the customer's AWS account. Illumio Segmentation for the Cloud 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.

Handling encrypted AWS VPC flow logs

If service-side encryption with KMS (SSE-KMS) keys is enabled for the S3 bucket, Cloud requires additional permissions for the log service to be added to the KMS key before enabling flow logs.

To allow the log service to write VPC Flow Logs in the designated S3 bucket, the AWS Logs Delivery System must be granted permission to the Encrypt, Decrypt, ReEncrypt, Generate-

DataKey*, and Describe key on the key that is used to encrypt the data in the S3 bucket. Below is an example policy showing the necessary permissions in place for the key policy.

```
{

"Sid": "Allow Log Delivery to use the key",

"Effect": "Allow",

"Principal": {

"Service": "delivery.logs.amazonaws.com"

},

"Action":

"kms:Encrypt",

"kms:Decrypt",

"kms:ReEncrypt*",

"kms:GenerateDataKey*",

"kms:DescribeKey"

],

 "Resource": "*"

"Condition": {

"StringEquals": {

"aws:SourceAccount": "<account-id>"

},

"ArnLike": {

"aws:SourceArn": "arn:aws:logs:<region>:<account-id>:*"

}

}

}
```

To read flows stored in encrypted buckets, the Assume Role requires access to the key used for encrypting the contents of the S3 bucket. This key decrypts the contents of the S3 bucket. The following is the policy document required to gain access to the key and decrypt the flow logs. Adding this permission automatically allows the Assume Role, created during on-boarding, to decrypt the contents of the bucket (In this case, the flow logs). No additional settings are required.

```
{

"Version": "2012-10-17",

"Statement":[

{

"Effect": "Allow",

"Action":

"kms:Decrypt"

],

"Resource": [

"arn:aws:kms:<region>:<account-id>:key/<key-id>" // Replace with your KMS
key ARN

]

}

]

}
```

The following CloudFormation Template gets the Assume Role ARN and the KMS Key ARN as input and grants the decrypt permission on the KMS Key to the Assume Role.

```
AWSTemplateFormatVersion: "2010-09-09"

Description: "Grant Decrypt permission on KMS key for CloudSecure's Assume
Role"

Parameters:

    IAMRoleName:

        Type: String

        Description: IAM Role name used by Cloud.

    KMSKeyARNs:

        Type: CommaDelimitedList

        Description: List of KMS Key ARNs.

Resources:

      IllumioKMSDecryptPolicy:

        Type: 'AWS::IAM::Policy'

        Properties:

            PolicyName: IllumioKMSDecrypt

            PolicyDocument:

                Version: 2012-10-17

                Statement:

                    - Effect: Allow

                    Sid: IllumioKMSKeyAccess

                    Action:

                    - 'kms:Decrypt'

                    Resource: !Ref KMSKeyARNs

        Roles

            - !Ref IAMRoleName
```

For more information, see the AWS website.

AWS flow logs

For a list of ports and IP addresses required for flow log access, see AWS Flow Log Access IP Addresses.

Supported Flow Log Fields

Illumio Segmentation for the Cloud 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} ${srcport} ${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 Illumio Segmentation for the Cloud fetches the flow logs depends on your configuration (e.g., a central account or multiple accounts)

- When flow log access is first enabled in Illumio Segmentation for the Cloud, there's a 15-minute latency until traffic flows are first displayed in Cloud Map, Traffic and Inventory pages. After this initial latency, traffic flows are periodically updated every two minutes.

Updating AWS permissions on the Assume Role

Illumio updates permissions required for the Assume Role on a continuous basis. Use these steps to provide permissions for the newly added resources.

1. Download the permissions that are provided in the first part of the wizard. Depending on whether you chose read-only or read and write, be sure to download the correct file below.
   - Read and write
   - Read-only
2. Run the CloudFormation Stack (CFT).
3. Login to the AWS console of account to which you need to update the permissions to run the CloudFormation stack.
4. Under services click **CloudFormation**.
5. Click **Create stack**.
6. In the Choose template page, select template ready and upload a template file option, and upload the downloaded template and click **Next**.
7. 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.

8. In the IAMRoleName box, enter the name of the assume role created in AWS when onboarding with Illumio Segmentation for the Cloud. By default, the name is IllumioCloudIntegrationRole. Click **Next**.
9. If you gave 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.)
10. 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.

AWS permissions background

When you start the onboarding process and begin creating IAM roles from the Illumio Segmentation for the Cloud 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 Illumio Segmentation for the Cloud

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 Illumio from synchronizing resources.

AWS 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.

Onboard an AWS Cloud organization

Learn to onboard an AWS organization. An AWS organization is a service AWS provides that allows you to consolidate multiple accounts into an organization and manage them centrally. Onboard an AWS organization to take advantage of Illumio Segmentation for the Cloud security features and minimize an attacker's lateral movement.

1. Review the prerequisites. Prerequisites for Onboarding AWS [94].
2. Onboarding a subscription with the wizard automatically provides the required permissions. See Permissions for Onboarding AWS [95].
3. Onboard your organization. See Onboard AWS Organizations in Illumio [108].

AWS organizations and root accounts

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 Illumio Segmentation for the Cloud, all the accounts under the root will be onboarded (provided the user runs the StackSet). Cloud 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.

Onboard AWS organizations in Illumio

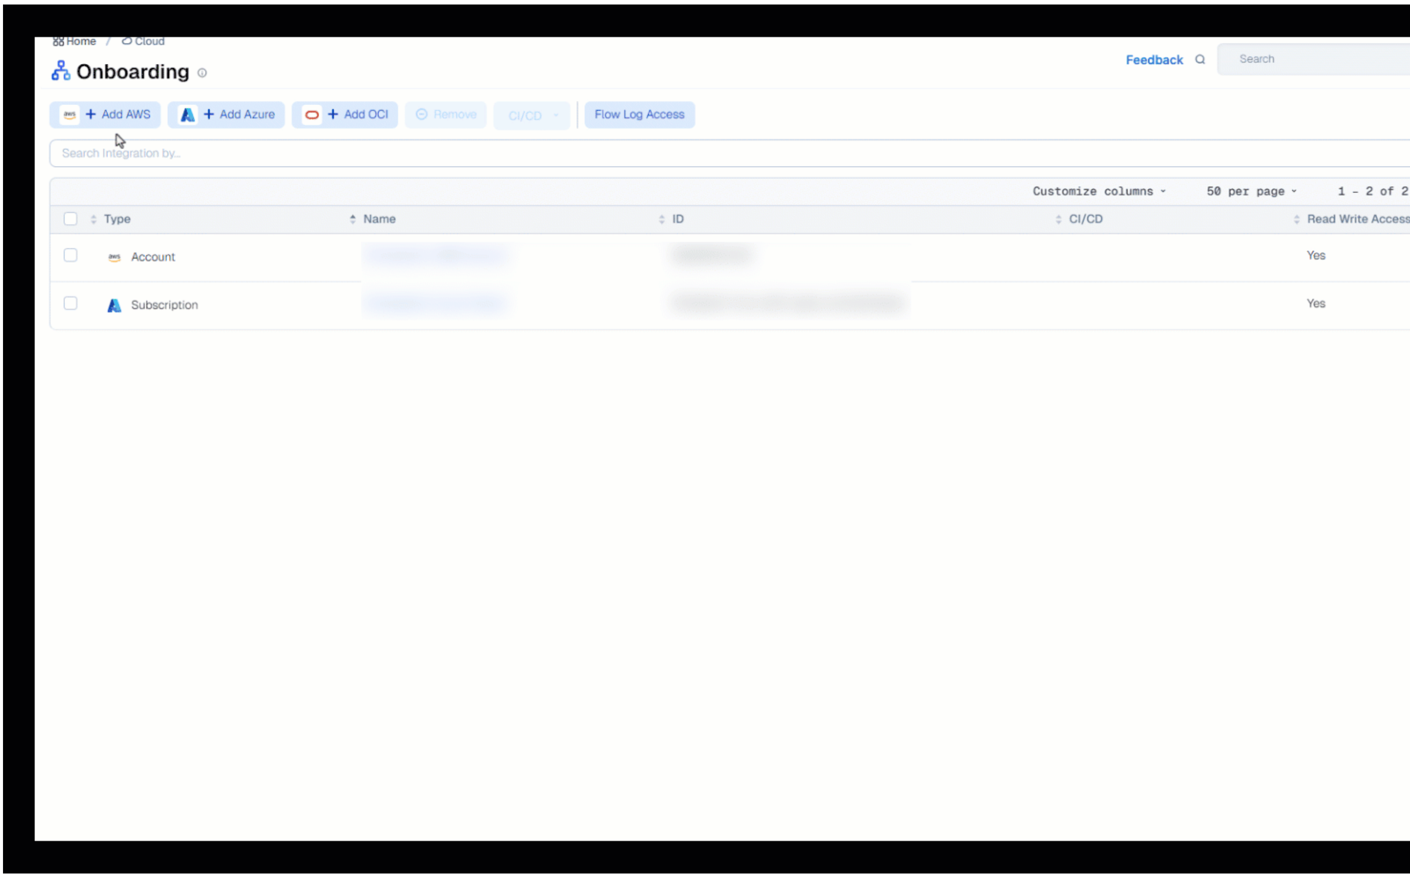

□	Type	Name	ID	CI/CD	Read Write Access
□	Account				Yes
□	Subscription				Yes

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 Illumio. The name should be descriptive so that you can easily identify it in Illumio.
   - The AWS ID of the root account you are onboarding into Illumio

   NOTE

   The page contains a toggle below the Account ID field to specify the type of access Illumio 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 Illumio 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 Illumio. Illumio 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 Cloud Tenant Id. It is 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 stackset in AWS. Cloud 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**.

   Cloud 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 [111].

In order to create the Assume role and provide Illumio Segmentation for the Cloud 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 Cloud YAML file provided by Illumio is a valid stack template file.)
5. Upload the Illumio Segmentation for the Cloud 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 CloudTenantId 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 Illumio Segmentation for the Cloud with the RoleARN, ExternalId, OrgId, and MasterAccountId. The RoleARN and ExternalId will be used by Illumio Segmentation for the Cloud to connect with the account and sync resources. The OrgId and MasterAccountId will be used by Illumio Segmentation for the Cloud 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 Cloud script will notify Cloud that the stack was successfully created and Cloud will detect that the organization was onboarded and begin synchronizing the organization resources with Cloud. 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 [110] before performing these steps.

In order to create the Assume role and provide Illumio Segmentation for the Cloud 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 Illumio Segmentation for the Cloud YAML file and click **Next**. (The Cloud 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 CloudTenantId 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 Illumio Segmentation for the Cloud 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 Illumio Segmentation for the Cloud with the RoleARN, ExternalId, OrgId, and MasterAccountId. The RoleARN and ExternalId will be used by Illumio Segmentation for the Cloud to connect with the account and sync resources. The OrgId and MasterAccountId will be used by Illumio Segmentation for the Cloud 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 Cloud script will notify Cloud that the stack was successfully created and Cloud will detect that the organization was onboarded and begin synchronizing the organization resources with Cloud. Clicking the organization in the Onboarding page will let you see the accounts under it.

What's next after onboarding your organization?

For the next steps after onboarding organization, see Onboarding AWS Cloud [94] and After onboarding your accounts [146].

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 AWS organization 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 Cloud and click **Delete**.

   Initially the stack deletion will fail. The CloudFormation template provided by Cloud 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. See the AWS website.

   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 Cloud.
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 [113]. Once that is done, repeat the instructions to Remove the stackset in AWS [113].

   Once the stacks are completely removed, select the stackset again and choose **Delete StackSet** from the Actions drop-down menu.

Onboard an AWS Cloud account

Onboard an AWS account to take advantage of Illumio Segmentation for the Cloud security features and minimize an attacker's lateral movement.

1. Review the prerequisites. Prerequisites for Onboarding AWS [94].
2. Onboarding a subscription with the wizard automatically provides the required permissions. See Permissions for Onboarding AWS [95].
3. Onboard your organization. See Onboard AWS by running a CloudFormation stack [114], Onboard AWS using a stack template [116], and Onboard AWS in concert with CodeDeploy [117].

Ways to onboard your AWS account

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 Cloud 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 Cloud to create and run the CloudFormation stack, Cloud 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 Cloud to run the stack.

If you wish, you can incorporate AWS CodeDeploy as described in the instructions when you onboard AWS accounts using any of the above methods. See Onboard AWS in concert with CodeDeploy [117].

Onboard AWS by running a 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 a stack template [116].

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 Illumio. The name should be descriptive so that you can easily identify it in Illumio.
   - The AWS account ID of the account you are onboarding into Cloud

   NOTE

   The page contains a toggle below the Account ID field to specify the type of access Cloud 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 Cloud to your AWS account, click **Download Permissions**.

> **NOTE**
>
> This page contains a CI/CD Integration toggle for enabling CodeDeploy. This is optional, but you will want to select **Yes** if you wish to make use of the AWS CodeDeploy feature as described in Onboard AWS in concert with CodeDeploy [117]. The toggle position will default to whatever setting you pick in Settings > CI/CD Integration. You can also select multiple, onboarded AWS accounts in the Onboarding page and click **CI/CD > Disable/Enable** to disable or enable CodeDeploy.

When done completing your 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 Cloud. Cloud 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. Cloud 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**. Cloud 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 Cloud. 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 Cloud 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 Cloud and a temporary Lambda function named `LambdaExecutionRoleIllumioCloudAPICall`. The Lambda function passes back to Cloud two credentials:

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

Now, Cloud 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 Cloud for your account, see Prerequisites for Onboarding AWS [94].

e. Leave the AWS console and return to Cloud.
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.

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

Onboard AWS using a 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**.

   Cloud 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 Illumio Segmentation for the Cloud 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 Cloud 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 Cloud script will notify Cloud that the stack was successfully created and Cloud will detect that account was onboarded and begin synchronizing the account resources with Cloud. A new row for that account appears in the Onboarding page.

Onboard AWS in concert with CodeDeploy

Illumio Segmentation for the Cloud uses your CodeDeploy configuration when onboarding an account. You can choose to opt-in to CodeDeploy auto-discovery at anytime.

AWS CodeDeploy is a deployment service that automates application deployments to Amazon services. The Illumio Segmentation for the Cloud integration with CodeDeploy automatically onboards all the applications and deployments defined in AWS into Illumio Segmentation for the Cloud. This allows you to onboard quickly in order to gain insights and visibility into the application and deployment traffic for analysis and segmentation.

Additionally, Illumio Segmentation for the Cloud lets you use AWS to:

- Auto-discover your AWS CodeDeploy applications that already exist, and include them as Illumio Segmentation for the Cloud applications
- Auto-discover your Illumio Segmentation for the Cloud environments (development, staging, production, etc.)
- Visualize application drift between your Illumio Segmentation for the Cloud environments for security review

Keep in mind the following notes:

- When new applications and deployments are pulled from CodeDeploy, Illumio Segmentation for the Cloud auto-approves these definitions by default
- Illumio Segmentation for the Cloud uses CodeDeploy as the source of truth for the application and deployment definitions. Any changes made in CodeDeploy to an application or deployment are synced and updated into Illumio Segmentation for the Cloud.
- Illumio Segmentation for the Cloud does not support bi-directional syncing. Illumio Segmentation for the Cloud pulls from CodeDeploy to maintain a consistent state across all applications and deployments.

- When the CodeDeploy integration is disabled, Illumio Segmentation for the Cloud stops syncing with AWS but the applications and deployments are not removed from Illumio Segmentation for the Cloud

Naming Conventions

When Illumio Segmentation for the Cloud syncs the application and deployment names from AWS, it converts the names into the following string format:

```
{ApplicationName}-{AWS_Account_ID}-{Region}
```

For example:

```
CodeDeployAppName-0123456789101-US-West-2
```

You can create Illumio Segmentation for the Cloud names for the applications and deployments that CodeDeploy creates. Once you create the Illumio Segmentation for the Cloud-specific name, you can search for the application or deployment using this name moving forward.

Tag Display

Illumio Segmentation for the Cloud displays AWS tags and group tags, including those associated with the CodeDeploy application or deployment. These are not editable through the Illumio Segmentation for the CloudUI. You can edit these tags in the AWS console.

Use CodeDeploy

Refer to the steps above in Onboard AWS by running a CloudFormation stack [114] and Onboard AWS using a stack template [116].

Remove the AWS 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 Cloud and click **Delete**.

   Initially the stack deletion will fail. The CloudFormation template provided by Cloud 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. See the Amazon website.

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

What's next after onboarding your AWS account?

For the next steps after onboarding an account, see Onboarding AWS Cloud [94] and After onboarding your accounts [146].

Onboarding GCP

Review the onboarding workflow for your cloud environment:

GCP onboarding workflow

If you have the permissions to run the onboarding Cloudshell script, follow the default setup for your GCP projects and organizations.

1. Before you begin, review the prerequisites and permissions.
   - See Prerequisites for onboarding GCP [119].
   - See Permissions for onboarding GCP [320].
2. Use the wizard to onboard projects, organizations, and folders.
   - See Onboard a GCP project [120].
   - See Onboard a GCP organization [122].
   - See Onboard a GCP folder [125].
3. Set up flow logs and grant access.
   - See Prerequisites for granting flow log access to your CSPs [140].
   - See Set up Flow Logs [132].
   - See Grant flow log access to your CSPs [141].

After you onboard your GCP organizations or projects, you can visualize your resources, define your public cloud environments, and create policies. See After onboarding your accounts [146].

Prerequisites for onboarding GCP

Review these prerequisites before you begin onboarding your GCP organizations or projects.

Before you begin onboarding GCP

Once you review these prerequisites, return to Onboarding GCP [119] for next steps.

☐ Log into a GCP account. The onboarding wizard flow assumes that you are already logged into a GCP account.

☐ The default installation assumes that you have enabled the APIs for all the services in your GCP projects, irrespective of whether you onboard them separately or as part of an organization.

☐ The default installation assumes that you have Read/ReadWrite permissions for the following: See Permissions for onboarding GCP [320].
   - Assigning the following IAM roles:

  - roles/iam.securityReviewer (This read only role can be truncated.)
  - roles/compute.viewer (The full read only role is required.)
  - roles/browser (This read only role is required for organization and folder onboarding only. It reads organization projects and folders.)
  - roles/cloudasset.viewer (The full read only role is required.)
- Assigning custom roles:
  - IllumioPubSubFlowLogAccess (The full role is required.)
  - illumio_write_role (The full role is required.)
  - illumio_api_enable_role (The full role is required.)
- Creating a GCP service account and assigning it impersonation permissions

☐ Know your organization ID, project ID, and Role Name.

☐ If you are restricting public access to flow logs, you need to make certain ports and IP addresses available to Illumio Segmentation for the Cloud. See GCP Flow Log Access IP Addresses [319].

☐ If you are restricting public access to flow logs, make certain ports and IP addresses available to Illumio Segmentation for the Cloud.

Required GCP permissions

See Permissions for onboarding GCP [320].

Onboard a GCP project

Onboard a Google Cloud Platform (GCP) project to take advantage of Illumio Segmentation for the Cloud security features and minimize an attacker's lateral movement.

1. Review the prerequisites. See Prerequisites for onboarding GCP [119].
2. Onboard your project. Onboarding a project with the wizard automatically provides the required permissions. See Permissions for onboarding GCP [320].

Onboard a GCP project with Illumio

NOTE

Enabling APIs is optional, but Illumio Segmentation for the Cloud functionality is affected if you don't enable APIs. See Permissions for onboarding GCP [320] for a list of supported services and enable their corresponding service APIs.

NOTE

Note the following about service accounts and projects:

- You do not need to specify a project to assign a service account to because the service account will be assigned to the project that you provided.
- To confirm that a service account has been created, after you run the onboarding script, navigate to the Service Accounts page for the project that you assigned the service account to.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/GCP+Onboarding+Project.mp4

1. If this is the first time logging in to Illumio Segmentation for the Cloud, click **+ GCP** on the Onboarding page to onboard your first account.

   If you've already onboarded other accounts, choose **Onboarding**. The Onboarding page appears. Click **+Add GCP**.

   The **Add GCP Cloud Project** wizard starts and displays the first step: **Connect to GCP**.
2. Provide the following information about your GCP account:
   - **Name:** Specify a name for the account; this name appears in Illumio Segmentation for the Cloud. Make the name descriptive so that you can easily identify it.
   - **Organization ID:** Paste this ID that you copied from GCP.
   - **Project ID:** Paste this ID that you copied from GCP.
   - **Onboarding toggle options:**
     - Read/Write Access:
       - Illumio has Read and Write access to ensure compliance (choose **Yes**)

         This grants the Illumio Cross Account Role permission to view your GCP organization resources and to apply policy to them, choose this option. To view the permissions you are granting Illumio Segmentation for the Cloud to your GCP organizations, click **Download Permissions**.
       - Illumio has Read and Write access to ensure compliance (choose **No**)

         This grants the Illumio Cross Account Role read-only access.
3. Click **Next**.

Next, 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 are needed for the Cloudshell script to return the GCP credentials back to Illumio Segmentation for the Cloud.
2. Enter the ServiceAccountToken in the appropriate field.
3. Provide a Role Name from the GCP console. The Illumio onboarding script creates an IAM role to add access permissions.

   The Project Deployment command field populates with a command to run the `gcp_onboarding_prod.sh` script in GCP. Illumio securely hosts the script so that it can run during the onboarding process. The command automatically appends the IDs, role, service account name, and secret from the first step of the wizard.

   In summary, the command does the following:

   - Creates a GCP service account
   - Enables APIs (script asks if you want to)

   NOTE

   All service APIs are enabled at the project level. This is because projects serve as "containers" for resources, so enabling APIs at the project level allows you to control which services are available, to assign certain IAM roles to the project, and to prevent unauthorized access to services that shouldn't be used.

- Creates an IAM role with the appropriate permissions based on Read/ReadWrite mode
- Binds the IAM role to the service account
- Binds pre-defined IAM roles:
  - roles/iam.securityReviewer (This read only role can be truncated.)
  - roles/compute.viewer (The full read only role is required.)
  - roles/browser (This read only role is required for organization and folder onboarding only. It reads organization projects and folders and allows Illumio Segmentation for the Cloud to view the account structure.)
  - roles/cloudasset.viewer (The full read only role is required.)
- Grants impersonation permission to the Illumio service account
- Sends the service account email to the Illumio endpoint
- Sends the IDs, role, service account name, key, and secret to Illumio Segmentation for the Cloud.

4. To the left of the 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 GCP console and paste the copied command in the Cloudshell prompt window to run it.

   The command provides Illumio the information and permissions necessary to onboard your project.
6. Leave your GCP console and return to Illumio Segmentation for the Cloud. 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 GCP, and click **Next**.
8. The final step of the wizard appears. This step displays a summary of the GCP project information you just specified for onboarding.
9. Click **Save and Confirm**.

Set up and enable flow logs after onboarding your GCP project

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

For the next steps after onboarding an project, set up and enable flow logs. See Set up flow logs in your CSP environment [132]. Once you set up and enable flow logs, see After onboarding your accounts [146].

Onboard a GCP organization

Onboard a Google Cloud Platform (GCP) organization to take advantage of Illumio Segmentation for the Cloud security features and minimize an attacker's lateral movement.

1. Review the prerequisites. See Prerequisites for onboarding GCP [119].
2. Onboard your organization. Onboarding an organization with the wizard automatically provides the required permissions. See Permissions for onboarding GCP [320].

Onboard a GCP organization with Illumio

> **NOTE**
>
> Enabling APIs is optional, but Illumio Segmentation for the Cloud functionality is affected if you don't enable APIs. See Permissions for onboarding GCP [320] for a list of supported services and enable their corresponding service APIs.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/GCP+Cloud+Organization.mp4

1. If this is the first time logging in to Illumio Segmentation for the Cloud, click **+ GCP** on the Onboarding page to onboard your first account.

   If you've already onboarded other accounts, choose **Onboarding**. The Onboarding page appears. Click **+Add GCP**.

   The **Add GCP Cloud Organization** wizard starts and displays the first step: **Connect to GCP**.
2. Provide the following information about your GCP account:
   - **Name:** Specify a name for the account; this name appears in Illumio Segmentation for the Cloud. Make the name descriptive so that you can easily identify it.
   - **Organization ID:** Paste this ID that you copied from GCP.
   - **Onboarding toggle options:**
     - Project Onboarding:
       - Onboard all Projects in this Organization (choose **Yes**)

         This onboards all member projects along with the organization.
       - Onboard all Projects in this Organization (choose **No**)

         This does not onboard any projects in the organization. Go to the Onboarding page to onboard projects individually.
     - Read/Write Access:
       - Illumio has Read and Write access to ensure compliance (choose **Yes**)

         This grants the Illumio Cross Account Role permission to view your GCP organization resources and to apply policy to them, choose this option. To view the permissions you are granting Illumio Segmentation for the Cloud to your GCP organizations, click **Download Permissions**.
       - Illumio has Read and Write access to ensure compliance (choose **No**)

         This grants the Illumio Cross Account Role read-only access.
3. Click **Next**.

Next, 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 are needed for the Cloudshell script to return the GCPcredentials back to Illumio Segmentation for the Cloud.
2. Enter the ServiceAccountToken in the appropriate field.
3. Provide a Project ID for the GCP service account. Even if you are not onboarding projects with the organization, Illumio requires a project ID to assign to the GCP service account.

The Organization Deployment command field populates with a command to run the `gcp_onboarding_prod.sh` script in GCP. Illumio securely hosts the script so that it can run during the onboarding process. The command automatically appends the IDs, role, service account name, and secret from the first step of the wizard.

In summary, the command does the following:

- Creates a GCP service account
- Enables APIs

> **NOTE**
>
> - If APIs have been consented to, when you onboard a GCP folder or organization, the Service Usage API is enabled. The Service Usage API allows to list enabled and disabled APIs and services. If the APIs for supported resources are not enabled, Illumio cannot fetch inventory.

- Creates an IAM role with the appropriate permissions based on Read/ReadWrite mode
- Binds the IAM role to the service account
- Binds pre-defined IAM roles:
  - roles/iam.securityReviewer (This read only role can be truncated.)
  - roles/compute.viewer (The full read only role is required.)
  - roles/browser (This read only role is required for organization and folder onboarding only. It reads organization projects and folders and allows Illumio Segmentation for the Cloud to view the account structure.)
  - roles/cloudasset.viewer (The full read only role is required.)
- Grants impersonation permission to the Illumio service account
- Sends the service account email to the Illumio endpoint
- Sends the Project ID, Organization ID, and GCP service account email to Illumio Segmentation for the Cloud.

4. To the left of the 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 GCP console and paste the copied command in the Cloudshell prompt window to run it.

   The command provides Illumio the information and permissions necessary to onboard your organization.
6. Leave your GCP console and return to Illumio Segmentation for the Cloud. 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 GCP, and click **Next**.
8. The final step of the wizard appears. This step displays a summary of the GCP organization information you just specified for onboarding.
9. Click **Save and Confirm**. Your child projects are now onboarded.

Set up and enable flow logs after onboarding your GCP organization

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

For the next steps after onboarding an organization, set up and enable flow logs. See Set up flow logs in your CSP environment [132]. Once you set up and enable flow logs, see After onboarding your accounts [146].

Onboard a GCP folder

Onboard a Google Cloud Platform (GCP) folder to take advantage of Illumio Segmentation for the Cloud security features and minimize an attacker's lateral movement.

1. Review the prerequisites. See Prerequisites for onboarding GCP [119].
2. Onboard your folder. Onboarding a folder with the wizard automatically provides the required permissions. See Permissions for onboarding GCP [320].

Onboard a GCP folder with Illumio

> **NOTE**
>
> Enabling APIs is optional, but Illumio Segmentation for the Cloud functionality is affected if you don't enable APIs. See Permissions for onboarding GCP [320] for a list of supported services and enable their corresponding service APIs.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/GCP+Onboarding+Folder.mp4

1. If this is the first time logging in to Illumio Segmentation for the Cloud, click **+ GCP** on the Onboarding page to onboard your first account.

   If you've already onboarded other accounts, choose **Onboarding**. The Onboarding page appears. Click **+Add GCP**.

   The **Add GCP Cloud Folder** wizard starts and displays the first step: **Connect to GCP**.
2. Provide the following information about your GCP account:
   - **Name:** Specify a name for the account; this name appears in Illumio Segmentation for the Cloud. Make the name descriptive so that you can easily identify it.
   - **Folder ID:** Paste this ID that you copied from GCP.
   - **Organization ID:** Paste this ID that you copied from GCP.
   - **Onboarding toggle options:**
     - Project Onboarding:
       - Onboard all Projects in this Folder (choose **Yes**)

         This onboards all member projects along with the folder.
       - Onboard all Projects in this Folder (choose **No**)

         This does not onboard any projects in the folder. Go to the Onboarding page to onboard projects individually.
     - Read/Write Access:
       - Illumio has Read and Write access to ensure compliance (choose **Yes**)

         This grants the Illumio Cross Account Role permission to view your GCP folder resources and to apply policy to them, choose this option. To view the permissions you are granting Illumio Segmentation for the Cloud to your GCP folders, click **Download Permissions**.
       - Illumio has Read and Write access to ensure compliance (choose **No**)

         This grants the Illumio Cross Account Role read-only access.

3. Click **Next**.

Next, 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 are needed for the Cloudshell script to return the GCPcredentials back to Illumio Segmentation for the Cloud.
2. Enter the ServiceAccountToken in the appropriate field.
3. Provide a Project ID for the GCP service account. Even if you are not onboarding projects with the folder, Illumio requires a project ID to assign to the GCP service account.

   The Folder Deployment command field populates with a command to run the `gcp_onboarding_prod.sh` script in GCP. Illumio securely hosts the script so that it can run during the onboarding process. The command automatically appends the IDs, role, service account name, and secret from the first step of the wizard.

   In summary, the command does the following:

   - Creates a GCP service account
   - Enables APIs (script asks if you want to)

   > **NOTE**
   >
   > If APIs have been consented to, when you onboard a GCP folder or organization, the Service Usage API is enabled. The Service Usage API allows Illumio to list enabled and disabled APIs and services. If the APIs for supported resources are not enabled, Illumio cannot fetch inventory.

   - Creates an IAM role with the appropriate permissions based on Read/ReadWrite mode
   - Binds the IAM role to the service account
   - Binds pre-defined IAM roles:
     - roles/iam.securityReviewer (This read only role can be truncated.)
     - roles/compute.viewer (The full read only role is required.)
     - roles/browser (This read only role is required for organization and folder onboarding only. It reads organization projects and folders and allows Illumio Segmentation for the Cloud to view the account structure.)
     - roles/cloudasset.viewer (The full read only role is required.)
   - Grants impersonation permission to the Illumio service account
   - Sends the service account email to the Illumio endpoint
   - Sends the Project ID, Folder ID, and GCP service account email to Illumio Segmentation for the Cloud.
4. To the left of the 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 GCP console and paste the copied command in the Cloudshell prompt window to run it.

   The command provides Illumio the information and permissions necessary to onboard your folder.
6. Leave your GCP console and return to Illumio Segmentation for the Cloud. 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 GCP, and click **Next**.
8. The final step of the wizard appears. This step displays a summary of the GCP folder information you just specified for onboarding.

9. Click **Save and Confirm**. Your child projects are now onboarded.

Set up and enable flow logs after onboarding your GCP folder

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

For the next steps after onboarding an folder, set up and enable flow logs. See Set up flow logs in your CSP environment [132]. Once you set up and enable flow logs, see After onboarding your accounts [146].

Onboarding OCI

Review the onboarding workflow for your cloud environment:

1. Before you begin, review the prerequisites and permissions.
   - See Prerequisites for onboarding OCI [127].
2. Use the wizard to onboard your OCI cloud tenants.
   - See Onboard an OCI tenant [129].
3. Set up flow logs and grant access.
   - See Set up Flow Logs [132].
   - See Grant OCI flow log access [131].

After you onboard your OCI tenants, you can visualize your resources, define your public cloud environments, and create policies. See After onboarding your accounts [146].

Prerequisites for onboarding OCI

Overview of OCI onboarding prerequisites

The following information is important to understanding how Illumio interacts with OCI.

For a list of ports and IP addresses required for flow log access, see OCI Flow Log Access IP Addresses [343].

OCI onboarding checklist

- Access to Cloud
- Access to the OCI Console
- The user must have an IAM management policy in OCI Cloud. (The Illumio Segmentation for the Cloudonboarding script runs Terraform to create a group, a user for Illumio Segmentation for the Cloud, and add permissions to the group.)
- The OCI tenant ID and home region of the OCI root tenant

Oracle Cloud Stack

The Oracle Cloud Stack is a feature that allows you to automate the creation of multiple cloud resources as a single unit, called the stack. Oracle Cloud lets you use Terraform to create stacks and manage resources. Illumio Segmentation for the Cloud makes use of this Stack feature to create the resources that are required to interact with Oracle cloud.

Oracle IAM users

A user is an identity created in OCI's Identity and Access Management (IAM) service that represents a person or an application that interacts with OCI services. Users allow for the authentication and authorization of individuals or entities to access and manage OCI resources in accordance with assigned permissions. API keys are created for a user, which can be used for API/SDK access over the resources in the OCI tenant.

Illumio Segmentation for the Cloud creates a new user when the stack is created, and adds an API key to the user. This API key is be used in to communicate with OCI tenant, synchronize the resources, and read flows.

Oracle IAM groups

An Oracle IAM group is a collection of users. Groups allow you to efficiently manage access permissions for multiple users at once, rather than needing to manage permissions for each user individually. By assigning users to groups, you can apply policies to the group as a whole, granting or revoking privileges to all members of the group simultaneously.

Illumio Segmentation for the Cloud creates an IAM group and adds the user to the group and write IAM policies.

Oracle IAM policies

OCI's IAM policies specify who has what type of access to your OCI resources. They play a crucial role in securing your OCI environment by granting precise permissions to users and groups, determining how they can interact with OCI resources. After creating the group, add the permissions mentioned in the Illumio-required policies [128] section to access the resources.

Illumio-required policies

Illumio Segmentation for the Cloud requires the following onboarding policies:

```
"Allow group <groupname> to inspect all-resources in tenancy",
```

```
"Allow group <groupname> to read network-security-groups in tenancy",
```

```
Allow group <groupname> to read security-lists in tenancy",
```

```
"Allow group <groupname> to read serviceconnectors in tenancy",
```

```
"Allow group <groupname> to read load-balancers in tenancy",
```

Illumio Segmentation for the Cloud requires the following flow policies:

```
Allow group <groupname> to read objects in tenancy where all {target.bucket.name = '<bucket>', any{request.permission='OBJECT_INSPECT', request.permission='OBJECT_READ'}}
```

Terraform-created resources

Terraform creates the following resources during onboarding:

- A group with the following name format<username>-group
- A policy document, adding it to the group
- A user, adding it to the group
- An API key with the public key appended to the script

During flow access enablement, Terraform creates a policy document allowing access to the destinations for the group created during onboarding.

Onboard an OCI tenant

This topic explains how to onboard an OCI tenant. Before onboarding, see Prerequisites for onboarding OCI [127].

Background for onboarding an OCI tenant

An OCI tenant is a service Oracle provides that allows you to consolidate multiple compartments and manage them centrally. The hierarchy of OCI is as follows:

- Tenant - The parent container for all accounts. It consists of compartments
- Compartment- The standard OCI account that contains the OCI resources

When the OCI tenant is onboarded into Illumio Segmentation for the Cloud, all the compartments (accounts) are onboarded, up to six parent-child levels deep. Illumio Segmentation for the Cloud supports onboarding tenants. It does not support onboarding individual compartments.

Onboarding of an OCI tenant is a two-step process.

1. Run a Terraform script on a root account.
2. Use the information to populate Illumio Segmentation for the Cloud onboarding dialog fields.

Onboard OCI Tenants in Cloud

The following instructions describe how to begin the tenant onboarding sequence in Illumio Segmentation for the Cloud.

Connect to OCI

The following instructions describe how to begin the tenant onboarding sequence in Illumio Segmentation for the Cloud .

1. Launch the onboarding wizard in either of the following ways:
   - Click **+ OCI** in the Onboarding page to onboard your first tenant when you sign in for the first time
   - From the left navigation, choose **Onboarding** and click **+ Add OCI** at the top of the page
2. Provide the following information about your OCI tenant:
   - Name for the tenant
   - This name is what appears in Illumio Segmentation for the Cloud. The name should be descriptive so that you can easily identify it.
   - The Root Tenancy/Compartment OCID of the root account you are onboarding. It might look something like `ocid1.tenancy.oc1..xxxxxxxyz1a2b3c...`.

- The home region

  This is the geographic area that applies to your tenant. Select one from the list.

3.

> **NOTE**
>
> The page contains a toggle to specify the type of access Illumio Segmentation for the Cloud has to your OCI tenant. To view the permissions you are granting Illumio Segmentation for the Cloud to your OCI tenant, click **Download Permissions**.
>
> The write feature is in BETA.
>
> By participating in the BETA program for OCI features you agree that your company's use of the BETA version of OCI features will be governed by Illumio's Beta Terms and Conditions.

4. Click **Next**.

   The wizard advances to step two: **Set up Access**.
5. Click **Download Terraform File** to get the .zip file containing the necessary terraform scripts.

   Before you proceed in the onboarding wizard, you first need to open the OCI console and perform some steps.

Running the Terraform Scripts in the OCI Console

1. Open the OCI Console at https://cloud.oracle.com. From the menu, navigate to **Developer Services > Resource Manager > Stacks** and click **Create Stack**.
2. Select **My configuration,** and in the stack, configuration click the **.Zip** file radio button, and upload the cs_connector.zip file.

   This will auto populate the Name for the stack.
3. Provide a description if needed, and make sure that the root compartment is selected under the Create in Compartment option. Leave the rest of the defaults if desired, and click **Next**.
4. In the Configuration variables page, all the values will be auto populated. If needed, the username can be changed. Click **Next**.'
5. Verify all the values in the review page and, in the Run apply on the created stack option, make sure to select the Run Apply check box and click **Create**. The stack will run and create the required resources in the OCI console.
6. Once the stack completes running, select the output page and copy the values from the following fields:

- User OCID. It might look something like `ocid1.user.oc1..xxxxxxxyz1a2b3c...`.
- Group Name. It might look something like `<username>-group`.
- API Fingerprint. It might look something like `12:34:56:78:90:ab:cd:ef:12:34:56:78:90:ab:...`.

Now you will return to the Illumio Segmentation for the Cloud onboarding wizard.

Set up Access

1. Click the **Terraform script was successfully run** check box.
2. Paste the outputs from your OCI console into the following fields and click **Next**:

- User OCID. It might look something like `ocid1.user.oc1..xxxxxxxyz1a2b3c....`
- Group Name. It might look something like `<username>-group`.
- API Fingerprint. It might look something like `12:34:56:78:90:ab:cd:ef:12:34:56:78:90:ab:....`

The final step of the onboarding wizard (Confirm and Save) appears.

Confirm and Save

1. 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.
2. To edit the account information, such as the name and read/write access, click the account in the Onboarding page and click **Edit**.

Next steps after onboarding your OCI tenant

For the next steps after onboarding your OCI tenant including enabling access to flow logs and viewing traffic, see Onboarding OCI [127] and After onboarding your accounts [146].

Remove the OCI tenant integration

You can delete the integration for a given organization by selecting the it in the Onboarding page and clicking **Remove > Remove**.

Once the OCI onboarding is removed from Illumio Segmentation for the Cloud, open the OCI console, navigate to the stack details, and click the **Destroy** button. Once the access is destroyed, select **More actions > Delete stack**. This will completely remove the resources created during the onboarding and granting flow access processes.

Grant OCI flow log access

Learn how to allow Illumio Segmentation for the Cloud access to your OCI cloud account flow logs.

To set up flow logs, which you need to do before you grant flow log access, see Set up Flow Logs [132].

To review your destinations before granting flow log access, which you should do before you grant flow log access, See Review destinations before granting flow log access [145].

See prerequisites for granting flow log access

Review the prerequisites for your CSP. See Prerequisites for granting flow log access to your CSPs [140].

Instructions for granting OCI flow log access

Illumio Segmentation for the Cloud uses flow logs to display the flows. Granting access to flow logs allows Illumio Segmentation for the Cloud to use these flow logs. For OCI you can enable VCN logs. For instructions on how to enable flow log access in Illumio Segmentation for the Cloud, see the in-application help. For instructions on how to set up flow logs, which you need to do before you grant flow log access, see Set up Flow Logs [132].

1. In the Flow log grant dialog box, download the cs_connector.zip file, which contains the Terraform files to grant access to the selected destinations.
2. After downloading the file, open the OCI Console and navigate to **Developer Services → Resource Manager → Stacks**. Select the stack created during the onboarding process. Selecting a different stack will result in failure to grant access to flows.
3. In the stack details page select **Edit > Edit Stack**.
4. In the edit page, upload the new zip file and click **Next**.

   The page shows the variables that were added during onboarding process.
5. Click **Next**. In the Review page, under the Run apply on the created stack option, make sure the Run Apply check box is *not* selected and click **Save Changes**. (For onboarding, the check-box must be selected. For enabling flows, it must not be selected.)
6. In the Stack Details page, click **Plan** and run the Plan.

   This will run a diff between the previous configuration and the new configuration. Once the plan is completed, it shows the new policies to be added to the group created during onboarding. The new policies are required to allow Illumio Segmentation for the Cloud to read flows from the specified destination.
7. After the Plan job has successfully completed, click **Apply**.
8. In the Apply dialog under the Apply job plan resolution, select **Automatically Approve** (selected by default) and click **Apply**. The stack will run granting the access to the destinations for Illumio.
9. Once the stack completes, return to Illumio Segmentation for the Cloud and click script run successfully and click **Save**.

Set up flow logs in your CSP environment

Learn how to setup flow logs for use by Illumio Segmentation for the Cloud.

Illumio Segmentation for the Cloud uses flow logs to provide visual diagrams and analysis of traffic between sources and destinations. In order for these features to work, after you set up flow logs, you need to grant Illumio Segmentation for the Cloud access to these flow logs through your CSP. The sequence is as follows:

1. Set up flow logs as described here.
2. Review destinations before granting flow log access. See Review destinations before granting flow log access [145].
3. Review prerequisites for granting flow log access. See Prerequisites for granting flow log access to your CSPs [140]
4. Grant flow log access. For AWS, Azure, and GCP, see Grant flow log access to your CSPs [141]. For OCI, see Grant OCI flow log access [131].

NOTE

Do you have central log storage or cross-account storage? This is where flow logs generated from one account (such as "account A1") are stored in a destination belonging to another account (such as "account A2").

You can onboard the accounts in any order but if you onboard "account A1" first, you will not see its flow logs until you onboard "account A2." See Onboarding AWS Cloud [94], Onboarding Azure [59], Onboarding GCP [119], and Onboarding OCI [127].

If you onboard the account storing the flow logs ("account A2") first, you will not see the traffic of "account A1" because it has not yet been onboarded. You also must onboard the account that is generating flow logs ("account A1"). Otherwise, Illumio Segmentation for the Cloud does not properly map the flow logs.

In short, you must onboard both.

Set up flow logs in AWS

You can set up flow logs in AWS using the console, a CloudFormation template, or the command line. You must do this before you Grant flow log access to your CSPs [141]. Note that for AWS, Illumio Segmentation for the Cloud can read flows from S3 buckets only, so it is important to configure these accordingly.

For multi-account AWS customers, Illumio recommends deploying a centralized log storage strategy. This allows you to configure a single S3 bucket destination for all cross-VPC and cross-account flow logs. For details on configuring the necessary roles and permissions, see the AWS website.

Set up AWS flow logs 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 v5 fields. For optimal performance and complete information, enable all v5 fields.

> **NOTE**
>
> Illumio supports only the Text as Log file format. It does not support the parquet format. You can optionally enable hive-compatible prefixes and partition logs by time. Use defaults for all other values.

7. After entering the required information click **Create flow log**.

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} ${vpc-id} ${subnet-id} ${instance-id} $
{tcp-flags} ${type} ${pkt-srcaddr} ${pkt-dstaddr} ${region} ${az-id} $
{sublocation-type} ${sublocation-id} ${pkt-src-aws-service} ${pkt-dst-aws-
service} ${flow-direction} ${traffic-path}"
      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 AWS website.

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 AWS website.

What's next for AWS flow logs

Now you can Grant flow log access to your CSPs [141].

Set up flow logs in Azure

You must do this before you Grant flow log access to your CSPs [141]. For VNETs and NSGs, see this Azure website.

Set up flow logs for Azure Firewalls

See this Azure website.

1. Open the Azure portal and browse to **Firewalls** and click one of your Azure Firewalls.
2. Click **Diagnostic settings** and click **+Add diagnostic setting.**
3. Enter a Diagnostic setting name.
4. Select all the logs you wish Azure to generate. Be sure to select at least **Azure Firewall Network Rule** for Illumio to be able to ingest flow logs. At time of writing, Illumio does not support any other type of Azure Firewall rule log, including any legacy diagnostics.
5. Select **Archive to a storage account**. For subscription and storage account selections, follow the same steps as for VNETs and NSGs.
6. Click **Save**. If you have not already done so, onboard the Azure account. See Onboarding Azure [59].

What's next for Azure flow logs

Now you can Grant flow log access to your CSPs [141].

Set up flow logs in GCP

You must do this before you Grant flow log access to your CSPs [141]. For VPC and Firewalls, use the following steps. In this example, you'll set up flow logs for a VPC.

NOTE

Illumio Segmentation for the Cloud supports the Compute Engine API. It does not support the Network Management API at time of writing.

First, create a flow log configuration.

1. Login to your Google Cloud shell and go to the **VPC networks** page.
2. Click on your VPC to open its details panel.
3. Click the Subnets tab, select the subnets for which you wish to enable flow logs, and click **Manage flow logs**.
4. In the dropdown menu that appears, click **Add new configuration**.
5. Select **Compute Engine**. Illumio recommends that you use the default settings for the rest of the fields.
6. Click **Save**.
7. Browse to the **VPC Flow Logs** page > **Subnet** table to verify that you created the flow log configuration.

   For more information on using flow logs and enabling subnets, see the Google website.

Now, set up a log router to tell the GCP log explorer to redirect the logs to Illumio Segmentation for the Cloud.

1. Type 'log' in the search bar and select the **Logs Explorer** result.
2. In the Logs Explorer, go to the left-hand navigation and click **Log Router**.
3. Click **Create sink**.
4. In **Sink details**, enter a name and optional description for your sink, and click **Next**.
5. In **Sink destination**, select **Cloud Pub/Sub topic** for your sink service, and select an existing topic or create a new one. Click **Next**.
6. In **Choose logs to include in sink** and **Choose logs to filter out of sink**, create your inclusion and exclusion filters. That way the log router will redirect only the logs that match the filter. For example, you might include resource.type="gce_subnetwork" or something similar. You can put 'NOT' in front of an inclusion filter entry to exclude it. Note that at time of writing, Illumio Segmentation for the Cloud processes only VPC and Firewall flow logs, so you may want to include only those or exclude all other logs.

NOTE

Illumio supports the following log sink filters. Note that at least one log sink for each of your topics must have at least one of these filters applied. Otherwise, Illumio will not display the topics and log sinks on the flow access page.

VPC Flow Logs:

- projects/<PROJ-ID>/logs/compute.googleapis.com%2Fvpc_flows
- folders/<FOLDER-ID>/logs/compute.googleapis.com%2Fvpc_flows
- organizations/<ORG-ID>/logs/compute.googleapis.com%2Fvpc_flows
- projects/<PROJ-ID>/logs/networkmanagement.googleapis.com%2Fvpc_flows
- folders/<FOLDER-ID>/logs/networkmanagement.googleapis.com%2Fvpc_flows
- organization/<ORG-ID>/logs/networkmanagement.googleapis.com%2Fvpc_flows

Firewall Logs:

- projects/<PROJ-ID>/logs/compute.googleapis.com%2Ffirewall
- folders/<FOLDER-ID>/logs/compute.googleapis.com%2Ffirewall
- organizations/<ORG-ID>/logs/compute.googleapis.com%2Ffirewall
- projects/<PROJ-ID>/logs/networkmanagement.googleapis.com%2Ffirewall
- folders/<FOLDER-ID>/logs/networkmanagement.googleapis.com%2Ffirewall
- organizations/<ORG-ID>/logs/networkmanagement.googleapis.com%2Ffirewall

7. Click **Update Sink**.

Now, verify that the flow logs are being redirected to Illumio Segmentation for the Cloud.

1. Go to the Logs Explorer as previously described.
2. Click **Log Router**, and look in the **Name** column for the name of the log router you created.
3. Look in the **Destination** column to verify that the destination contains the topic you selected or created when choosing your sink destination.
4. Type 'topic' in the search bar and select the **Topics Pub/Sub** result.
5. In the list of topics, click the topic that you just verified in the Log Router.
6. Click on your subscription, click the **Messages** tab, and click **Pull** to verify that your sink is sending messages to the subscription.

What's next for GCP flow logs

Now you can Grant flow log access to your CSPs [141].

Set up flow logs in OCI

You can set up flow logs in the console. You must do this before you Grant flow log access to your CSPs [141].

This guide provides a concise step-by-step process for you to enable Virtual Cloud Network (VCN) flow logs in Oracle Cloud Infrastructure (OCI) and create a service connector to store these logs in an Object Storage bucket. Observe the following prerequisites:

- OCI Access: Ensure that you have the necessary permissions to manage Networking, Logging, Service Connector Hub, and Object Storage services.
- Existing VCN: Identify the VCN for which you want to enable flow logs.
- Bucket: Create a bucket to store flow logs.

Enable VCN Flow Logs

1. Browse to **Networking > Virtual Cloud Networks**.
2. Select the compartment and choose your VCN.
3. Browse to **Networking > Resources > Flow Logs** on the OCI console.
4. Configure the Flow Log with the following:
   - Name: Enter a name for the flow log.
   - Compartment: Select the appropriate compartment.
   - Enablement Point: Select VCN.
   - Flow logs that need to be enabled: Select the appropriate VCN for which flow logs need to be enabled.
   - Flow Log Type: Choose **All Traffic** with 100% sampling rate.
   - Log Group: Select an existing log group or create a new one.
   - Log: Provide a name for the log.
5. Click **Create Flow Log**.

Store Logs

1. If an object storage bucket is not already available, browse to **Storage > Bucket**s.
2. Click **Create Bucket**, provide a name, and set the desired configurations.
3. Navigate to **Analytics and AI > Messaging > Service Connectors**.
4. Configure the Service Connector with the following:
   - Name: Enter a name for the connector.
   - Source:
     - Select **Logs**
     - Configure the source to use the Log Group and Log from your flow log
   - Target:
     - Select **Object Storage** and choose the bucket you created.
     - (Optional) Set the Object Name Prefix, Batch Rollover Size, and Batch Time Interval.
5. Click **Create Service Connector**.
6. Navigate to your Object Storage Bucket and confirm that log files are being generated and stored. This usually takes some time to show up, as connector will start streaming from the logging service.

What's next for OCI flow logs

Now you can Grant flow log access to your CSPs [141].

Prerequisites for granting flow log access to your CSPs

Review the prerequisites before you grant flow log access to your CSPs.

Set up flow log access before you grant flow log access to your CSPs. See Set up flow logs in your CSP environment [132].

To grant flow log access to AWS, Azure, and GCP see Grant flow log access to your CSPs [141].

NOTE

Granting OCI flow log access is different than granting flow log access to other CSPs. See Grant OCI flow log access [131].

CSP-specific prerequisites for granting flow log access

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

- To grant AWS flow log access, you 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 Onboarding AWS Cloud [94].
- To grant Azure flow log access, you 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 you copied from the Grant Access... dialog box. See Onboarding Azure [59].
- To grant GCP flow log access, you 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 Google Cloud shell open in a browser window, so that you can run the script you copied from the Grant Access... dialog box. Alternatively, you can use your local command line interface. See Onboarding GCP [119].
  - Permission to create an IAM role and bind it to projects and topics using service account as the principal.
- To grant OCI flow log access, you need:
  - Your Oracle portal open in a browser window, so that you can run the Terraform file. See Onboarding OCI [127].

Grant flow log access to your CSPs

Learn how to allow Illumio Segmentation for the Cloud access to your AWS and Azure cloud account flow logs.

> **NOTE**
>
> Granting OCI flow log access is different than granting flow log access to other CSPs. See Grant OCI flow log access [131].

To set up flow logs before you grant flow log access, see Set up Flow Logs [132].

To review your destinations before granting flow log access, see Review destinations before granting flow log access [145].

Grant flow log prerequisites

Review the prerequisites for your CSP. See Prerequisites for granting flow log access to your CSPs [140].

Grant AWS and Azure flow log access in Illumio Segmentation for the Cloud

Granting access to flow logs is done using the onboarding page. For AWS you can enable SG flow logs, and for Azure you can enable Azure Firewall, NSG, and VNET flow logs. For a digest of instructions on how to enable flow log access in Illumio Segmentation for the Cloud, see the in-application help.

First review your flow log access to determine whether you see the flow logs that you expect. See Review destinations before granting flow log access [145].

AWS:

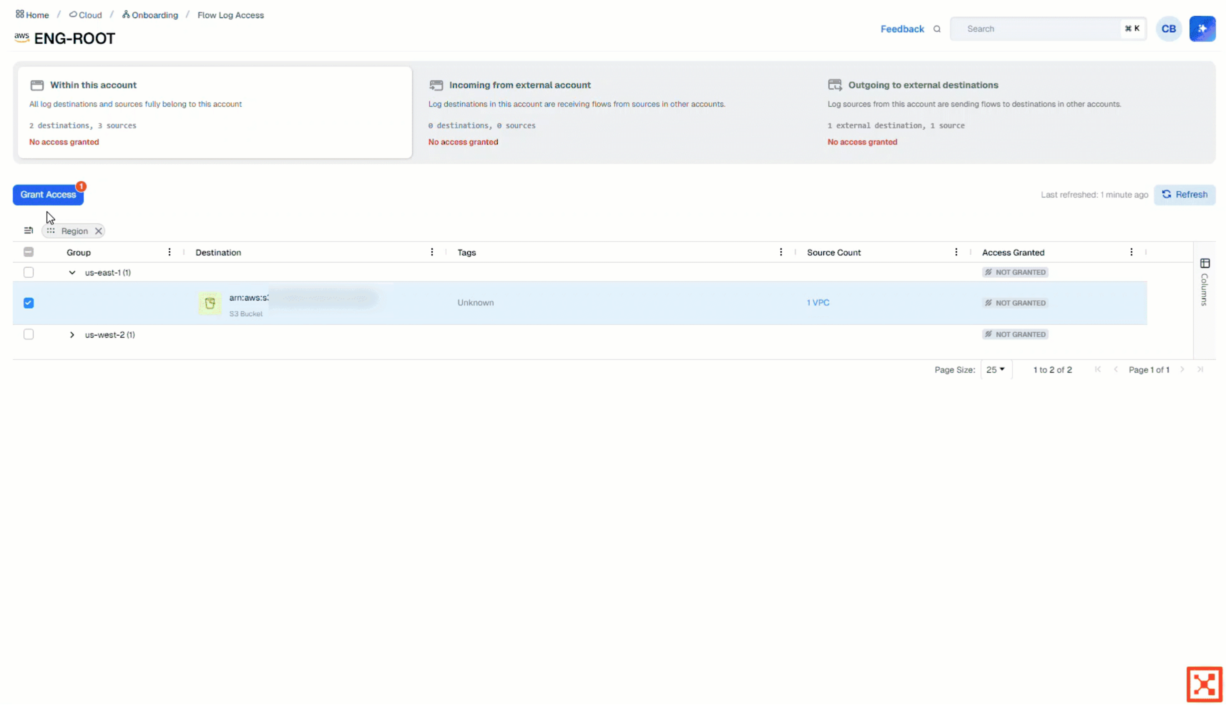

Home / Cloud / Onboarding / Flow Log Access

aws ENG-ROOT

Within this account
All log destinations and sources fully belong to this account
`2 destinations, 3 sources`
No access granted

Incoming from external account
Log destinations in this account are receiving flows from sources in other accounts.
`0 destinations, 0 sources`
No access granted

Outgoing to external destinations
Log sources from this account are sending flows to destinations in other accounts.
`1 external destination, 1 source`
No access granted

	Group	Destination	Tags	Source Count	Access Granted
□	us-east-1 (1)				NOT GRANTED
☑		arn:aws:s[redacted] S3 Bucket	Unknown	1 VPC	NOT GRANTED
□	us-west-2 (1)				NOT GRANTED

Page Size: 25 ▾ 1 to 2 of 2 Page 1 of 1

Azure:

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Flow+Log+Access+Azure.mp4

Steps for granting Azure and AWS flow log access

1. Click **Flow Log Access** on the Onboarding page to open the Flow Log Access page.
2. Find the account name you want.
3. Grant access by first selecting individual or grouped accounts.
   1. Confirm that the flow log destinations you want are selected. If you wish to grant access to flow log destinations within this account or incoming from external accounts:
   2. Click **Grant Access** and use the above prerequisites information in the Grant Access… dialog box, as explained in the in-application help.

NOTE

If you wish to grant access to external flow log destinations, use this set of steps instead.

1. Click the **Outgoing to external destinations** tile.
2. Browse to **Onboarding > Flow Log Access** and find the account that contains the destination resource.
3. Click **Grant Access** and use the above prerequisite [141] information in the Grant Access… dialog box, as explained in the in-application help.

Grant GCP flow log access in Illumio Segmentation for the Cloud

Granting access to flow logs is done using the onboarding page. For GCP you can enable VPC and Firewall flow logs.

For a digest of instructions on how to enable flow log access in Illumio Segmentation for the Cloud, see the in-application help. First review your flow log access to determine whether you see the flow logs that you expect. See Review destinations before granting flow log access [145].

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/GCP+Flow+Log+Access.mp4

When you execute the script, Illumio Segmentation for the Cloud enables the following permissions:

- Project level permissions:
  - pubsub.subscriptions.consume
  - pubsub.subscriptions.create
  - pubsub.subscriptions.delete
- Topic level permissions:
  - pubsub.topics.attachSubscription
  - Read access to the topic paths for your topics

    When granting flow log access, Illumio Segmentation for the Cloud creates two roles with the following bindings and permissions:
    - Role: IllumioFlowAccessRole
      - Binding: topic level permission: pubsub.topics.attachSubscription
    - Role: IllumioPubSubAccessRole
      - Binding: project level permission: pubsub.subscriptions.consume, pubsub.subscriptions.create, pubsub.subscriptions.delete

Use the following steps to grant GCP flow log access:

1. Click **Flow Log Access** on the Onboarding page to open the Flow Log Access page.
2. Find the account name you want.
3. Grant access by first selecting individual or grouped accounts.
   1. Confirm that the flow log destinations you want are selected. If you wish to grant access to flow log destinations within this account or incoming from external accounts:
   2. Click **Grant Access** and use the above prerequisite [141] information in the Grant Access... dialog box, as explained in the in-application help.

NOTE

If you wish to grant access to external flow log destinations, use this set of steps instead.

1. Click the **Outgoing to external destinations** tile.
2. Browse to **Onboarding > Flow Log Access** and find the account that contains the destination resource.
3. Click **Grant Access** and use the above prerequisite [141] information in the Grant Access... dialog box, as explained in the in-application help.

Limitations for granting flow log access

To get traffic flow visibility in the Illumio Segmentation for the Cloud Map and Traffic pages, you need to provide access to flow logs. Once the flow logs are configured in the cloud console, the flow details will be displayed in the flow log access page of Illumio Segmentation for the Cloud. By granting access to flow logs, you will allow Illumio Segmentation for the Cloud to read the flows and provide details about network traffic in the traffic page.

NOTE

To grant access to flow logs stored in a different account than the one you onboarded, you must also onboard the account containing those flow logs. See Onboarding AWS Cloud [94], Onboarding Azure [59], and Onboarding OCI [127].

For example, do you have central log storage or cross-account storage? This is where flow logs generated from one account are stored in a destination belonging to another account.

In this case, you must onboard the account with the destinations first, and then onboard the account with flow logs. Otherwise, Illumio Segmentation for the Cloud does not properly map the flow logs. This may prevent you from granting access to the flow logs despite enabling them.

For example, first onboard the account with the storage account that contains the flow logs (such as "account A1"). Then onboard the accounts that are sending flow logs to the storage accounts in "account A1."

NOTE

Regarding AWS flow log access: Once you select your S3 buckets and Illumio generates a Cloud Formation Template, the template is available to download or run for only 15 minutes. After 15 minutes, you have to re-start the grant flow log access process to generate the template again.

NOTE

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

- In AWS, Illumio Segmentation for the Cloud supports reading flow logs that are stored in S3 buckets only. Currently, other storage destinations are not supported.
- 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 Illumio Segmentation for the Cloud. If the account that owns the S3 bucket is not onboarded, Illumio Segmentation for the Cloud will not be able to fetch the flow logs of that S3 bucket.
- When flow log access is first enabled in Illumio Segmentation for the Cloud, there's a 15-minute latency until traffic flows are first displayed in Cloud Map, Traffic and Inventory pages. After this initial latency, traffic flows are periodically updated every two minutes.

Review destinations before granting flow log access

Review your destinations before allowing Illumio Segmentation for the Cloud access to your cloud account flow logs.

Set up flow logs before you grant flow log access, see Set up Flow Logs [132].

Review flow log access in Illumio Cloud

Review your flow log access to determine whether you see the flow logs that you expect. Click **Flow Log Access** on the Onboarding page.

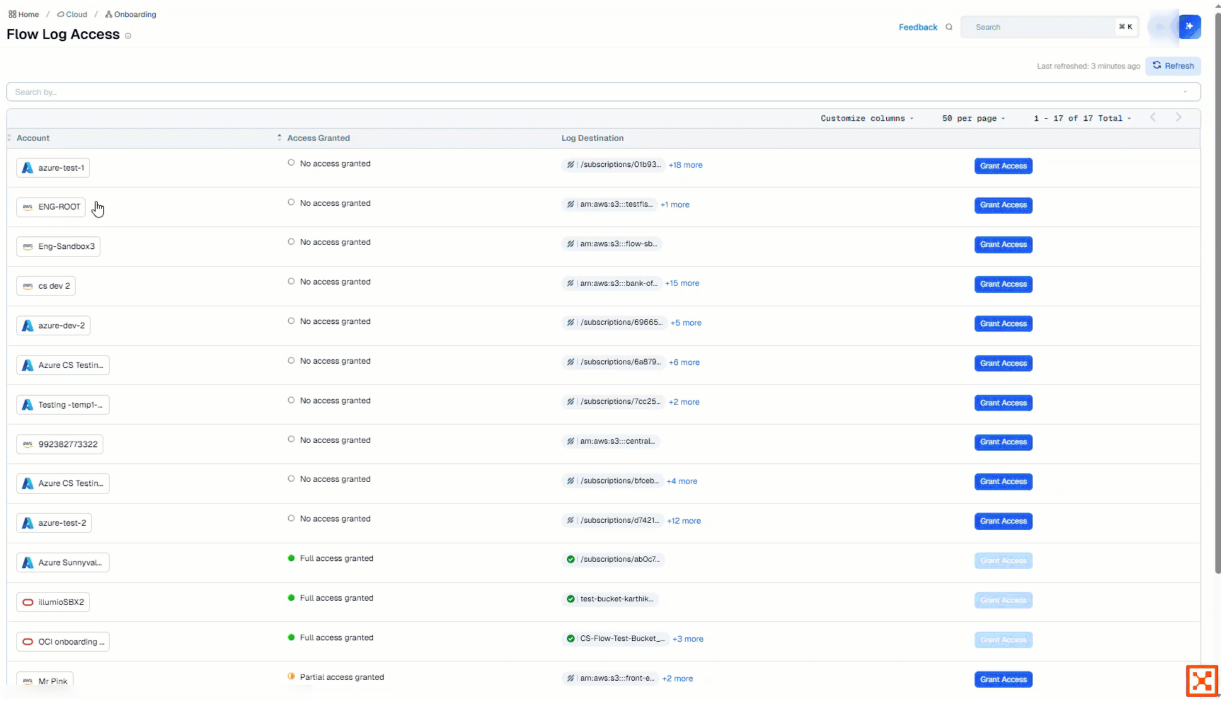

Account	Access Granted	Log Destination	
azure-test-1	○ No access granted	/subscriptions/01b93... +18 more	Grant Access
ENG-ROOT	○ No access granted	arn:aws:s3:::testfls... +1 more	Grant Access
Eng-Sandbox3	○ No access granted	arn:aws:s3:::flow-sb...	Grant Access
cs dev 2	○ No access granted	arn:aws:s3:::bank-of... +15 more	Grant Access
azure-dev-2	○ No access granted	/subscriptions/69665... +5 more	Grant Access
Azure CS Testin...	○ No access granted	/subscriptions/6a879... +6 more	Grant Access
Testing -temp1-...	○ No access granted	/subscriptions/7cc25... +2 more	Grant Access
992382773322	○ No access granted	arn:aws:s3:::central...	Grant Access
Azure CS Testin...	○ No access granted	/subscriptions/bfceb... +4 more	Grant Access
azure-test-2	○ No access granted	/subscriptions/d7421... +12 more	Grant Access
Azure Sunnyval...	● Full access granted	✓ /subscriptions/ab0c7...	Grant Access
illumioSBX2	● Full access granted	✓ test-bucket-karthik...	Grant Access
OCI onboarding ...	● Full access granted	✓ CS-Flow-Test-Bucket_... +3 more	Grant Access
Mr Pink	◐ Partial access granted	arn:aws:s3:::front-e... +2 more	Grant Access

Review flow log access details

1. Find the Account name you want and click on that row.
2. Organize your log destinations and sources if you have several.
   - Select from the tiles to show only resources with the following flows. Each tile indicates the access granted status and lists the number of sources and destinations.
     - Within this account tile. Both the flow sources and flow destinations are in the same account.
     - Incoming from external account tile. The flow destinations receive flows from sources in other accounts in your environment.
     - Outgoing to external destinations tile. The flow sources send flows to destinations in other accounts. This can include accounts outside your environment.
   - Group your destinations by Region or Access Granted
     - Drag and drop the column heading to the 'Drag here to set row groups' text.
     - Click checkboxes to select and deselect destinations. For example, for one tile you may want to group all destinations that don't have access, so that you can grant access to all of them simultaneously.
3. Review the details for each account. Click on a Source Count entry to see which sources are sending flow logs to this destination, and to see details like the region, tags, and so on for each of the sources.
   - 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 Full access or Partial access Granted in the Access Granted column
   - Review your cloud account settings for any child flow logs that are listed as not granted in the access details
   - Refresh the Flow Log Access Details page immediately after granting access to make sure that the updated status
   - You can elect multiple destinations from the Within this account and the Incoming from external account tiles, but not the Outgoing to external destinations tile
   - Checkboxes for grouping by access are available only to users with write access.
4. Review prerequisites [140] and grant flow log access. For AWS, Azure, and GCP, see Grant flow log access to your CSPs [141]. For OCI, see Grant OCI flow log access [131].

After onboarding your accounts

This topic explains what to do next after onboarding your cloud service provider (CSP) accounts.

1. Visualize your resources

After you've onboarded your CSP accounts, your resources are discovered and ingested.

Before defining your environment, Illumio recommends you review your ingested resources.

- About the Cloud Map
  The Map displays a view of your cloud inventory as a network topology map for the your 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.

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

Illumio

See Map [168].

- About the Inventory Page

  View your cloud resources from accounts that were discovered in your environment.

  Search and filter cloud resources on different parameters. View preset filters and set custom columns.

  See Inventory [157].
- Review Your Traffic Flows

  Before you write policy rules to either allow or block traffic, Illumio recommends you determine if there are traffic flows between resources. Filter your resources by flow status, source labels and addresses, destination labels and addresses/ports, and so forth.

  See Traffic [219].

2. Define your public clouds in Illumio Segmentation for the Cloud

Defining your accounts is a multi-step process:

1. **Define your deployment stacks:**

   With Illumio Segmentation for the Cloud, you may decide to create deployment stacks as part of specifying which applications in your cloud account to protect.

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

   See Define a deployment [228].
2. **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.

   See Define an application automatically [231].
3. **Approve your application definitions:**

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

   Stakeholders are in the loop to approve your application definitions.

   See View and approve an application [238].

3. Create policy in Illumio Segmentation for the Cloud

Now that you've reviewed your ingested resources and defined your environment, you are ready to create security policy.

1. **Create your organization polices:**
   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.

   See Writing Organization Policy [254].
2. **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.
   See Writing application policy [255].

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

- Organization policies
- Application policies

Agentless Containers overview

Illumio Segmentation for the Cloud Agentless Containers is a solution that monitors your containerized Kubernetes workloads without requiring agents on each node. It's designed to provide visibility across large, complex, and diverse cloud and hybrid environments, giving Kubernetes administrators full insights into container traffic and workloads.

The intended audience for this topic includes Kubernetes administrators, security professionals, devops engineers, and application developers looking to:

- Extend security across the container lifecycle, including reporting, and auditing

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Agentless_Containers.mp4

Supported platforms

Agentless Containers provides visibility into Kubernetes workloads running on the following:

Cloud

- Amazon EKS

- Azure AKS
- Google GKE

On-Premise

- Self-managed Kubernetes
- OpenShift

Supported container network interfaces (CNI)

Agentless Containers is a solution that requires a Container Network Interface (CNI). It uses the CNI to ingest pod-level network flows and provide visibility into Kubernetes traffic. Agentless Containers currently supports the following configurations:

- Cilium with Hubble Relay (recommended option)

  When Cilium is used as the CNI, enabling Hubble allows Agentless Containers to natively ingest network flows for deep visibility and identity-aware traffic analysis.

  For best results, Illumio recommends viewing videos in Chrome.

  https://product-docs-repo.illumio.com/Tech-Docs/Containers/Videos/Remove_AWS_CNI_Install_Cilium_HLS.mp4
- OpenShift OVN (Open Virtual Networking)-Kubernetes

  When OVN-Kubernetes is used as the CNI, the Illumio Segmentation for the Cloud operator becomes an IPFIX collector, allowing you to export from your OBM cluster to the Illumio service. See Configure OpenShift OVN-Kubernetes [151].
- Falco Plugin (alternative option)

  As an alternative to Cilium or OpenShift, you can enable Falco to ingest network flows. See Onboard and Offboard Kubernetes Clusters [151].

Agentless Containers benefits

This solution allows rapid onboarding of Kubernetes clusters to Illumio Segmentation for the Cloud, and reduces time and complexity. You can:

- Simplify onboarding by eliminating the need for agents on nodes
- Integrate Illumio Segmentation for the Cloud with your existing cloud-native tools and infrastructure on large, distributed, multi-cloud environments
- Eliminate the dependency on iptables and underlying Kubernetes infrastructure

Onboarding and managing your Kubernetes clusters

For directions on using Illumio Segmentation for the Cloud to onboard and offboard your Kubernetes clusters, see Onboard and Offboard Kubernetes Clusters [151].

Viewing and managing your Kubernetes inventory

Illumio Segmentation for the Cloud lets you identify and protect clusters in your environment. You can filter for Kubernetes resources by resources, regions, and other parameters.

See Kubernetes Resources Inventory [165].

Using the Kubernetes Map View

See Navigating the Map Kubernetes View [184].

Viewing your Kubernetes traffic visibility

See Search traffic [220].

Agentless Containers limitations

- Network policies are not supported at the time of writing
- Cluster and node-level policy enforcement is not available for on-premise environments

Upgrade the Kubernetes Operator in Agentless Containers

This topic describes how to upgrade the Kubernetes Operator in Agentless Containers via a Helm upgrade command. Upgrading allows uptaking a patch or release to benefit form bug fixes or new functionality.

Before you begin

This process requires the same onboarding credential file used to initially onboard the agentless container.

Upgrade Kubernetes Operator with a Helm Chart Command

1. Locate the onboarding credential used to initially onboard.
2. Upgrade the Kubernetes Operator by running the Helm Chart `upgrade` command.

   ```
   helm upgrade illumio -f <file location> --namespace illumio-cloud oci://
   ghcr.io/illumio/charts/cloud-operator --version <ver#> --create-namespace
   ```

   The option `--version <ver#>` is the version you want to upgrade to, and `-f <file location>` indicates the location of your onboarding credential.

   For example, if the version you want to upgrade to is `v1.3.3` and the onboarding credential is located in `illumio-cloud-operator-values-f73eba32-638b-47f8-b92f-e4852b8cf7fb.yaml`, the command would be:

   ```
   helm upgrade illumio -f illumio-cloud-operator-values-f73eba32-638b-47f8-
   b92f-e4852b8cf7fb.yaml --namespace illumio-cloud oci://ghcr.io/illumio/
   charts/cloud-operator --version v1.3.3 --create-namespace
   ```
3. Verify the upgrade by running the Helm Chart `version` command:

   ```
   helm version
   ```

Configure OpenShift OVN-Kubernetes

To retrieve network flows from OVN-Kubernetes, the cloud operator acts as an IPFIX collector. This is a standard format collector that OVN-Kubernetes natively supports for exporting.

The preferred collector is the Service IP. OVN-Kubernetes cannot use a Service IP within the same cluster as a network flow collector. When configured with a Service IP, flow data will not reach the collector.

This document offers the cloud operator's pod IP address as an alternative solution until the service IP address functionality is resolved.

Prerequisites

- An installation of OpenShift CLI (oc).
- Access to a kubernetes cluster onboarded with a cloud operator.
- Log into the cluster with a user with cluster-admin privileges.
- Obtain the pod IP address for the cloud operator by running `kubectl get pods -n illumio-cloud -o wide`

Steps to configure OVN

1. Create a patch yaml file that looks like the following. Replace <ip_address> with your `cloud-operator`'s pod IP address.

   ```
   spec:
     exportNetworkFlows:
       ipfix:
         collectors:
           - <ip_address>:4739
   ```
2. Use the following command to apply this patch file to openshift-ovn-kubernetes on each node.

   ```
   oc patch network.operator cluster --type merge -p "$(cat
   <file_name>.yaml)"
   ```
3. To validate this worked, run the following command:

   ```
   oc get network.operator cluster -o jsonpath="{.spec.exportNetworkFlows}"
   ```

   This will show you the following if your patch file worked as expected.

   ```
   {"netFlow":{"collectors":[<ip_address>:4739]}}
   ```

Onboard and Offboard Kubernetes Clusters

This section describes how to onboard and offboard Kubernetes clusters to and from Illumio Segmentation for the Cloud. For an overview of Agentless Containers, see Agentless Containers overview [148].

Kubernetes Administrators can onboard containerized infrastructure to Illumio Segmentation for the Cloud with little time and effort. The onboarding wizard improves consistency and reduces repetitive manual configuration tasks for new Kubernetes clusters. Agentless Containers gives you visibility into your inventory, applications, network traffic, and Kubernetes resources, providing deeper visibility and stronger security controls for containerized workloads.

The Kubernetes Clusters tab shows managed clusters organized by regions. An Illumio region is a designated cloud region where onboarded Kubernetes clusters connect for enhanced visibility and control. Select the nearest Illumio Region for each cluster to optimize performance and security.

NOTE

Prerequisites:

- Onboard your cloud account. This automatically ingests your cloud-managed Kubernetes clusters. Note that at this point, Illumio can only see the cluster, not what is inside the cluster.
- If the account has cloud-managed Kubernetes clusters, click the Kubernetes Clusters tab on the Onboarding page.

CAUTION

You can only download a given credential (.yaml file) **once**.

Create a New Onboarding Credential

After your clusters are automatically ingested during cloud account onboarding, follow these steps to create credentials. These credentials connect the cluster to the Illumio data plane in the selected region, enabling visibility into the cluster’s contents.

1. If you have cloud-managed clusters, browse to the Onboarding page and click a region tile.

   NOTE

   If you don't have any cloud-managed clusters in a specific region, but want to onboard clusters to an Illumio region, click the **Onboard other Illumio regions** tile instead.
2. To onboard credentials, click **Add** in the Onboarding credentials section. This launches the Kubernetes cluster onboarding wizard. The first part is credential creation, and the second part is cluster onboarding.
3. In the credential part of the wizard, fill out the required fields and any optional fields and click **Save**. This creates the onboarding credential and takes you to the onboard cluster instructions part of the wizard. You can use a single onboarding credential to onboard multiple clusters within the selected Illumio region. You can only download a given credential **once**. Credential creation fields include the following:

- Onboarding Credential Name (required)
- Description (optional)
- Illumio Region (required) this is the region for the tile you clicked in step 1.

4. Click **Save**.
5. In the cluster onboarding step of the wizard, be sure to download the Helm values (.yaml) file. Once the wizard is closed, you won't be able to download it again. The .yaml files are available for download only when creating new credentials.
6. From your Helm deployment environment, configure Helm to connect to the Kubernetes cluster. Refer to Helm documentation for details.
7. Copy the Illumio Segmentation for the Cloud deployment command and run it using Helm from outside the cluster, making sure it targets the desired Kubernetes cluster.
8. Copy the command for confirming the deployment and pairing, and run it in your Helm deployment environment. When you have confirmed the deployment and pairing, click **Done** in the wizard.

   You have now onboarded your Kubernetes cluster to Illumio Segmentation for the Cloud. To see statistics for it, browse to the Inventory page. See Kubernetes Resources Inventory [165]. To see a visualization of it, browse to the Cloud Map. See Navigating the Map Kubernetes View [184].

> **NOTE**
>
> If you are onboarding an Amazon EKS cluster using the AWS VPC CNI and require visibility into network flows, deploying Falco is required. As per the vendor's recommendation, Falco should be used alongside the AWS CNI to enhance security monitoring and network flow analysis in your EKS environment.
>
> Before adding configuration values to your credentials file, add the namespace "falco" using the following kubcutl command:
>
> ```
> kubectl create ns falco
> ```
>
> To ensure seamless integration, include the following configurations in your credentials file:
>
> ```
> illumio-cloud-operator-values-xxxxxx.yaml
>
> falco:
>   enabled: true
>
> onboardingSecret:
> ....
> ```

9. Click **Done**. The cluster status changes to onboarded, and the last connection time and cluster status change. Illumio can now see the Kubernetes resources and traffic for the cluster.

Onboard a Kubernetes Cluster Using an Existing Credential

To onboard additional clusters using a previously created credential, follow these steps. For example, if you initially onboarded 10 clusters during your cloud account setup, you might now want to onboard 15 more in the same region. These steps allow you to reuse the same .yaml credential file that you previously downloaded. The wizard provides a Helm command that uses this credential to onboard the new clusters.

Task-focused example:

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Kubernetes_Cluster_Onboarding.mp4

Detailed example:

https://product-docs-repo.illumio.com/Tech-Docs/Containers/Videos/Self+Managed+Kubernetes+Clusters.mp4

1. To onboard all the clusters for a region, browse to the Onboarding page, click the region tile, and click **Onboard**. To onboard individual clusters for a region, click the region tile, click a cluster, and click **Onboard**.

   This launches a Kubernetes cluster onboarding wizard, that is slightly different from the one used to add new credentials. The first part is credential selection, and the second part is cluster onboarding.
2. In the credential part of the wizard, select an existing credential from the dropdown menu and click Next.
3. Configure Helm to connect to the Kubernetes cluster.
4. Copy the command for deploying the Illumio Segmentation for the Cloud and use the command in Helm, outside of the cluster, and have it pointed to the desired cluster.
5. Copy the command for confirming the deployment, and run it in your Helm deployment environment. When you have confirmed the deployment, click **Done** in the wizard.

   You have now onboarded your additional Kubernetes cluster to Illumio Segmentation for the Cloud using the same same credentials (.yaml) file you downloaded. To see statistics for the cluster, browse to the Inventory page. See Kubernetes Resources Inventory [165]. To see a visualization of the cluster, browse to the Cloud Map. See Navigating the Map Kubernetes View [184].

Configuring Proxy with Cloud Operator

If your Kubernetes cluster does not have direct internet access and operates in a proxy-restricted environment, you can configure the Cloud Operator to route outbound traffic through an HTTPS proxy. See Configure Proxy with Cloud Operator [156].

Disable a Kubernetes Cluster

1. Browse to the **Onboarding > Kubernetes** Clusters tab.
2. Click the tile for the region containing the cluster you want to disable.
3. In the Clusters section, click the checkbox next to the cluster and click **Disable**. To reenable it, click the checkbox and click **Enable**.

Offboard a Kubernetes Cluster

Use the following steps to offboard a Kubernetes cluster, which removes all its workloads from Inventory but not completely remove it from Illumio Segmentation for the Cloud .

To completely remove a Kubernetes cluster after you have offboarded it, see Remove an Offboarded Kubernetes Cluster [155].

https://product-docs-repo.illumio.com/Tech-Docs/Containers/Videos/Kubernetes+cluster+offboarding.mp4

1. Browse to the **Onboarding > Kubernetes Clusters** tab.
2. Click the tile for the region containing the cluster you want to offboard.
3. In the Clusters section, click the checkbox next to the cluster and click **Off-board**. In the dialog that appears, click **Off-board**.

Remove an Offboarded Kubernetes Cluster

Offboarding a cluster removes it from Inventory, but does not completely remove it. Illumio recommends that you remove offboarded clusters. To remove offboarded clusters:

Open your Helm console (or a different environment console if you used one to onboard the cluster originally) and run the following command:

```
helm uninstall illumio -n illumio-cloud
```

Remove an Onboarding Credential

After you have created credentials to onboard clusters, you can remove them for security purposes. Use the following steps to remove a credential:

1. Browse to the **Onboarding > Kubernetes Cluster** tab.
2. Click the tile for the region containing the cluster that has the credential you wish to remove.
3. In the Onboarding Credentials section, click the checkbox next to the credential and click **Remove**. In the dialog that appears, click **Remove**.

   This removes the credential without affecting any clusters you previously onboarded with this credential.

Other Agentless Container Solutions documentation

For upgrading the Kubernetes Operator for Agentless Containers, see Upgrade the Kubernetes Operator in Agentless Containers [150]

For Kubernetes Clusters in Inventory, see Kubernetes Resources Inventory [165].

For Kubernetes Clusters in Cloud Map, see Navigating the Map Kubernetes View [184].

For Kubernetes Clusters in Traffic, see Search traffic [220].

Configure Proxy with Cloud Operator

If your Kubernetes cluster does not have direct internet access and operates in a proxy-restricted environment, you can configure the Cloud Operator to route outbound traffic through an HTTPS proxy.

This is done by setting the HTTPS_PROXY environment variable in the Helm values file. The value is passed into the container at runtime, and all outbound connections made by the Cloud Operator will route through the configured proxy.

Prerequisites

Ensure that the proxy server is accessible from the environment where the Cloud Operator is running. URIs for cloud.illum.io and subdomains *.cloud.illum.io must be allowed.

Steps to configure the proxy environment variable

1. Open the `illumio_values.yaml` file.
2. Add the env section with the httpsProxy variable specifying your proxy URL.
   Example:

   ```
   env:
     httpsProxy: "https://example.com:8888"
   ```
3. Deploy or update the Cloud Operator using Helm.

Visualize

The content in this category introduces you to the Illumio Segmentation for the Cloud visualization tools and explains how to use key features of Illumio Segmentation for the Cloud, such as inventory, traffic, and the Cloud Map.

Inventory

This topic describes the Illumio Segmentation for the Cloud 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 GUI.

For information about the Inventory Details pages for your resources, see Inventory Details [162].

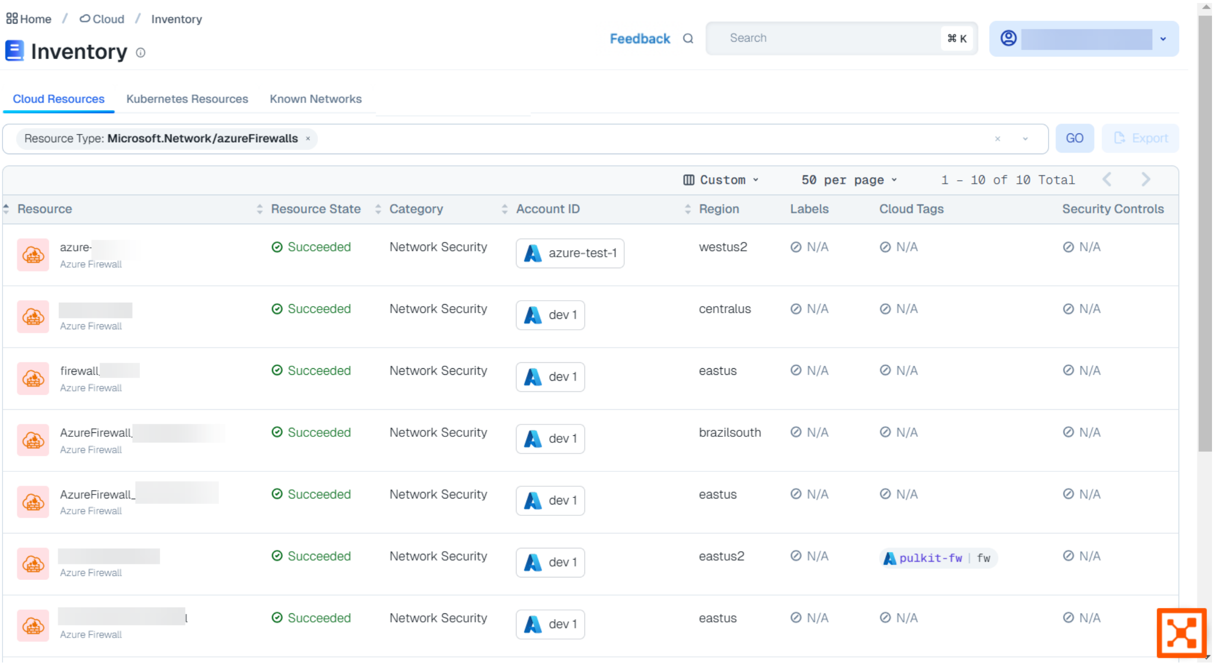

Resource Type: **Microsoft.Network/azureFirewalls** ×

1 – 10 of 10 Total

Resource	Resource State	Category	Account ID	Region	Labels	Cloud Tags	Security Controls
azure-[redacted] Azure Firewall	✓ Succeeded	Network Security	azure-test-1	westus2	⊘ N/A	⊘ N/A	⊘ N/A
[redacted] Azure Firewall	✓ Succeeded	Network Security	dev 1	centralus	⊘ N/A	⊘ N/A	⊘ N/A
firewall[redacted] Azure Firewall	✓ Succeeded	Network Security	dev 1	eastus	⊘ N/A	⊘ N/A	⊘ N/A
AzureFirewall[redacted] Azure Firewall	✓ Succeeded	Network Security	dev 1	brazilsouth	⊘ N/A	⊘ N/A	⊘ N/A
AzureFirewall_[redacted] Azure Firewall	✓ Succeeded	Network Security	dev 1	eastus	⊘ N/A	⊘ N/A	⊘ N/A
[redacted] Azure Firewall	✓ Succeeded	Network Security	dev 1	eastus2	⊘ N/A	pulkit-fw | fw	⊘ N/A
[redacted]l Azure Firewall	✓ Succeeded	Network Security	dev 1	eastus	⊘ N/A	⊘ N/A	⊘ N/A

Supported Resource Types

See Illumio visibility for resource types [193].

For a list of resources against which you can write policy, see Policy enforcement and resource types [262].

Searching Your Inventory Resources

This feature provides a quick way to quickly search your inventory resources. Some parameters show only those values that are relevant to your other parameter selections. Selecting specific parameters (cloud, region, resource type, regions and categories) helps you craft effective queries. For example, if you select "cloud = AWS", and then open the region parameter, you see only AWS regions listed. This gives you regions in the context of your cloud selection.

You can search your inventory by the following parameters:

- Cloud
- Account ID
- Account Name
- Region
- Resource Type (for example, Azure Firewall, EC2, Subnet, OCI, compute instance, and more
- Resource Group
- Resource Name
- Resource State (This reflects updates from the CSP. Frequent state changes may experience delayed refreshes, but updating any of these other parameters triggers an immediate refresh.)
- VPC/VNET ID
- Subnet ID
- Cloud Tags
- Labels (for example, <application name>, VPC deployments, and more.)
- Categories (for example, databases, containers, and more.)
- 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. Only valid IPs are returned, so users do not have to type full length valid addresses. This is limited to IPv4 and IPv6.)

Resources Use Cases and Example

Illumio Segmentation for the Cloud discovers your resources when cloud onboarding is done. This feature lets you search through a table of your discovered resources. Such a search lets you confirm general expectations of what resources you have and what is in a given region, or display information about a specific type of resource.

Filtering by an AWS EC2 Instance

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 Resource 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
   - Resource name
   - Resource state
   - Account ID

   In addition the the above, other general properties of the Inventory table may display, again depending on your column customization. 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 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 a heading shows details for that entry such as its state or creation date.

Similarly, selecting a database category and then a CSP, like AWS, would give you a list of list of all the AWS databases in the Illumio Segmentation for the Cloud environment.

Filtering by an Azure Firewall

NOTE

Illumio Segmentation for the Cloud does not support Classic Azure Firewall.

For this example, suppose you are interested in reviewing a particular Azure Firewall. The following steps illustrate how you would do that.

1. The first part of the sequence might be to filter by Resource Type and select **Microsoft.Network/azureFirewalls**:

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

   - Resource
   - Resource state
   - Region
2. The next step in the sequence would be to click one of the entries in the Resource column.

   The Detail page opens on the Summary tab. In the case of an Azure Firewall, you see additional information beyond Name, ID , Cloud and so on, such as:

   - Resource Group
   - Sku Tier
   - Threat Intel Mode
3. The next step might be to click the Attached Resources tab. That tab displays the following information:
   - Virtual Networks
   - Subnets
   - Firewall Policy
4. Next, you may want to see the firewall policies that you have. Click on the **Policy Rules** tab, which shows the top-down Firewall Policy, with the policy listed at the top with its parent policy listed. The first level below that is the Rule Collection Groups.

   The first level below that is the Rule Collection Groups.

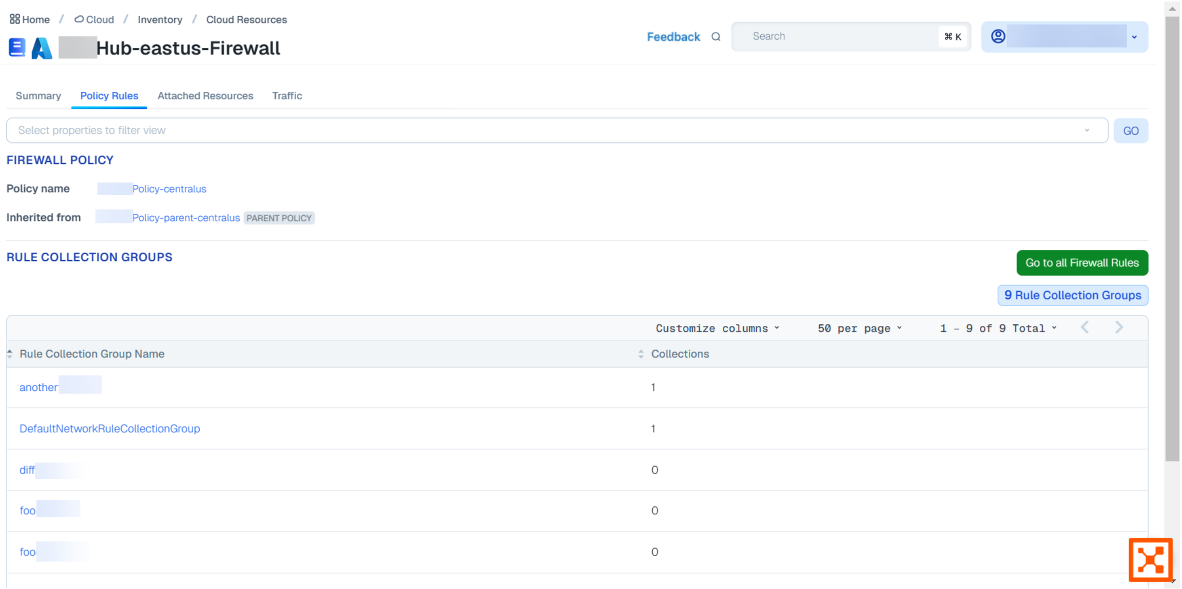

Home / Cloud / Inventory / Cloud Resources

[redacted]Hub-eastus-Firewall

Feedback | Search ⌘ K | [redacted]

Summary | Policy Rules | Attached Resources | Traffic

Select properties to filter view | GO

FIREWALL POLICY

Policy name [redacted]Policy-centralus

Inherited from [redacted]Policy-parent-centralus PARENT POLICY

RULE COLLECTION GROUPS

Go to all Firewall Rules

9 Rule Collection Groups

Customize columns | 50 per page | 1 – 9 of 9 Total

Rule Collection Group Name	Collections
another[redacted]	1
DefaultNetworkRuleCollectionGroup	1
diff[redacted]	0
foo[redacted]	0
foo[redacted]	0

5. Click a rule collection group to display its details and show the the second level, which is the Rule Collections.

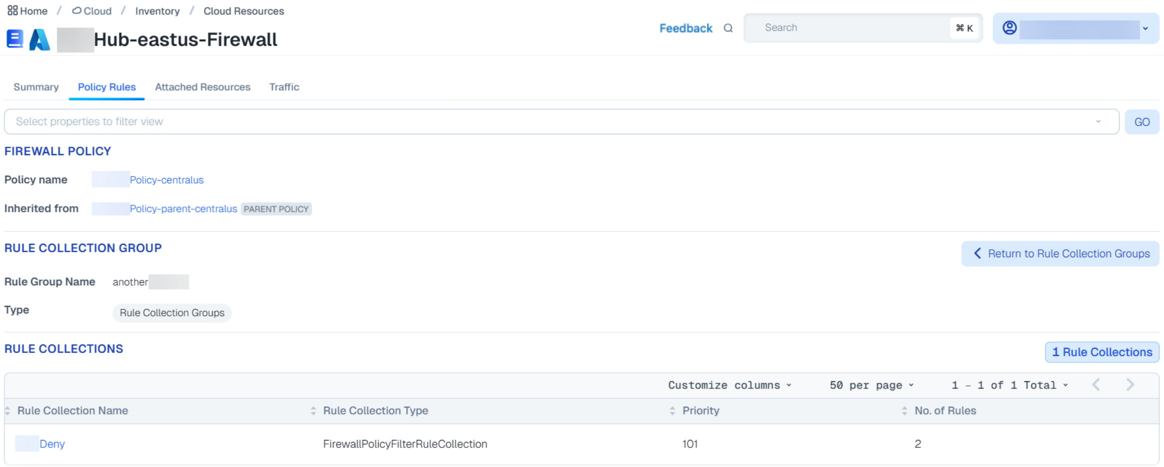

Home / Cloud / Inventory / Cloud Resources

[redacted]Hub-eastus-Firewall

Feedback | Search ⌘ K | [redacted]

Summary | Policy Rules | Attached Resources | Traffic

Select properties to filter view | GO

FIREWALL POLICY

Policy name [redacted]Policy-centralus

Inherited from [redacted]Policy-parent-centralus PARENT POLICY

RULE COLLECTION GROUP

Return to Rule Collection Groups

Rule Group Name another[redacted]

Type Rule Collection Groups

RULE COLLECTIONS

1 Rule Collections

Customize columns | 50 per page | 1 – 1 of 1 Total

Rule Collection Name	Rule Collection Type	Priority	No. of Rules
[redacted]Deny	FirewallPolicyFilterRuleCollection	101	2

Click a rule collection to display its details and show the third level, which is the Rules. You can then click on a rule, and so on. At each level, you can click **Return...** to go back up a level.

6. To get an all in one view, from the Rule Collection Groups page, click **Go to All Firewall Rules** and select one of the rule tabs. In this example, you select the **Network Rules** tab.

The columns show the rules with their parent Rule Collection Groups and Rule Collections.

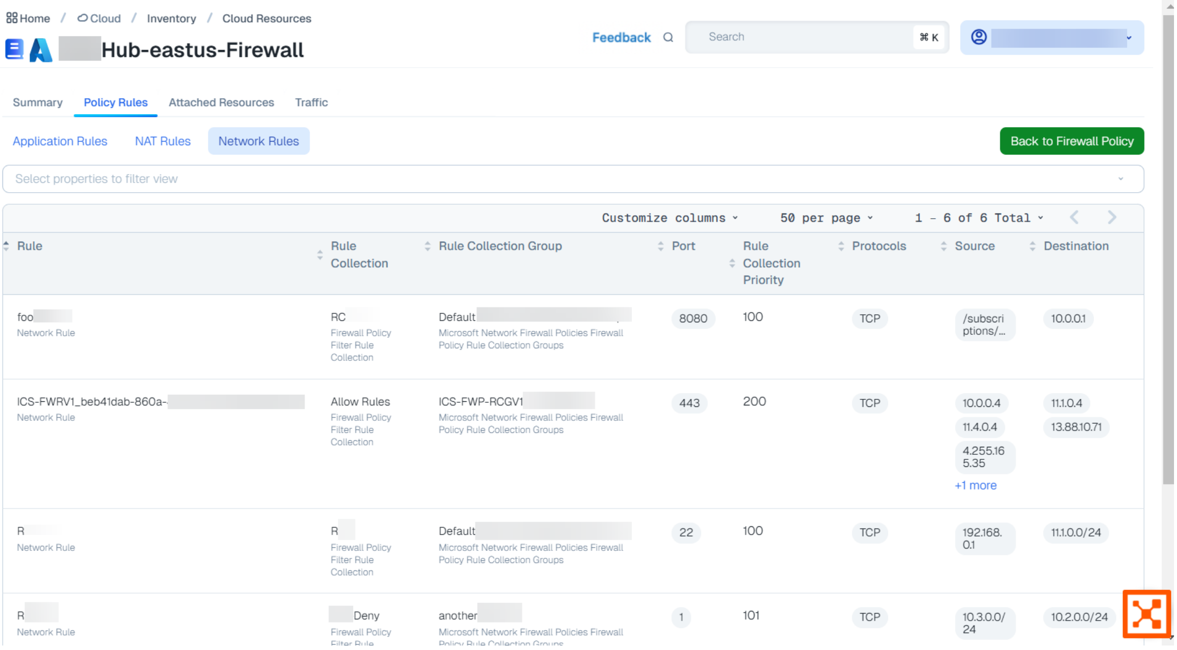

Home / Cloud / Inventory / Cloud Resources

[redacted]Hub-eastus-Firewall

Rule	Rule Collection	Rule Collection Group	Port	Rule Collection Priority	Protocols	Source	Destination
foo[redacted] Network Rule	RC[redacted] Firewall Policy Filter Rule Collection	Default[redacted] Microsoft Network Firewall Policies Firewall Policy Rule Collection Groups	8080	100	TCP	/subscriptions/...	10.0.0.1
ICS-FWRV1_beb41dab-860a-[redacted] Network Rule	Allow Rules Firewall Policy Filter Rule Collection	ICS-FWP-RCGV1[redacted] Microsoft Network Firewall Policies Firewall Policy Rule Collection Groups	443	200	TCP	10.0.0.4 11.4.0.4 4.255.165.35 +1 more	11.1.0.4 13.88.10.71
R[redacted] Network Rule	R[redacted] Firewall Policy Filter Rule Collection	Default[redacted] Microsoft Network Firewall Policies Firewall Policy Rule Collection Groups	22	100	TCP	192.168.0.1	11.1.0.0/24
R[redacted] Network Rule	[redacted]Deny Firewall Policy Filter Rule	another[redacted] Microsoft Network Firewall Policies Firewall Policy Rule Collection Groups	1	101	TCP	10.3.0.0/24	10.2.0.0/24

For more information on Inventory page searching, see the Context-based search section of Search and filtering [190].

Exporting an Inventory Report

Use this feature to export a list of your resources subject to your search query.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Export+an+Inventory+Report.mp4

1. Click **Export** to export the inventory data.
2. Edit the report name and select the format.
3. Click the **Scheduling Section** toggle to the on position to schedule the export unless you want to export the report immediately.
4. If you choose to schedule your report, select your recurrence and time.
5. Click **Save** when done.
6. Go to the Generated reports [212] page to download the exported report.

Known Networks

The Known Networks tab displays a list of known networks (IP lists). This list populates the options for the Known Networks filter option on the Traffic page. See Search traffic [220]. To add one, follow the in-application help directions.

The Cloud Map also displays known networks as a type of resource.

Tooltips

Hover over items in the following columns to see tooltips summarizing information about them:

- Resource (Name, Resource ID, Type, Account ID, CSP, Region, Labels, Last Updated)
- Account ID (Account ID, Name, both of which you can copy)
- Labels (Type, which you can copy)
- Cloud Tags (Cloud Tag, Illumio Label)
- Security Controls (Properties of the resource you have, and Policy Sync and Policy Last Applied.)

NOTE

Last Updated refreshes only when there is a change in the CSP resource. Policy Sync and Policy Last Applied refresh only when there is a change to existing policy.

Inventory Details

When you click on a resource in the Inventory page, that open details for that inventory item. There, you see the details for that inventory. Details include a summary comprising properties of the inventory item, such as general properties , which are also seen in the inventory table's default columns. Other details include additional properties that are not a part of general properties. These include IP addresses, K8 versions in the case of an ECS, auto scaling definitions in the case of autoscaling groups, and so forth. The Illumio labels applied, and the cloud tags on the resource, are also shown in the summary of the inventory details. All of these together provide you with a clear understanding of your inventory.

Every resource may have additional details tabs depending on the type of resource. This includes:

- **Attached resources:** All the resource types in relation to the resource that is currently viewed
- **Traffic:** (If the resource has traffic activities)
- **Resource Graph:** Each resource has a graph that shows the resource and its related resources in a graphical structure. (This feature is in preview and is subject to limitations listed toward the end of this topic).

For information about the Inventory page itself, see Inventory [157].

This page describes some types of details that may benefit from further description.

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.
- You can see in the attributes whether a VPC is shared.
- 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 Illumio Segmentation for the Cloud 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.

Route Table Resource Details

The Details page for AWS route tables includes a Routes tab with the following information:

- Destination: This is the IP address or CIDR block that a route specifies

- Origin: This describes how the route was created, whether automatically, manually, or by route propagation
- State: This is the state of the route
- Target ID: This is the ID of the gateway, network interface, or connection that will receive the destination traffic

The Details page for Azure route tables includes a Routes tab with the following information:

- Address Prefix: This represents a range of IP addresses (in CIDR notation) to which this route applies
- Next Hop IP Address: This is the specific IP address of the device (like a virtual network gateway or a network virtual appliance) that will receive the traffic for further routing
- Next Hop Type: This specifies where traffic should be directed based on a route
- Route Name: This is a user-defined label for easy identification of the route within the table

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.

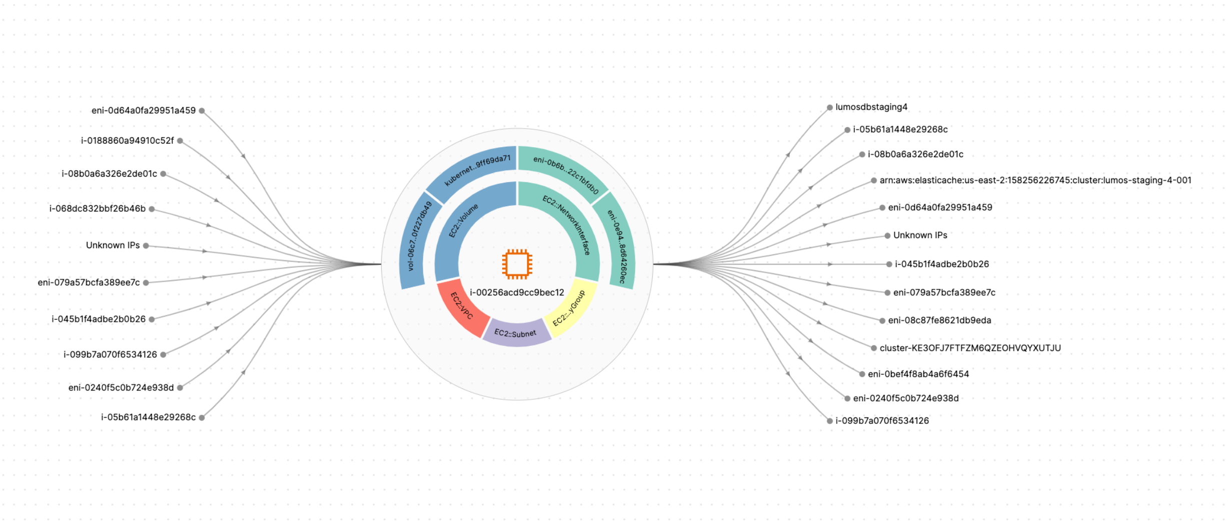

Details Resource Graph limitations

The following are limitations of the resource graph:

- If there is no traffic going in or out of the resource, you do not see traffic flows
- The resource graph does not show rules
- If a resource does not have attached resources, it does not have a resource graph

Kubernetes Resources Inventory

Learn about the Illumio Segmentation for the Cloud Kubernetes Resources Inventory. Kubernetes administrators benefit from advanced search and filter capabilities to quickly locate Kubernetes resources such as clusters, nodes, namespaces, or workloads for both cloud-managed and self-managed clusters. The Kubernetes Resource Inventory helps you quickly and efficiently identify resources that may pose security risks.

See Agentless Containers overview [148].

For information about Inventory Details pages for your resources, see Inventory Details [162].

For general information about the Inventory page, see Inventory [157].

Key capabilities

The Kubernetes Resources Inventory provides a comprehensive list of Kubernetes resources and visualizes connections across clusters and workloads. You can drill down to display Kubernetes resources detailed information on an inventory page, as described in Kubernetes Resources [166] on this page.

Use the following capabilities to identify Kubernetes resources that may be at risk:

- **Discover clusters:** Agentless Containers lets you search for cloud-managed clusters in your cloud service source (CSP). This lets you proactively identify clusters that are not yet protected by Illumio Segmentation for the Cloud. For cloud-managed clusters, it helps you determine the total number of protected clusters within each CSP. For example, you can search by the following across different environments:
  - Region
  - Kubernetes clusters
  - Combinations of whole or partial cluster names and CSPs
- **Identify workload security gaps:** Identify and address security gaps across Kubernetes environments by searching according to workload types, including:
  - Deployment
  - DaemonSet
  - Job
  - CronJob
  - Application
  - Labels
- **Identify resource type security gaps:** In addition to searching by workload, you can search by resource types including:
  - Namespace
  - NetworkPolicy
  - Node
  - Services

Kubernetes resources

The Kubernetes Resources tab displays a list of known Kubernetes resources, including both cloud-managed and self-managed clusters. Illumio Segmentation for the Cloud supports all Kubernetes resources, including:

- Clusters
- Nodes
- Namespaces
- Deployments
- Network Polices
- Kubernetes Labels

You can filter them by the following parameters (columns):

- Resource (name)
- Cloud (self-managed clusters have empty entries)
- Region (self-managed clusters have empty entries)
- Cluster/Namespace (if the resource is not itself a namespace)
- Kubernetes Labels
- Last Updated On

Click a Kubernetes entry in the Resource column to see its details, including Resource state and the above column information. This is useful when screen space limits the number of columns visible to you. When you click on a namespace or cluster resource, the details panel displays the following tabs:

- General (similar details as seen for other Kubernetes resource types. Note that for non-cluster or namespace resources, you will get only the General tab)
- Attached Resources (resources associated with the namespace, with parameters shown in columns similar to those on the main Kubernetes Resources tab)
- Map (Kubernetes-focused map where you can view traffic lines and drill down into regions, clusters, and the like)

You can hover over the entries in the Resource, and Cluster/Namespace columns to see their details as well. The hover details may be more concise than those in the details panel.

The Map also displays Kubernetes resources in a separate view. See Map [168].

Limitations for Kubernetes Resources Inventory

Note that cloud-managed Kubernetes resources have Accounts, Regions, and Clouds associated with them. These three categories are shared across Illumio Segmentation for the Cloud. This means that in addition to showing cloud resources like EC2 instances, filtering on terms like Account also shows cloud-managed Kubernetes resources like nodes. However, it does not show self-managed Kubernetes resources because they are not associated with any of these three categories. Filtering on Kubernetes-specific categories like K8s Resource Name shows both cloud- and self-managed resources.

Note that for EKS, Kubernetes Inventory supports only nodes that are part of a node group.

Azure Firewalls Overview

You can gain visibility into your Azure Firewalls and enforce policy on them using Illumio Segmentation for the Cloud. This visibility includes a clear view of your Azure Firewall inventory, with details about your firewalls and their current policies. The Map provides Azure firewall visualization within a hub-spoke architecture. This provides you with insights into firewall relationships with virtual hubs and VNETs (differentiated by hub and spoke), firewall policy details, and traffic flows passing through the firewall. You can also see your firewall network flows using the Traffic page. These combined capabilities make it easier to see which VNets talk to each other at a glance, as well as to see which VNets need to have firewalls applied.

NOTE

Illumio Segmentation for the Cloud does not support Classic Azure Firewall.

Benefits of Using Azure Firewall

This visual inspection feature lets you:

- Understand all your onboarded Azure subscriptions' firewall deployments within the Azure network topology
- Gain deeper insights into your Azure accounts' firewall configuration regarding network topology, policy details, and traffic from source to firewall and firewall to source

With Azure Firewalls you can:

- Open the Cloud Map configuration panel to select the Firewall Topology view. Highlighting lets you see which VNets and virtual hubs have firewalls.
- Distinguish between hubs and spokes immediately. (Hub VNets contain firewalls, and VNets peered to hub VNets are spoke VNets).
  - Hubs are labeled as such and specify whether they are virtual hubs or VNet hubs
  - Spokes are labeled as such
- Determine immediately whether a virtual hub or VNet has a firewall deployed (firewall icon) or has a peering connection to a firewall (dashed peering line). In peering mode, the line will be purple. In the Firewall topology view the line will be orange.
- Apply filters by application, region, resource, and more to display associated Azure firewalls, resources, and traffic flows. Displayed traffic flows include:
  - Traffic passing through the firewalls,
  - Traffic between resources within the region,
  - Traffic across the region, cloud, and the internet

A Typical Use Use Case

For security operations administrators, gaining complete visibility into the firewalls associated with critical applications, is essential to minimizing exposure to risks. This visibility ensures that payments applications are protected from potential security threats. If any VNets lack firewalls, follow this workflow to put the firewalls in place:

1. Use the Illumio Segmentation for the Cloud Azure Firewalls feature to see which of your applications' VNets have firewalls associated with them. See Navigating Azure Firewalls [179].
2. Identify VNets. See Inventory [157].
   - Compile a list of all VNets associated with payments applications
   - Determine whether each VNet has an active firewall
   - Evaluate the risks posed by VNets without firewalls, considering application sensitivity and compliance requirements
3. Define Firewall Rules. See Writing Azure Firewall policy [259].
   - Configure rules in Illumio Segmentation for the Cloud to meet application requirements while adhering to the principle of least privilege.
   - Set up alerts in Azure to monitor traffic anomalies
4. Validate and Monitor. See Navigating Azure Firewalls [179].
   - After deployment, validate the configuration and monitor traffic to ensure the firewall operates as intended.

Map

This topic explains how to work with the Map, found in the Cloud > Explore menu. For information on navigating the Map, see Cloud Map navigation [171].

For a list of resources you can view on the Map, see Illumio visibility for resource types [193]. For a list of resources against which you can write policy, see Policy enforcement and resource types [262].

What is the 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. Cloud is designed to handle this challenge. Cloud analyzes these relationships to provide a view of assets with proper cloud hierarchy.

The Map displays a view of your cloud inventory as a network topology map for the your 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.

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

Illumio

How the Map is organized

Using the Infrastructure View

Cloud organizes the Infrastructure View first by cloud — AWS, Azure, etc. 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

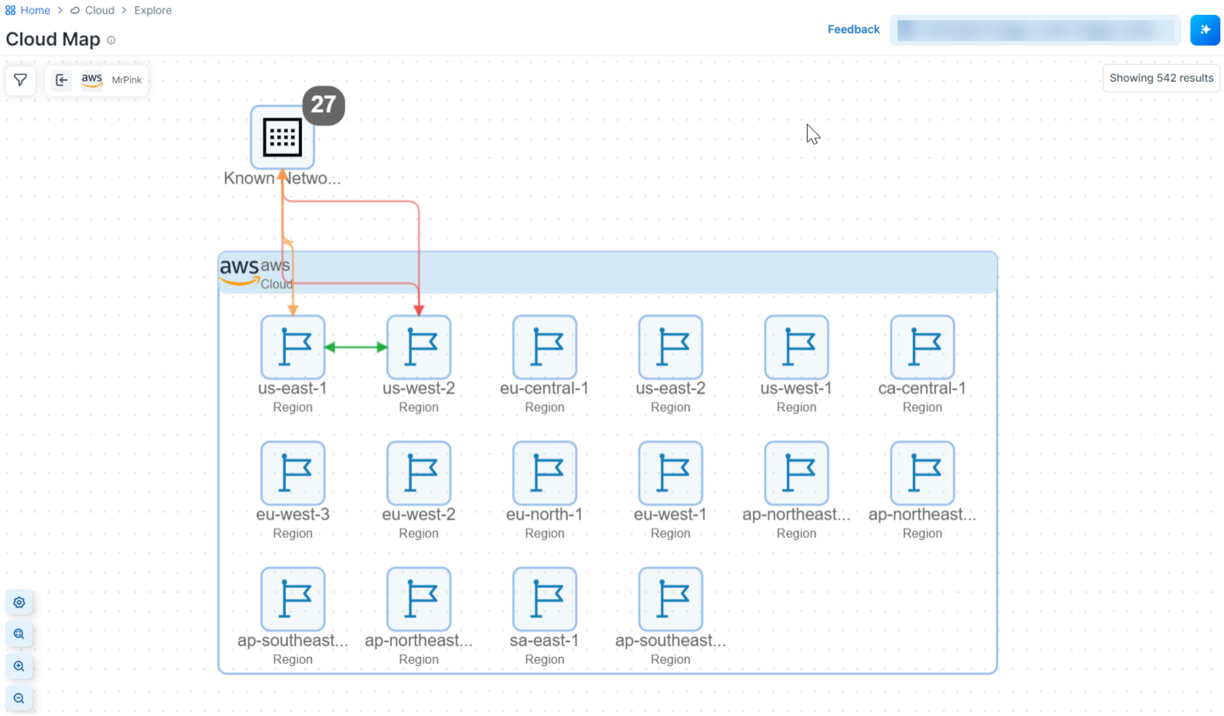

The Map displays your resources within the regions. This example shows us-west-2 region in your AWS **13##########** account.

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

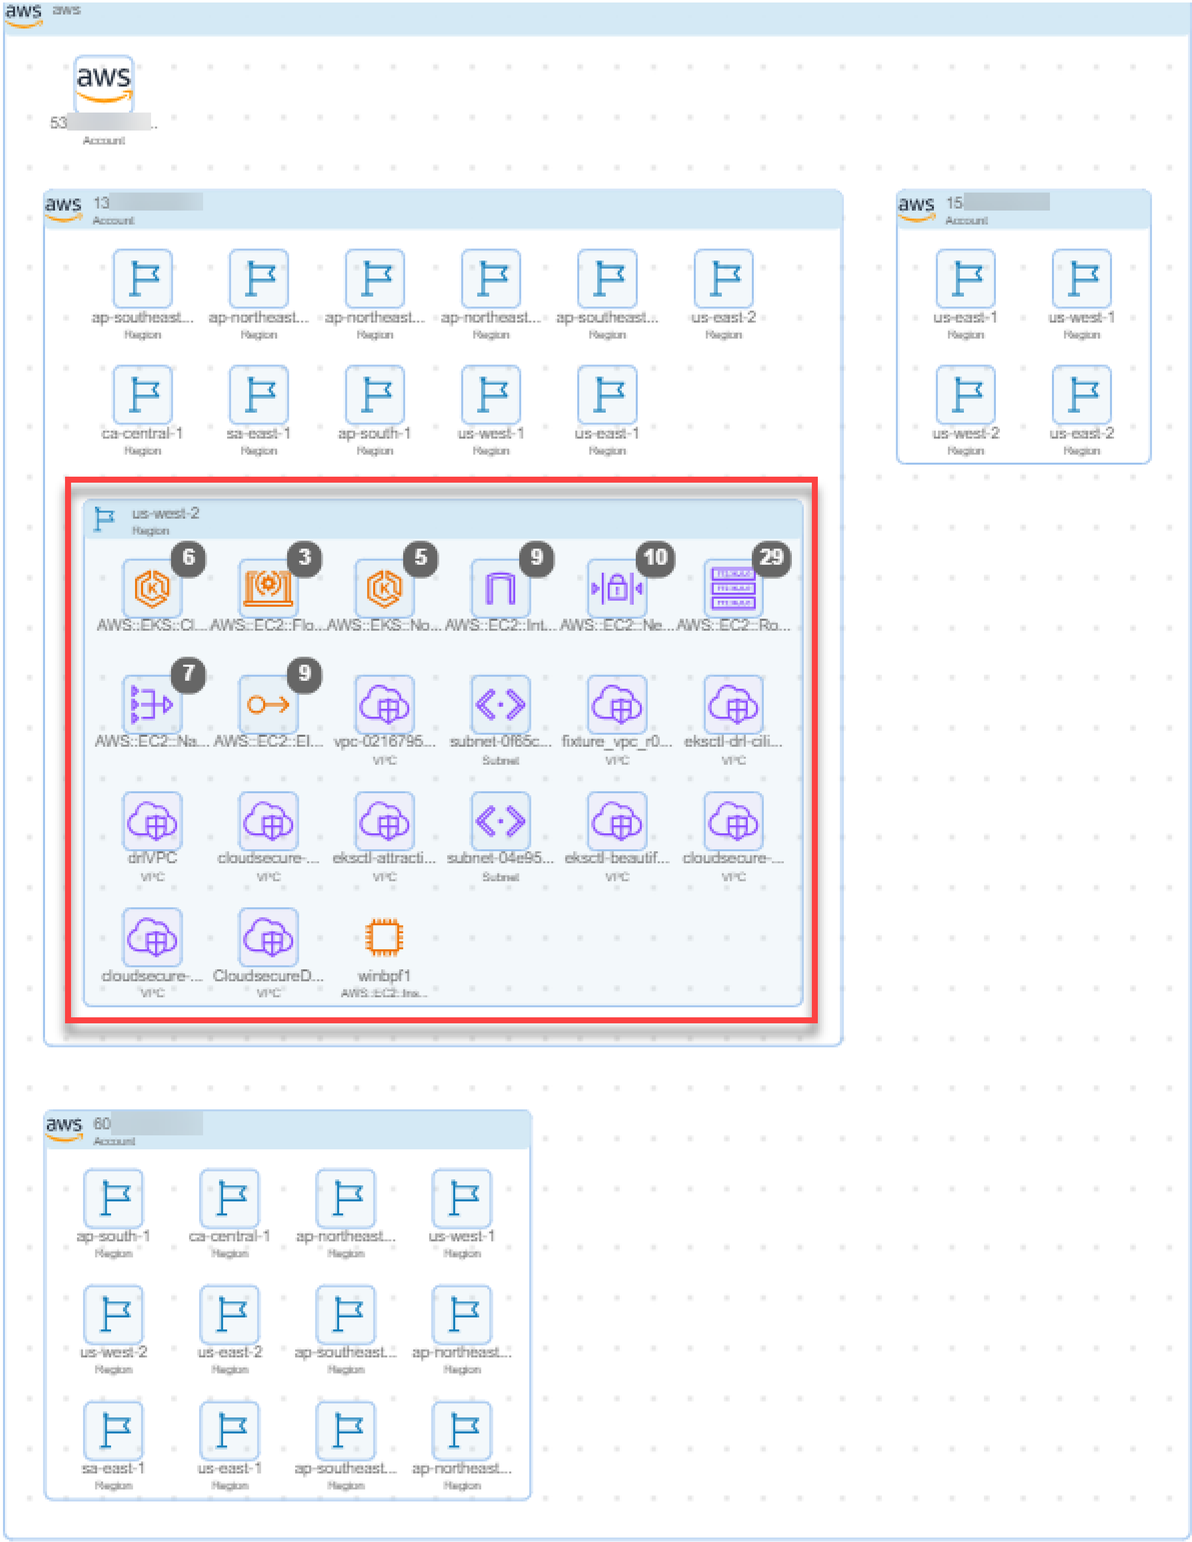

Each region of the Map contains the following types of objects:

- **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.

Limitations for using the Map

When the Map loads, there are limits on the number of objects it displays.

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

These display limitations are not configurable. After you onboard your cloud accounts, Cloud discovers all their resources. To provide optimal 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 Navigating the Cloud Map Infrastructure View [173] 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.

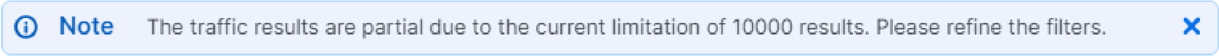
ⓘ Note The traffic results are partial due to the current limitation of 10000 results. Please refine the filters. ✕

When the Map has more than 100 objects, by default it collapses the view for usability rather than displaying all objects.

Traffic flow ingestion limitations

When flow log access is first enabled in Illumio Segmentation for the Cloud, there's a 15-minute latency until traffic flows are first displayed in Cloud Map, Traffic and Inventory pages. After this initial latency, traffic flows are periodically updated every two minutes.

Cloud Map navigation

This topic provides additional details about navigating the Map's Infrastructure View. For details specific to AKS, EKS, and GKE clusters, see Navigating AKS, EKS, and GKE clusters [178].

Interacting with Cloud Map

You have these ways to navigate the map:

- Use the filters at the top of the page to locate and zoom in to specific areas or resources; see Filtering your map resources [173].
- 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:

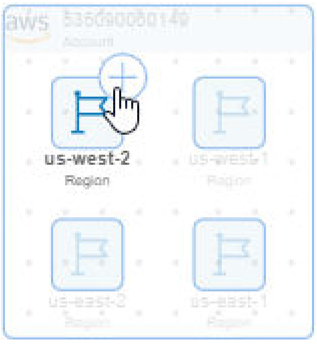

To collapse a group so that you it's not expanded and you see the resources within it, click the minus (-) icon.

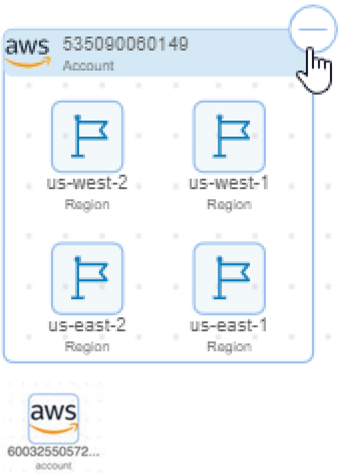

  - Click the white space within a group to zoom in:

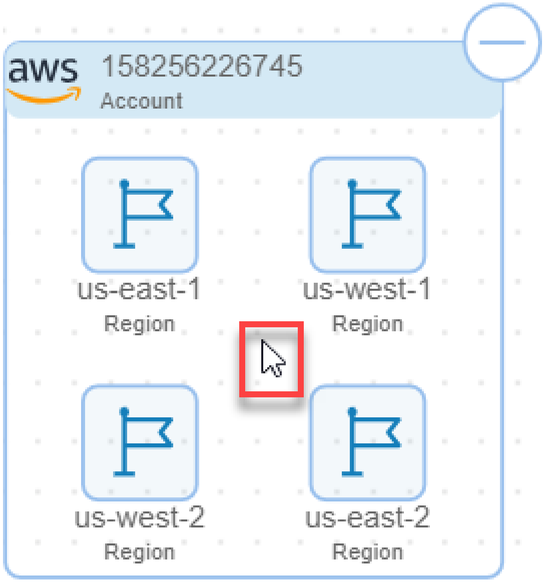

  - Use the map tools to more easily investigate the traffic and resources of most interest. For example, configure the view so that it's isolated to specific resources and relationships, such as only resources with a peered VPC/VNet. Or, use the map tools to zoom to the appropriate level for easier viewing.

⚙	Map Configurations. See Configuring your View [174].
⌖	Zoom to fit entire map within view
⊕	Zoom in
⊖	Zoom out

Filtering your map resources

At the top of the page, the Map includes a **Resource** filter. To make it appear, click the magnifying glass icon:

🔍

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 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:

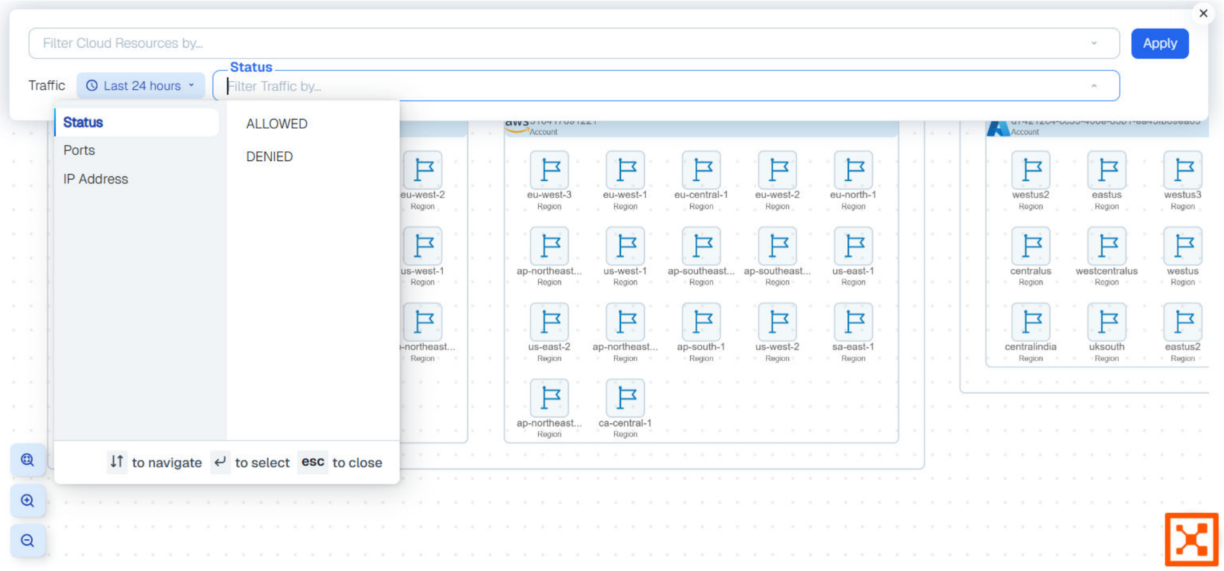

IMPORTANT

As you use the filters to manipulate the Map display and display details about accounts and the resources in them, a message may display saying 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.

See Limitations for Using Map [171].

When filtering by IP addresses, use CIDR blocks to include a range of IP addresses. For example, adding "/16" to an IP address will search for flows with IP addresses starting with the same first 16 bits as the specified IP address, such as 10.104.XXX.XXX. Similarly, adding "/24" or "/30" will search for flows with IP addresses starting with the same first 24 or 30 bits as the specified IP address, respectively. Note that the number after the slash specifies the prefix length.

NOTE

In GCP, under a network you have subnets. The subnet hierarchy is cloud > projects > global resources (like VPC) > region > regional resources. However, when there is no filter applied, the region is not rendered on the map. This means that all the regional resources are displayed rather than grouped by region. To display regional resources grouped by region, use a filter (such as cloud, account, or any feature).

When considering filtering, note that GCP virtual machines are zonal but display at the regional level. You can see the zone in the resource properties panel.

Note that if your GCP account is associated with multiple NICs under various subnets and VPCs, the map displays the instance and traffic at the region level instead of the project level. For easier navigation in GCP environments, Illumio recommends using filters on the map, such as filter `cloud:gcp`, and using additional filters as you navigate to narrow your view as needed.

Configuring your 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.

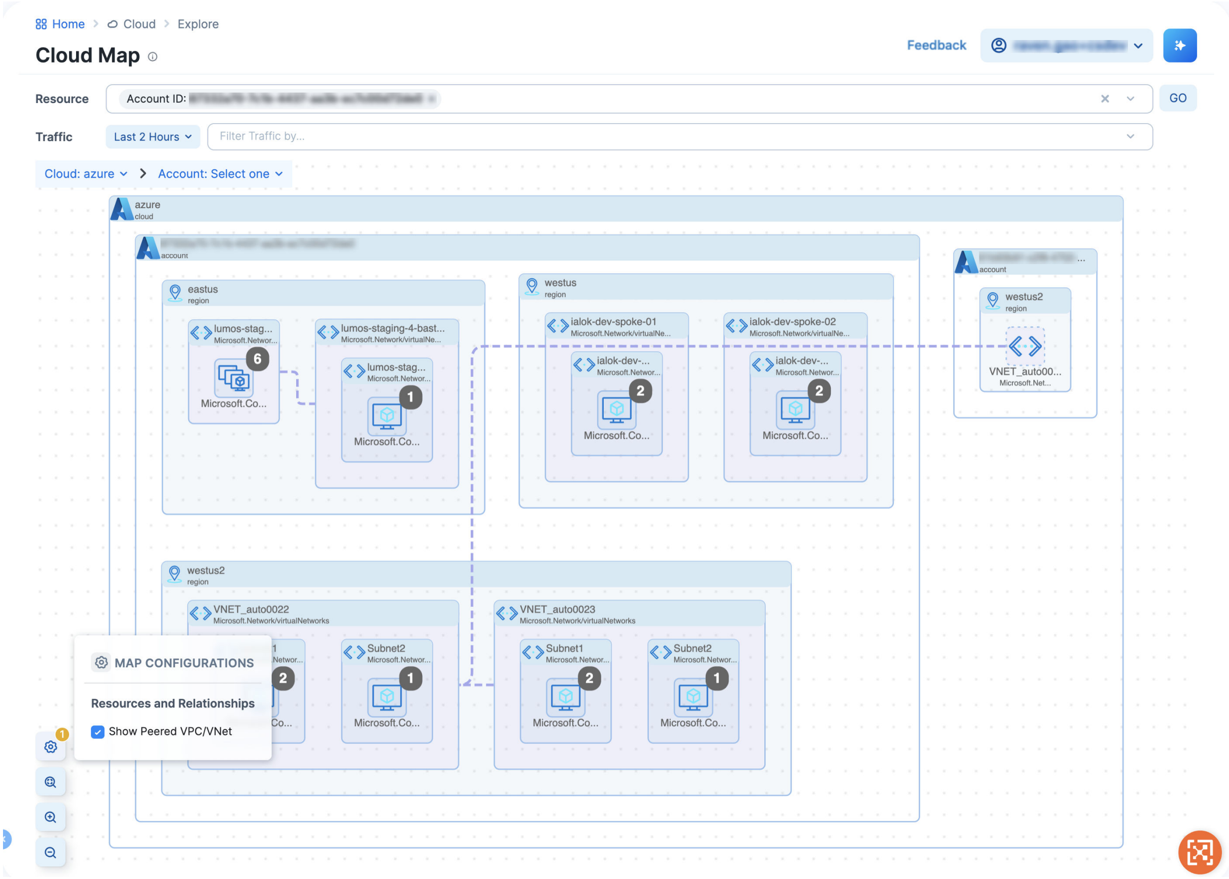

Home > Cloud > Explore

Cloud Map

Resource: Account ID: [redacted]

Traffic: Last 2 Hours

Cloud: azure > Account: Select one

MAP CONFIGURATIONS

Resources and Relationships

☑ Show Peered VPC/VNet

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

Display resource side panel

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

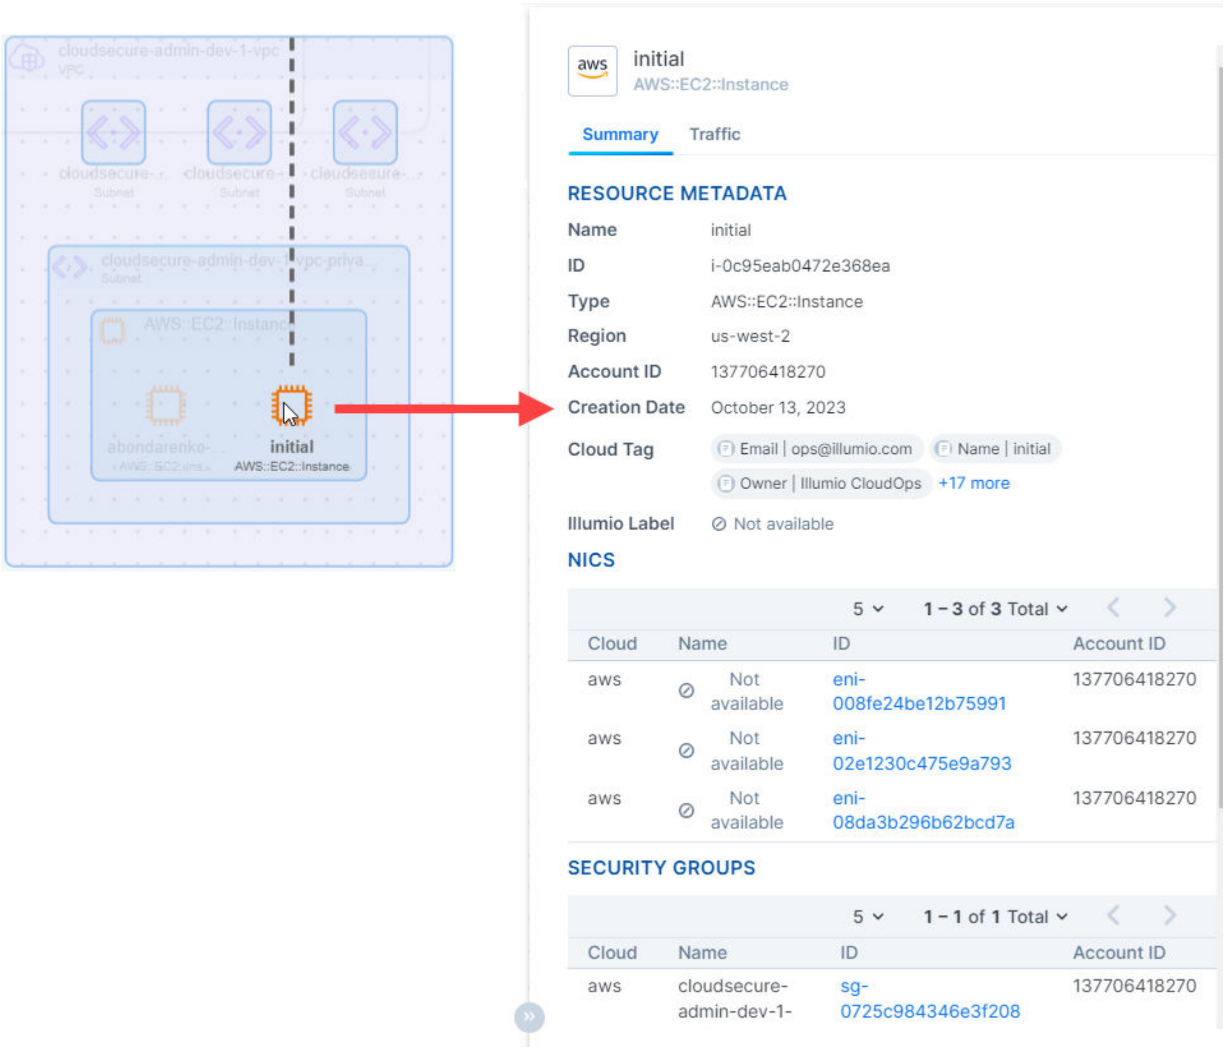

initial
AWS::EC2::Instance

Summary Traffic

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

5 ⌄ 1 – 3 of 3 Total ⌄ ‹ ›

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

5 ⌄ 1 – 1 of 1 Total ⌄ ‹ ›

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 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.

Map traffic lines

The 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.

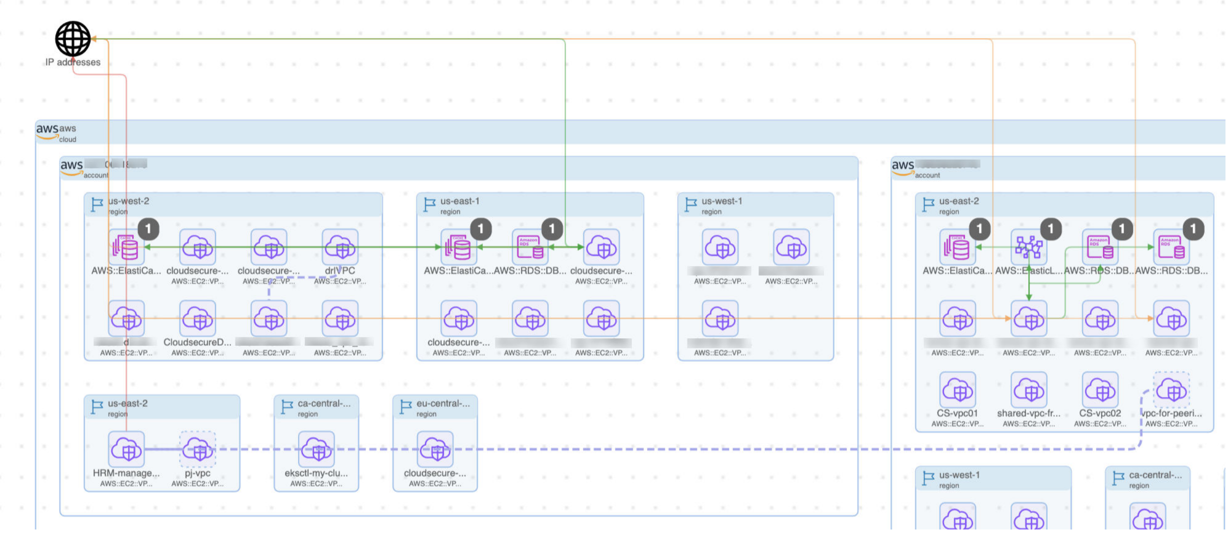

Orange lines indicate mixed state (both denied and allowed) traffic. Green lines indicate allowed traffic. Red lines indicate denied traffic. These traffic lines are displayed from the lowest level node selected. For example, you may have green lines between two regions, indicating that strictly regional 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.

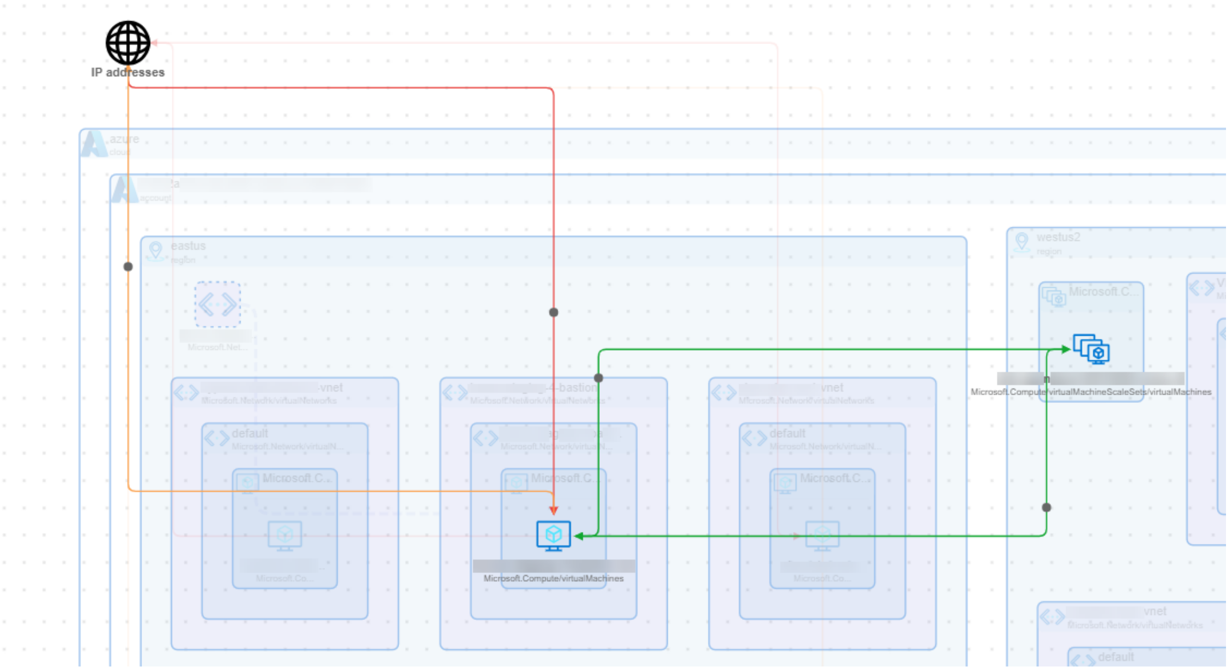
IP addresses
azure cloud
account
eastus region
westus2 region
vnet
bastion
default
Microsoft.Network/virtualN...
Microsoft.C...
Microsoft.Compute/virtualMachines
Microsoft.Compute/virtualMachineScaleSets/virtualMachines

Navigating AKS, EKS, and GKE clusters

Learn how to navigate your AKS, EKS, and GKE clusters in the Map. For information on general Map navigation, see Cloud Map navigation [171].

About AKS, EKS, and GKE Clusters

AKS, EKS, and GKE clusters have two deployment structures:

- Node Groups: These contain EC2 instances that comprise the nodes. This is the structure assumed for this topic.
- Serverless Option (Fargate Profiles): These contain Fargate nodes. Fargate Profiles are shown for visibility purposes but do not have flow support or Map support at this time.

Drilling into AKS, EKS, and GKE Clusters

Click on a cluster to learn more about it. This opens a details panel, containing the following tabs:

- Summary: This tab contains information such as name, ID, cloud, region, category, private and public addresses, and so forth.
- Attached Resources: This tab contains information about things like any attached Network Interfaces, Subnets, Security Groups, and so forth. Resources that do not have flow support, such as Fargate profiles, will have information in only the details panel.
- Traffic: This tab contains traffic information, much like that shown in the Traffic page, but filtered for the selected luster in this example.
- Resource Map: This tab shows the structure comprising the cluster. In this example, assume it shows one node group that contains two nodes (EC2 instances) that are container hosts. However, it could potentially contain many more.

  The Map displays only the traffic flowing to and from the AKS, EKS, or GKE cluster as a whole. However, when you open the cluster in the resource map, it expands to reveal the traffic within the cluster. For example, the cluster shows the node groups contained within it. If you further expand a node group in the resource map, it displays the EC2 instances (nodes) within the group. The traffic between these nodes is displayed on the Traffic page. The image shows an example of an EKS cluster.

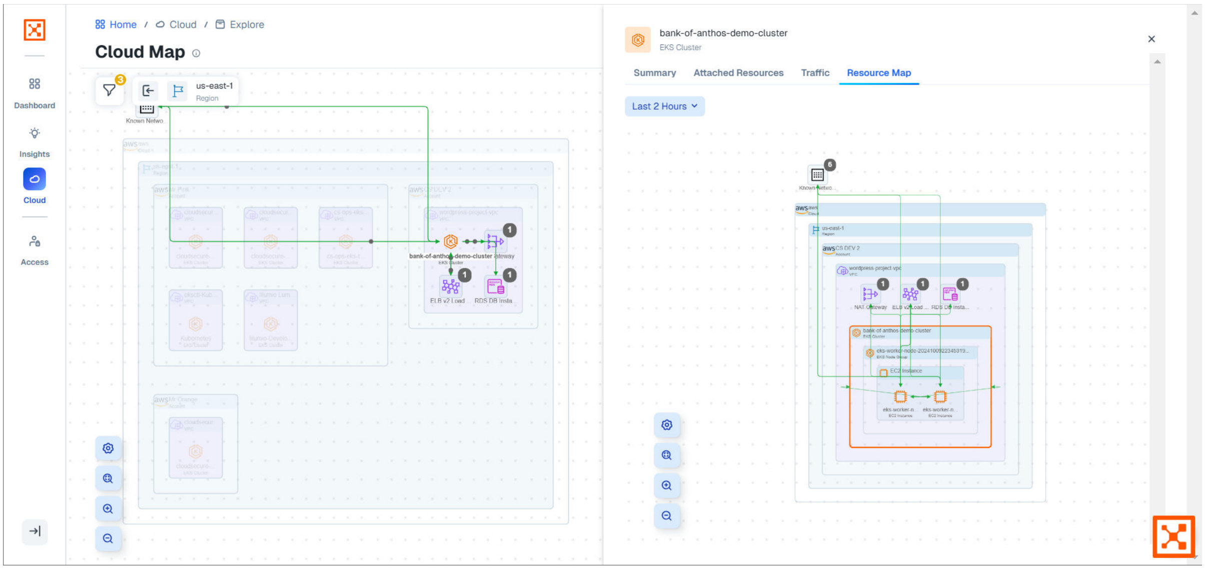

Additionally, communication between the Kubernetes control plane and its compute instances is reflected as AKS-to-node, EKS-to-node, or GKE-to-node traffic on the Traffic page.

- See the Amazon website.
- See the Azure website.
- See the Google website.

NOTE

Cluster VMs appear only under the GKE cluster node pool and not as standalone instances.

Navigating Azure Firewalls

To gain insights into firewall relationships with VNETs (differentiated by hub and spoke), firewall policy details, and traffic flows passing through the firewall, you will need to apply the correct filters and drill down through the results displayed on the Map. This makes it easier to see which VNets talk to each other at a glance, as well as to see which VNets need to have firewalls applied. You can also browse from an application to its Details Page Map tab to visualize Azure Firewalls specific to that application.

NOTE

Illumio Segmentation for the Cloud does not support Classic Azure Firewall.

For navigating the Infrastructure View, see Cloud Map navigation [171].

For an overview of Azure Firewalls, see Azure Firewalls Overview [167].

Here are guidelines and instructions for navigating Azure Firewalls in the Map and from an application. There are three scenarios described:

- Filtering the Map by Azure Firewall Example [179]
- Filtering the Map by Account Example [180]
- Browsing from an Application Example [182]

Filtering the Map by Azure Firewall Example

Follow these steps to get a broad view of your environments firewalls.

1. In the Map, use the filter to select Azure Firewalls (Resource Type = Microsoft.Network/azureFirewalls) and click **Apply**.

   The Map shows VNets containing Azure Firewalls but does not show which ones are hubs or spokes. An example VNet with an Azure Firewall:

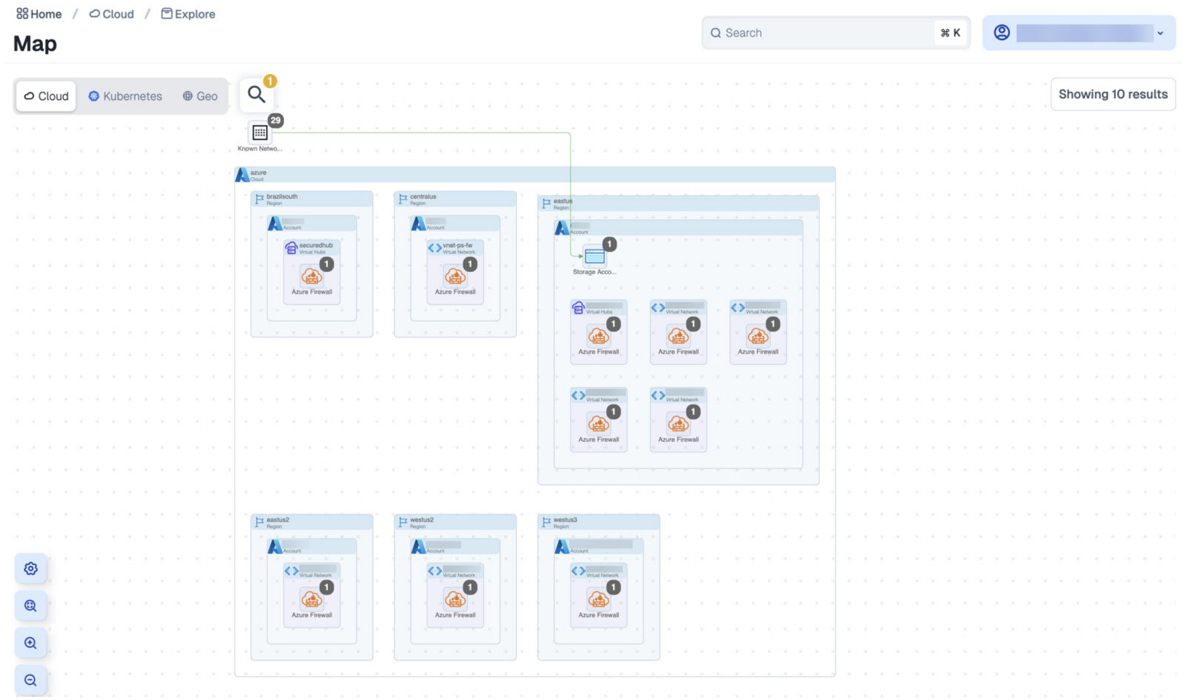

You will need to show the firewall topology to see which ones are hubs vs. spokes, as described in the next step.

2. Click the gear icon to open the Map Configuration dialog and select **Show Firewall Topology**.

   The Map displays your VNets containing firewalls and reveals the Firewall Topology, showing peering traffic and labeling them as hubs and spokes. To see how to expand hub VNets and hover over them for more details, see Filtering by Account Example [180].

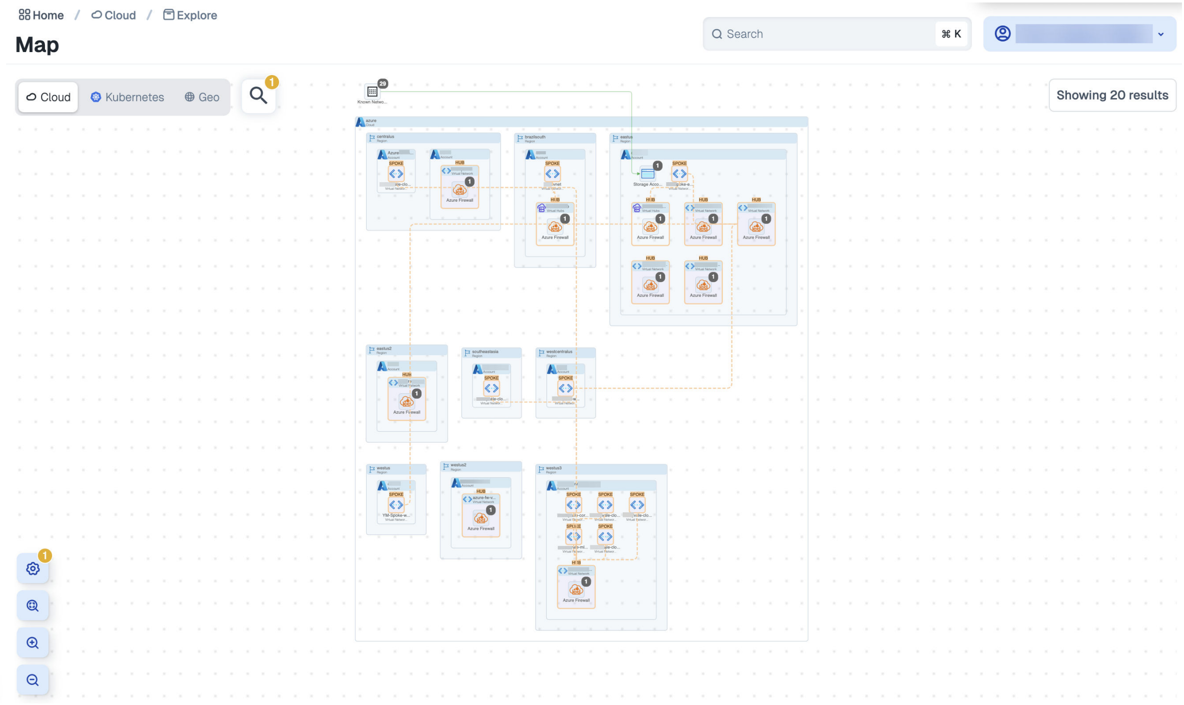

Filtering the Map by Account Example

Follow these steps to filter the Map by an account that you want to check for firewalls.

1. In the Map, use a filter like Account = <account name> and click **Apply**. It shows the Illumio Segmentation for the Cloud Map for that account. You can also filter by region, resource type, etc.
2. Click the gear icon to open the Map Configuration dialog and select **Show Firewall Topology**.

   If none of the VNets associated with the filter query have firewalls, a message displays to that effect.
3. To use the filter to select for different query, click the back arrow icon in the Illumio Segmentation for the Cloud Map (not the browser) and use a filter like Region = eastus and click **Apply**.

   If the VNets associated with the filter query do have firewalls, the Illumio Segmentation for the Cloud Map displays the regions and the accounts that match the filter, including those accounts that have Azure firewalls, as indicated border highlighting. Click an account with firewall topology to expand it. In this example you see this:

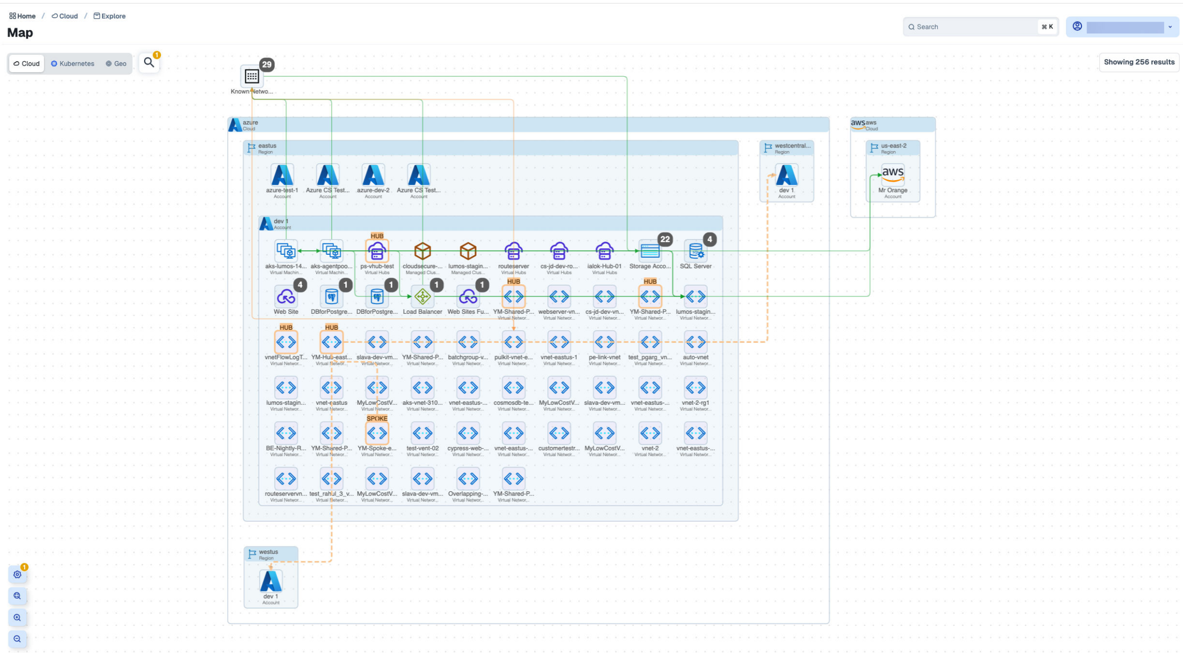

4. When you see hubs and spokes, which are labeled as such, look for things like the following:
   - Peering connections between hubs and spokes, as indicated by dashed lines
   - Traffic entering or leaving the firewall, as indicated by green lines. Note that if a VNet or virtual hub is collapsed you see traffic from it, and if it is open you see traffic from the particular resource.
5. Expand the hub VNet.

   The expanded hub is "pulled out" to facilitate the enlarged view and shows an Azure Firewall indicated by the flame icon, along with peering traffic lines.

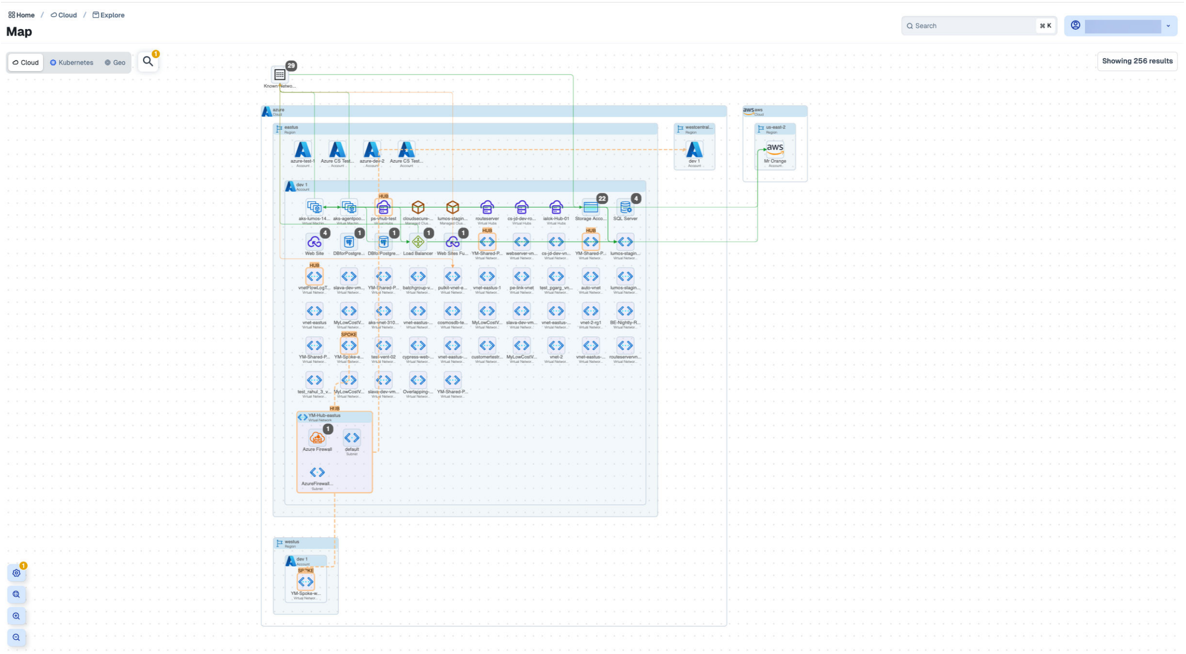

6. Hover on a hub VNet to see its details. In this example you see this:

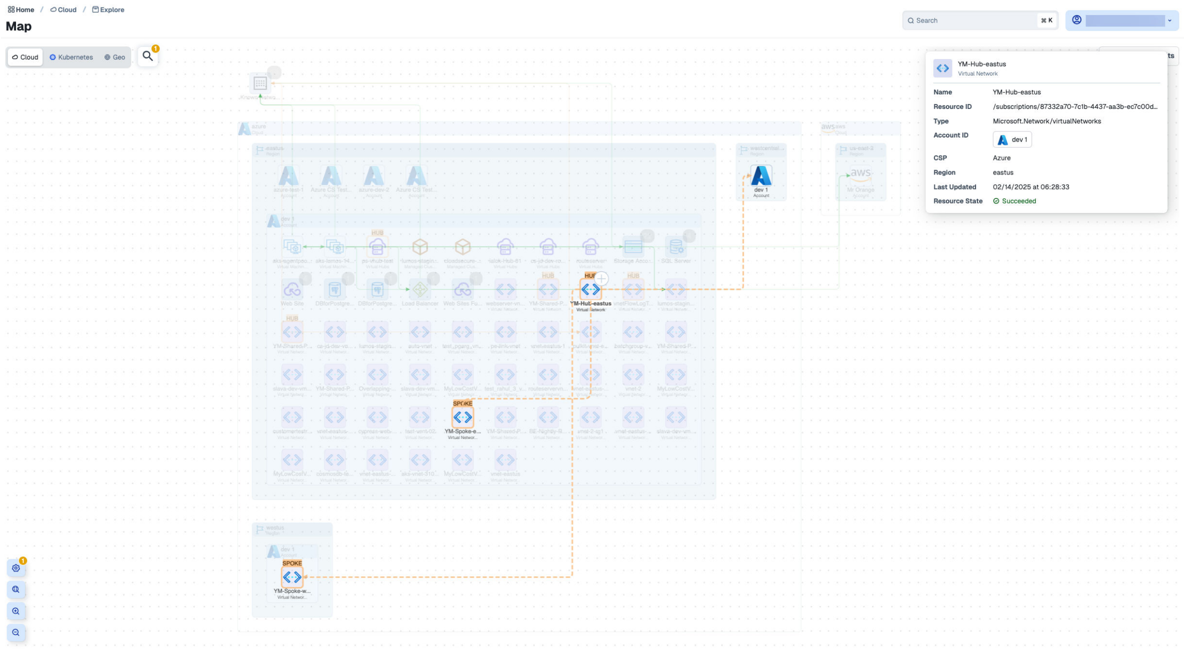

YM-Hub-eastus
Virtual Network

Name	YM-Hub-eastus
Resource ID	/subscriptions/87332a70-7c1b-4437-aa3b-ec7c00d...
Type	Microsoft.Network/virtualNetworks
Account ID	dev 1
CSP	Azure
Region	eastus
Last Updated	02/14/2025 at 06:28:33
Resource State	Succeeded

Browsing from an Application Example

Follow these steps to navigate from an application to its application-specific map to gain visual insights for firewalls associated with that application.

1. Instead of starting from the Map, browse to **Explore > Applications** and select an application that you know to have a firewall.

   The application's Details page opens to the the Summary tab by default.
2. Click the **Map** tab in the Details page.

   The Map shows the application's VNets that contain firewalls but does not show which ones are hubs or spokes.

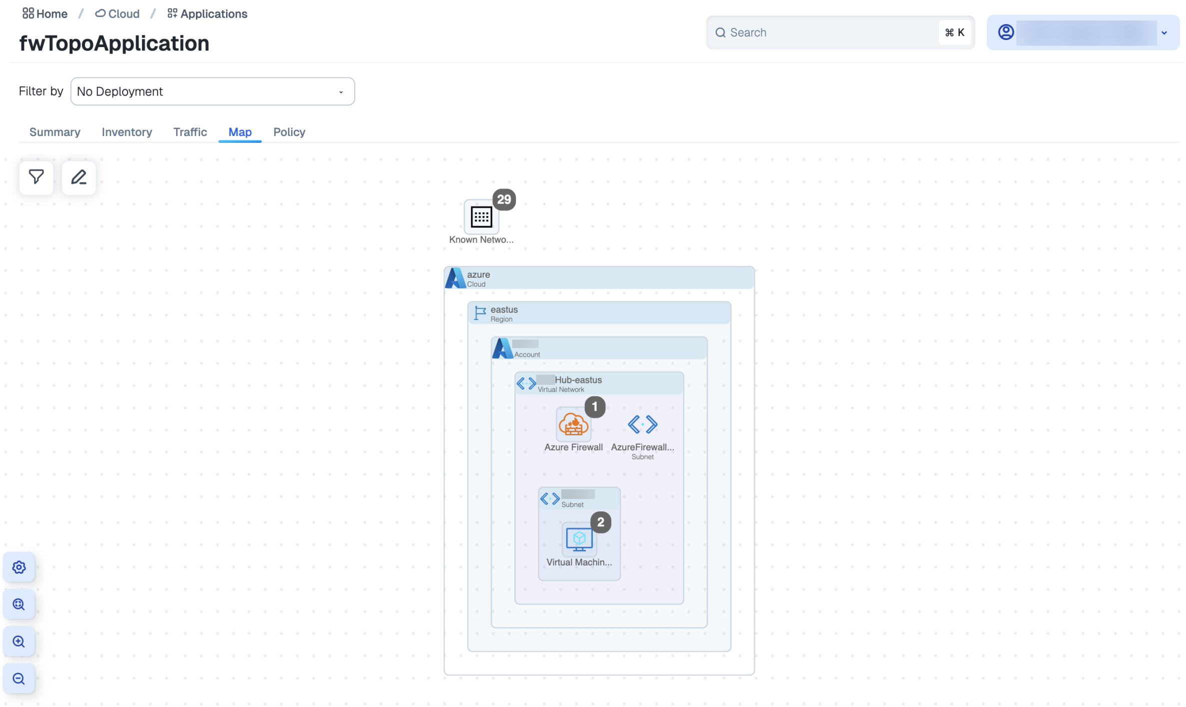

You will need to show the firewall topology to see which ones are hubs vs. spokes, as described in the next step.

3. Click the gear icon to open the Map Configuration dialog and select **Show Firewall Topology**.

   The map displays your application VNets containing firewalls and reveals the Firewall Topology, showing peering traffic and labeling them as hubs and spokes. To see how to expand hub VNets and hover over them for more details, see Filtering by Account Example [180].

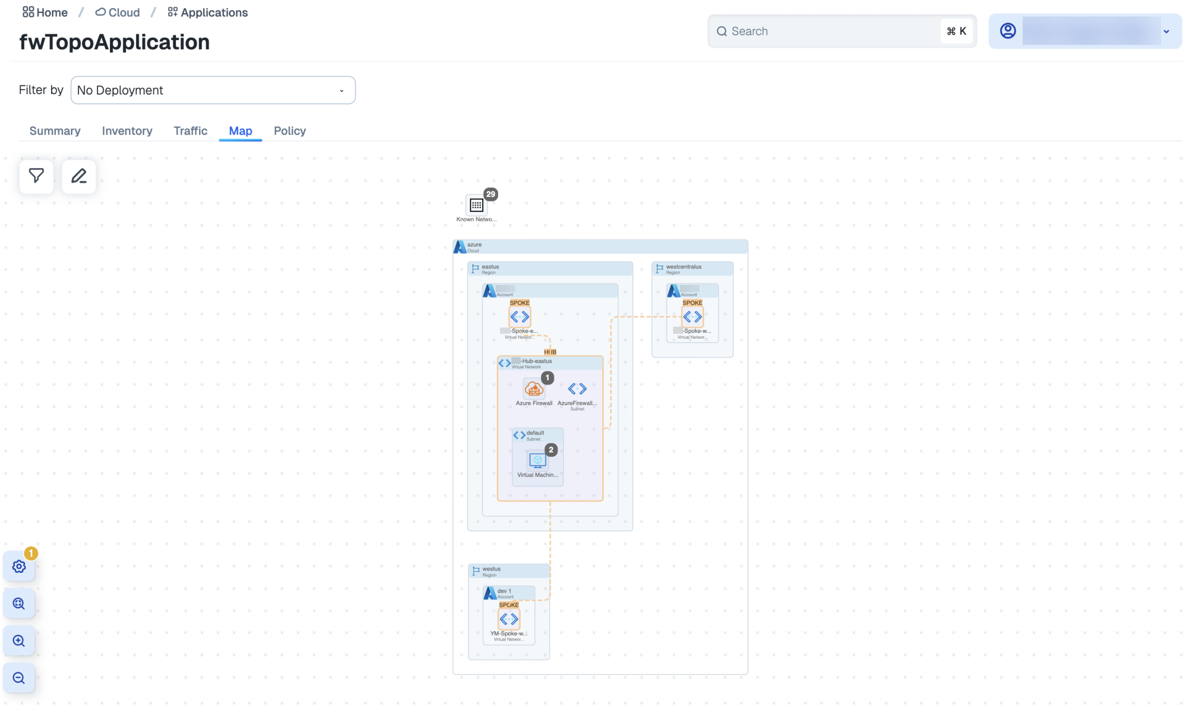

Navigating the Map Kubernetes View

This topic provides details about navigating the Kubernetes View. Kubernetes administrators benefit from advanced filter and map capabilities to quickly locate specific resources like clusters, nodes, namespaces, or workloads for both cloud-managed and self-managed clusters.

See Agentless Containers overview [148].

For general information on the Map, see Map [168].

To see the Kubernetes View on the Map, navigate to Cloud > Explore > Map and click the **Kubernetes** button. For general Map navigation instructions, see Navigating the Map Infrastructure View [171].

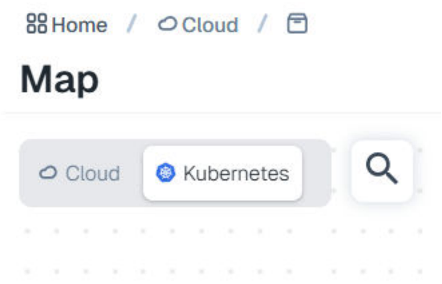

Filtering Your Map Resources

At the top of the page, the Map includes a **Resource** filter. The **Resource** filter includes several options, including Cloud, Account, Region, Resource Type, Cluster, Namespace, and others. Note that the Kubernetes View supports only resource-based filtering, so it is not refined by source or destination, for example. To make the filter appear, click the magnifying glass icon:

By default, when you open the Kubernetes view, the Map displays all onboarded Kubernetes clusters that are either cloud-managed such as EKS and AKS, or self-managed. Only the cloud-managed clusters are rendered under cloud account, region, and more. Self-managed clusters are rendered in a self-managed box, which means they do not have cloud-managed hierarchies such as accounts and regions.

If you have no filters selected in the Kubernetes View, you will see only clusters and traffic at the cluster level. This high-level default view provides context. To leave this default view and see traffic details and what is inside your clusters:

- Drill down on a traffic line
- Drill down on a cluster
- Apply filters (recommended)

For example, you could filter by namespaces.

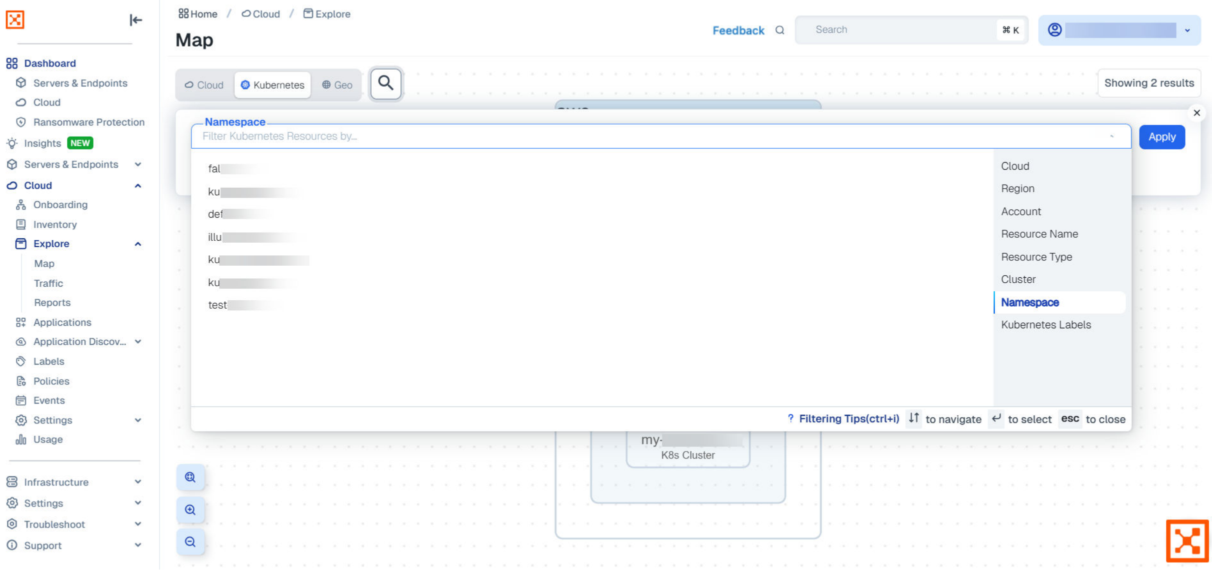

The Map shows your Kubernetes resources as filtered by namespaces, displaying the contents and traffic inside the cluster.

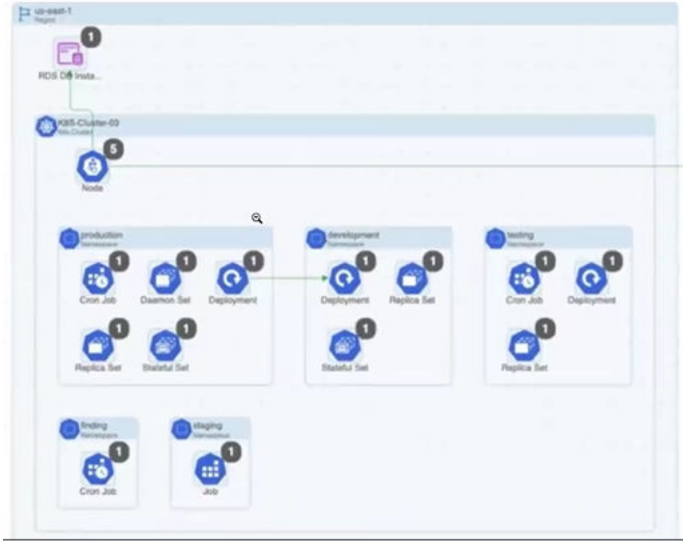

IMPORTANT

As you use the filters to manipulate the Map display and display details about accounts and the resources in them, a message may display saying that the Map 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.

See Limitations for Using Map [171].

Display Resource Side Panel

When you click a Kubernetes cluster in the Map, it opens a right-side panel that displays the resource metadata. For example, you can click a Kubernetes cluster to see:

- Summary information: Name, ID, Created, Type, Cloud, Region
- Attached Resources: Custom Resource Definition, Daemon Set, Deployment, Endpoint, Namespace, and more.
- Cluster map: Zooms in to the cluster you selected
- Cloud/Kubernetes button: Toggle from Kubernetes to Cloud to switch between a Kubernetes cluster and its Illumio Segmentation for the Cloud equivalent such as an EKS. In the Kubernetes view, this toggle appears for Kubernetes clusters and Kubernetes nodes. In the Map Infrastructure view, this toggle appears for EKS clusters and EC2 instances inside EKS clusters.

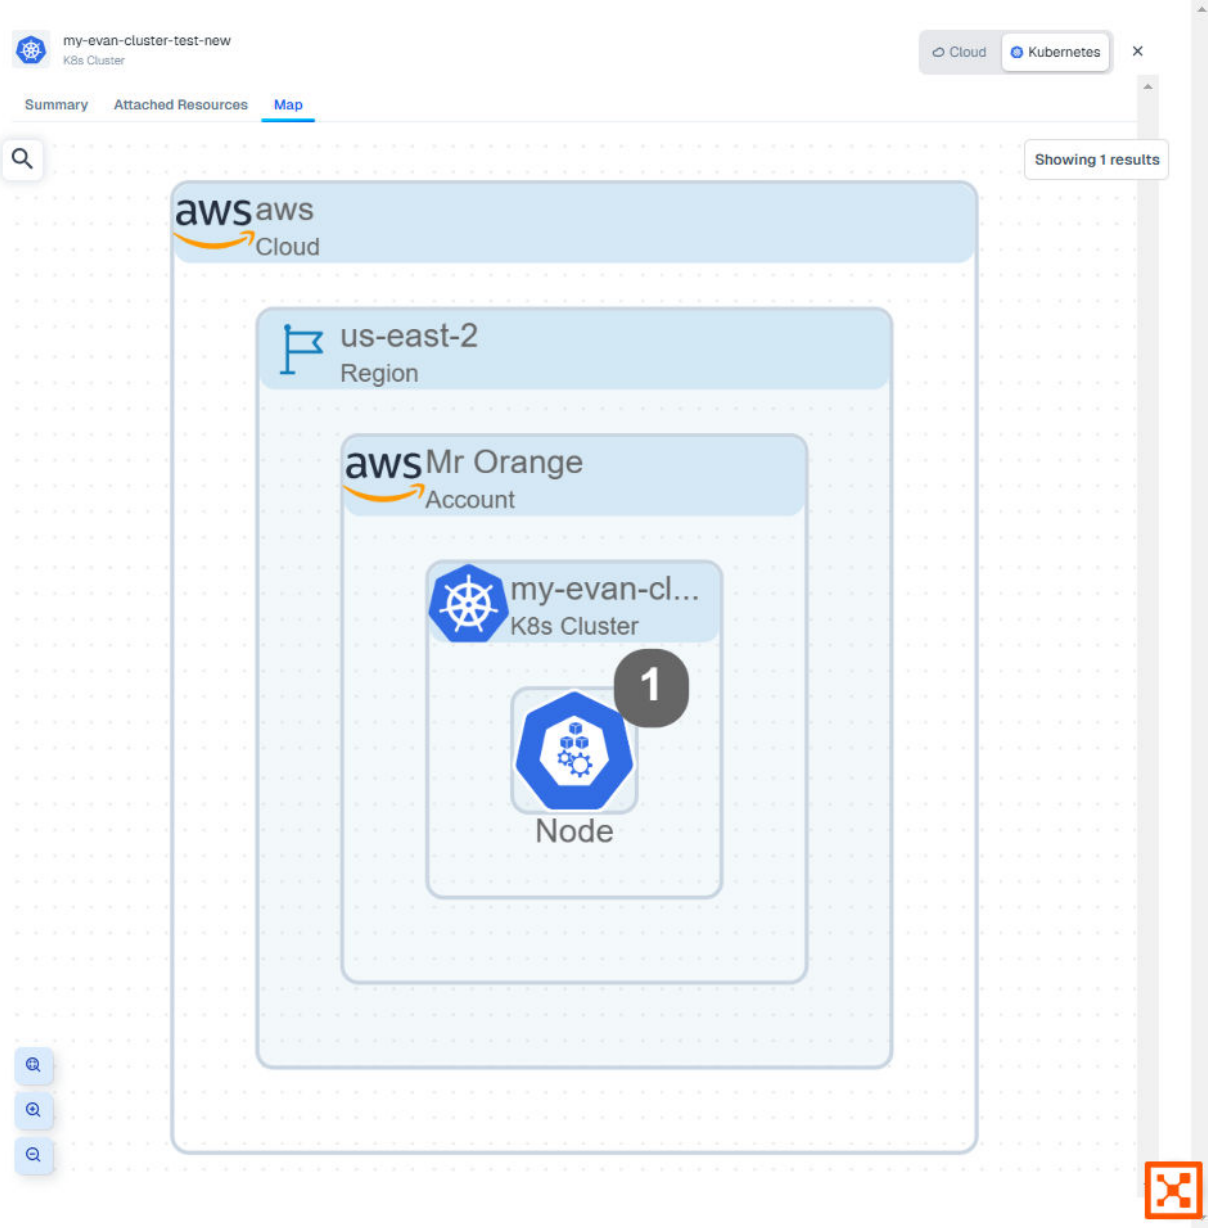

At this time, the Map only supports displaying traffic data for Azure Firewalls, VM (Azure), and EC2 instance (AWS) resources. For resources that don't support displaying traffic flows, the panel includes a **Summary** tab only.

Map Traffic Lines

The 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.

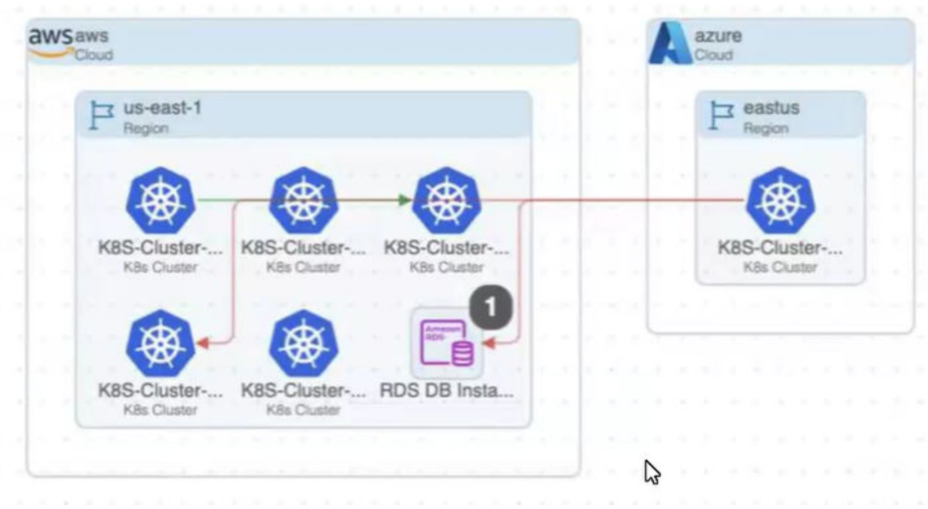

Orange lines indicate a mixed state (both denied and allowed) traffic. Green lines indicate allowed traffic, while red lines indicate denied traffic. These traffic lines are displayed from the lowest level node selected. For example, you may have green lines between two regions, indicating that strictly 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 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 the hover effect 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.

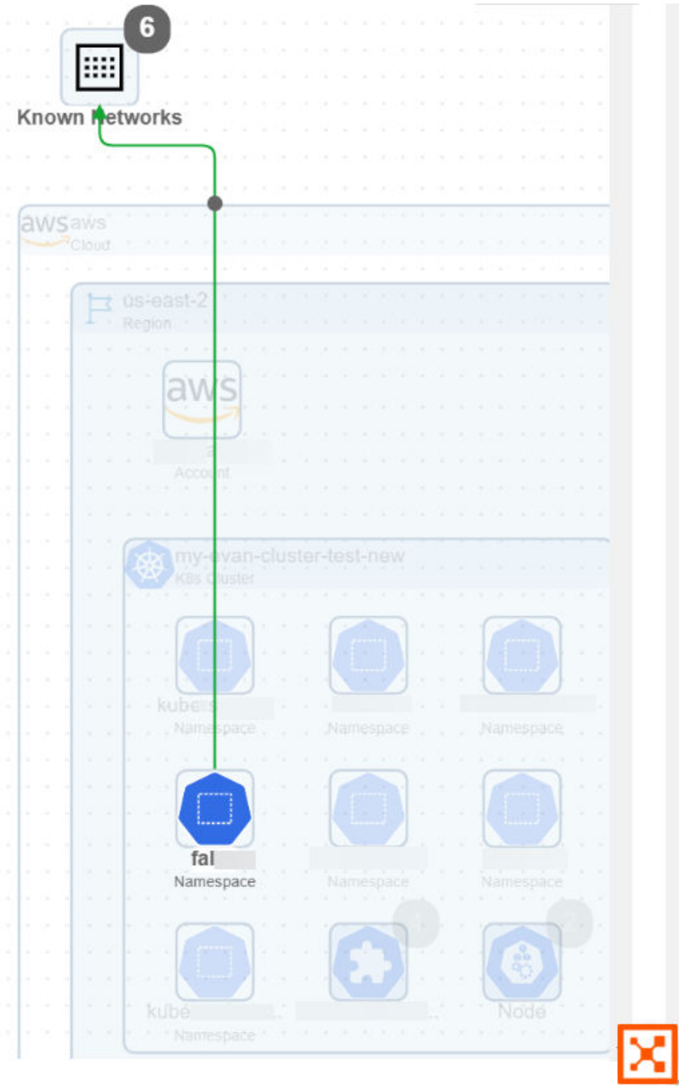

Dashboard overview

The Dashboard provides a quick way to understand your cloud presence at a glance.

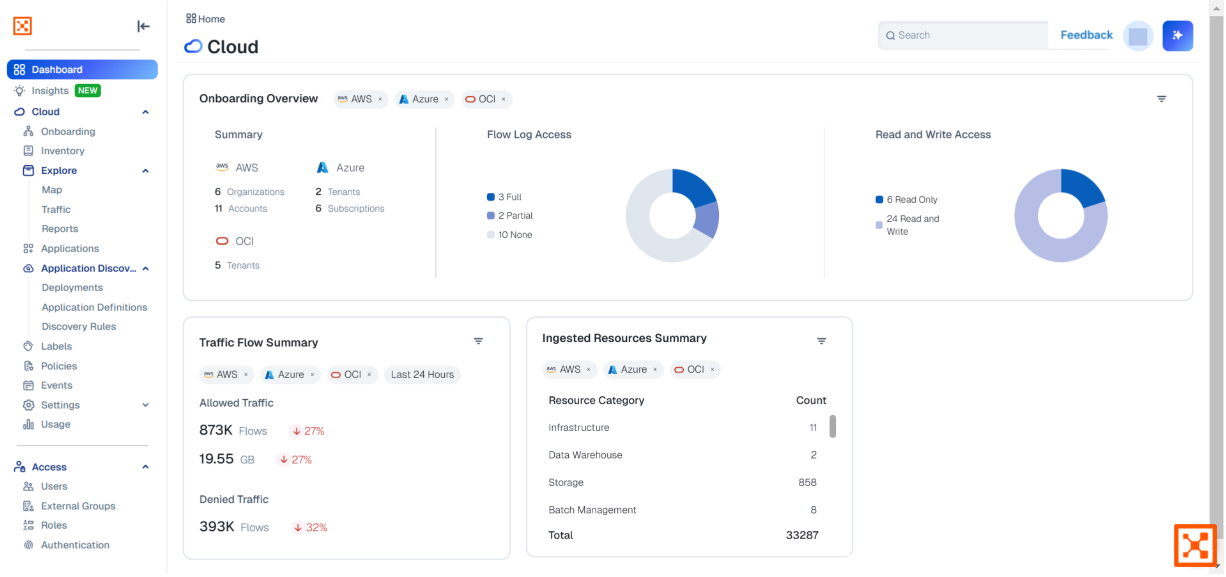

Resource Category	Count
Infrastructure	11
Data Warehouse	2
Storage	858
Batch Management	8
Total	33287

Dashboard sections

The dashboard provides the following sections:

Onboarding overview

This has three sub-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.
- Click the Flow Log Access section to go to the Flow Log Access page. See Grant flow log access to your CSPs [141].
- 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 section's filters

Traffic Flow summary

This has two sub-sections, Allowed Traffic and Denied Traffic.

- Click the count of Allowed Traffic to go to the Traffic [219] 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 section 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 section'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 includes two main sub-sections: Resource Category and Count, displaying resource types that support visibility.

- Click the filter icon in the upper right-hand corner of the section to change the section'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 section.
- Click any of the following in the Resource Category column to see both billable and non-billable 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 section. Hover over the numbers to see the percentage of the total that each represents.

The Illumio Segmentation for the Cloud dashboard gets updated over time, so check here as new sections are added.

Search and filtering

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

- 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. 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 your search on a given page, a subsequent search would give results for only your AWS cloud resources but not your Azure cloud resources.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Context+Based+Search.mp4

Pages with the feature

- Inventory
- Traffic
- Map
- Application > Inventory tab
- Application > Traffic tab

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):
  - Cloud
  - Account Name
  - Account ID
  - Region
  - Resource Type
  - Resource Group
  - Resource Name
  - Resource State
  - VPC/VNET ID
  - Subnet ID
  - Cloud Tags
  - Labels
  - Categories
- `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.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Using+Filter+Based+Search.mp4

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.

Events

This topic describes the purpose of the Illumio Segmentation for the Cloud 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, removes system events older than 30 days, and removes audit events older than 365 days.

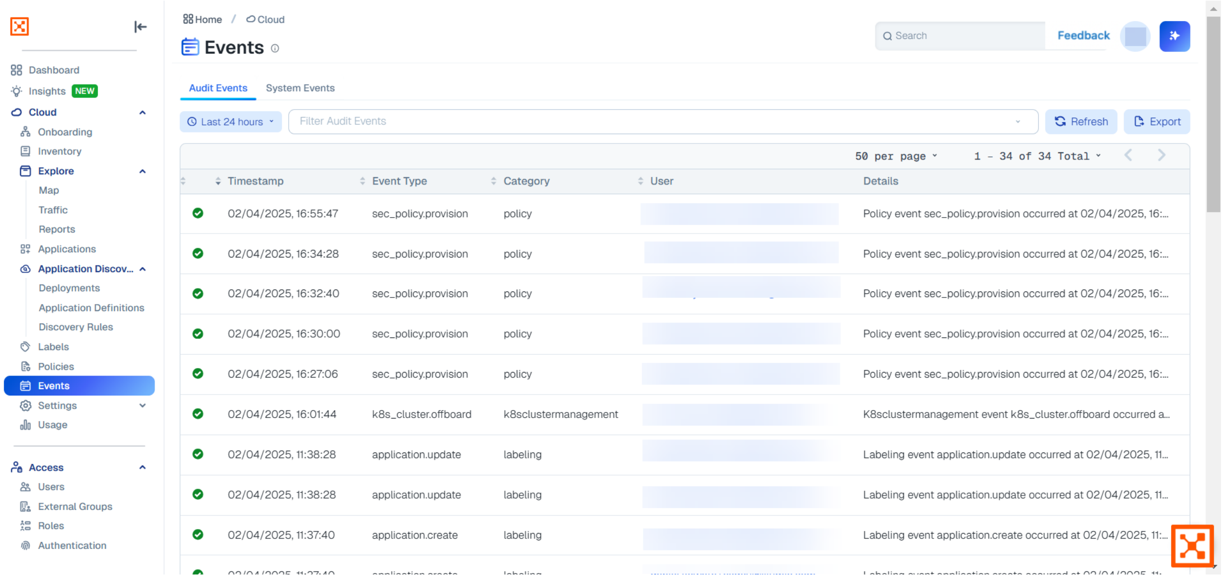

Events

50 per page   1 – 34 of 34 Total

	Timestamp	Event Type	Category	User	Details
✓	02/04/2025, 16:55:47	sec_policy.provision	policy	[redacted]	Policy event sec_policy.provision occurred at 02/04/2025, 16:...
✓	02/04/2025, 16:34:28	sec_policy.provision	policy	[redacted]	Policy event sec_policy.provision occurred at 02/04/2025, 16:...
✓	02/04/2025, 16:32:40	sec_policy.provision	policy	[redacted]	Policy event sec_policy.provision occurred at 02/04/2025, 16:...
✓	02/04/2025, 16:30:00	sec_policy.provision	policy	[redacted]	Policy event sec_policy.provision occurred at 02/04/2025, 16:...
✓	02/04/2025, 16:27:06	sec_policy.provision	policy	[redacted]	Policy event sec_policy.provision occurred at 02/04/2025, 16:...
✓	02/04/2025, 16:01:44	k8s_cluster.offboard	k8sclustermanagement	[redacted]	K8sclustermanagement event k8s_cluster.offboard occurred a...
✓	02/04/2025, 11:38:28	application.update	labeling	[redacted]	Labeling event application.update occurred at 02/04/2025, 11...
✓	02/04/2025, 11:38:28	application.update	labeling	[redacted]	Labeling event application.update occurred at 02/04/2025, 11...
✓	02/04/2025, 11:37:40	application.create	labeling	[redacted]	Labeling event application.create occurred at 02/04/2025, 11:...

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

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Export+an+Events+Report.mp4

1. Click **Export** on either the Audit Events or System Events tab.
2. Edit the report name and select the format.
3. Click the **Scheduling Section** toggle to the on position to schedule the export unless you want to export the report immediately.
4. If you choose to schedule your report, select your recurrence and time.
5. Click **Save** when done.
6. Go to the Reports [212] page to download the exported report.

Illumio visibility for resource types

This page lists the resource types that Illumio Segmentation for the Cloud supports for metadata searches and for displaying in inventory once you onboard your Cloud Service Provider (CSP) accounts. In addition to appearing on the Inventory page, these resource types also appear on the Map and the Traffic pages. Policies can be written for a subset of these resources. See Policy enforcement and resource types [262].

Illumio Segmentation for the Cloud supports the AWS, Azure, and OCI cloud resource types listed below along with the attached resources that appear in the Inventory resource details panel. A resource type is defined as the instances of the object types that Illumio Segmentation for the Cloud retrieves from the CSP SDKs or APIs. These attached resource types provide you an overview of the cloud infrastructure and help you visualize the resource types on the Map. Resource types not listed below are not supported.

Resource types come from the CSP (e.g., a resource type could be AWS::EC2::VPC is a VPC in the EC2 category. Illumio also defines resource categories (e.g., VPCs are placed under the Network Management category as seen in the AWS table below). A resource type can be defined as all the types that build up to a resource.

Note that some resource types that don't come directly from the CSP also appear on the Inventory page.

Illumio Segmentation for the Cloud's support for network traffic flows is tied to resource types, because categories such as "database" can include many different types of resources. Keep in mind that although traffic may appear on the Traffic page, not all flows are "decorated." Decoration means that flows are labeled with resource type details. For example, traffic between IPs such as 1.1.1.1 and 2.2.2.2 may be shown, but if Illumio hasn't associated those IPs with specific resource types, the flow appears undecorated. When Illumio does recognize and label the source and destination IPs with their resource types, it might show something like: 1.1.1.1 (vm1) → 2.2.2.2 (db1). See the image.

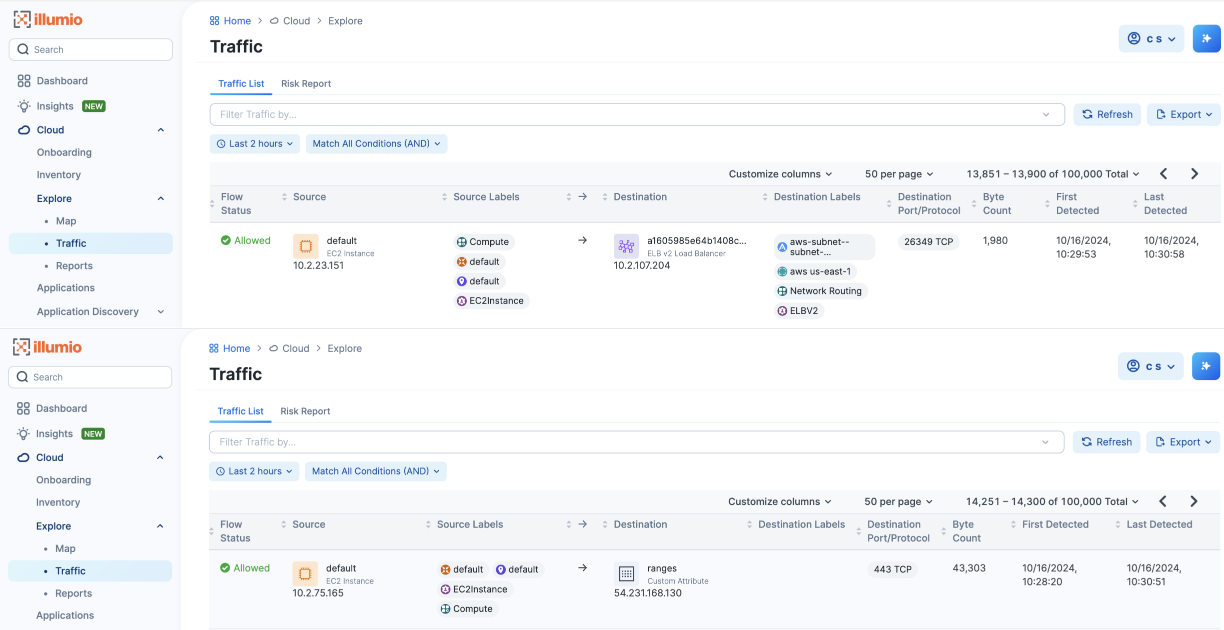

Flow Status	Source	Source Labels	→	Destination	Destination Labels	Destination Port/Protocol	Byte Count	First Detected	Last Detected
Allowed	default EC2 Instance 10.2.23.151	Compute default default EC2Instance	→	a1605985e64b1408c... ELB v2 Load Balancer 10.2.107.204	aws-subnet--subnet-... aws us-east-1 Network Routing ELBV2	26349 TCP	1,980	10/16/2024, 10:29:53	10/16/2024, 10:30:58

Flow Status	Source	Source Labels	→	Destination	Destination Labels	Destination Port/Protocol	Byte Count	First Detected	Last Detected
Allowed	default EC2 Instance 10.2.75.165	default default EC2Instance Compute	→	ranges Custom Attribute 54.231.168.130		443 TCP	43,303	10/16/2024, 10:28:20	10/16/2024, 10:30:51

Decorated (top) vs. undecorated (bottom) flows

For AWS, Illumio Segmentation for the Cloud supports network traffic flows visibility to all supported resource types within the boundary of a VPC, if VPC flow logging is enabled and onboarded in Cloud. However this support is subject to the limitations of the CSP. See the Amazon Virtual Private Cloud Flow Log Limitations link below for information.

Illumio Segmentation for the Cloud's policy feature supports AWS resource types that have Security Group and NACL attachments. The Amazon AWS Security Documentation link below details whether Security Group and NACL enforcement is available for each resource type in the resource type page, under the infrastructure security section.

For Azure, Illumio Segmentation for the Cloud provides visibility into network traffic flows for supported resource types that have NSG attachments with NSG flow logs enabled, or when the resource types are within a VNET boundary with VNET flow logs enabled. Some limitations regarding NSG and VNET flow logs apply, as detailed in the Microsoft Azure Network Watcher Virtual Flow Logs and NSG Flow Logs below.

Illumio Segmentation for the Cloud policy supports all resource types in Azure with NSG attachments. For details, see the Microsoft Azure Security Baseline link.

For GCP, Illumio Segmentation for the Cloud supports IP-based policy on firewalls. For details, see the Google VPC Firewall Rules link.

- Amazon Virtual Private Cloud Flow Log Limitations
- Amazon AWS Security Documentation
- Microsoft Azure Network Watcher Virtual Flow Logs
- Microsoft Azure Network Watcher NSG Flow Logs Overview
- Microsoft Azure Security Baseline
- Google VPC Firewall Rules

AWS

Resource Types by Category	Resource Type Relationships (Attached resource types on Details Page)	Flow Support	Display on Map	Display in Inventory
Compute				
EC2 Instance	ENI, Subnet, VPC, Security Group (SG), Elastic IP, Elastic Block Storage (EBS) Volume, Target Group, Load Balancer, Route Table, EKS Node Group, EKS Cluster	Yes	Yes, on Subnet level	Yes
EKS Cluster	Subnet, VPC, EKS Node Group, EKS Addon, EKS Fargate Profile, Security Group, ENI	Yes	Yes, on VPC level	Yes
Databases				
RDS DB Cluster	ENI, Subnet, VPC, SG, KMS Key	Yes	Yes, on VPC level	Yes
ElastiCache CacheCluster	ENI, Subnet, VPC, Security Group, KMS Key	Yes	Yes	Yes
MemoryDB Cluster	ENI, Subnet, VPC, Security Group, KMS Key	Yes	Yes	Yes
RDS DB Instance	ENI, Subnet, VPC, SG, KMS Key	Yes	Yes on RDS DB Cluster level	Yes
Network Routing				
ElasticLoadBalancingV2 Load Balancer	ENI, Subnet, VPC, SG, Target Group	Yes	Yes	Yes
NAT Gateway	Route Table, Elastic IP, ENI, VPC, Subnet	Yes	Yes	Yes
Serverless				
Lambda Function	Subnet, VPC, SG, Key Management Services (KMS) key, ENI	Yes, see the note at the bottom of this page	Yes, on VPC level	Yes
Storage				
S3 Bucket	Bucket Policy, VPC Endpoint, VPC, ENI, SG, Subnet, Flow Log	Yes, on VPC Endpoint	Yes	Yes

Azure

Resource Types by Category	Resource Type Relationships (Attached resource types on Details Page)	Flow Support	Display on Map	Display in Inventory
Compute				
Virtual Machine	NIC, NSG , Subnet, VNet, VM ScaleSet	Yes	Yes, on Subnet level	Yes
VirtualMachineScaleSet Virtual Machine	VM, VM ScaleSet VM	Yes	Yes, on VM ScaleSet level	Yes
AKS Cluster	Network Public IP, Agent Pool, Private Link Service, Agent Pool Machine, Virtual machine Scale Set, Virtual Machine Scale Set Virtual Machine, Network Interface, Loadbalancer, Subnet, Virtual network, Route table, NSG, NSG Flow Log, Storage Account	Yes	Yes	Yes
Databases				
DBforPostgreSQL Flexible Server	Private Endpoint, NIC, Subnet, VNet, NSG, DBforPostgreSQL Flexible Server Database	Both NSG and VNET flows for Private Endpoint Configuration. See the note at the bottom of this page.	Yes	Yes
DBforPostgreSQL Server	Private Endpoint, NIC, Subnet, VNet, NSG, DBforPostgreSQL Server Database	Both NSG and VNET flows for Private Endpoint Configuration. See the note at the bottom of this page.	Yes	Yes
DBforPostgreSQL ServerGroup V2	DBforPostgreSQL ServerGroup V2 Server	Both NSG and VNET flows for Private Endpoint Configuration. See the note at the bottom of this page.	Yes	Yes
DocumentDB Database Account	Private Endpoint, NIC, Subnet, VNet, NSG, SQL Database, DocumentDB Table, Document DB Gremlin Database, DocumentDB Cassandra	Both NSG and VNET flows for Private Endpoint Configuration.	Yes	Yes

Resource Types by Category	Resource Type Relationships (Attached resource types on Details Page)	Flow Support	Display on Map	Display in Inventory
	Keyspace, DocumentDB Mongo Database	See the note at the bottom of this page.		
DocumentDB Mongo Cluster	Private Endpoint, NIC, Subnet, VNet, NSG, DocumentDB Database Account	Both NSG and VNET flows for Private Endpoint Configuration.	Yes	Yes
		See the note at the bottom of this page.		
Redis Cache	VNet, Subnet, Private Endpoint	Both NSG and VNET flows for Private Endpoint Configuration.	Yes	Yes
		See the note at the bottom of this page.		
SQL Managed Instance	SQL Managed Instance Private Endpoint Connection, Subnet, Private Endpoint, VNET, NIC, NSG	Yes	Yes	Yes
SQL Server	Private Endpoint, NIC, Subnet, VNet, NSG, SQL Server Database	Both NSG and VNET flows for Private Endpoint Configuration.	Yes	Yes
		See the note at the bottom of this page.		
Network Security				
Azure Firewall	Azure Firewall Policy, Diagnostic Settings, Subnet, Network Public IP	Yes	Yes, on VNET level	Yes
Security Infrastructure				
Key Vault	Key Vault Private Endpoint Connection, Subnet, Private Endpoint, VNET, NIC, NSG	Both NSG and VNET flows for Private Endpoint Configuration.	Yes	Yes
		See the note at the bottom of this page.		

Resource Types by Category	Resource Type Relationships (Attached resource types on Details Page)	Flow Support	Display on Map	Display in Inventory
App Service (Web App, Function App)	Function App Function, App Service Private Endpoint Connection, Subnet, Private Endpoint, VNET, NIC, NSG	Both NSG and VNET flows for Private Endpoint Configuration. See the note at the bottom of this page.	Yes	Yes
Storage				
Storage Account	Private Endpoint, NIC, Subnet, VNet, NSG	Both NSG and VNET flows for Private Endpoint Configuration. See the note at the bottom of this page.	Yes	Yes

GCP

Resource Types by Category	Resource Type Relationships (Attached resource types on Details Page)	Flow Support	Display on Map	Display in Inventory
Compute				
Instance	Network, Subnetwork, Instance Template, Instance Group Manager, Instance Group, Autoscaler, Disk, Routes, Target Pool, Cluster, Node Pool	Yes	Yes	Yes
Containers				
Cluster	Network, Subnetwork, Node Pool, Instance group, Instance	Yes	Yes	Yes
Database				
Cloud Sql Instance	Cloud Sql Database, Cloud Sql SSL Cert, Cloud Sql Users	Yes	Yes, at the project level	Yes
Network Management				
Subnetwork	Network, Instance, Disk, Instance Template, Instance Group, Cluster, Node Pool, Address, Network Attachment, Network Endpoint Group, Router	Yes	Yes	Yes

OCI

Resource Types by Category	Resource Type Relationships (Attached resource types on Details Page)	Flow Support	Display on Map	Display in Inventory
Compute				
Compute Instance	VNIC, VNIC Attachment, Subnet, VCN, Container node pool, Container cluster	Yes	Yes on Subnet level	Yes
Containers				
Container Cluster	Container node pool, VCN, Subnet, Instance	Yes	Yes	Yes
Container Node Pool	Container cluster, VCN, Subnet, Instance	Yes	Yes	Yes
Database				
Autonomous Database	Subnet, NSG, VCN, Database Tools Connection, Private Endpoint, VNIC	Yes, for private IPs	Yes, on Subnet level	Yes

NOTE

Because Illumio Segmentation for the Cloud 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
- AWS Lambda Functions

Notes about resource type visibility

NOTE

If an ENI is associated with a single Lambda function, the flow logs will clearly identify the corresponding Lambda. However, due to the design of AWS Lambda architecture, complete flow visibility may not be achievable in Illumio Segmentation for the Cloud under the following conditions:

- Multiple ENIs that share the same subnet and VPC security group are associated with a single Lambda function
- Multiple Lambda functions that share the same subnet and VPC security group are associated with a single ENI
- Multiple ENIs associated with different Lambda functions that share the same subnet and VPC security group

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.

NOTE

Azure PaaS offerings in the Data, Database, Storage, Serverless, and Security Infrastructure categories log both NSG and VNET traffic only when a private endpoint is configured. Traffic to private endpoints can only be captured at the source VM. See Microsoft Q&A documentation.

The traffic is recorded with source IP address of the VM and destination IP address of the private endpoint. Illumio Segmentation for the Cloud cannot record traffic at the private endpoint itself due to platform limitations. See the Microsoft Azure documentation: NSG Flow Logs Overview - Azure Network Watcher

NOTE

Azure network IP configurations are no longer a visibility-supported resource type. However, they appear in the Microsoft network interface Inventory details tab. It may take time for existing IP configuration resources to stop displaying, so they may temporarily display as resources.

NOTE

For GCP, Illumio supports egress from a virtual machine's NIC to the load balancer's proxy infrastructure, which is captured in the VPC flow logs.

Global Map

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

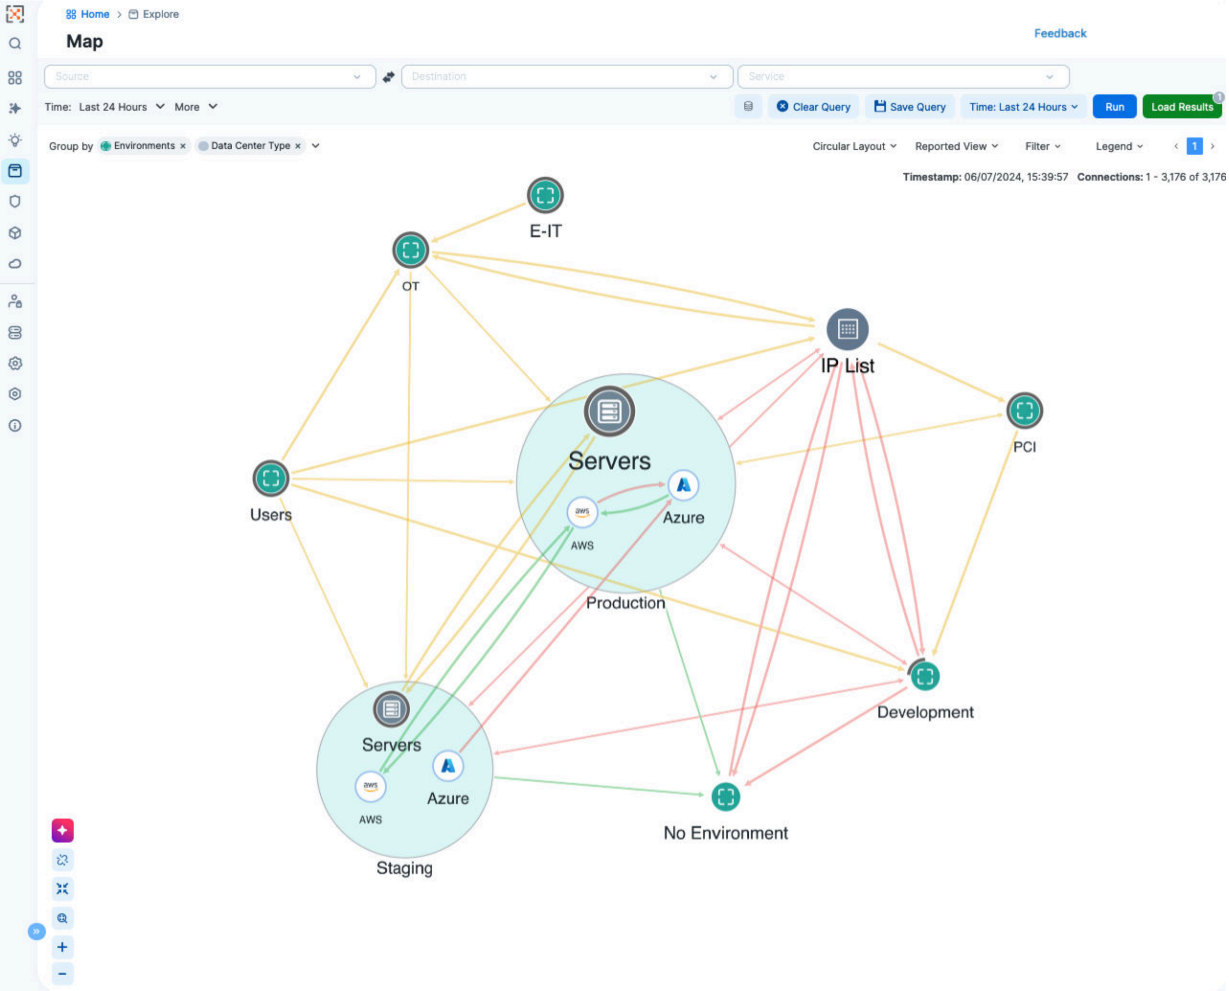

- You can hover your mouse over a cloud item, such as a region. Illumio will display information about it such as the number of resources and applications. Right-click items to see additional details.
- Left-click items to write policy for them. See Writing Organization Policy [254].

Grouping in the Global Map

Groups 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 Global 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 Global 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 displays a pop-up dialog box with the list of labels and the number of workloads using the labels.

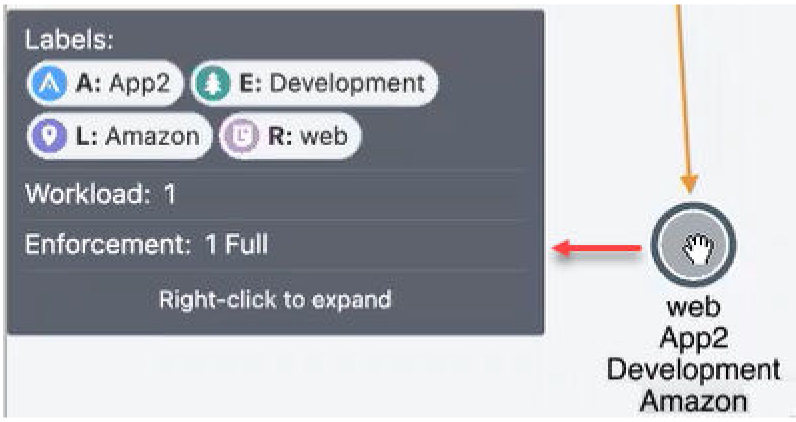

- In the **Group by** drop-down list, you can drag and drop labels in the list to re-order how group display. Labels at the top of the list control the prominence of those groups .
- The UI displays the groups using the colors you've selected for your labels. Use these colors to help orient yourself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

Global Map Layout Options

You can choose how the UI displays the Global Map:

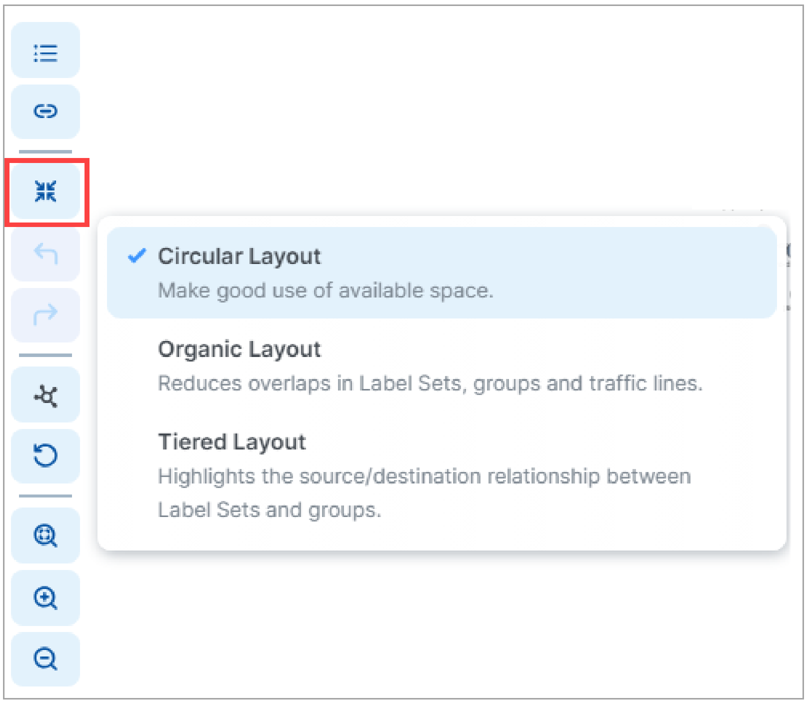
✓ Circular Layout
Make good use of available space.
Organic Layout
Reduces overlaps in Label Sets, groups and traffic lines.
Tiered Layout
Highlights the source/destination relationship between Label Sets and groups.

Not every layout choice is good for your 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.

Panels in the Global 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 Global 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 Global 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.

Reading the Global 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.

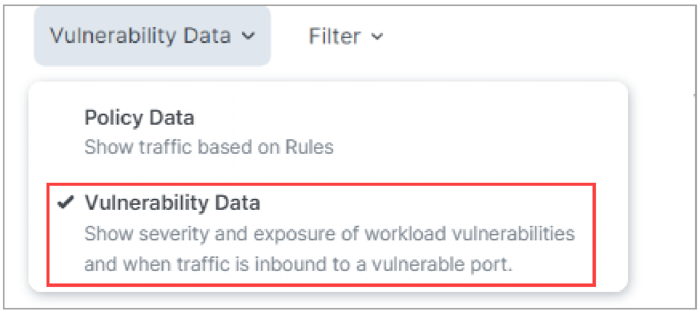

Policy Data
Show traffic based on Rules

✓ Vulnerability Data
Show severity and exposure of workload vulnerabilities and when traffic is inbound to a vulnerable port.

Legend - Policy Data

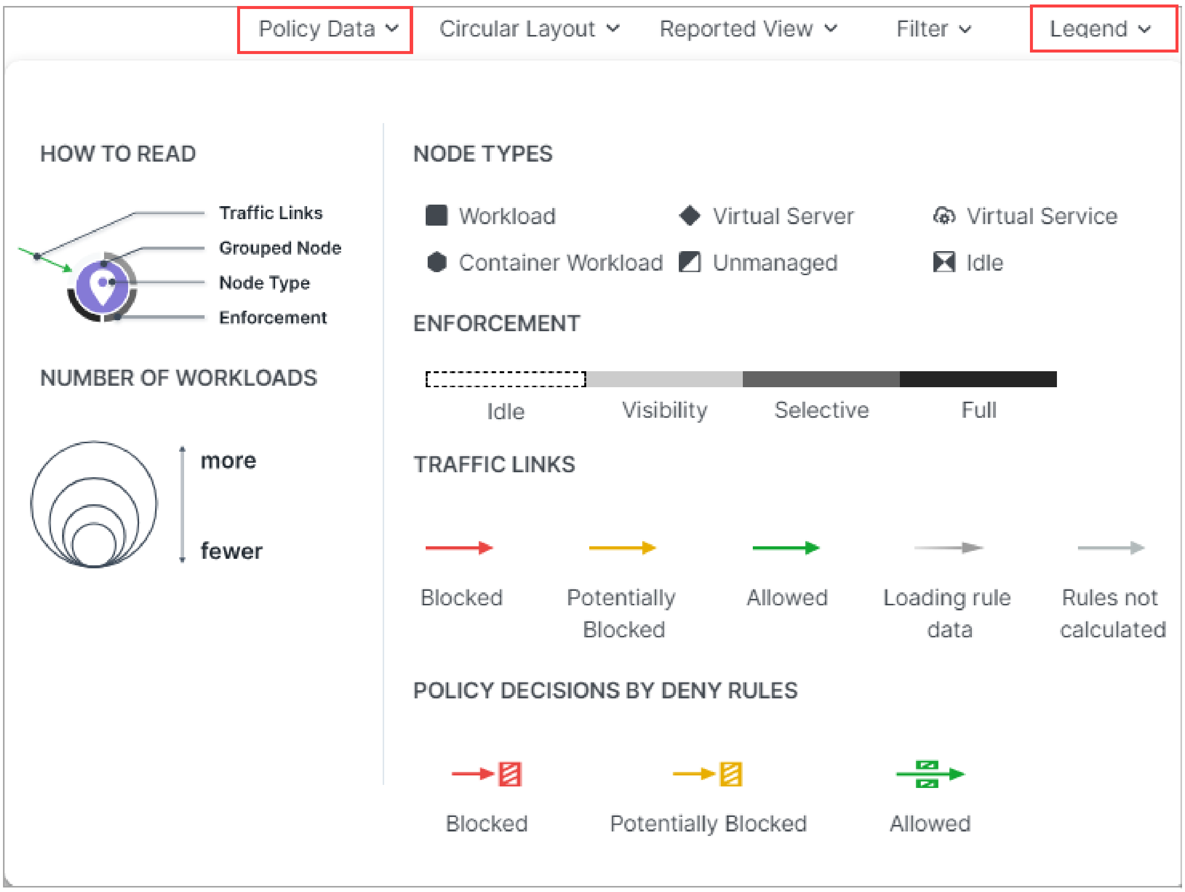

Symbols Explained

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

Reported View

The Illumio UI displays the traffic on the Global 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.

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 , 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 Global Map

Connections Menu

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

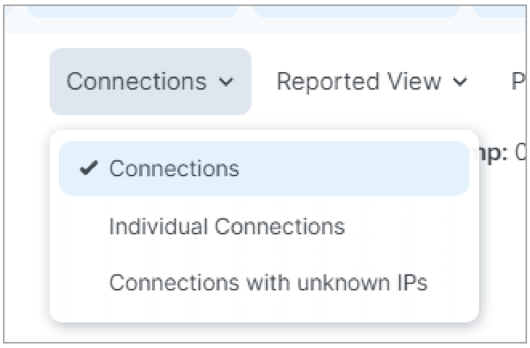

Filter drop-down

Options in the Filter drop-down allow you to control which traffic information is displayed on the Global 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.

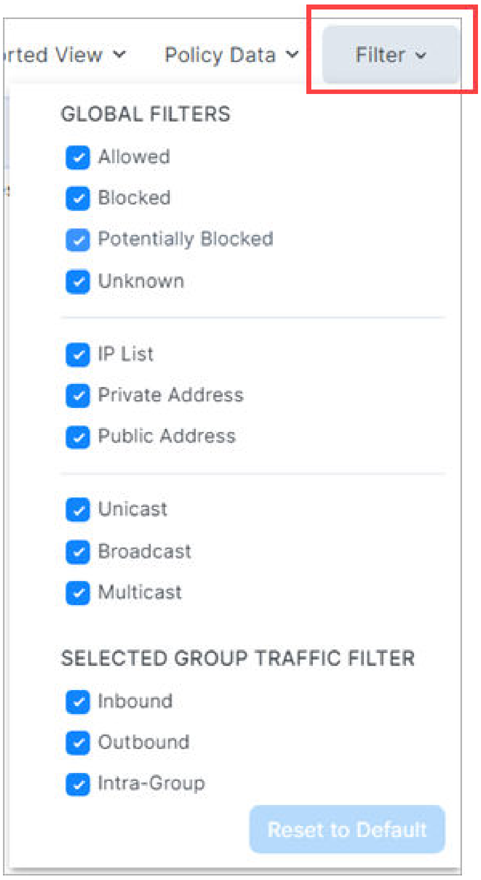

GLOBAL FILTERS

☑ Allowed
☑ Blocked
☑ Potentially Blocked
☑ Unknown

☑ IP List
☑ Private Address
☑ Public Address

☑ Unicast
☑ Broadcast
☑ Multicast

SELECTED GROUP TRAFFIC FILTER

☑ Inbound
☑ Outbound
☑ Intra-Group

The Filter dropdown presents two types of filters:

Global Filters

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

Selected Group Filters

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

Usage

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

Displaying Usage

In the 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)
- Illumio Workload Hours by date, with the following types of workload hours:
  - Total Illumio Workload Hours
  - Compute Workload Hours
  - Database Workload Hours
  - Container Hosts
  - Serverless Containers
  - Serverless Functions
- Data Processed by date and Daily Totals of Data Processed (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.

NOTE

Not all resource types ingested within a given category are billable. Only resource types with policy support are considered billable and therefore reflected in the usage data. As a result, the total resource count displayed in the Ingested Resources summary dashboard tile for a category may differ from the total count shown on the usage page for the same category. See Policy enforcement and resource types [262].

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.

Exporting a Usage Report

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Exporting+Usage+Reports.mp4

1. Click **Export** to export the usage data.
2. Edit the report name and select the format.
3. Click the **Scheduling Section** toggle to the on position to schedule the export unless you want to export the report immediately.
4. If you choose to schedule your report, select your recurrence and time.
5. Click **Save** when done.
6. Go to the Reports [212] page to download the exported report.

Generated reports

The Generated Reports page lets you view all the reports generated in Illumio Segmentation for the Cloud, including ad-hoc reports, scheduled reports and report exports from across the product.

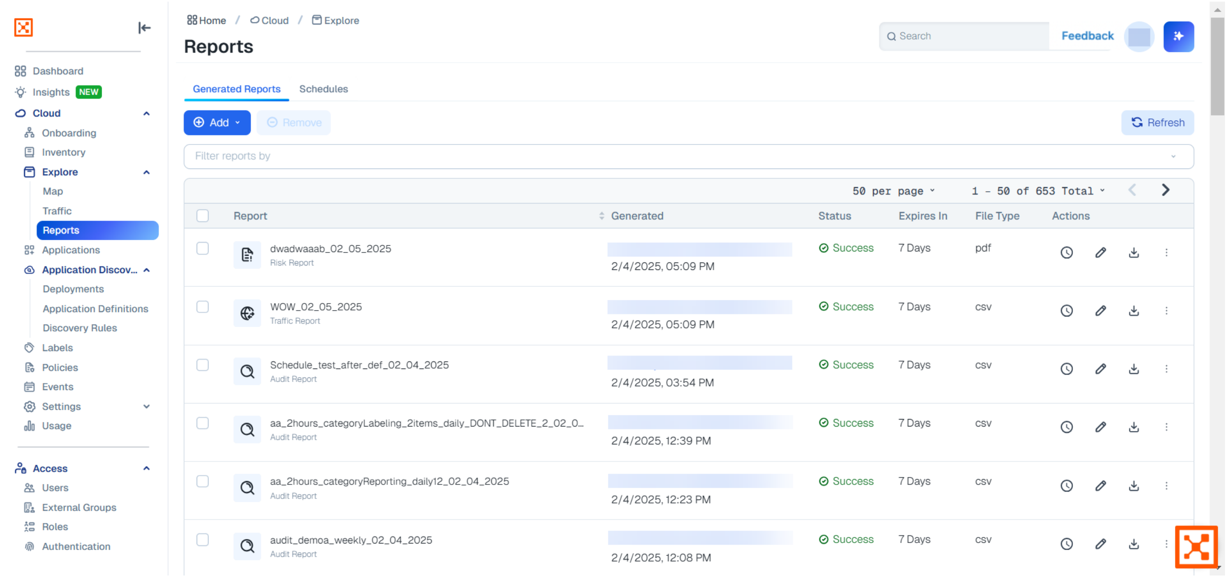

Reports

□	Report	Generated	Status	Expires In	File Type	Actions
□	dwadwaaab_02_05_2025 Risk Report	[redacted] 2/4/2025, 05:09 PM	Success	7 Days	pdf	
□	WOW_02_05_2025 Traffic Report	[redacted] 2/4/2025, 05:09 PM	Success	7 Days	csv	
□	Schedule_test_after_def_02_04_2025 Audit Report	[redacted] 2/4/2025, 03:54 PM	Success	7 Days	csv	
□	aa_2hours_categoryLabeling_2items_daily_DONT_DELETE_2_02_0... Audit Report	[redacted] 2/4/2025, 12:39 PM	Success	7 Days	csv	
□	aa_2hours_categoryReporting_daily12_02_04_2025 Audit Report	[redacted] 2/4/2025, 12:23 PM	Success	7 Days	csv	
□	audit_demoa_weekly_02_04_2025 Audit Report	[redacted] 2/4/2025, 12:08 PM	Success	7 Days	csv	

Use reports

Use these guidelines when you want to download or remove 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.

Add reports

Illumio Segmentation for the Cloud supports generation of the following reports from the Generated Reports page:

- Audit: This is generated from the Events page Audit Events tab. See Events [192].
- System: This is generated from the Events page System Events tab. See Events [192].
- Traffic: This is generated from the Traffic page. See Traffic [219].
- Risk: This is a traffic risk report generated from Illumio Segmentation for the Cloud your traffic. See Risk Reports [216].
- Inventory: This is generated from the Inventory page. See Inventory [157].
- Usage: This is generated from the Usage page. See Usage [211].
- Application: This is generated from the Application page. See View and approve an application [238].
- Application Definition: This is generated from the Application Definitions page. See Define an application individually [235] and Define an application automatically [231].

Reports and schedules

You can run reports either once or at regular intervals.

Generate a Report

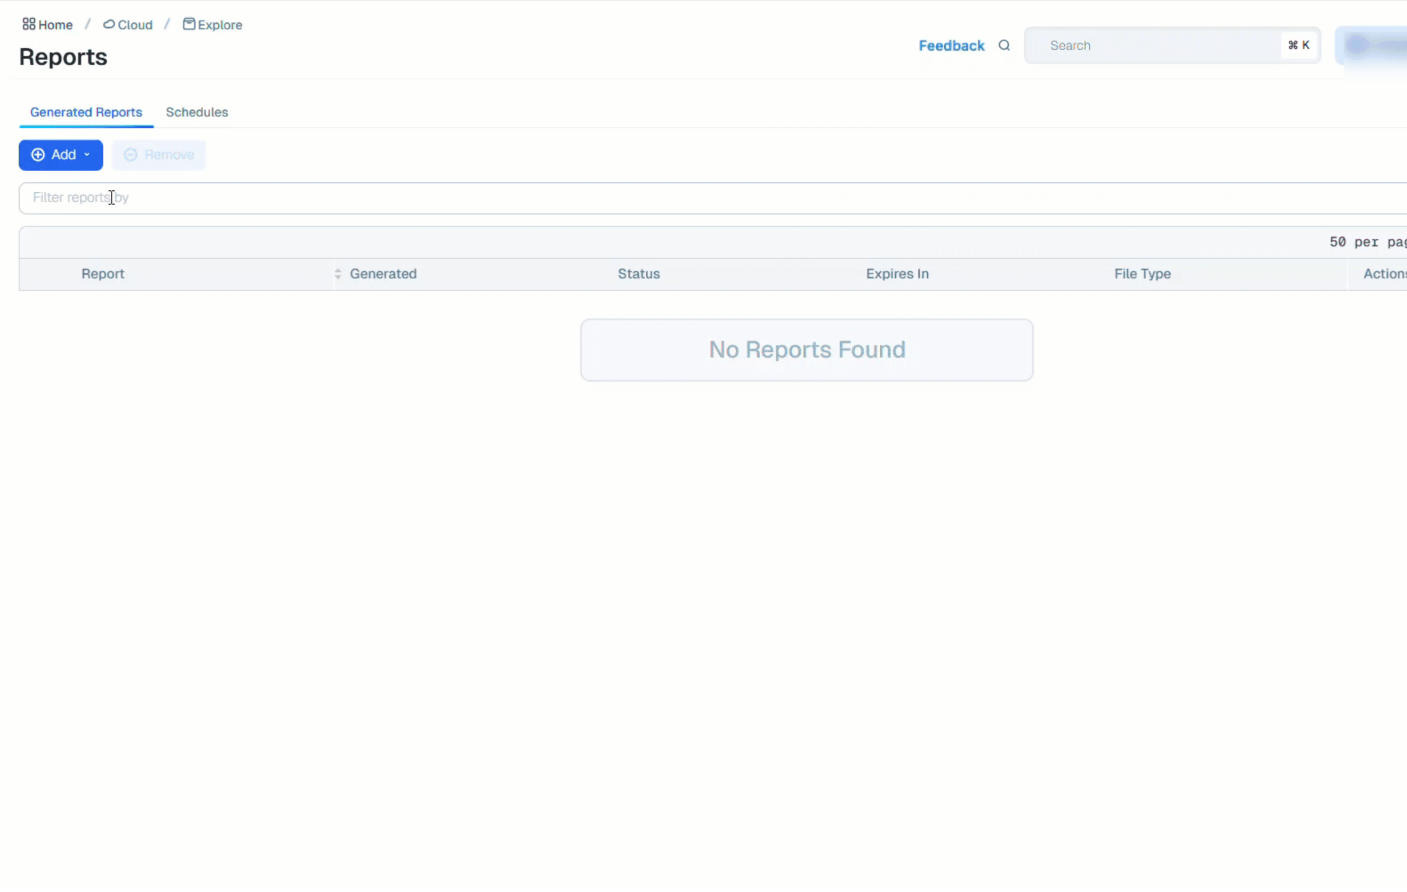

Report	Generated	Status	Expires In	File Type	Action

No Reports Found

1. Click **Add** to select the type of report from the drop-down menu.
2. Enter a name (without spaces), select the format, and specify the parameters.
   a. **Audit**
      - Configuration
        - Time range: A variety of ranges, including custom
        - Filters: Cloud, User, Account ID, Status, Category, Event Type
      - Export Format: .JSON or .CSV

   b. **Traffic**
      - Configuration
        - Time range: A variety of ranges, including custom
        - Filters: Source/Destination, Category (Account ID, Account Name, Resource Name, etc.), and Operator (= and !=). See Search traffic [220].
        - Conditions: Match Any Condition (OR), and Match Any Condition (AND). See Search traffic [220].
      - Export Format: .JSON or .CSV

   c. **Risk**
      - Configuration
        - Time range: A variety of ranges, including custom
        - Details to Include: Top Sources & Destinations, Top Conversations
        - Sort: Byte Count, Flow Count
      - Export Format: .PDF

   d. **System**
      - Configuration
        - Time range: A variety of ranges, including custom

    - Filters: Cloud, Account ID, Status, Category, Event Type
  - Export Format: .JSON or .CSV

  e. **Inventory**
  - Export Format: .JSON or .CSV
  - Filters: Cloud, Account, IP Address, Region, Resource Type, Resource Group, Resource Name, Resource State, VPC/VNET ID, Subnet ID, Cloud Tags, Labels, and Categories

  f. **Usage**
  - Export Format: .CSV
  - Time range: A variety of ranges, including custom

  g. **Application**
  - Export Format: .JSON or .CSV
  - Filters: Name, Deployments, and Cloud Accounts

  h. **Application Definition**
  - Export Format: .JSON or .CSV
  - Filters: Application Label, Deployments, Associated Labels, Approval Status, Source

  i. **All report types**
  - Retention duration: (seven days at time of writing)
  - Scheduling: See Create a report schedule [215] on this page.

3. Select your time range.
4. Click **Save**.

Create a report schedule

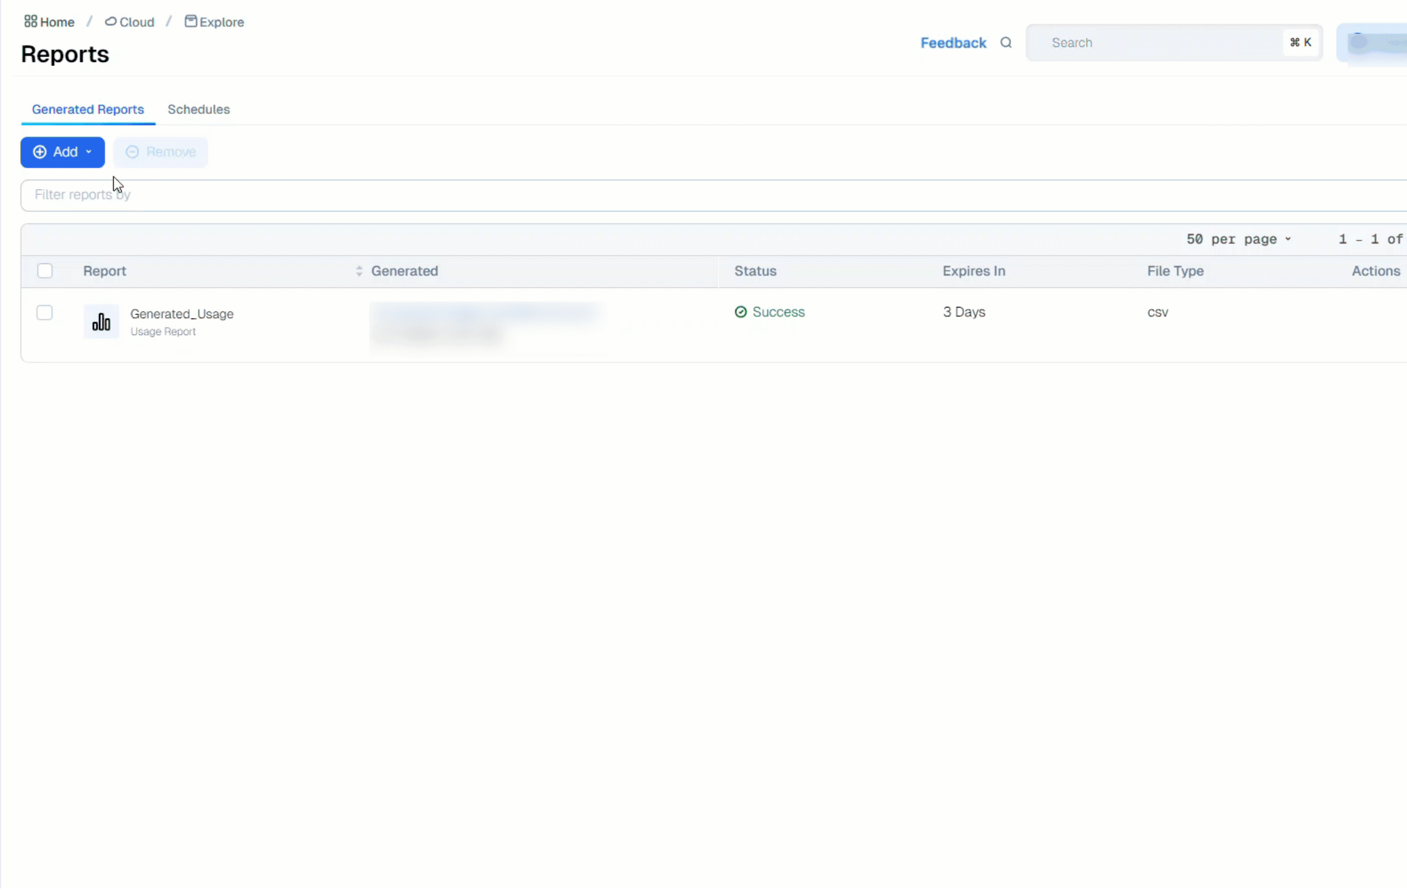

☐	Report	Generated	Status	Expires In	File Type	Actions
☐	Generated_Usage Usage Report	[redacted]	Success	3 Days	csv	

1. Click **Add** to select the type of report from the drop-down menu.
2. Enter a name (without spaces) , select the format, and specify the parameters as described in Generate a report [214].

3. Click the **Scheduling Section** toggle to the on position if you want to schedule a report.
4. Select your recurrence and time range.

> **NOTE**
>
> If you select a custom time range for a scheduled report, you will generate an identical report each time, effectively duplicating your reports. This is because the custom time range is always fixed. For this reason, use only the non-custom time ranges ("last X days") when creating a report schedule. Use the custom time range for generating non-scheduled reports only.

5. Click **Save**.

Generate ad hoc reports from your schedules

1. Go to the Schedules tab of the Reports page and find your Schedule.
2. Click the **Generate Report Now** icon in the Actions column.
3. Go to the Generated Reports page, find your newly generated ad-hoc scheduled report, and click the **Download** icon. Click **Download**.

Edit schedules

1. Click the pencil icon for the schedule you want to edit.
2. Make your changes and click **Save**.

> **NOTE**
>
> You cannot edit the schedule for an Insights report.

Enable and disable schedules

Within the **Schedules** tab, select the row for the schedule and click **Enable** or **Disable**. The **Status** column reflects whether a schedule is enabled or disabled.

Risk Reports

This is an overview of the Risk Report feature. For instructions on generating a Risk Report, see Generated reports [212]. For a list of services that Illumio considers to be at risk, see Risky Services [217].

The Generated Reports page 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, Illumio Segmentation for the Cloud 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 Illumio Segmentation for the Cloud 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 with the following details:
  - Heading row, containing the following columns:
    - Severity
    - Service
    - Port
    - Protocol

Risky Services

Learn about services that Illumio considers to be at risk. For information on Risk Reports, see Risk Reports [216].

Ransomware Risky 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-POPV3	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

This topic describes the purpose of the Illumio Segmentation for the Cloud traffic feature, found in the 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 [220].

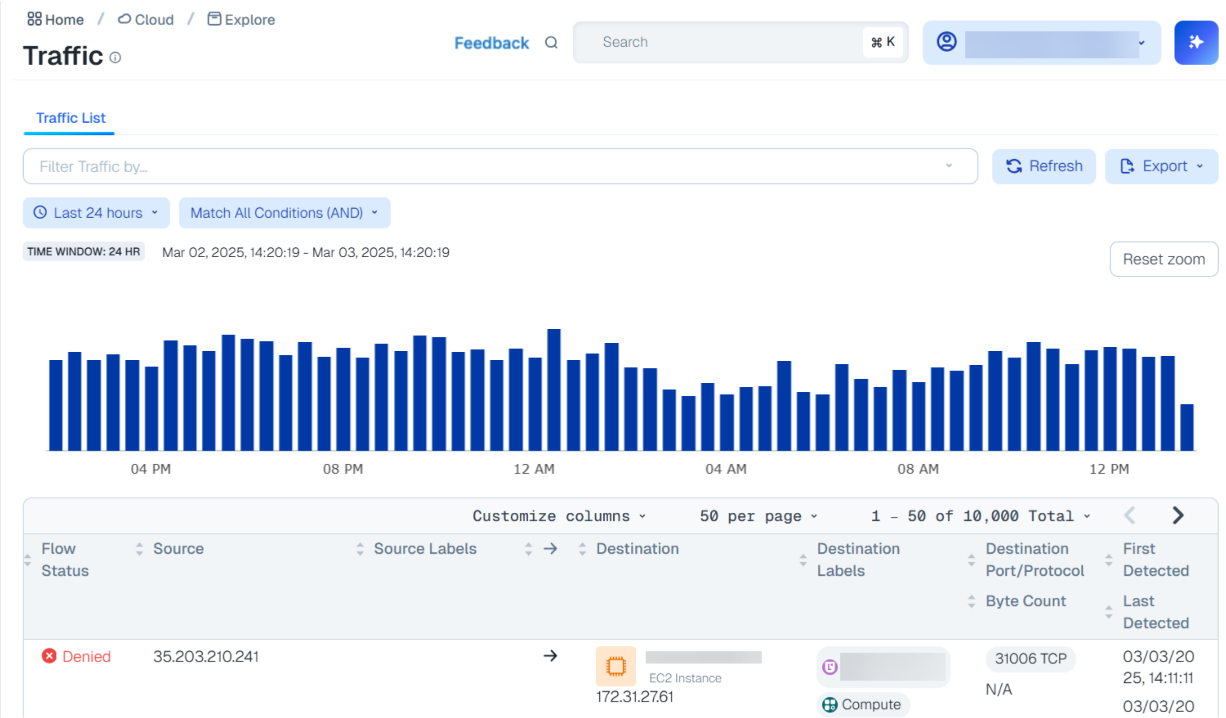

Last 24 hours · Match All Conditions (AND)

TIME WINDOW: 24 HR Mar 02, 2025, 14:20:19 - Mar 03, 2025, 14:20:19

1 – 50 of 10,000 Total

Flow Status	Source	Source Labels	→	Destination	Destination Labels	Destination Port/Protocol Byte Count	First Detected Last Detected
Denied	35.203.210.241		→	[redacted] EC2 Instance 172.31.27.61	[redacted] Compute	31006 TCP N/A	03/03/2025, 14:11:11 03/03/20

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.

Exporting Traffic Lists

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Export+Traffic+Lists.mp4

1. Click **Export** to export the traffic data.
2. Edit the report name, select a time range, and select the format.
3. Click the **Scheduling Section** toggle to the on position to schedule the export unless you want to export the report immediately.
4. If you choose to schedule your report, select your recurrence and time.
5. Click **Save** when done.
6. Go to the Generated reports [212] page to download the exported report.

Click the down arrow and select **View All Reports** to go to the Generated reports [212] page and 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. Illumio 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. Note that traffic records older than 90 days will be automatically removed.

Illumio Segmentation for the Cloud doesn't display GCP traffic flowing through load balancers because VPC flow logs don't capture load-balancer-specific details. For supported resources, see Illumio visibility for resource types [193].

Search traffic

This topic describes the steps for searching the Illumio Segmentation for the Cloud traffic, found in Cloud > Explore, and provides a general example of how you would use it. For an overview of the Traffic page, including Risk Report generation, see Traffic [219]. For information on how to use the search function, see the in-application pop-up on the Traffic page.

Navigate to the Traffic page to get a view into your traffic patterns over a specified time period.

Searching Traffic Guidelines

Use Time Slider and the Filter to view traffic patterns during specific time periods.

About the Traffic Time Slider

The Traffic Time Slider provides a visual representation of traffic spikes over the last 24 hours, 7 days, 14 days, or a custom-defined time range. It displays the total number of flows within the selected period, allowing for in-depth traffic analysis. You can zoom into a specific timeframe to examine flow activity in greater detail. For instance, selecting a single bar reveals traffic data for that moment, while selecting a range of bars enables deeper analysis over the chosen period.It also provides detailed flow information, including flow status, source and destination, and the timestamps when flows were detected.

Using the Traffic Time Slider

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Traffic+Time+Slider.mp4

1. Apply any filters. See About the Filter [221].
2. Click and drag to select one or more bars to zoom.

   The Traffic Time Slider view and the traffic list table update simultaneously.
3. Hover over a bar to see the time frame and the number of flows. Right click to zoom or download a .png of all the visible bars.

4. Click **Reset Zoom** to reset the view.

About the Filter

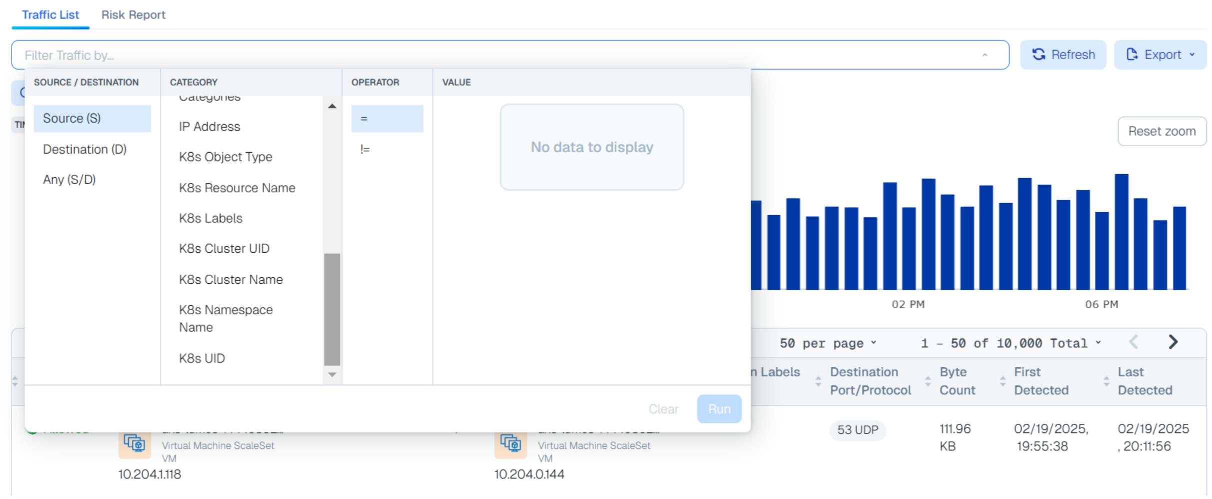

Use these guidelines to define filters and search for traffic patterns. Click in the filter field to start narrowing your search parameters.

- Use operators such as '!=' and '='
- Select the **Match All Conditions (AND)** or **Match Any Conditions (OR)** dropdown to switch the automatically inserted joiners from OR to AND. You can add additional search terms without deleting existing terms.
- Filter by:
  When filtering by IP addresses, use CIDR blocks to include a range of IP addresses. For example, adding "/16" to an IP address will search for flows with IP addresses starting with the same first 16 bits as the specified IP address, such as 10.104.XXX.XXX. Similarly, adding "/24" or "/30" will search for flows with IP addresses starting with the same first 24 or 30 bits as the specified IP address, respectively. Note that the number after the slash specifies the prefix length.

  - Source/Destination (can change depending on Category selection)
  - Category
    Categories include the following:
    - Cloud
    - Account
    - Label
    - Flow Status
    - IP Address
    - Port
    - Subnet
    - VPC
    - Known Network
    - Resource Type

    Kubernetes categories include the following:
    - K8s Object Type
    - K8s Resource Name
    - K8s Labels

    - K8s Cluster UID
    - K8s Cluster Name
    - K8s Namespace Name
    - K8s UID
  - Operator (can change depending on Category selection, such as the Label category, which has '=' but not'!=')
  - 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.
- When filtering by Known Networks, the Inventory page Known Networks tab provides the values that populate your filter. See Inventory [157].
- Click **Run** to apply your filter and display the results
- Hover over your result entries to see a brief list of details. Click on your result entries to see more details. For example, click on a resource to go to its Inventory details page, or click on labels, ports, and the like to see a details panel about the flow.
- Click **Refresh** to refresh the traffic data.

Define

The content in this category explains how you define your cloud environment in Illumio Segmentation for the Cloud; in particular, this category explains how to work with and set up deployments and applications, and cloud tag to label mapping. The deployments and application section contains topics that explain how to define deployments and applications, as well as how to view and approve an application.

Deployments and Applications

This topic explains defining deployments and defining applications in Illumio Segmentation for the Cloud. 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 Illumio Segmentation for the Cloud?

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 Cloud, 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.

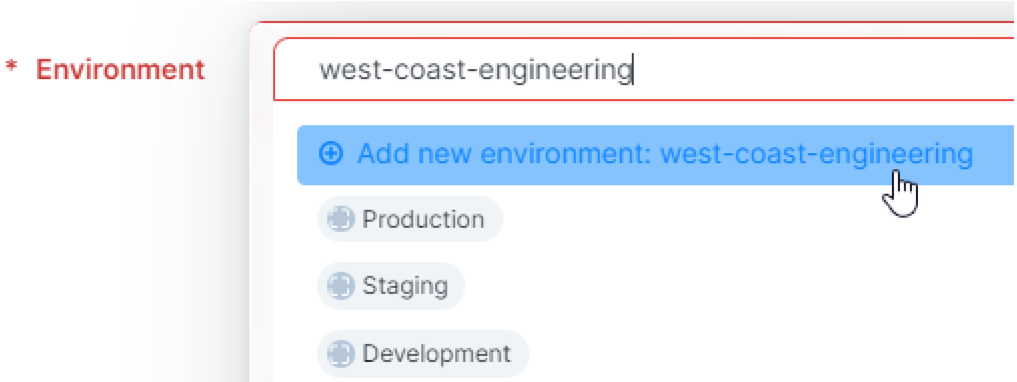

* Environment: west-coast-engineering

For example, you might realize that you want a deployment stack in Illumio Segmentation for the Cloud 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.

With Illumio Segmentation for the Cloud, you may decide to create deployment stacks as part of specifying which applications in your cloud account to protect.

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

Relationship between Deployments and Applications

To more fully use Illumio Segmentation for the Cloud, you need to understand the relationship between applications and their deployments.

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.

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 [241] for information.

Illumio Segmentation for the Cloud 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.

This 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.

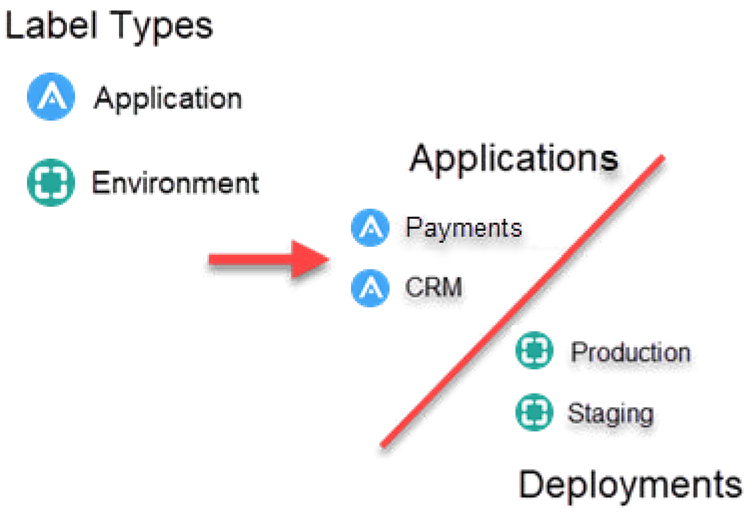

To break these concepts down further, see how Illumio Segmentation for the Cloud utilizes Application and Environment labels.

Application Labels

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

For example, you have two applications — a Payment application and a CRM application. In Cloud, 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.

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.

Illumio Segmentation for the Cloud Discovers Your Application Environments

When you define a deployment, it doesn't discover anything about your applications. You defined your deployment stacks separately. Then, after you defined each application, Illumio Segmentation for the Cloud analyzes them by reviewing the associated cloud metadata, such as account, VPC, subnet, tags, etc. Illumio Segmentation for the Cloud 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 definition (CRM in this example) and save it. Illumio Segmentation for the Cloud begins the process of discovering the environments in which it's running and the resources that it has. The *Application Definitions* page refreshes to include the new application definition, and the Details page for that application definition displays the message "DISCOVERY STATUS: QUEUED."

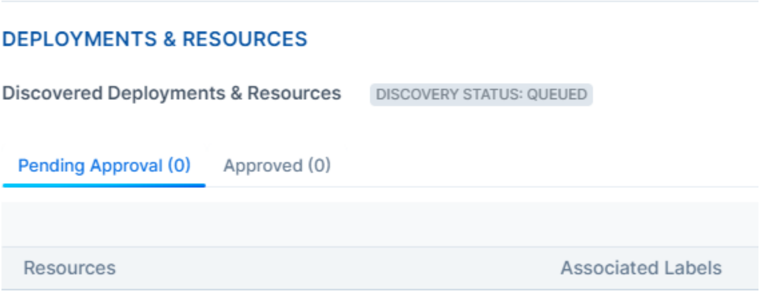

After discovering all the associated environments and resources, the Details page for that application definition displays the message "DISCOVERY STATUS: COMPLETE," and the Application Definitions page lists all the discovered deployments and resources matching your selected cloud metadata. If it discovers, deployments or resources, the approval status will show as pending.

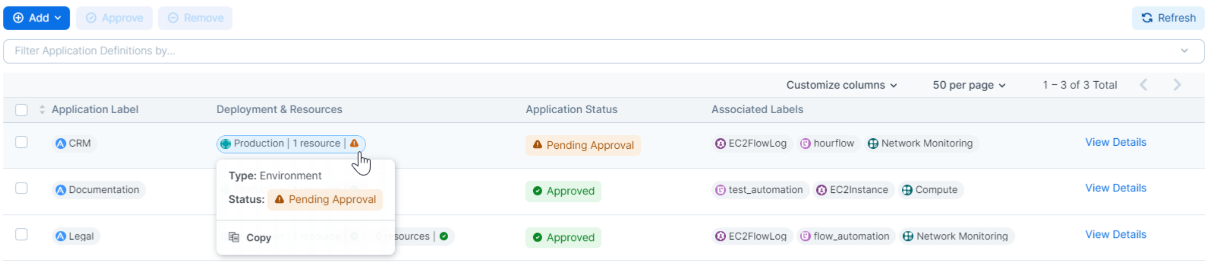

□ Application Label	Deployment & Resources	Application Status	Associated Labels	
□ CRM	Production | 1 resource | ⚠	⚠ Pending Approval	EC2FlowLog hourflow Network Monitoring	View Details
□ Documentation		✓ Approved	test_automation EC2Instance Compute	View Details
□ Legal	sources | ✓	✓ Approved	EC2FlowLog flow_automation Network Monitoring	View Details

Type: Environment
Status: ⚠ Pending Approval
Copy

Illumio will not populate the Deployments & Resources column if none are discovered or if existing application definitions already feature the ones you selected for your new application definition.

Why is Environment Discovery Important?

Illumio 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.

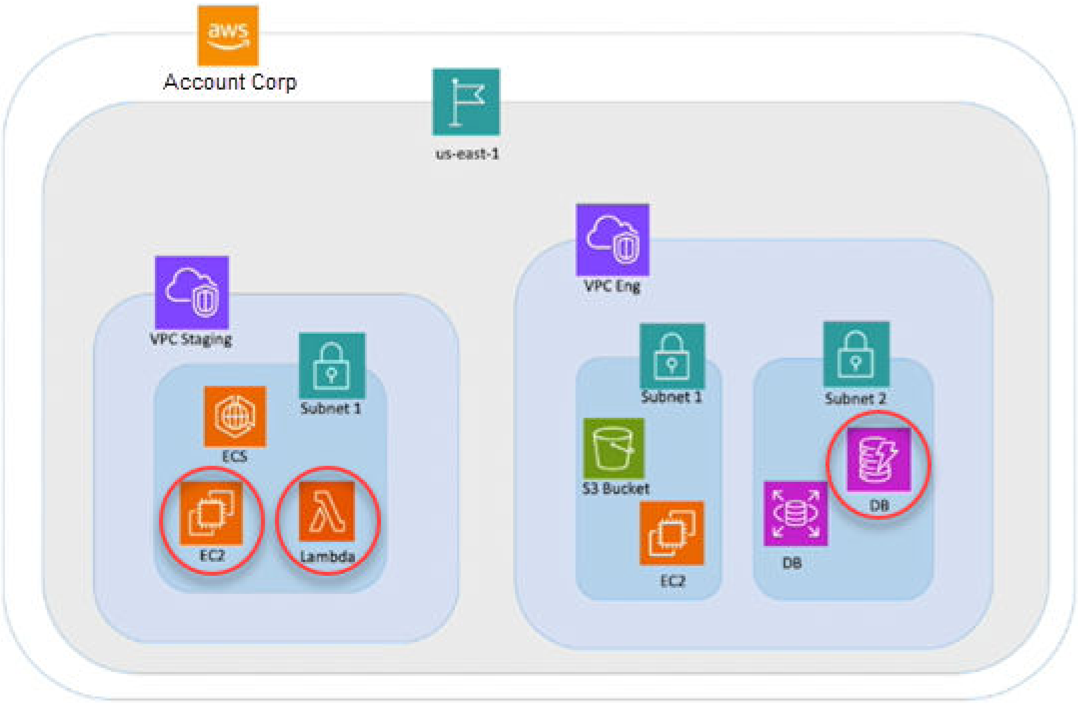

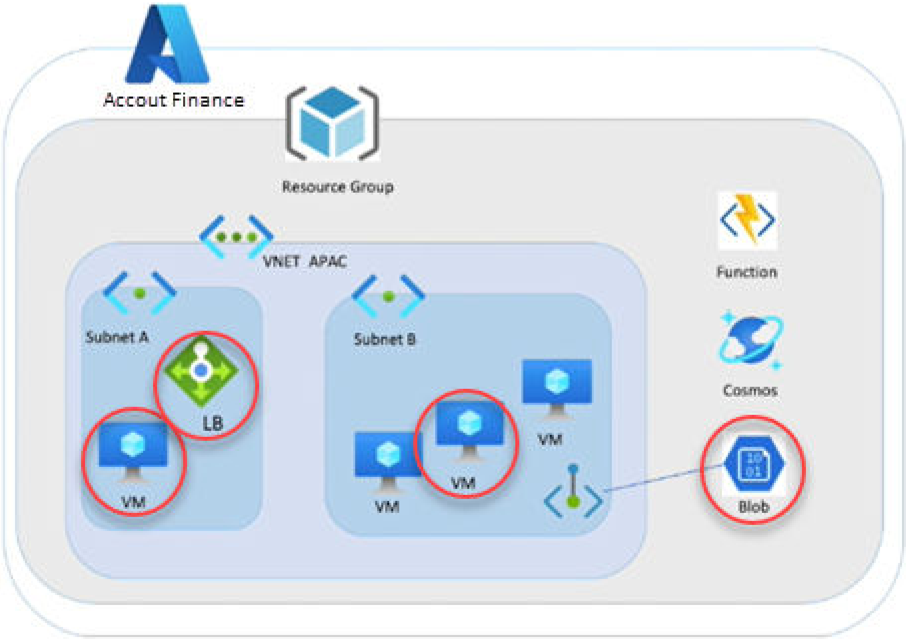

In Illumio Segmentation for the Cloud, the company defines the resources that are part of this application.

They begin by defining all three deployments— development, staging, and production. See Define a Deployment [228]. Then, they are ready to define the application in Illumio Segmentation for the Cloud. When they define the application, they specify the scope for what comprises the application. See Define an Application [231].

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. Illumio Segmentation for the Cloud can now determine that the application has resources in the AWS development and staging environments and resources in the Azure production environment.

In Illumio Segmentation for the Cloud, 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 Illumio Segmentation for the Cloud.

For an explanation of how Illumio Segmentation for the Cloud uses deployments and why they are helpful despite not being required, see Deployments and Applications [223]

Prerequisites

Before you define a deployment, you must have onboarded one or more public cloud accounts and give Illumio Segmentation for the Cloud time to synchronize with them so that you can configure the correct boundaries for the deployment.

- See Onboarding AWS Cloud [94].
- See Onboarding Azure [59].
- See Onboarding OCI [127].

Before defining a deployment, Illumio recommends that you review your Inventory and Map topology to gain an understanding of how your cloud resources are utilized and how they are communicating. See Map [168] and Inventory [157] for information about using these features.

Define a Deployment (Optional)

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Define+a+Deployment.mp4

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, Illumio Segmentation for the Cloud includes Environment labels for “Production,” “Staging,” and “Development.” If you select a label that already has a deployment defined for it, Illumio Segmentation for the Cloud 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.

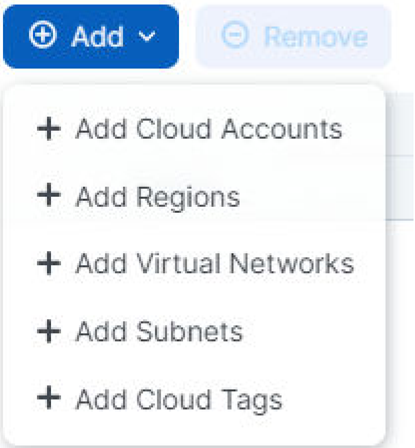

When you select an item from the resource drop-down menu, the **Add Deployment Stacks** dialog box opens.

6. In the **Includes** field, select the resource to use from the pre-populated drop-down list. The list includes resource that Cloud 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

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Edit+a+Deployment.mp4

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, Cloud includes Environment labels for “Production,” “Staging,” and “Development.” If you select a label that already has a deployment defined for it, Cloud 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.

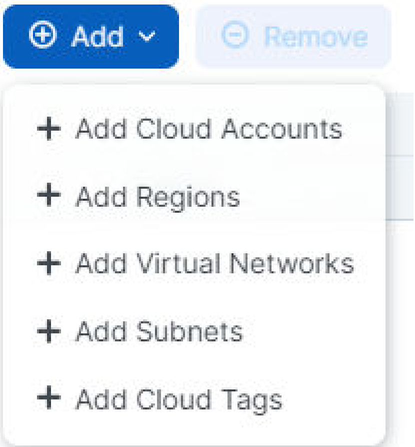

When you select an item from the resource drop-down menu, the **Add Deployment Stacks** dialog box opens.

6. To change the resource to use, select the **Includes** field to open the pre-populated drop-down list. The list includes resource that Cloud 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. Cloud 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**.

> **NOTE**
>
> Illumio supports deployment creation including using GCP labels.
>
> Illumio recognizes GCP labels under Illumio cloud tags. This means that when you use the tag to label mapping feature for GCP, cloud tags appear in the dropdown menu with the relevant prefix that indicates it is a GCP tag or label. For example, cloud tags for GCP may have values like `label/gcp-key:gcp-value`.
>
> Illumio supports GCP resource manager tags and labels at this time. Because GCP label values are optional, you may occasionally see empty tag values.

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 [231].

Define an application automatically

This topic explains how to define an application in Illumio Segmentation for the Cloud using Application Discovery Rules. To define an application individually, see Define an application individually [235].

For an explanation of application definitions and how they relate to deployments, see Deployments and Applications. [223]

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 [228].

Define Applications Automatically

Although Illumio Segmentation for the Cloud 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.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Define+Automatically.mp4

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 [238], *unless* you click the Auto Approve Setting toggle to **ON**. Do this if you want Illumio Segmentation for the Cloud 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 [238] 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`

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 Discovery Rule-Created Application Definitions

When you delete applications that are pending approval, Illumio Segmentation for the Cloud simply deletes the application definitions.

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

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:

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Delete+Application+Discovery.mp4

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.

Exporting an Application Definition Report

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/Videos/Export+App+Definitions+Report.mp4

1. Click **Export** on the Application Definitions tab.
2. Edit the report name and select the format.
3. Click the **Scheduling Section** toggle to the on position to schedule the export unless you want to export the report immediately.
4. If you choose to schedule your report, select your recurrence and time.
5. Click **Save** when done.
6. Go to the Reports [212] page to download the exported report.

What's Next

Approve your application. (Each instance of the application in different deployments requires approval.) See View and Approve an Application [238] for information.

Begin creating policy for your application. See Writing Application Policy [255] for information.

Define an application individually

This topic explains how to manually define an individual application in Illumio Segmentation for the Cloud.

TIP

To automatically define an application with application discovery rules, edit application definitions, or export application definition reports, see Define an Application Automatically [231].

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 [228].

Define Applications Individually (Manually)

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Define+Manually.mp4

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 Illumio Segmentation for the Cloud. The name should be descriptive so that you can easily identify it.

   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. Cloud 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 Cloud labels. Your applications are defined for Cloud 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 your chose) that defines the application.

> **TIP**
>
> The list is pre-populated with values that Cloud 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.

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

8. 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.

> **IMPORTANT**
>
> When adding multiple tags, be aware that matching of tags between resources and application definitions is done by logical OR. Consequently, a resource matches an application definition when it shares one or more tags.

9. When done fully defining all resources for the application, click **Confirm Selection**. The dialog box closes, and your selected tags or metadata appears in the Selected section.

   If you restart the steps to add a definition, the existing definition is cleared.

10. Click the Auto Approve Setting toggle to **ON** if you want Illumio Segmentation for the Cloud 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 [238] for information.

11. 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 Illumio Segmentation for the Cloud is discovering any defined deployments that host this application.

When the discovery process finishes, the list includes any deployments where Illumio Segmentation for the Cloud discovered matching cloud tags or metadata.

Illumio Segmentation for the Cloud does not populate the Deployments column if you choose not to define any for that application.

When it 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 Illumio Segmentation for the Cloud 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.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Edit+an+Application+Definition.mp4

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 Individually Created Application Definitions

When you delete applications that are pending approval, Illumio Segmentation for the Cloud simply deletes the application definitions.

When you delete approved applications, Illumio Segmentation for the Cloud deletes the application definitions and the rulesets (policies) associated with the application definitions and the application instances. Illumio Segmentation for the Cloud 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.

> **NOTE**
>
> Illumio recognizes GCP labels under Illumio cloud tags. This means that when you use the tag to label mapping feature for GCP, cloud tags appear in the dropdown menu with the relevant prefix that indicates it is a GCP tag or label. For example, cloud tags for GCP may have values like `label/gcp-key:gcp-value`.
>
> Illumio supports GCP resource manager tags and labels at this time. Because GCP label values are optional, you may occasionally see empty tag values.

What's Next

To understand application definitions and how they relate to deployments, see Deployments and Applications. [223]

To approve your application, where each instance of the application in different deployments requires approval, see View and Approve an Application [238] .

To begin creating policy for your application, see Writing Application Policy [255].

View and approve an application

This topic explains how to approve an application definition after you've created it. See Define an Application [231] for information.

Prerequisites

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

Why is Approval Required?

After you define an application, it appears in the Application Definitions list. First, if you have defined a deployment, Illumio discovers any environments where the application is running. See Illumio Discovers Your Application Environments [225] for information.

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

NOTE

The Application Definition page lets you toggle whether you want Illumio Segmentation for the Cloud to automatically approve all discovered applicable deployments and resources. Similarly, the Application Discovery Rule page lets you toggle whether you want Illumio Segmentation for the Cloud to automatically approve all discovered application definitions, as well as any updates made to their deployments and resources. See the Define an Application [231] documentation on the portal.

NOTE

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, Illumio ensures other key stakeholders are in the loop to approve your application definitions.

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

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

Approve a Given Application Definition

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Approve+a+Given+Application+Definition.mp4

1. From the left navigation, choose **Application Discovery** > **Application Definitions**.
2. The list of defined applications appears.
3. 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.
4. The **Approve** button becomes enabled.
5. Click **Approve**. A confirmation dialog box appears displaying the application you are approving.
6. 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. Illumio Segmentation for the Cloud lets you approve two or more such application deployments in bulk.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Approve+Application+Deployments.mp4

1. From the left navigation, choose **Application Discovery** > **Application Definitions**.
2. The list of defined applications appears.
3. Select the application that you want to review and/or approve.
4. Click **Approve**. A confirmation dialog box displays the application's associated deployments and/or resources..
5. 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.
6. 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.

Illumio Segmentation for the Cloud does not let you bulk approve application definitions associated with different applications as their basis.

Viewing Application Information

Once you have approved an application, you can view various information about the application beyond what the Application page lists in the table. When you click on an application listed on the Application page, you will see the following tabs for that application:

- Summary: This gives you general information about the application, such as the following:
  - Name
  - All Owners
  - All Cloud Accounts
  - Created With (This indicates whether the application was created manually or with a particular discovery rule.)
  - Associated Labels
  - Resources by Deployment (This circle graph indicates service categories, service roles, resources, security controls, and firewall rules. Click on the graph to see details, such as the security control count. The security control count for each resource is the total number of Security Groups, Network Security Groups, and Network ACLs associated directly with the resource with the subnet of which the resource is a part .)
- Inventory: This gives you an application-specific view of what you would see on the Inventory page. See Inventory [157].

- Traffic: This gives you an application-specific view of what you would see on the Traffic page. See Traffic [219].
- Map: This gives you an application-specific view of what you would see on the Cloud Map page. See Cloud Map [168].
- Policy: This gives you a list of active policies for the application. See Cloud Policy Model [246].

Exporting an Application Report

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Export+Application+Report.mp4

1. Click **Export** on the Applications page.
2. Edit the report name and select the format.
3. Click the **Scheduling Section** toggle to the on position to schedule the export unless you want to export the report immediately.
4. If you choose to schedule your report, select your recurrence and time.
5. Click **Save** when done.
6. Go to the Reports [212] page to download the exported report.

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 [241].

See Writing Application Policy [255].

Cloud Tag to Label Mapping

Learn about the purpose of the Illumio Segmentation for the Cloud cloud tag to label mapping feature, and see a general example of how you would use it. To learn about viewing system-created labels with the Category Labels and Service Role Labels tabs, see View System Labels [243].

> **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 Cloud 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`, and so forth, 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. In 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:

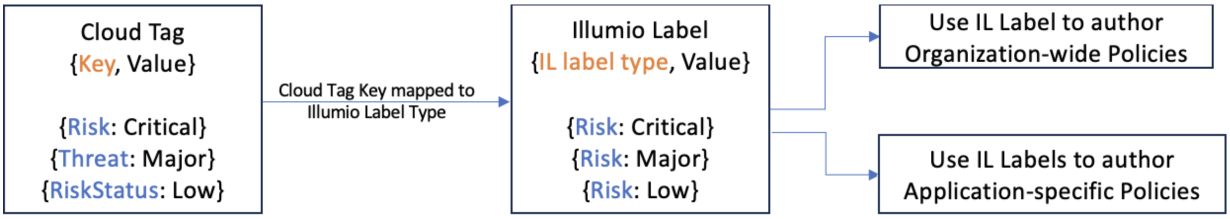
Cloud Tag
{Key, Value}
{Risk: Critical}
{Threat: Major}
{RiskStatus: Low}
Cloud Tag Key mapped to Illumio Label Type
Illumio Label
{IL label type, Value}
{Risk: Critical}
{Risk: Major}
{Risk: Low}
Use IL Label to author Organization-wide Policies
Use IL Labels to author Application-specific Policies

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 mapping:
   - Cloud tag key `Risk` mapped to Illumio Label `Risk`

   For example, if your resource has a cloud tag like `Risk:Critical`, you would map it to corresponding Illumio labels by specifying the tag key in the tag to label mapping. If you created a mapping using the tag key `Risk`, the resource would have the Illumio label `Risk:Critical`.

   You can also map multiple cloud tag keys to one Illumio label type, such as 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 Illumio.

   The following example supposes that you have an application that you wish to define using resources that you have associated with tag to label mappings.
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 polices 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`.

NOTE

Illumio recognizes GCP labels under Illumio cloud tags. This means that when you use the tag to label mapping feature for GCP, cloud tags appear in the dropdown menu with the relevant prefix that indicates it is a GCP tag or label. For example, cloud tags for GCP may have values like `label/gcp-key:gcp-value`.

Illumio supports GCP resource manager tags and labels at this time. Because GCP label values are optional, you may occasionally see empty tag values.

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 [250]. For information on tag to label mapping, see Cloud Tag to Label Mapping [241].

Category Labels

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

Use AI Labeling

Overview

Through machine learning, Illumio recommends day-one labels for select resources. The following are points to consider:

- At time of writing, Illumio Segmentation for the Cloud supports AI labeling on AWS resources only and is available for US cloud resources only. Contact your Customer Success or Account team to request access to this feature.
- AI label-based recommendations draw from a pool of over 300 labels, including role and application labels.
- The first time you use this feature, you need to acknowledge the dialog that appears.
- Optionally, you can accept, deny, or ignore the label recommendation, which can be in an Accepted, Denied, or Pending state.

Use AI Labeling

To get started from the left navigation pane, browse to **Labels > AI labels** tab.

The tab presents tiles listing your application and non-application resources that have pending AI recommended labels. The tile for application resources provides context about which resources belong to the applications that have the recommendation. The tile for non-application resources suggests labels for resources that do not yet belong to applications, which can help with visibility and rule-writing. The recommendations update every 24 hours.

The tab also presents a table listing the individual recommendations. You can filter by application to narrow the list of visible recommendations. You can also filter by Pending, Approved, or Denied, by clicking the respective buttons above the filter dropdown menu.

1. For the application of interest, hover over the label to view the label details, including the reason for the recommendation. You can also hover over the resource to see its details. You can copy the details of the resource or the label from the hover dialog. Note that a listed resource may have more than one entry, with each one recommending a different label.
2. Select the application/resource listings with the recommended labels that you want to apply.
3. For bulk approval or denial, click **Approve** or **Deny** above the filter dropdown menu. For individual approval or denial, you do not need to select the application/resource listing. Instead, click the vertical breadcrumb icon at the end of the row and click **Approve** or **Deny**.
4. If you click **Approve**, a message notifies you that the label has been applied to the resource, and the recommendation moves to the end of the list. Note that if you approve a label, Illumio Segmentation for the Cloud will cease recommending that specific label for that resource.

   If you have an existing label on the resource, a confirmation dialog lets you approve overriding the existing label.
5. If you click **Deny**, a message notifies you that the label has been denied for the resource, and the recommendation moves to the end of the list.

Guidelines

Keep the following guidelines in mind when you use the AI labeling feature:

- You can revise your decisions
  - If you approved a label that you later decide to deny, click **Approved** above the filter dropdown menu. Then select the label and click click **Deny**.
  - If you denied a label that you later decide to approve, click **Denied** above the filter dropdown menu. Then select the label and click **Approve**.
- You can hover over AI recommended labels to see details such as who denied or approved it, when they did so, why it may have been recommended, etc.
- • Recommend update: continues to suggest updated labels for this resource
  - Do not recommend until date: pauses label recommendations for this resource until a date you specify in a dropdown calendar that appears
  - Never recommend: permanently stops recommending labels for this resource
  - Report as incorrect labeling: Select this checkbox if the labeling suggestion is substantially incorrect
- If a resource already has a role label created by Tag to Label Mapping, and Illumio Segmentation for the Cloud recommends a new role label for it, a small exclamation point displays for that recommendation. This exclamation point provides the context for the existing and recommended Role label value to inform your approval decisions. (Before making the approval decision, you must take into account whether existing policies are using the existing label.) If you choose to approve the recommendation, a dialog tells you that approving the recommendation removes the existing label in the process of applying the recommended one. Once the approval action is taken, the AI will not recommend the role label for the resource again.

Policy

The content in this category introduces you to the Illumio Segmentation for the Cloud policy model and explain key features of Illumio Segmentation for the Cloud, such as policy attributes and the differences between organization and application policies, as well as how to write both of those varieties of policies.

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 Segmentation for the Cloud 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 [241] 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 Segmentation for the Cloud. 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 [247] 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 policy

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

- Writing Application Policy [255]
- Writing Organization Policy [254]

Overview of policy attributes

Illumio Segmentation for the Cloud uses the following policy attributes that help you write your security policy:

- Labels [250]
- Services [251]
- IP Lists [247]

Policy attributes

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.

See IP Lists [247].

See Labels [250].

See Services [251].

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.

Illumio allows use of '!' to exclude IP addresses.

Adding IP Lists

1. To add an IP List, browse to **Policies > Policy Objects > IP Lists** and click **Add > Manually** or **Add > Bulk Import**.
2. If you clicked Bulk Import, choose a .CSV file and click **Add to IP Lists**. Note that the .CSV headers include "name,description,attribute,ranges" and must not include spaces. Attributes are optional.
3. Enter a name, a description, and an attribute type. Use attribute types to more easily filter your IP lists. You can use default attributes or create your own.
4. Type in the IP addresses field to add the first IP address, IP address ranges, or CIDR blocks. Press **Enter** to do a carriage return and type another address if needed.

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

5. Type in the FQDN field to add your fully qualified domain names. Group FQDNs in an IP list to enable FQDN-based policy enforcement on security controls.

   > **NOTE**
   > Illumio Cloud does not support FQDNs at this time.

6. Click **Save**.

   Your IP List appears in the table.
7. Select your new IP List in the table and click **Provision**.

Examples of Different IP List Entries

Single IP

You can create IP lists that use IPv4 or IPv6, and Illumio will decorate them, but you cannot search for the traffic of an IP list that contains IPv6.

Examples:

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

CIDR Block

Use a slash to indicate a CIDR Block. You can indicate CIDR Blocks that use IPv4 or IPv6, and Illumio will decorate them, but you cannot search for the traffic of an IP list that contains a CIDR block.

Examples:

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

IP Ranges

Use a hyphen to indicate an IP range. Use cases include when you can see traffic decorated with an IP list that contains a range, but you cannot search for the traffic of an IP list that contains a 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 Notes

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 Segmentation for the Cloud 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 Cloud to discover applicable deployments for applications.
Service-Category	This label type allows you to describe key resources by their categories, 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 Illumio visibility for resource types [193]. To view your ServiceCategory labels, see View System Labels [243]. Note that this label type is system-generated and cannot be edited or removed. Also note that Illumio Segmentation for the Cloud 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 [255].
Service-Role	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 youre ServiceRole labels, see View System Labels [243]. Also note that Illumio Segmentation for the Cloud 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 [255].
Other labels	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.

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 [241] 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 Cloud 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 automatically [231] for information.

AI Recommended Labels

Through machine learning, Illumio recommends day-one labels for critical resources. At time of writing, this is for role labels only. See Use AI Labeling [244].

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.

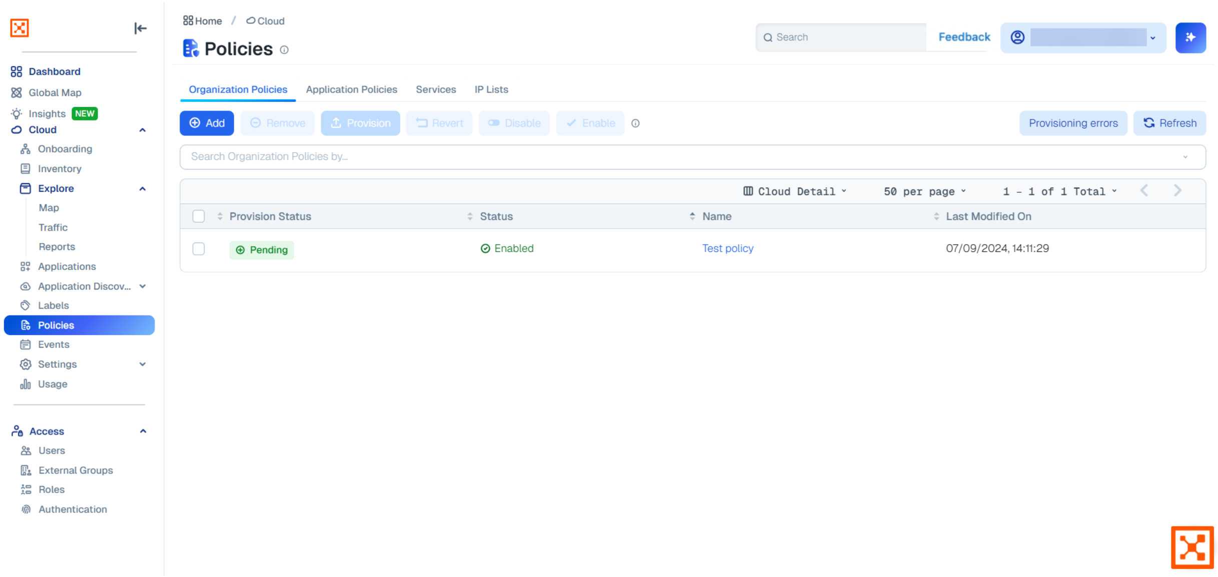

□	Provision Status	Status	Name	Last Modified On
□	Pending	Enabled	Test policy	07/09/2024, 14:11:29

For information about creating these types of policies, see Writing Organization Policy [254] and Writing Application Policy [255].

About Illumio Segmentation for the Cloud Policies

The Policies page lists all the different policies you have created. 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 [254].

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 may want to first define services and IP lists by going to the Policies menu and selecting the **Services** and **IP List** tabs. See Services [251] and IP Lists [247] for information. This is optional.

You may also 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 [241] for information. This is optional, as the application definition workflow itself also creates labels.

To write application policies, go to **Applications > your application > Policy** tab. See Writing Application Policy [255] for information.

Writing Organization Policy

This topic provides an overview of using rules to write organization Illumio Segmentation for the Cloud 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 Segmentation for the Cloud policy model, see Cloud Policy Model [246].

For a list of resources against which you can write policy, see Resources that Support Policy [262].

In order to write policy, you must create rules for the policy. Illumio Segmentation for the Cloud 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 Illumio 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 [252] 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 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 [255].

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 [255] 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 Segmentation for the Cloud policy model, see Policy model [246].

For a list of resources against which you can write policy, see Policy enforcement and resource types [262].

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

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

- 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.

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 [252].

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, Illumio Segmentation for the Cloud 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, Illumio Segmentation for the Cloud 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 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 [256].
- 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 [257].
- 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, Illumio Segmentation for the Cloud 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 Segmentation for the Cloud 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.).

If Source is	And Service is	Destination can be
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 Segmentation for the Cloud 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 Segmentation for the Cloud 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, consider previewing what its impact will be when it's provisioned. This can help you gauge how the policy will map to destinations, security group rules, 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 security controls from the drop-down menu:

- All security controls
- Azure NSGs
- AWS Security Groups
- Network Access Control Lists
- Azure Firewall Policies
- GCP Firewall Rules

Each of these show you the following:

- CSP
- Resource
- Account ID
- Number of Protected Resources

Rules

If you select any particular affected AWS SG, AWS NACL, Azure NSG, or Azure Firewall, you may see rules that come from other applications and/or policies. Note that only Illumio-written rules will display. The 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 Illumio Segmentation for the Cloud 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.

Application policy 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, an Illumio 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 Policy enforcement and resource types [262].

> **NOTE**
>
> Because Illumio Segmentation for the Cloud 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
> - AWS Lambda Functions

Writing Azure Firewall policy

Security teams can drive segmentation policies to control network traffic using Illumio labels, and services. They can also use IP/IP lists to define what can talk to applications, what data can be transferred from an organization's network, and more. Creating policies, including Azure Firewall policies, is critical to minimizing an attacker's lateral movement. Use your Azure firewalls as enforcement points in your cloud environment and minimize an attacker's lateral movement.

> **NOTE**
>
> Illumio Segmentation for the Cloud does not support Classic Azure Firewall.

Illumio allows or denies traffic between Azure Firewalls using policies that you write. In order to write policies, you must create rules for the policy. Illumio Segmentation for the Cloud has the following types of rules for Azure Firewall 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.

Before you begin writing Azure Firewall policies

Review these topics to get an understanding of Illumio policies.

- For an overview of the policy model and how it helps you protect your resources, see Policy model [246].
- For general policy writing information such as guidelines, permitted rule writing combinations, provisioning, and caveats see Writing application policy [255] and Writing Organization Policy [254].
- For a list of resources against which you can write policy, see Policy enforcement and resource types [262].
- Verify that you have at least one Azure Firewall policy resource already attached to your Azure Firewall. See Verify Azure Firewall Policy resource [261].

Understanding Azure Firewalls

An Azure Firewall functions as a central enforcement point for network security within your Azure environment. This allows you to define and apply network and application level security rules across multiple virtual networks and subscriptions. This also allows you to effectively manage all traffic filtering from a single point of control through its centralized policy management capabilities.

In the Azure console, firewall policies have rule collection groups that contain rule collections. When you create a policy in Illumio, it creates a rule collection group with the prefix "ICS." The rule collection group that it creates always attempts to get the highest priority available. Illumio maintains this rule collection group on its own. If you modify this rule collection group directly in the Azure console, Illumio will overrule the modification. See Tamper Protection [272].

Rule collections have the following default priorities, where Illumio Segmentation for the Cloud writes the rules inside each collection based on type:

- Override Deny Rules: 100
- Allow Rules: 200
- Deny Rules: 300

In Illumio Segmentation for the Cloud, an Azure Firewall is a VNet-level enforcement point. If a VNet has a firewall, it is a hub. If a VNet has a peering relationship with a hub, it is a spoke.

At time of writing, Illumio Segmentation for the Cloud supports writing Azure Firewall policy only for network rules.

Write a policy to allow traffic between spokes on the same hub

Suppose that you want to write a policy to allow HTTP traffic between a pair of spokes peered to the same hub. Let's assume they're called AzSpoke1 and AzSpoke2.

1. Select the Applications menu and click one of your applications with one of the two the associated Azure Firewall spokes. Let's choose AzSpoke1.

2. Add an Allow rule.
3. Select the AzSpoke2 application from the Source drop-down list.
4. Select the AzSpoke1 application from the Destination drop-down list.
5. Select HTTP as your destination service and save the rule.

   Your policy is now ready to be provisioned. Before you provision the policy, preview the impact after it is provisioned to gauge how the policy will map to destinations, security group rules, enforcement points, and so forth.
6. To see these mappings on your Azure Firewall policy before you provision it, click **Show Impact**. Then choose a security control from the drop-down menu, like **All security controls** or **Azure Firewall Policies**.

   Show Impact shows the priority of Azure Firewall rule collection groups, rule collections, and rules.
7. Provision your policy.

   The rule is Enabled. Illumio Segmentation for the Cloudcreates a rule collection group with the prefix "ICS." It also creates rules in the Azure console for that firewall's policy and its parent policy.

NOTE

To troubleshoot any Illumio system messages you see regarding Azure Firewall rules, see Troubleshoot system-generated Azure Firewall messages [347].

Verify Azure Firewall Policy resource

Before you proceed with writing Azure Firewall policies in Illumio Segmentation for the Cloud, verify that at least one Azure Firewall Policy resource is already attached to your Azure Firewall.

NOTE

Use Case: You onboarded an Azure subscription or tenant and wish to enforce segmentation using Illumio policies.

Verify that an Azure Firewall Policy resource is attached

1. In Illumio, browse to the **Inventory > Cloud Resources** tab and filter the search field for Azure Firewalls.
2. Click your Azure Firewall and click the **Attached Resources** tab in the **Details** panel that appears.

   The tab displays your attached Azure Firewall Policy resource with its details.
3. If at least one Azure Firewall Policy resource is attached, you can proceed with writing Azure Firewall policies in Illumio. See Writing Azure Firewall policy [259].

   If you do not have an Azure Firewall resource attached, you need to create one in the Azure console. See Microsoft's website.

Policy enforcement and resource types

Illumio Segmentation for the Cloud 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 resource types that appear in the Inventory page, and additional details such as flow support, map support, and attached resources, see Illumio visibility for resource types [193].

AWS

Category	Resource Type
Compute	EC2 Instance
Containers	EKS Cluster
Databases	ElastiCache CacheCluster
Databases	MemoryDB Cluster
Databases	RDS DB Cluster
Databases	RDS DB Instance
Data Warehouse	Redshift Cluster
Network Routing	ElasticLoadBalancingV2 Load Balancer
Serverless	Lambda Function

Azure

Category	Resource Type
Compute	Virtual Machine (inclusive of "spot" VM)
Compute	Virtual Machine ScaleSet Virtual Machine
Containers	AKS Cluster
Databases	CosmosDB
Databases	DocumentDB Database Account
Databases	DBforPostgreSQL Flexible Server
Databases	DBforPostgreSQL Server
Databases	SQL Managed Instance
Databases	SQL Server (Microsoft.Sql/servers)
Network Management	Private Link Service: The following resources are attached to a subnet via a private link service. See note below table. • Load Balancer
Network Routing	Load Balancer
Network Routing	Private Endpoint: The following resources are attached to a subnet via a private endpoint. See note below table. • App Service (Web App, Function App) • DocumentDB/MongoDB Cluster • DocumentDB/Database Account • Cosmos DB • SQL Managed Instance • SQL Server (Microsoft.Sql/servers) • Storage Account • Key Vault
Network Security	Azure Firewall
	NOTE Illumio Segmentation for the Cloud does not support Classic Azure Firewall.

GCP

Category	Resource Type
Compute	Cloud SQL Instance. See note at bottom of page.
Container	GKE Cluster
Database	SQL Instance. See note at bottom of page.

OCI

By participating in the BETA program for OCI features you agree that your company's use of the BETA version of OCI features will be governed by Illumio's Beta Terms and Conditions.

Category	Resource Type
Compute	Instance. See note at bottom of page.

AWS ENI Considerations

NOTE

Because Illumio Segmentation for the Cloud 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
- AWS Lambda Functions

Azure, OCI, and GCP Considerations

NOTE

Illumio Segmentation for the Cloud applies policies for the different resources listed under Private Endpoints and Private Link Services by applying rules to the NSG on the subnet that hosts them.

NOTE

Azure PaaS databases must have private endpoint connectivity for Illumio Segmentation for the Cloud to enforce policies on private endpoint NSGs.

NOTE

Illumio enforces policies on OCI instances only if they have NSGs attached.

NOTE

When enforcing policy on GCP Cloud SQL instances, it must be done through forwarding rules. You can enforce policies on SQL traffic by targeting forwarding rules that route traffic to SQL-based services over Private Service Connect (PSC). By modeling the application at the VPC or subnet level, you can apply fine-grained controls to SQL endpoints behind load balancers or PSC endpoints, ensuring secure and auditable access.

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 [254].
- 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 [255].

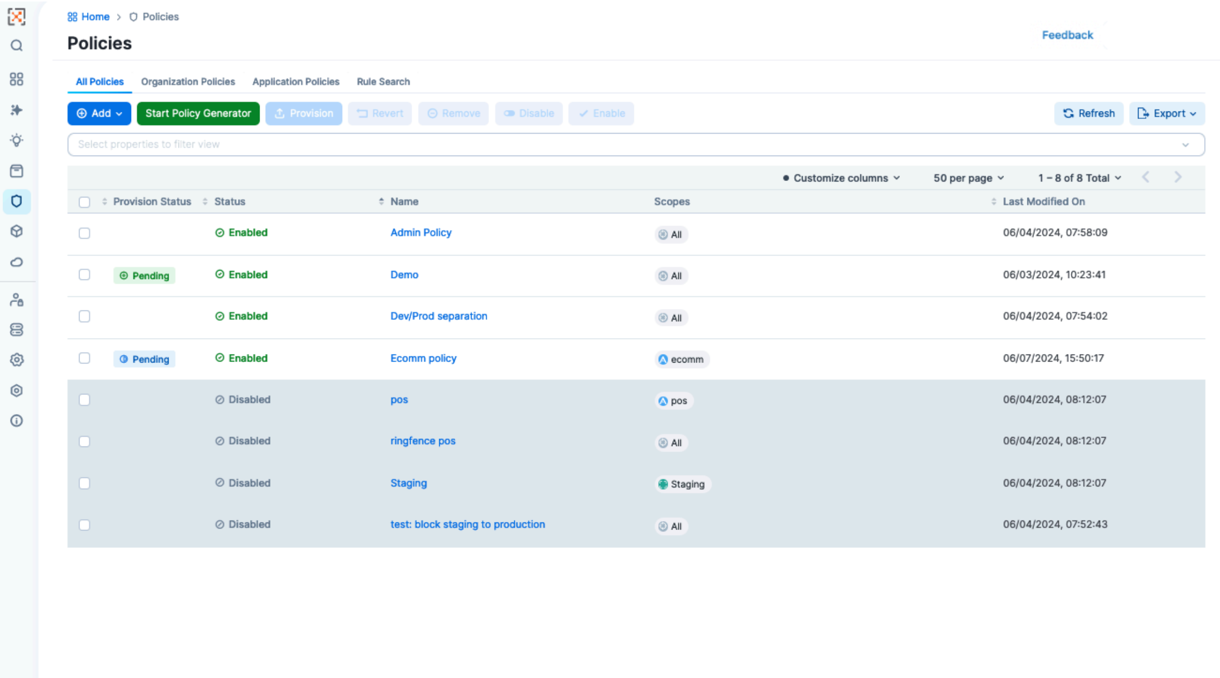

Provision Status	Status	Name	Scopes	Last Modified On
	Enabled	Admin Policy	All	06/04/2024, 07:58:09
Pending	Enabled	Demo	All	06/03/2024, 10:23:41
	Enabled	Dev/Prod separation	All	06/04/2024, 07:54:02
Pending	Enabled	Ecomm policy	ecomm	06/07/2024, 15:50:17
	Disabled	pos	pos	06/04/2024, 08:12:07
	Disabled	ringfence pos	All	06/04/2024, 08:12:07
	Disabled	Staging	Staging	06/04/2024, 08:12:07
	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 [252].

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:

- Custom IPtables Rules

  You can write rules for Linux workloads.

- 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.

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 [254].

Security reviews

Enable Azure and AWS policies using security reviews. Illumio Segmentation for the Cloud security reviews ensure that users review policy enforcement on Azure subscriptions and AWS accounts, reducing the risk of implementing ineffective rules. First, onboard your Azure subscription, Azure Tenant, AWS account, or AWS organization with either read or read and write permissions. Then, perform a security review to check for existing organization policy issues and review subscription security controls to correct ineffective policies. Conduct a security review before you write application or additional organization policies for your subscription.

> **NOTE**
> Use Case: You onboarded an Azure subscription with read and write permissions. You have now decided that you want to write organization policies for Illumio Segmentation for the Cloud to enforce on your subscription. You must review and approve your subscription policies before you enforce any application policies on your subscription.

Perform the security review

1. Onboard your Azure subscription, Azure Tenant, AWS account, or AWS organization with read and write permissions. If your subscription has only read permissions, see Change Azure permissions from read to read and write [93] and Enable read-write permissions [269]. If your account or organization has only read permissions, also see Enable read-write permissions [269].

   The Onboarding page Cloud Accounts tab shows that the security review is pending approval for your subscription.
2. Navigate to **Cloud > Security Review** and select your subscription.
3. Click **Review**.

   Using Azure as an example, the Security Review page displays your subscription's organization policies. You can filter by 'Has Policy Conflicts' or 'No Policy Conflicts,' which display warning icons or check marks respectively in the Policy Effectiveness column.

   A warning icon means the policy has one or more rules requiring your attention before you can effectively implement it. Review each organization policy with a warning icon to ensure it has the appropriate security controls and applications, as described in the next step. Illumio strongly recommends reviewing at least the policies with a warning icon.

Rule effectiveness can be defined as able to meet the flow condition, able to successfully provision so that it can act by denying or allowing a given flow, and able to be evaluated. If an existing rule is found to be effective, any equivalent rules of lower priority are not evaluated, and would be found ineffective by default. For example, a customer-defined rule that allows certain traffic, followed by an Illumio-written rule that denies the same traffic, cannot be effective due to order. Note that Illumio Segmentation for the Cloud attempts to give Illumio-written rules higher priority than non Illumio-written rules, although this may not always be possible.

4. Click **Review** for an organization policy.

   The Policy Impact tab displays the policy security controls. Reasons for ineffectiveness can include:

Reason	Description
Max Rule Limit	Because all new rules are non-effective, this warning doesn't return any particular rule. For example, SGs and NSGs can accommodate only a given number of rules. The rule limit varies depending on your setup.
Resource locked	For example, an NSG could be locked, preventing Illumio Segmentation for the Cloud from implementing rules.
Broader Rule Exists	For example, if you have an existing rule that allows traffic from a broad range of IPs on a given port, any new rules that are provisioned covering a subset of this IP range, *and* are given a lower rule priority (in other words, evaluated after the broader rule), it is marked as ineffective, because there is no scenario in which it can be evaluated.
Conflicting Rule Present	Illumio Segmentation for the Cloud does not implement conflicting new rules if an existing one is more permissive. For example, if you have a customer-defined Security Group that allows all inbound and outbound traffic, and a new, lower priority Illumio Segmentation for the Cloud rule that allows such traffic on port 80 only, the new rule is not evaluated, and is considered ineffective. For rules to conflict, there must be a difference in action take, like Allow vs. Deny, for example. If rules have the same actions, there may instead be a broader rule that makes another rule ineffective.
Unsupported Protocol	For example, this warning could be about Azure NSG policies not supporting ICMPv6 or IGMP protocols.

5. Click **View Rules** to see details about the rules in your security controls and investigate any ineffective rules.

   A side panel displays the security control, its outbound rules, and its inbound rules. It lists which rules are Illumio-authored, modified, deleted, added, and ineffective, by default. Ineffective rules display a warning icon. To view all rules, not just Illumio-authored, modified, and so forth, remove the filter.

6. Click **<number> Ineffective** to assess the ineffective rules.

   The side panel lists the rule and information such as source, destination, port, protocol, and whether it is cloud or Illumio-managed. Click **>** to list individual issues, such as conflicting or broader organization policy rules. For example, your ineffective rule may be a deny rule that conflicts with existing allow rules.

NOTE

Click **Export** to save a .csv of the ineffective rule and the reasons it is ineffective.

7. Resolve issues such as Conflicting Rule Present, Resource locked, and so forth.

Approve the security review

NOTE

If you have not enabled read and write permissions for your Azure subscription, you'll need the following:

- Permissions to run the provided read access script. See Prerequisites for Onboarding Azure [60] and Permissions for Onboarding Azure [61]. If you don't have permissions, see Change Azure permissions from read to read and write [93] and Enable read-write permissions [269].
- A service account and its token

NOTE

If you have not enabled read and write permissions for your AWS account or organization, you'll need the following:

- Permissions to run the provided read access script. See Prerequisites for Onboarding AWS [94] and Permissions for Onboarding AWS [95]. If you don't have permissions, see Enable read-write permissions [269].

1. Resolve any issues you encountered during your security review.
2. Navigate to **Cloud > Security Review,** select your subscription and click **Approve Security Review**.

   If you did not enable read and write permissions when you onboarded your subscription, the approval dialog prompts you to do so. See Onboard an Azure Cloud subscription - default setup [65]. Click the **Azure Deployment Complete** checkbox once you have completed the steps as prompted.
3. Click **Approve**.

Enable read-write permissions

Enable read and write permissions for Illumio to enforce Azure subscription and AWS account policies after onboarding.

NOTE

Use Case: You onboarded an Azure subscription with read only permissions. You have now decided that you want to write application policies for Illumio Segmentation for the Cloud to enforce on your subscription. You must enable read and write permissions.

NOTE

If you have not enabled read and write permissions for your subscription, you'll need the following:

- Permissions to run the provided read access script. See Prerequisites for Onboarding Azure [60] and Permissions for Onboarding Azure [61]. If you don't have permissions, see Change Azure permissions from read to read and write [93].
- A service account and its token

NOTE

If you change your AWS account from read to read and write by downloading the CloudFormation Stack, you must use the same role as you did during the initial onboarding for that account. This restriction does not apply if you instead click the link recommended in the wizard.

If you do not run the original CloudFormation Stack you may see an error like the following:

"Read to Read Write cannot be completed. Please delete and reonboard the AWS integration"

If you see such an error, re-onboard the AWS account as though you had never run the CloudFormation template as a stack in the first place.

If you wish to change an organization's child accounts to read and write, you must first run the CloudFormation template on the organization to update its permissions and then run the CloudFormation template as a stackset so that the update occurs in the child AWS accounts. See Onboard an AWS Cloud organization [107].

Enable read-write steps

1. If your subscription has read permissions only, browse to **Onboarding** and select your account.

2. Click **Enable Read Write**.
   The Read-Write Access Setup dialog displays with the following options:

   - Start with Security review before Enable Read Write Access (recommended)
     This option allows you to review and approve your subscription policies before you enforce any application policies on your subscription. See Security reviews [267].
   - Enable Read Write Access and skip security review
     This option automatically approves the security review without giving you the benefit of reviewing your subscription policy enforcement, reducing the risk of implementing ineffective rules.
3. Click **Acknowledge and Approve**.

Drift Detection

Drift Detection overview

Illumio Segmentation for the Cloud automatically monitors the state of security controls on your CSP for any changes.

For non-Illumio Segmentation for the Cloud-written rules in the security control, Drift Detection generates drift system events when rules are modified, added, or deleted from the security controls. Systems events will show drift to the security control.

In other words, any changes to customer-owned rules are considered to be drift.

NOTE
The feature is currently supported for the following:

- AWS SGs, NACLs
- Azure NSGs
- Azure Firewalls
  Drift Detection for Azure Firewalls is performed only for Illumio Segmentation for the Cloud-managed rule collection groups, rule collections, and rules.
- GCP Firewall Rules

Drift Detection vs. Tamper Protection

Drift Detection and Tamper Protection both alert you to changes made to security control rules. They address different scenarios:

- Drift Detection displays system events on the Events page System Events tab if non-Illumio-written rules are modified, deleted, or added, but takes no action

- Tamper Protection displays alerts on the Events page System Events tab if Illumio-written rules are modified or deleted, and automatically reinforces the original rule. See Tamper Protection [272].

Drift Detection example

In this example, assume you have an AWS SG and you recently created a rule for it in Illumio Segmentation for the Cloud. Suppose that you later add a customer-owned rule in the AWS console.

Drift Detection generates an event that makes note of this change. In the Events page System Events tab, use the filter to select **Event Type: network_security.drift_addition**. You would see an event message like this:

A new rule ExampleRule1 was added to ExampleSG

If you removed a customer-owned rule in the AWS console, that would also be considered drift. In the Events page System Events tab, use the filter to select **Event Type: network_security.drift_removal**. You would see an event message like this:

Rule ExampleRule1 from ExampleSG was removed

Tamper Protection

Tamper Protection overview

By default, Illumio Segmentation for the Cloud automatically monitors for any changes to the state of security controls on your CSP.

For non-Illumio Segmentation for the Cloud-written rules in the security control, Illumio Segmentation for the Cloud generates tamper system event(s) when rules are modified, but does not provide Tamper Protection (in other words, original rules are not enforced back to the security controls). System events show what has changed in each modified rule.

For Illumio Segmentation for the Cloud-written rules in security controls, Tamper Protection generates system events when rules are modified or deleted. Tamper Protection automatically re-enforces correct rules to the security control. System events show what has changed in each rule, if modified, or whether that rule was deleted.

NOTE

The feature is currently supported for the following:

- AWS SGs, NACLs
- Azure NSGs
- Azure Firewalls

  Tamper Protection for Azure Firewalls is performed only for Illumio Segmentation for the Cloud-managed rule collection groups, rule collections, and rules.
- GCP Firewall Rules

Tamper Protection vs. Drift Detection

Tamper Protection and Drift Detection both alert you to changes made to security control rules. They address different scenarios:

- Tamper Protection displays alerts on the Events page System Events tab if Illumio-written rules are modified or deleted, and automatically reinforces the original rule
- Drift Detection displays system events on the Events page System Events tab if non-Illumio-written rules are modified, deleted, or added, but takes no action. See Drift Detection [271].

Tamper Protection Example

In this example, assume you have an AWS SG and you recently created a rule for it in Illumio Segmentation for the Cloud. Suppose you later changed this rule in the AWS console.

Tamper Protection generates an event that makes note of this change, which would be considered tampering. In the Events page System Events tab, use the filter to select **Event Type: network_security.tampering_modification**. You would see an event message like this:

Illumio Cloud authored rule ExampleRule1 from ExampleSG was modified. Correct rule will be enforced again.

If you removed an Illumio Segmentation for the Cloud-generated rule in the AWS console, that would also be considered tampering. In the Events page System Events tab, use the filter to select **Event Type: network_security.tampering_removal**. You would see an event message like this:

Illumio Cloud authored rule ExampleRule1 from ExampleSG was removed. Correct rule will be enforced again.

Administer

The content in this category explains how you administer Illumio Segmentation for the Cloud for your cloud environment; in particular, this category explains how to set up a connector with Illumio Segmentation for the Cloud to receive notifications, add and manage users, and review events.

Connector

This topic describes the purpose of the Illumio Segmentation for the Cloud Connector feature, and provides a general example of how you would use it. For instructions on how to connect a specific a workflow and incident management tool, such as Slack, using the Connector page, see the applicable pop-up help in the user interface.

Apps Use Case and Example

This feature lets you connect workflow and incident management tools, such as messaging applications or others, to Illumio Segmentation for the Cloud. For example, you might want to receive a notification in your messaging application when a policy changes, when a deployment is removed, or any other audit event.

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-up help will give you instructions.
2. Depending on your application, you may need to provide the following:
   - Channel Name (Illumio Segmentation for the Cloud does not verify the name, so make sure it is correct.)
   - Webhook URL (This would be how Illumio Segmentation for the Cloud knows where to deliver the message.)
3. The dialog may have fields for other characteristics, depending on the application.
4. When a channel is configured, it automatically starts receiving any subsequent policy provisioning alert. For any other alert type, the automation rule needs to be created.
5. Policy provisioning 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.
6. 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 help in the user interface that is specific to that particular one.

Automation Use Case and Example

This feature lets you automate messaging after you have performed the above steps to connect workflow and incident management tools, such as messaging applications or others, to Illumio Segmentation for the Cloud.

1. The first step is to browse to the **Settings > Connector > Automation** tab.

2. Click **Add Rule** and enter a name in the dialog that appears. In this example, you want to have Slack notify you of system audit events. You might name it 'Successful Policy Update to Slack.'
3. Select one or more triggers by clicking **Add Trigger** and then selecting a trigger in the dropdown menu. In this case, you might pick something like 'policy is provisioned.'
4. Select one or more actions by clicking **Add Action** and then selecting an action in the dropdown menu. In this example it might be 'send a Slack message.'
5. Under Slack Channel, select a Slack channel, and under Message, enter a message.
6. Under Date and Time, select either **Send immediately**, or **Send later** and specify a time and frequency.
7. You may click **Send Test** to verify your system event Slack notification before you click **Add Rule** in the Add Rule dialog.

S3 Bucket Use Case and Example

This feature lets you connect Illumio Segmentation for the Cloud to your AWS S3 buckets so that you can export Illumio Segmentation for the Cloud traffic flows to your S3 buckets.

Exporting traffic flow logs sends enriched flow logs from Illumio Segmentation for the Cloud to the storage destination in your account. Flows are sent in batches every 60 minutes as the flow logs are collected. This feature is not limited to the 10,000 flows maximum allotment in the Traffic dashboard. The flow logs sent to the S3 bucket will include all the Illumio Segmentation for the Cloud-processed flows based on the filters set while configuring the export.

During the flow processing, Illumio Segmentation for the Cloud enriches these raw flow logs from the cloud providers with labels on the source and destination. This adds greater enrichment to the traffic flows which provides greater context for users to understand when viewing and investigating traffic through these flow logs.

In some cases, cloud providers send only partial flow data to Illumio Segmentation for the Cloud. When this happens, Illumio Segmentation for the Cloud makes periodic requests back to the cloud provider for the full flow log. This process can take up to forty minutes to retrieve the full flow log. To reduce issues related to partial flow data or duplicate flows being sent, Illumio Segmentation for the Cloud batches flow logs once it ensures that the full flow data is captured from the cloud provider.

Onboarding an S3 Bucket

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Connector.mp4

Use these steps to onboard an S3 bucket:

1. Log into the Illumio Console and navigate to the **Settings > Connector** tab.
2. Click the **S3 Bucket** tile.
3. Click **Connect S3 Bucket**.
4. In the dialog that appears, choose the radio button for either onboarded or unknown AWS accounts and select entries for the following credentials:

- Account ID
- S3 Bucket ARN
- Region

5. Click **Next**.
6. Select a Service Account.
7. Select your preferred type of Integration. Illumio recommends creating a cloud formation stack. Create the appropriate roles in the AWS console, and when you are done, click **Next** in the Illumio Segmentation for the Cloud dialog.
8. Click **Save**. This completes the connection and takes you to a list of added S3 buckets.

You can delete S3 bucket connections by selecting one S3 bucket at a time from the list and clicking **Remove**.

Testing the Onboarded S3 Bucket Connection

Test your connection to ensure that Illumio Segmentation for the Cloud exports traffic to your S3 bucket with the following steps:

1. For the desired S3 bucket in the list, click **Test Connection**. You will get either a 'Connection Successful' or a 'Connection Failed' message.
2. If you got a failure message, click **Configure** to change your selections as needed to successfully connect.
   - Verify that the provided account ID, bucket ARN, and region are correct
   - Once verified, grant access again by running the cloud formation template to grant Illumio Segmentation for the Cloud access to the bucket
   - Save the changes and the test connection again. If the cloud formation template succeeded, the connection should work.
3. If you got a success message, there is nothing more you need to do for that connection.

Exporting Traffic Data to an Onboarded S3 Bucket

Use the following steps to export data to your S3 buckets.

1. On the Traffic page, filter your traffic as desired and click **Export > Export to Connector**.
2. In the dialog that appears, choose the following selections:
   - Export Format: CSV or JSON
   - Connector: S3 Bucket
   - S3 Bucket: The S3 bucket of current interest
   - S3 Bucket Prefix Name: An optional prefix with meaning to you that will assist in sorting your exported collection of data
3. If you wish, click **Test Connection**.
4. When you are satisfied with your selections, click **Save**.
5. After saving, view your Illumio Segmentation for the Cloud traffic query export statuses in the Settings > Connector > S3 Bucket tab, under the specified bucket. Traffic data begins to appear in the AWS console S3 bucket on an hourly basis.

   Note that you can enable or disable traffic data export for a connector. To do this, find it in the connector list and click on it. In the panel that appears, select an export and click **Enable** or **Disable**, as appropriate. Disabling prevents export of any more flows for that specific configuration.

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 Segmentation for the Cloud capabilities for their zero-trust segmentation programs. Some users have owner privileges, meaning that those users can perform the same tasks . All Illumio Segmentation for the Cloud users are assigned a security administrator role that provides them access to all the capabilities in the product, including the ability to invite additional users. Servers & Endpoints 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 [58].

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.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Add+Users.mp4

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 [58].

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 Servers & Endpoints 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 [279].

My Profile

You can view and change the following in your profile:

- Email Address/Username (not editable)
- First Name
- Last Name
- Time Zone
- Color Mode (Normal Vision, Color Vision Deficiency)
- Reset Password

My Roles

The upper right-hand corner has a dropdown menu where you see your login name. The My Roles menu item, Cloud customers view current roles by the following columns:

- Scopes - Illumio Segmentation for the Cloud values may include: All
- Roles - Illumio Segmentation for the Cloud values may include: Owner, Viewer, and other Illumio Segmentation for the Cloud roles. A user may have a combination of any of these roles.

The following should be noted:

- Only unscoped roles are applicable to Illumio Segmentation for the Cloud
- Global roles other than Owner are not applicable to Illumio Segmentation for the Cloud
- Users with both Illumio Segmentation for the Cloud and Servers & Endpoints subscriptions may see roles and scopes that apply to only one of those products.

Role-based access control

This section describes how to manage user roles for your organization. For user management topics, see User Management [276]. For more information on roles, see My Roles [278].

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.

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Add+Roles.mp4

1. Click **Add Role> Add Unscoped Role**. (Scoped roles are not applicable to Illumio Segmentation for the Cloud.)
2. Select one or more of the roles listed below and click **Grant Access**.

The following roles are available for Illumio Segmentation for the Cloud:

- Multi-product Roles

The user has these roles for Illumio Segmentation for the Cloud and other Illumio products on the Illumio Console:

- Owner: All read and write permissions.
- Viewer: Read permissions only.
- Illumio Segmentation for the Cloud: -specific Roles

The user has these roles for Illumio Segmentation for the Cloud only:

- Cloud Security Onboarding Administrator: Grants permissions to add and delete integrations and flow log access, and is used by AWS Lambda and Azure PowerShell callbacks.
- Cloud Security Policy Author: Allows policy creation for all the workloads and applications in all the accounts that have been onboarded. This role includes permissions to view applications, labels, maps, inventory, and traffic.
- Cloud Security Label Administrator: Access to Tag-Label Mapping only. This governance is useful when authoring policies based on Illumio Segmentation for the Cloud labels.
- Cloud Security Auditor: Grants read-only permissions to view Cloud Map, Inventory, Traffic, and Policies (Organization and Application).
- Cloud Security Incident Responder: View access to Traffic, Inventory, and Cloud Map functions only.

- 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.

Service accounts

This topic describes the purpose of the Illumio Segmentation for the Cloud Service Accounts page, and provides a general example of how you would use it.

Use case and example

For best results, Illumio recommends viewing videos in Chrome.

https://product-docs-repo.illumio.com/Tech-Docs/CloudSecure/Videos/Service+Account+Use+Case.mp4

This feature lets you manage your service accounts that Illumio uses to interact with its own services (Cloud services) rather than directly with your AWS services. You will use the generated secret for the user account later, when you create a cloud formation stack.

The following steps illustrate how you might add a service account.

1. The first part of the sequence would be to browse to the **Settings > Service Accounts**.
2. Click **Add**.
3. Enter a name and description and click **Save**.
4. Click on the new service account in the table and click **Add**, under the SECRETS heading.
5. Enter a name and click **Save**.
6. Click **Copy** to copy the credential secret. This is very important. To see it, click **Show**.

   The next step would be to paste this copied secret into the credential secret field in the Onboard AWS by Running CloudFormation Stack section of Onboard an AWS Cloud Account [113].

If you wish to change the description for a service account, click the account in the table and click **Edit** (you cannot edit the name). If you want to remove secrets from a service account, click the account in the table, select the secret, and click **Remove**.

Illumio Terraform source

Terraform is an infrastructure-as-code tool that enables teams and organizations to safely provision and manage infrastructure in any cloud.

The Illumio Segmentation for the Cloud Terraform source is the first-class API support for the Illumio Segmentation for the Cloud platform. Illumio supports the creation of applications, deployments, and tag to label mappings with Terraform. In addition to the resources managed in the source, Illumio includes Terraform modules for Illumio Segmentation for the Cloud use. This includes:

- Onboarding of AWS and Azure Cloud Providers
  See Create a Terraform Illumio Onboarding Application for Azure [75].
  See Onboard an Azure Subscription using a Terraform Illumio Onboarding Application [78].
- Onboarding of network flow logs from AWS and Azure
- Creating deployments
- Creating applications
- Creating application policy rules
- Creating tag to label mappings
- Creating IP lists

Illumio frequently updates the Illumio Segmentation for the Cloud Terraform source, so refer to this documentation periodically.

Access, subscribe, and contribute to the Illumio Segmentation for the Cloud Terraform source using the links below.

Provider and modules in Terraform Registry

- Illumio Terraform Provider Configuration Documentation
- Illumio Terraform Modules

Github

- Illumio Terraform Provider on GitHub
- Illumio Terraform Modules on GitHub

Terraform module examples

- AWS onboarding
- AWS onboarding with flow logs
- Azure onboarding
- Azure onboarding with flow logs

Examples and schemas

Application schemas

- Terraform Application Schema
- Terraform Application AWS Resources Schema
- Terraform Application Azure Resources Schema

Deployment and tag to label schemas

- Terraform Deployment Schema
- Terraform Tag to Label Mapping Schema

Policy schemas

- Terraform Application Policy Rules Schema
- IP List Schema

Policy preferences

Learn about setting your policy preferences. This feature lets you define your security control preferences at the tenant level. The default for Azure is to have both Subnet and NIC Network Security Groups (NSGs) selected. The default value for AWS is to have Security Groups (SGs) selected.

For Azure NSGs, you can choose to apply rules at the NIC-level, subnet-level, or Azure Firewall-level (beta-only).

NOTE
Illumio Segmentation for the Cloud does not support Classic Azure Firewall.

For AWS environments, you can choose between configuring Security Groups, Network Access Control Lists (NACLs), or both. For example, if you switch from programming rules at both the Subnet and NIC levels (i.e., NACLs and SGs) to the NIC-level (SGs), Illumio Segmentation for the Cloud removes all the written rules from the NACLs. However, the Security Group rules remains intact and continues to get updated whenever there are changes to the policy or inventory resources. Conversely, if you switch from a NIC configuration to both NIC and subnet-level security controls, the NACLs are reprogrammed with Illumio Segmentation for the Cloud-written rules to reflect the updated policy.

An error is displayed if the rule limits are exceeded. In such cases, Illumio Segmentation for the Cloud does not apply the updated policy, and the last enforced policy remains active.

Set your enforcement points

1. To set your preferences for enforcement points, browse to **Settings > Policy Preferences** in the left-hand navigation panel. As each cloud environment can vary, this feature lets you choose a setting that covers your cloud environments best.
2. Click **Edit** to set your enforcement points to include different settings as described following. When you select a setting for a given CSP, explanatory text appears next to that selection. If you choose a default value, a message displays, saying that those are recommended. If you choose a non-default value, a different message displays, saying that there may be a potential effect on traffic flows until your changes take effect.

   Azure

   - All Azure Enforcement Points
   - NIC NSGs
   - Subnet NSGs

   AWS

   - Both NACLs and SGs
   - Security Groups
   - Network Access Control List (NACLs)
3. Click **Save** when you are done. This exits the editing mode, with only the current values displayed.

For information on policies, see Policy model [246], Writing application policy [255], Writing Organization Policy [254], and Writing Azure Firewall policy [259].

Licensing

Licensing and Usage for Illumio Segmentation

Illumio Segmentation license usage is measured using the number of Segmentation Workloads (SWL). SWLs are calculated based on the number and type of resources in your environment that are supported for segmentation.

Each resource is mapped to an Illumio Resource Type and grouped using these categories:

- Cloud Resources
- Data Center Resources
- Endpoint Resources

Cloud Resources

Illumio Segmentation has 7 resource types for cloud resources. Each resource type groups similar resources from different Cloud Service Providers (CSPs) together. As each CSP follows its own naming conventions, this table shows how their resource names match up with Illumio's resource types for Segmentation.

Illumio Resource Types	AWS	Azure	GCP
Cloud virtual machine	EC2 Instances (exclude VMs running containers)	Virtual Machines, Scale Sets Virtual Machines (exclude VMs running containers)	Compute Instances
Cloud container	EC Instances running Containers, EKS Cluster	VMs running Containers, Managed Clusters (AKS)	
Cloud database	Memory DB Cluster, RDS DB Cluster and DB Instance, ElastiCashe Cluster, Redshift Cluster, etc.	SQL Servers and ManagedInstances, Mongo Clusters, DatabaseAccounts, PostgresQL servers, etc.	
Cloud storage		Storage Accounts, etc.	
Cloud firewall		Azure Firewalls	
Serverless function	Lambda Functions	Sites, Functions	
Network routing and resource management	Load Balancers	Load Balancers	
Network security and management		KeyVault	

Data Center Resources

Illumio Resource Types	Definition
Standard Server OS	Number of Servers running non-legacy operating system
Legacy OS	Servers running supported legacy operating systems: • AIX 7.1 Technology Level 4 or greater • AIX 7.2 Technology Level 3 or greater • AIX 7.3 Technology Level 1 or greater • AIX 6.1 Technology Level 9 • Red Hat Enterprise Linux (RHEL) 5 • Solaris 10 Update 8 or greater on Solaris x86 (64 bit) or SPARC (64 bit) architecture • Solaris 11.1 or greater on Solaris x86 (64 bit) or SPARC (64 bit) architecture • Windows Server 2003 • Windows Server 2008 pre-R2
Oracle Exadata Database	Oracle Exadata database instances
Container Host	All the nodes in a Kubernetes cluster when using Illumio Containerized VEN (C-VEN)
Switch Port	Number of managed IPs on a switch
Load Balancer	Number of virtual IPs connected to a load balancer
IBM i-series LPAR	Number of IBM i-series LPAR (Logical Partition)
IBM zLinux IFL	Number of IBM zLinux Integrated Facility for Linux (IFL)
IBM zOS LPAR	Number of IBM zOS LPAR (Logical Partition)

Endpoint Resources

Illumio Resource Type	Definition
Endpoint	Count of virtual or physical desktops or laptops

Determining Segmentation Resource Count

To estimate resource count by resource type in your environment, run the resource count script provided by your technical contact at Illumio. If you have onboarded your cloud subscription, you can view the actual resource count under the Usage menu (Usage > Illumio Segmentation tab > Usage Details).

Daily, Weekly, and Monthly Averages

Cloud resource counts are taken every hour.

- To calculate daily usage, the average of the hourly counts over the past 24 hours is used. The monthly average resource count is then calculated by averaging the daily usage values across all days in the month.
- Since cloud resources can scale up or down throughout the day, using averages helps mitigate short-term usage spikes and reduces the risk of overage charges.

Data Center and Endpoint resources are metered daily. Their monthly average is calculated by summing the daily counts for the month and dividing them by the number of days.

If your Illumio Segmentation implementation is on-premises, you will be asked to provide usage information to Illumio on a quarterly basis at a minimum.

Converting Resource Count to Segmentation Workloads

After you have determined the number of resources in your environment, Segmentation for the Cloud Workloads (SWL) are calculated using these conversion ratios:

Use the workload calculator on Illumio's website or find it listed in the Usage menu (Usage > Workload Calculator > Illumio Segmentation tab) to convert your resource count into Insights Workloads.

Cloud	Illumio Resource types	Conversion Ratio
	Cloud Virtual Machine	1 Resource = 1 SWL
	Cloud Container	1 Resource = 2 SWLs
	Cloud Database	1 Resource = 1 SWL
	Cloud Storage	10 Resources = 1 SWL
	Cloud Firewall	1 Firewall = 100 SWLs
	Serverless Function	50 Functions = 1 SWL
	Network Routing and Resource Management	1 Resource = 10 SWLs
	Network Security and Management	10 Resources = 1 IWL
	Additional Data Processing *	50 MB = 1 SWL
* For Cloud resources, each Segmentation Workload (SWL) includes 50MB of data processing per day. Exceeding this allowance is rare for typical workloads. If your usage exceeds the allowance, you can purchase additional SWLs for additional data processing.		
Data Center	**Illumio Resource types**	**Conversion Ratio**
	Standard Server OS	1 OS = 1 SWL
	Legacy OS	1 OS = 2 SWLs
	Oracle Exadata Database	1 DATABASE = 10 SWLs
	Container Host	1 HOST = 6 SWLs
	Switch Port	1 MIP = 1 SWL
	Load Balancer	1 VIP = 1 SWL
	IBM i-series LPAR	1 LPAR = 10 SWLs
	IBM zLinux IFL	1 IFL = 15 SWLs
	IBM zOS LPAR	1 LPAR = 20 SWLs
	Additional Data Traffic Retention**	50 MB = 1 SWL
** For Data Center and Endpoint resources, each Segmentation Workload (SWL) includes 50MB data traffic retention. You can purchase additional data traffic.		
Endpoint	**Illumio Resource types**	**Conversion Ratio**
	Endpoint	5 ENDPOINTS = 1 SWL
	Additional Data Traffic Retention**	50 MB = 1 SWL

Sample Calculation for Segmentation

Consider an environment with the following number of cloud resources on a monthly average:

Cloud Resources	
Cloud Virtual Machine	100 virtual machines
Cloud Container	50 hosts
Cloud Database	20 databases
Cloud Storage	10 storage
Cloud Firewall	1 firewall
Serverless Function	50 functions
Network Security and Management	5 resources
Data Center Resources	
Standard Server OS	1000 OS
Container Host	50 hosts
Load Balancer	20 VIPs

Use the Illumio Workload Calculator to determine the number of Segmentation Workloads you need.

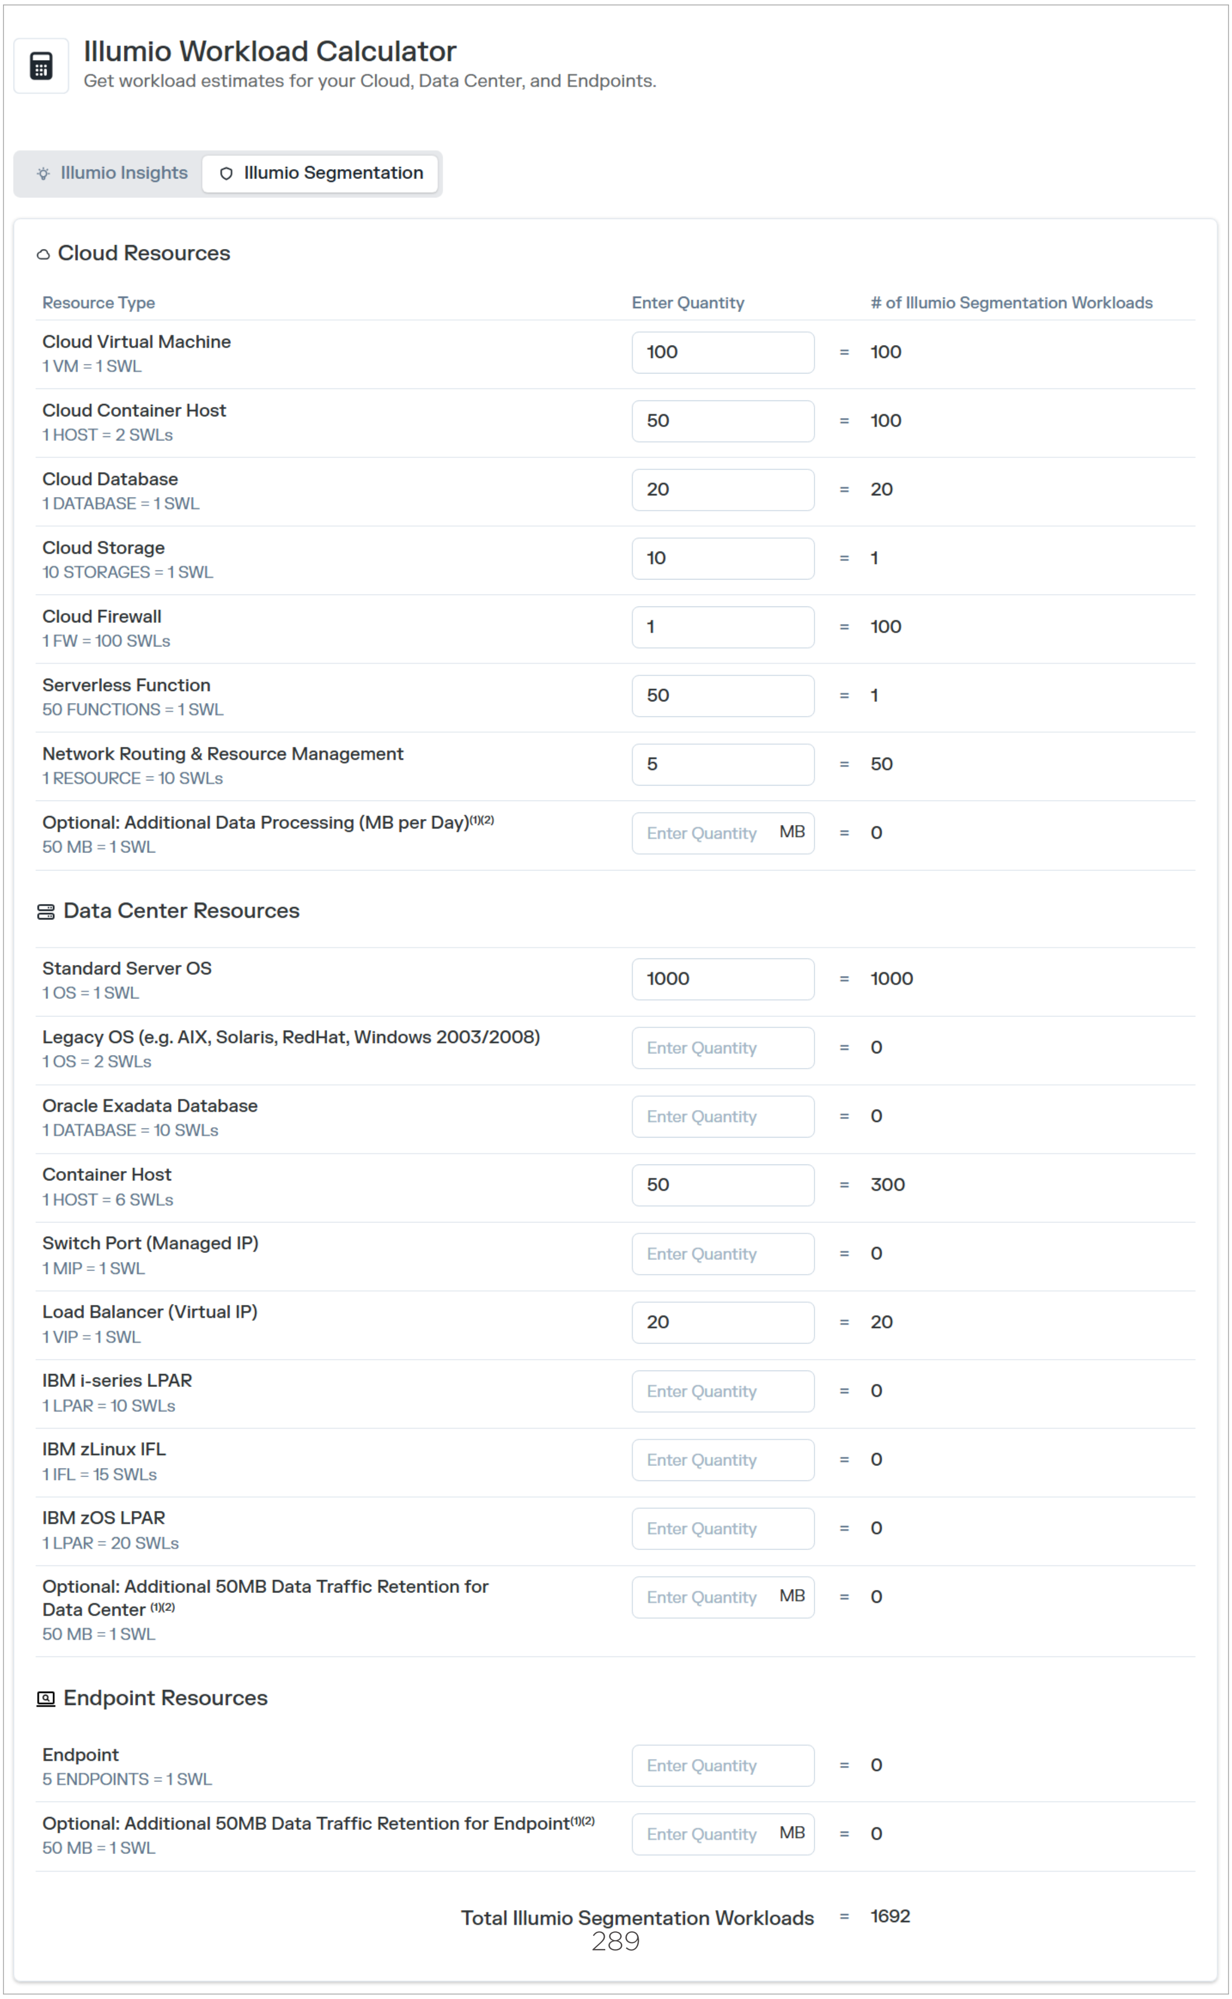

Illumio Workload Calculator

Get workload estimates for your Cloud, Data Center, and Endpoints.

Cloud Resources

Resource Type	Enter Quantity		# of Illumio Segmentation Workloads
Cloud Virtual Machine 1 VM = 1 SWL	100	=	100
Cloud Container Host 1 HOST = 2 SWLs	50	=	100
Cloud Database 1 DATABASE = 1 SWL	20	=	20
Cloud Storage 10 STORAGES = 1 SWL	10	=	1
Cloud Firewall 1 FW = 100 SWLs	1	=	100
Serverless Function 50 FUNCTIONS = 1 SWL	50	=	1
Network Routing & Resource Management 1 RESOURCE = 10 SWLs	5	=	50
Optional: Additional Data Processing (MB per Day)(1)(2) 50 MB = 1 SWL	Enter Quantity MB	=	0

Data Center Resources

Standard Server OS 1 OS = 1 SWL	1000	=	1000
Legacy OS (e.g. AIX, Solaris, RedHat, Windows 2003/2008) 1 OS = 2 SWLs	Enter Quantity	=	0
Oracle Exadata Database 1 DATABASE = 10 SWLs	Enter Quantity	=	0
Container Host 1 HOST = 6 SWLs	50	=	300
Switch Port (Managed IP) 1 MIP = 1 SWL	Enter Quantity	=	0
Load Balancer (Virtual IP) 1 VIP = 1 SWL	20	=	20
IBM i-series LPAR 1 LPAR = 10 SWLs	Enter Quantity	=	0
IBM zLinux IFL 1 IFL = 15 SWLs	Enter Quantity	=	0
IBM zOS LPAR 1 LPAR = 20 SWLs	Enter Quantity	=	0
Optional: Additional 50MB Data Traffic Retention for Data Center (1)(2) 50 MB = 1 SWL	Enter Quantity MB	=	0

Endpoint Resources

Endpoint 5 ENDPOINTS = 1 SWL	Enter Quantity	=	0
Optional: Additional 50MB Data Traffic Retention for Endpoint(1)(2) 50 MB = 1 SWL	Enter Quantity MB	=	0

Total Illumio Segmentation Workloads = 1692

Resource Coverage Expansion

Illumio continuously enhances its security coverage by adding support for additional resource types. When new resources are added to capabilities, they will be automatically included in your security coverage, which may increase your usage metrics and associated billing.

TIP

As a best practice, Illumio recommends that you review your usage dashboard regularly to manage your consumption effectively. Contact your Illumio account team with any questions about usage trends or billing impacts.

Usage

Learn how to view your organization's workload usage details and model usage of your workload resources.

View your total workload usage and license information in the **Usage** page. View your current workload usage statistics and the number of licenses you are currently entitled to. Use the workload calculator to estimate the additional workloads that might be required for your organization's growth.

On the Usage page, select the Illumio product you want to view: Illumio Insights or Illumio Segmentation. After you select an Illumio product, you'll see a summary of the average billable workloads used over the last 30 days out of the number of licenses currently active.

Next, view your organization’s license entitlement:

- License Name
- Status (Active, Future, or Expired)
- Start and End Date
- Number of billable workloads you are entitled to

Usage Details

Drill into details of your usage by clicking the “Usage Details” button below the Summary.

The **Usage Details** page displays graphs of daily usage based on the time range selected.

- **Total Workloads (Billable):** The value next to the chart title represents the daily workload usage averaging across the time period you select. Move your cursor to a specific date on the chart to see the average number of workloads used on that date.
- **Total Flow Data Processed:** The value next to the chart title represents the daily flow data processed averaging across the time period you select. Moving your cursor to a specific date on the chart gives you the average flow data processed on that date.

- **Resource Count:** The remaining charts show the raw count for each resource type. These charts are automatically sorted by the resource type with highest workload consumption.
  The value next to the chart title represents the daily resource count averaging across the time period you select. When you move your cursor to a specific date on the chart, you will see the average number of resources on that date.
  View the **Illumio Segmentation for Cloud** usage details.

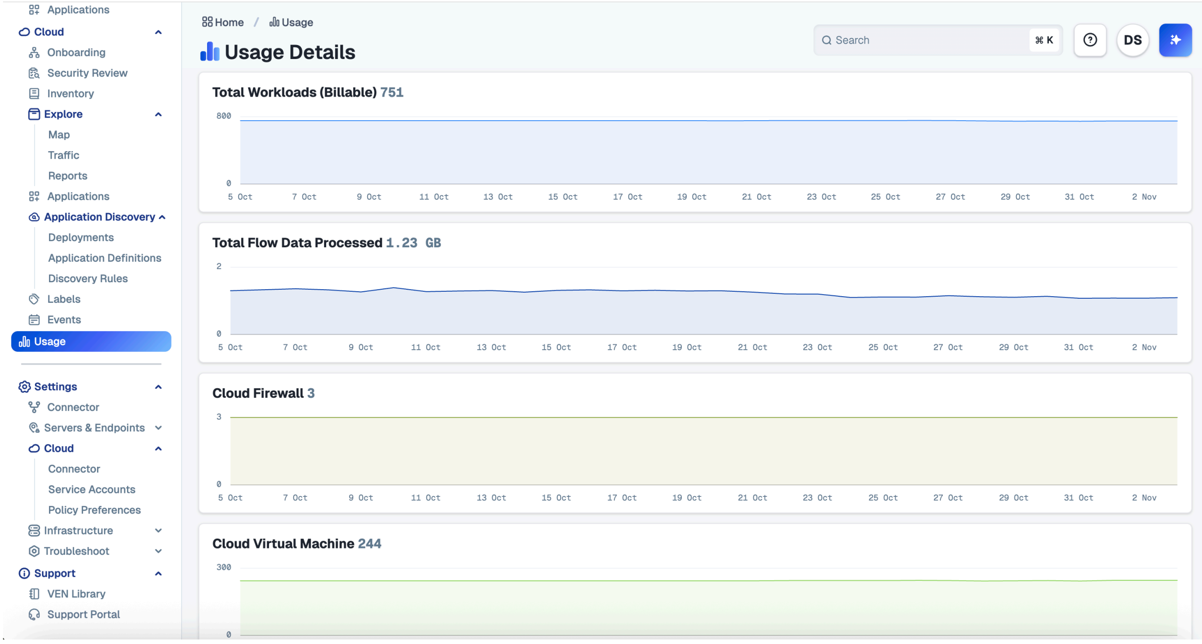
Applications
Cloud
Onboarding
Security Review
Inventory
Explore
Map
Traffic
Reports
Applications
Application Discovery
Deployments
Application Definitions
Discovery Rules
Labels
Events
Usage
Settings
Connector
Servers & Endpoints
Cloud
Connector
Service Accounts
Policy Preferences
Infrastructure
Troubleshoot
Support
VEN Library
Support Portal
Home / Usage
Usage Details
Search ⌘ K
DS
Total Workloads (Billable) 751
800
0
5 Oct 7 Oct 9 Oct 11 Oct 13 Oct 15 Oct 17 Oct 19 Oct 21 Oct 23 Oct 25 Oct 27 Oct 29 Oct 31 Oct 2 Nov
Total Flow Data Processed 1.23 GB
2
0
5 Oct 7 Oct 9 Oct 11 Oct 13 Oct 15 Oct 17 Oct 19 Oct 21 Oct 23 Oct 25 Oct 27 Oct 29 Oct 31 Oct 2 Nov
Cloud Firewall 3
3
0
5 Oct 7 Oct 9 Oct 11 Oct 13 Oct 15 Oct 17 Oct 19 Oct 21 Oct 23 Oct 25 Oct 27 Oct 29 Oct 31 Oct 2 Nov
Cloud Virtual Machine 244
300
0

For a list of Resource Types and their definitions, see the License and Usage section.

Reference

The content in this category explains the requirements for granting Illumio Segmentation for the Cloud access permissions to your cloud environment.

AWS flow log access IP addresses

AWS data planes, regions, and IPs

Illumio Segmentation for the Cloud uses TCP port 443 to access your flow logs, so open that port for the IP addresses listed in this section.

> **NOTE**
> You don't need to whitelist these IPs under most circumstances. However, there may be conditions where a service control policy or something similar denies IP access to AWS S3 buckets outside certain ranges. If this is the case, try whitelisting the IPs for the data planes listed.

Illumio Control Plane (For all AWS Regions)

The Illumio Segmentation for the Cloud control and data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound/outbound allowed list:

- 35.167.22.34
- 52.88.124.247
- 52.88.88.252

Illumio US West Data Plane for AWS

The Illumio Segmentation for the Cloud US West data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound allowed list for the AWS regions listed following.

- 35.163.224.94
- 44.226.137.227
- 54.190.103.0

AWS Regions Requiring above IPs for US West Data Plane

- af-south-1
- ca-central-1
- ca-west-1

- sa-east-1
- us-east-1
- us-east-2
- us-gov-east-1
- us-gov-west-1
- us-west-1
- us-west-2

Illumio EU West (UK) Data Plane for AWS

The Illumio EU West (UK) data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound allowed list:

- 18.169.5.9
- 13.41.233.77
- 18.169.6.17

AWS Regions Requiring above IPs for EU West (UK) Data Plane

- eu-central-1
- eu-central-2
- eu-north-1
- eu-south-1
- eu-south-2
- eu-west-1
- eu-west-2
- eu-west-3

Illumio APAC Data Plane for AWS

The Illumio APAC data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound allowed list for the AWS regions listed following:

- 13.54.140.138/32
- 52.63.108.169/32
- 52.64.120.98/32

AWS Regions Requiring above IPs for APAC Data Plane

- ap-east-1
- ap-northeast-1
- ap-northeast-2
- ap-northeast-3

- ap-south-1
- ap-south-2
- ap-southeast-1
- ap-southeast-2
- ap-southeast-3
- ap-southeast-4
- ap-southeast-5
- ap-southeast-6
- ap-southeast-7
- cn-north-1
- cn-northwest-1

Illumio Middle East Data Plane for AWS

The Illumio Middle East data plane uses the following public IP addresses to reach customer networks; add them to your firewall inbound allowed list for the AWS regions listed:

- 40.172.14.68
- 3.28.82.254
- 51.112.162.113

AWS Regions Requiring above IPs for Middle East Data Plane

- il-central-1
- me-central-1
- me-south-1

AWS Resource Type Permissions

When you add permissions to the Illumio AssumeRole as described in Permissions for Onboarding AWS, you will need to ensure that your permissions for resource types match those seen in the .json file example. Make sure to frequently check with your Support or Customer Success teams to make sure you have updated your .json file. Each time Illumio Segmentation for the Cloud supports a new resource type, you will need to update your resource permissions.

Permissions for Onboarding AWS

This page describes the set of required permissions that are created when onboarding AWS as described in Onboard an AWS Cloud account [113] and Onboard an AWS Cloud organization [107].

AWS IAM Permissions

To onboard your AWS account, you will need to use the CloudFormation Stack to create an IAM role within your AWS account, which Illumio assumes to make API calls. This role must be granted permissions to specific AWS resources for Cloud to provide visibility and manage policies for those resources. It is important to note that this 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.

Read and Write Permissions by Service and Category

Service	Category	Resource Types
Read and Write (IllumioCloudAWSProtectionPolicy)		
EC2	Network Security	DBSecurity Group, Network ACL, Security Group, Security Group Rule
RDS	Network Security	DB Security Group
Read (IllumioCloudAWSIntegrationPolicy)		
CodeDeploy	Infrastructure	Application, Deployment Group
DirectConnect	Network Routing	Connection, Gateway, Lag, Virtual Interface
DocumentDB	Database	Cluster
DynamoDB	Database	Table
EC2	Compute	Instance, Spot Fleet Request, Spot Instance Request
EC2	Network Management	EIP, Network Interface, Subnet, VPC, VPC Peering
EC2	Network Monitoring	Flow Log
EC2	Network Routing	Carrier Gateway, Customer Gateway, Egress Only Internet Gateway, Instance Connect Endpoint, Internet Gateway, Nat Gateway, Route Table, Transit Gateway, Transit Gateway Attachment, Transit Gateway Route Table, Transit Gateway Multicast Domain, VPC Endpoint, VPC Endpoint Service, VPN, VPN Connection, VPN Gateway
EC2	Network Security	Security Group
EC2	Storage	Volume
ECS	Containers	Cluster, Container Instance
EKS	Containers	Addon, Cluster, Fargate Profile, Node Group,
Elasticache	Database	Cache Cluster
ElasticLoadBalancingV2	Network Routing	Load Balancer
Glacier	Storage	Vault
Lambda Function	Serverless	Function
IAM	Account Management	Account, User
KMS	Security Infrastructure	Key
MemoryDB	Database	Cluster
Network Manager	Network Routing	Global Network, Core Network, Connect Attachment, VPC Attachment, Site To Site VPN Attachment, Transit Gateway Route Table Attachment, Transit Gateway Peering, Transit Gateway Registration

Service	Category	Resource Types
RAM	Resource Management	Resource Share
RDS	Database	DB Cluster, DB Instance, DBSecurityGroup
Redshift	Data warehouse	Cluster
S3	Storage	Bucket, Bucket Policy
Target Groups	Network Routing	Target Group

IAM Role Configuration

To facilitate access to your AWS environment, you must create an IAM role within your AWS account. This role must be assigned the following policies:

- **SecurityAudit (managed by AWS) and IllumioCloudAWSIntegrationPolicy:** Permissions in these policies are required to read the resources in your AWS account.
- **IllumioCloudAWSProtectionPolicy:** Permissions in this policy are required 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:

```
READ ONLY Policy

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*'
            - 'docdb-elastic:GetCluster'
            - 'docdb-elastic:ListTagsForResource'
            - 'dynamodb:List*'
            - 'dynamodb:Describe*'
            - 'ec2:Describe*'
            - 'ec2:SearchTransitGatewayMulticastGroups'
            - 'ecs:Describe*'
            - 'ecs:List*'
        - 'eks:DescribeAddon'
            - 'eks':ListAddons'
            - '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'
            - 'networkmanager:ListCoreNetworks'
            - 'networkmanager:GetCoreNetwork'
            - 'networkmanager:ListAttachments'
            - 'networkmanager:GetVpcAttachment'
            - 'networkmanager:GetSiteToSiteVpnAttachment'
            - 'networkmanager:GetConnectAttachment'
            - 'networkmanager:GetTransitGatewayRouteTableAttachment'
            - 'networkmanager:ListPeerings'
            - 'networkmanager:GetTransitGatewayPeering'
            - 'networkmanager:GetTransitGatewayRegistrations'
```

Write Policy

The following items are AWS IAM write permissions that you will need to grant to the Illumio AssumeRole.

```
READ ONLY Policy

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*'
            - 'docdb-elastic:GetCluster'
            - 'docdb-elastic:ListTagsForResource'
            - 'dynamodb:List*'
            - 'dynamodb:Describe*'
            - 'ec2:Describe*'
            - 'ec2:SearchTransitGatewayMulticastGroups'
            - 'ecs:Describe*'
            - 'ecs:List*'
            - 'eks:DescribeAddon'
            - 'eks':ListAddons'
            - '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'
            - 'networkmanager:ListCoreNetworks'
            - 'networkmanager:GetCoreNetwork'
            - 'networkmanager:ListAttachments'
            - 'networkmanager:GetVpcAttachment'
            - 'networkmanager:GetSiteToSiteVpnAttachment'
            - 'networkmanager:GetConnectAttachment'
            - 'networkmanager:GetTransitGatewayRouteTableAttachment'
            - 'networkmanager:ListPeerings'
            - 'networkmanager:GetTransitGatewayPeering'
            - 'networkmanager:GetTransitGatewayRegistrations'

WRITE Policy
- PolicyName: IllumioCloudAWSProtectionPolicy
  PolicyDocument:
    Version: 2012-10-17
    Statement:
      - Effect: Allow
        Resource:
          - 'arn:aws:ec2:*:*:security-group-rule/*'
          - 'arn:aws:ec2:*:*:security-group/*'
          - 'arn:aws:ec2:*:*:network-acl/*'
        Action:
          - 'ec2:AuthorizeSecurityGroupIngress'
          - 'ec2:RevokeSecurityGroupIngress'
          - 'ec2:UpdateSecurityGroupRuleDescriptionsIngress'
          - 'ec2:AuthorizeSecurityGroupEgress'
          - 'ec2:RevokeSecurityGroupEgress'
          - 'ec2:UpdateSecurityGroupRuleDescriptionsEgress'
          - 'ec2:ModifySecurityGroupRules'
          - 'ec2:DescribeTags'
          - 'ec2:CreateTags'
          - 'ec2:DeleteTags'
          - 'ec2:DescribeNetworkAcls'
          - 'ec2:CreateNetworkAclEntry'
          - 'ec2:ReplaceNetworkAclEntry'
          - 'ec2:DeleteNetworkAclEntry'
          - 'ec2:ModifyNetworkInterfaceAttribute'
          - 'ec2:CreateSecurityGroup'
          - 'ec2:DeleteSecurityGroup'
          - 'ec2:DescribeSecurityGroups'
```

AWS Resource Type Permissions

When you add permissions to the Illumio AssumeRole as described in Permissions for On-boarding AWS, you will need to ensure that your permissions for resource types match those seen in the .json file example. Make sure to frequently check with your Support or Customer Success teams to make sure you have updated your .json file. Each time Illumio Segmentation for the Cloud supports a new resource type, you will need to update your resource permissions.

FLOW READ Policy

```
's3:ListBucket'

's3:ListBucketVersion'

's3:GetBucketLocation'

's3:GetObject'
```

Service Accounts and IAM Roles for AWS

The following information is important to understanding how Illumio interacts with AWS.

Service Accounts in the Illumio Segmentation for the Cloud Context

Within the Illumio Segmentation for the Cloud platform, a "service account" refers to an account used by Illumio Segmentation for the Cloud to interact with its own services (Illumio Segmentation for the Cloud services) rather than directly with your AWS services. This account is primarily used for internal operations within Illumio Segmentation for the Cloud, such as making API calls to the Illumio Segmentation for the Cloud 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, Illumio Segmentation for the Cloud requires the creation of an identification and access management (IAM) role within the customer's AWS account. Illumio Segmentation for the Cloud 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.

Handling encrypted AWS VPC flow logs

If service-side encryption with KMS (SSE-KMS) keys is enabled for the S3 bucket, Cloud requires additional permissions for the log service to be added to the KMS key before enabling flow logs.

To allow the log service to write VPC Flow Logs in the designated S3 bucket, the AWS Logs Delivery System must be granted permission to the Encrypt, Decrypt, ReEncrypt, GenerateDataKey*, and Describe key on the key that is used to encrypt the data in the S3 bucket. Below is an example policy showing the necessary permissions in place for the key policy.

```
{

"Sid": "Allow Log Delivery to use the key",

"Effect": "Allow",

"Principal": {

"Service": "delivery.logs.amazonaws.com"

},

"Action":

"kms:Encrypt",

"kms:Decrypt",

"kms:ReEncrypt*",

"kms:GenerateDataKey*",

"kms:DescribeKey"

],

 "Resource": "*"

"Condition": {

"StringEquals": {

"aws:SourceAccount": "<account-id>"

},

"ArnLike": {

"aws:SourceArn": "arn:aws:logs:<region>:<account-id>:*"

}

}

}
```

To read flows stored in encrypted buckets, the Assume Role requires access to the key used for encrypting the contents of the S3 bucket. This key decrypts the contents of the S3 bucket. The following is the policy document required to gain access to the key and decrypt the flow logs. Adding this permission automatically allows the Assume Role, created during on-boarding, to decrypt the contents of the bucket (In this case, the flow logs). No additional settings are required.

```
{
"Version": "2012-10-17",
"Statement":[
{
"Effect": "Allow",
"Action":
"kms:Decrypt"
],
"Resource": [
"arn:aws:kms:<region>:<account-id>:key/<key-id>" // Replace with your KMS
key ARN
]
}
]
}
```

The following CloudFormation Template gets the Assume Role ARN and the KMS Key ARN as input and grants the decrypt permission on the KMS Key to the Assume Role.

```
AWSTemplateFormatVersion: "2010-09-09"

Description: "Grant Decrypt permission on KMS key for CloudSecure's Assume Role"

Parameters:

    IAMRoleName:

        Type: String

        Description: IAM Role name used by Cloud.

    KMSKeyARNs:

        Type: CommaDelimitedList

        Description: List of KMS Key ARNs.

Resources:

      IllumioKMSDecryptPolicy:

        Type: 'AWS::IAM::Policy'

        Properties:

            PolicyName: IllumioKMSDecrypt

            PolicyDocument:

                Version: 2012-10-17

                Statement:

                    - Effect: Allow

                    Sid: IllumioKMSKeyAccess

                    Action:

                    - 'kms:Decrypt'

                    Resource: !Ref KMSKeyARNs

        Roles

            - !Ref IAMRoleName
```

For more information, see the AWS website.

AWS flow logs

For a list of ports and IP addresses required for flow log access, see AWS Flow Log Access IP Addresses.

Supported Flow Log Fields

Illumio Segmentation for the Cloud 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} ${srcport} ${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 Illumio Segmentation for the Cloud fetches the flow logs depends on your configuration (e.g., a central account or multiple accounts)

- When flow log access is first enabled in Illumio Segmentation for the Cloud, there's a 15-minute latency until traffic flows are first displayed in Cloud Map, Traffic and Inventory pages. After this initial latency, traffic flows are periodically updated every two minutes.

Updating AWS permissions on the Assume Role

Illumio updates permissions required for the Assume Role on a continuous basis. Use these steps to provide permissions for the newly added resources.

1. Download the permissions that are provided in the first part of the wizard. Depending on whether you chose read-only or read and write, be sure to download the correct file below.
   - Read and write
   - Read-only
2. Run the CloudFormation Stack (CFT).
3. Login to the AWS console of account to which you need to update the permissions to run the CloudFormation stack.
4. Under services click **CloudFormation**.
5. Click **Create stack**.
6. In the Choose template page, select template ready and upload a template file option, and upload the downloaded template and click **Next**.

7. 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.
8. In the IAMRoleName box, enter the name of the assume role created in AWS when onboarding with Illumio Segmentation for the Cloud. By default, the name is IllumioCloudIntegrationRole. Click **Next**.
9. If you gave 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.)
10. 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.

AWS permissions background

When you start the onboarding process and begin creating IAM roles from the Illumio Segmentation for the Cloud 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 Illumio Segmentation for the Cloud

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 Illumio from synchronizing resources.

AWS 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 Flow Log Access IP Addresses

Illumio Segmentation for the Cloud uses TCP port 443 to access your flow logs, so open that port for the IP addresses listed in this section.

Illumio Control Plane (for all Azure Regions)

The Illumio Segmentation for the Cloud control and data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound/outbound allowed list:

- 35.167.22.34
- 52.88.124.247
- 52.88.88.252

Illumio US East Data Plane for Azure

The Illumio Segmentation for the Cloud US East data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound allowed list for the Azure regions listed following:

- 13.68.238.145
- 13.68.232.36
- 13.68.236.178

Azure Regions Requiring above IPs for US East Data Plane

- southeastasia
- centralus
- southafricanorth
- centralindia
- eastasia
- japaneast
- koreacentral
- canadacentral
- uaenorth
- brazilsouth
- centraluseuap
- eastus2euap
- qatarcentral
- centralusstage
- eastusstage
- eastus2stage
- northcentralusstage
- southcentralusstage
- westusstage
- westus2stage
- asia asiapacific
- australiaeast
- canada
- global
- india
- japan
- korea
- singapore
- southafrica
- uae
- unitedstates
- unitedstateseuap
- eastasiastage
- southeastasiastage

- brazilus
- eastusstg
- northcentralus
- westus
- jioindiawest
- devfabric
- westcentralus
- southafricawest
- australiacentral
- australiacentral2
- australiasoutheast
- japanwest
- jioindiacentral
- koreasouth
- southindia
- westindia
- canadaeast
- uaecentral
- brazilsoutheast
- chinaeast
- chinaeast2
- chinaeast3
- chinanorth
- chinanorth2
- chinanorth3

Illumio EU West (Germany) Data Plane for Azure

The Illumio Segmentation for the Cloud EU West (Germany) data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound allowed list for the Azure regions listed following:

- 4.185.170.43
- 4.185.170.168
- 4.185.170.165

Azure Regions Requiring above IPs for EU West (Germany) Data Plane

- germanynortheast
- polandcentral
- swedensouth
- eastus
- northeurope
- swedencentral
- uksouth
- westeurope
- francecentral
- germanywestcentral
- norwayeast
- switzerlandnorth

- europe
- france
- germany
- norway
- switzerland
- uk
- francesouth
- germanynorth
- norwaywest
- switzerlandwest
- ukwest
- italynorth
- germanycentral

Illumio US West 2 Data Plane for Azure

The Illumio Segmentation for the Cloud US West data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound allowed list for the Azure regions listed following:

- 52.183.73.40
- 52.148.136.248
- 20.191.118.116

Azure Regions Requiring above IPs for Illumio Cloud US West 2 Data Plane for Azure

- eastus
- eastus2
- westus

Permissions for Onboarding Azure

This section describes the set of permissions that you grant to the Illumio Segmentation for the Cloud App that is registered in Azure Active Directory.

These permissions are required, irrespective of whether you use the default method provided by the wizard or the guided method.

- If you are onboarding using the default method described in Onboard an Azure Cloud tenant - default setup [79] and Onboard an Azure Cloud subscription - default setup [65], it automatically provisions the permissions described here.
- If you are onboarding using the guided method, which does not involve running the PowerShell script provided in the wizard, or if you lack Owner access, described in Onboard an Azure Cloud Tenant - Guided Setup [84]) and Onboard an Azure Cloud Subscription - Guided Setup [67], you need to set these permissions via the Azure Console.

Azure permission descriptions

Permission Type	Permission Name	Notes
Read	Reader - role	This role gives Illumio Segmentation for the Cloud the permissions to read data or resources from your subscription or tenant. This role allows the viewing of all resources, but does not allow modification.
Write	Writer - role	This role gives Illumio Segmentation for the Cloud the permissions to modify data or resources in your subscription or tenant. This role allows the modification of resources.
NSG, Azure Firewall	Multiple, see below.	Use these permissions to create custom roles. Define any custom roles with elevated permissions, as part of the PowerShell script that is run when you onboard an Azure subscription. If the user onboarding Azure has Owner permissions, these permissions are automatically assigned to the "Illumio Network Security Administrator" custom role that is created when the onboarding PowerShell script is run. However, if the user onboarding Azure does *not* have Owner permissions, you must create the"Illumio Network Security Administrator" custom role with these NSG and Azure Firewall permissions *before* the onboarding PowerShell script is run.
Flow	Storage Blob Data Reader – role	

Azure read and write policy

When you grant read and write permissions to Illumio Segmentation for the Cloud, the following roles are created in the Azure tenant.

```
Reader Role - Built In Role
{
  "assignableScopes": [
    "/"
  ],
  "description": "View all resources, but does not allow you to make any
changes.",
  "id": "/providers/Microsoft.Authorization/roleDefinitions/
acdd72a7-3385-48ef-bd42-f606fba81ae7",
  "name": "acdd72a7-3385-48ef-bd42-f606fba81ae7",
  "permissions": [
    {
      "actions": [
        "*/read"
      ],
      "notActions": [],
      "dataActions": [],
      "notDataActions": []
    }
  ],
  "roleName": "Reader",
  "roleType": "BuiltInRole",
  "type": "Microsoft.Authorization/roleDefinitions"
}
Illumio Network Security Administrator Role - Custom Role
{
    "properties": {
        "roleName": "Illumio Network Security Administrator",
        "description": "Illumio Network Administration Role",
        "assignableScopes": [
            "/"
        ],
        "permissions": [
            {
                "actions": [
                    "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.Network/networksecuritygroups/providers/
Microsoft.Insights/diagnosticSettings/read",
                    "Microsoft.Network/networksecuritygroups/providers/
Microsoft.Insights/diagnosticSettings/write",
                    "Microsoft.Network/networksecuritygroups/providers/
Microsoft.Insights/logDefinitions/read",
                    "Microsoft.Network/networkWatchers/securityGroupView/
action",
                    "Microsoft.Network/networkSecurityGroups/*",
                    "Microsoft.Network/networkInterfaces/read",
                    "Microsoft.Network/networkInterfaces/write",
                    "Microsoft.Network/virtualNetworks/read",
                    "Microsoft.Network/virtualNetworks/subnets/write",
                    "Microsoft.Authorization/locks/*",
                    "Microsoft.Compute/virtualMachines/read"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
}
Illumio Firewall Administrator Role - Custom Role
{
    "properties": {
        "roleName": "Illumio Firewall Administrator",
        "description": "Illumio Firewall Administrator role",
        "assignableScopes": [
            "/"
        ],
        "permissions": [
            {
                "actions": [
                    "Microsoft.Network/azurefirewalls/read",
                    "Microsoft.Network/azurefirewalls/learnedIPPrefixes/
action",
                    "Microsoft.Network/azureFirewalls/
applicationRuleCollections/write",
                    "Microsoft.Network/azureFirewalls/
applicationRuleCollections/delete",
                    "Microsoft.Network/azureFirewalls/
applicationRuleCollections/read",
                    "Microsoft.Network/azurefirewalls/providers/
Microsoft.Insights/logDefinitions/read",
                    "Microsoft.Network/azureFirewalls/natRuleCollections/
write",
                    "Microsoft.Network/azureFirewalls/natRuleCollections/
read",
                    "Microsoft.Network/azureFirewalls/natRuleCollections/
delete",
                    "Microsoft.Network/azureFirewalls/
networkRuleCollections/read",
                    "Microsoft.Network/azureFirewalls/
networkRuleCollections/write",
                    "Microsoft.Network/azureFirewalls/
networkRuleCollections/delete",
                    "Microsoft.Network/azureFirewallFqdnTags/read",
                    "Microsoft.Network/azurefirewalls/providers/
Microsoft.Insights/metricDefinitions/read",
                    "Microsoft.Network/firewallPolicies/read",
                    "Microsoft.Network/firewallPolicies/write",
                    "Microsoft.Network/firewallPolicies/join/action",
                    "Microsoft.Network/firewallPolicies/certificates/
action",
                    "Microsoft.Network/firewallPolicies/delete",
                    "Microsoft.Network/firewallPolicies/
ruleCollectionGroups/read",
                    "Microsoft.Network/firewallPolicies/
ruleCollectionGroups/write",
                    "Microsoft.Network/firewallPolicies/
ruleCollectionGroups/delete",
                    "Microsoft.Network/firewallPolicies/ruleGroups/read",
                    "Microsoft.Network/firewallPolicies/ruleGroups/write",
                    "Microsoft.Network/firewallPolicies/ruleGroups/delete",
                    "Microsoft.Network/ipGroups/read",
                    "Microsoft.Network/ipGroups/write",
                    "Microsoft.Network/ipGroups/validate/action",
                    "Microsoft.Network/ipGroups/updateReferences/action",
                    "Microsoft.Network/ipGroups/join/action",
                    "Microsoft.Network/ipGroups/delete"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
}
```

Azure flow log support

Illumio Segmentation for the Cloud supports NSG Flow logs version 2 (includes flow state and byte counts), but does not support version 1. It also supports VNet flow logs and Azure Firewall flow logs.

See Set up flow logs in your CSP environment.

Update Service Principals for Onboarded Azure Subscriptions and Tenants

This topic describes how to update the service principal used by Illumio Segmentation for the Cloud for accessing customer Azure Resources and how to rotate the secret for an existing client when the secret expires.

NOTE

Use Case: You onboarded an Azure tenant a long time ago, and the service principal is about to expire. You can update the existing Service Principal, which is used by Illumio Segmentation for the Cloud, with a new Service Principal by using the PowerShell script provided by Illumio Segmentation for the Cloud during the onboarding process. Alternatively, you can create it on your own and make use of the custom PowerShell script to send the credentials back to Cloud. You can also run a command to create a new secret if your secret expires after one year.

Create a New Service Principal with the Cloud Onboarding Script

You can use the onboarding PowerShell script provided by Illumio Segmentation for the Cloud to create a new Service Principal and update it in Cloud. Use the following commands, depending on your needs:

Azure Subscription Onboarding

If you want to use Read-only Mode, use this command:

```
Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -sid <azure_subscription_id> -serviceAccountKey <service_account_key> -serviceAccountToken
<service_account_token> -csTenantId <cs_tenant_id>; -url https://cloud.illum.io
```

If you want to use Read/Write Mode, use this command:

```
Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -sid <azure_subscription_id> -serviceAccountKey <service_account_key> -serviceAccountToken
<service_account_token> -csTenantId <cs_tenant_id> -url https://cloud.illum.io -nsg
```

Azure Tenant Onboarding

If you want to use Read-only Mode, use this command:

```
Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -tid <azure_tenant_id>
-serviceAccountKey <service_account_key> -serviceAccountToken <service_account_token> -csTenantId <cs_tenant_id>-url https://cloud.illum.io
```

If you want to use Read/Write Mode, use this command:

```
Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-
west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -tid <azure_tenant_id>;
-serviceAccountKey <service_account_key> -serviceAccountToken <service_ac-
count_token> -csTenantId <cs_tenant_id> -url https://cloud.illum.io -nsg
```

Rotate Secrets for an Existing Service Principal

The secrets for Azure Service Principal are set with an expiry of 365 days (1 year) when created using Cloud. After the expiry, use the following commands to create a new secrets and update them in Cloud.

Subscription Onboarding

```
Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-
west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -sid <subscription_id>
-serviceAccountKey <service_account_key> -serviceAccountToken <service_ac-
count_token> -csTenantId <cloudsecure_tenant_id>; -clientId <service_princi-
pal_client_id> -url https://cloud.illum.io -rotateSecret
```

To obtain the <service_principal_client_id>, from within the Azure Portal, do the following:

1. Go to Microsoft Entra ID and select **App registrations** from the menu.
2. Select the app registration that starts with `Illumio-CloudSecure-Access` and copy the client Id.

Tenant Onboarding

```
Invoke-WebRequest -Uri https://cloudsecure-onboarding-templates.s3.us-
west-2.amazonaws.com/cloudsecure/illumio-init.ps1 -OutFile (Join-Path
$PWD.Path "illumio-init.ps1"); ./illumio-init.ps1 -tid <azure_tenant_id>;
-serviceAccountKey <service_account_key> -serviceAccountToken <service_ac-
count_token> -csTenantId <cloudsecure_tenant_id> -clientId <service_princi-
pal_client_id> -url https://cloud.illum.io -rotateSecret
```

Once you have created the new secrets, then you can follow the steps in Send Secrets Back to the Cloud [317].

Send Secrets Back to Cloud

If you have a new service principal with the required permissions, based on the type of onboarding, you can run the following PowerShell script to send the secrets back to Cloud. Copy and save the contents to a file with the .ps1 extension, e.g., "web_request.ps1" or something similar.

Before running the script the following information in the file should be changed to your actual values:

- <YourServiceAccountKeyId> - Cloud's service account key id. Service account can be created under Settings

- <YourServiceAccountToken> - Token of the service account being used
- <Your ClientSecret> - New Service Principals secret
- <CloudSecureTenantId> - Customer's Illumio Segmentation for the Cloud Tenantid
- <ClientId> - New Service Principal's client id
- <SubscriptionId> - Azure subscription Id. This is required only for subscription onboarding. If the onboarding type is an Azure tenant, remove the entire line.
- <AzureTenantId> - Azure Tenant Id of the customer

PowerShell Script

```
# Set your service account key ID, token, and client secret
$serviceAccountKeyId = "<YourServiceAccountKeyId>"
$serviceAccountToken = "<YourServiceAccountToken>"
$clientSecret = "<YourClientSecret>" # The actual client secret to be
encoded

# Combine the key ID and token with a colon and base64 encode for the
Authorization header
$authString = "$($serviceAccountKeyId):$($serviceAccountToken)"
$encodedAuthString =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($authString))

# Base64 encode the client secret separately
$encodedClientSecret =
[Convert]::ToBase64String([Text.Encoding]::ASCII.GetBytes($clientSecret))

# Construct the headers with the encoded Authorization header
$headers = @{
  "X-Tenant-Id"   = "<CloudsecureTenantId>"
  "Content-Type"  = "application/json"
  "Authorization" = "Basic $encodedAuthString"
}

# Construct the request body with the encoded client secret
$body = @{
  "type"             = "AzureRole"
  "client_id"        = "<ClientId>"
  "client_secret"    = $encodedClientSecret  # Use the base64 encoded client
secret
  "subscription_id" = "<SubscriptionId>" # remove this and use
azure_tenant_id if onboarding the entire tenant.
  "azure_tenant_id" = "<AzureTenantId>" # both azure tenant id and
subscription_id should be present for subscription onboarding.
} | ConvertTo-Json -Depth 10

# Send the POST request
$response = Invoke-RestMethod -Uri 'https://cloud.illum.io/api/v1/
integrations/cloud_credentials' -Method Post -Headers $headers -Body $body


# Output the response
Write-Host $response
```

GCP Flow Log Access IP Addresses

Illumio Segmentation for the Cloud uses TCP port 443 to access your flow logs, so open that port for the IP addresses listed in this section.

Illumio Control Plane (For all GCP Regions)

The Illumio Segmentation for the Cloud control pane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound/outbound allowed list:

- 35.167.22.34
- 52.88.124.247
- 52.88.88.252

Illumio US Region Data Plane for GCP

The Illumio Segmentation for the Cloud US Region data plane uses the following public IP address to reach customer networks. Add this IP to your firewall inbound allow list for the GCP regions listed below.

- IP Address
  - 57.154.51.200
- GCP Regions requiring this IP for US Region Data Plane:
  - us-east1
  - africa-south1
  - northamerica-northeast1
  - northamerica-northeast2
  - northamerica-south1
  - southamerica-east1
  - southamerica-west1
  - us-central1
  - us-east4
  - us-east5
  - us-south1
  - us-west1
  - us-west2
  - us-west3
  - us-west4

Illumio Australia East Region Data Plane for GCP

The Illumio Segmentation for the Cloud Australia East data plane uses the following public IP address to reach customer networks. Add this IP to your firewall inbound allow list for the GCP regions listed below.

- IP Address:
  - 4.197.250.30

- GCP Regions requiring this IP for East Data Plane:
  - asia-south2
  - asia-south1
  - asia-northeast3
  - asia-northeast2
  - asia-northeast1
  - australia-southeast1
  - australia-southeast2
  - asia-east2
  - asia-east1
  - asia-southeast1
  - asia-southeast2

Illumio North Europe Region Data Plane for GCP

The Illumio Segmentation for the Cloud North Europe data plane uses the following public IP addresses to reach customer networks. Add these IPs to your firewall inbound allow list for the GCP regions listed below.

- IP Addresses:
  - 52.169.78.75
  - 52.169.75.70
  - 52.169.76.44
- GCP Regions requiring these IPs for North Europe Data Plane:
  - europe-west9
  - europe-west8
  - europe-west6
  - europe-west4
  - europe-west3
  - europe-west2
  - europe-west12
  - europe-west10
  - europe-west1
  - europe-southwest1
  - europe-north1
  - europe-central2
  - me-central1
  - me-central2
  - me-west1

Permissions for onboarding GCP

This section describes the set of permissions that you grant to Illumio Segmentation for the Cloud.

GCP permission descriptions

Permission Type	Permission Name	Notes
Read	roles/iam.securityReviewer (link lists all resources)	This role gives Illumio Segmentation for the Cloud the permissions to read data or resources from your organization or project. This role allows the viewing of all resources, but does not allow modification. This role can be truncated to the following minimum resource permissions. (Limiting permissions to cloud services currently supported by Illumio): • https://cloud.google.com/iam/docs/roles-permissions/batch#batch.jobsViewer • https://cloud.google.com/iam/docs/roles-permissions/logging#logging.viewer • https://cloud.google.com/iam/docs/roles-permissions/cloudsql#cloudsql.viewer • https://cloud.google.com/iam/docs/roles-permissions/pubsub#pubsub.viewer • https://cloud.google.com/iam/docs/roles-permissions/resourcemanager#resourcemanager.folderViewer • https://cloud.google.com/iam/docs/roles-permissions/resourcemanager#resourcemanager.organizationViewer • https://cloud.google.com/iam/docs/roles-permissions/resourcemanager#resourcemanager.tagAdmin • https://cloud.google.com/iam/docs/roles-permissions/resourcemanager#resourcemanager.tagViewer Although Illumio Segmentation for the Cloud does not currently support every resource listed in the full securityReviewer role, if you do not truncate the role, Illumio Segmentation for the Cloud assigns all the resource permissions so that you don't have to update permissions each time Illumio supports additional resources.
Read	roles/compute.viewer (link lists all resources)	This role gives Illumio Segmentation for the Cloud the permissions to read compute resources from your organization or project. This role allows the viewing of all resources, but does not allow modification. The full role is required, so truncation is not supported.
Read	roles/browser (link lists all resources)	Required for organization and folder onboarding. This role gives Illumio Segmentation for the Cloud the permissions to read organization projects and folders. The full role is required, so truncation is not supported.
Read	roles/cloudasset.viewer (link lists all resources)	This role gives Illumio Segmentation for the Cloud the permissions to read cloud asset resources from your organization or project. This role allows the viewing of all resources, but does not allow modification. The full role is required, so truncation is not supported.

Permission Type	Permission Name	Notes
Read Flows	IllumioPubSubFlowLogAccess	This custom role gives Illumio Segmentation for the Cloud the permissions to read flows by creating and attaching subscriptions. The permission is also sought for offboarding flow access workflows. The full role is required, so truncation is not supported. This role assigns the following resource permissions: • pubsub.topics.attachSubscription • pubsub.subscriptions.consume • pubsub.subscriptions.create • pubsub.subscriptions.delete
Read Flows	IllumioPubSubAccessRole	This custom role gives Illumio Segmentation for the Cloud the permissions to read flows by creating and attaching subscriptions. The permission is also sought for offboarding flow access workflows. This role assigns the following resource permissions: • pubsub.topics.attachSubscription
Write	illumio_api_enable_role	This custom role gives Illumio Segmentation for the Cloud the permissions to enable APIs. The full role is required, so truncation is not supported. This role assigns the following resource permissions: • serviceusage.services.enable • serviceusage.services.list • serviceusage.services.get
Write	illumio_write_role	This custom role gives Illumio Segmentation for the Cloud the permissions to modify data or resources in your organization or project. This role allows the modification of resources. Enable read-write mode from in the onboarding wizard to create the role. The full role is required, so truncation is not supported. This role assigns the following resource permissions: • compute.firewalls.create • compute.networks.updatePolicy • compute.firewalls.update • compute.firewalls.get • compute.firewalls.delete

GPC permissions script example

When you grant read and write permissions to Illumio Segmentation for the Cloud, the gcp_onboarding_prod.sh script below creates roles and sets permissions. The script prompts you for permission to enable APIs. Enabling APIs is optional, but Illumio Segmentation for the Cloud functionality is affected if you don't enable APIs.

NOTE

If billing is not enabled for your projects, Illumio Segmentation for the Cloud cannot enable APIs.

```
#!/bin/bash

#Constants
RESOURCE_TYPE_PROJECT="project"
RESOURCE_TYPE_ORGANIZATION="organization"
RESOURCE_TYPE_FOLDER="folder"
DEFAULT_ROLE_NAME="illumio_role_$(date +%s)"
DEFAULT_SA_NAME="illumio-sa-$(date +%s)"
DEFAULT_SA_DISPLAY_NAME="Illumio Service Account"
DEFAULT_POST_URL="https://cloud.illum.io"
PREDEFINED_ROLES="roles/iam.securityReviewer,roles/compute.viewer,roles/
cloudasset.viewer"
DEFAULT_WRITE_ROLE_NAME="illumio_write_role_$(date +%s)"
WRITE_PERMISSIONS="compute.firewalls.create,compute.firewalls.delete,compute
.firewalls.get,compute.firewalls.update,compute.networks.updatePolicy"
DEFAULT_API_ENABLE_ROLE_NAME="illumio_api_enable_role_$(date +%s)"
API_ENABLE_PERMISSION="serviceusage.services.enable,serviceusage.services.li
st,serviceusage.services.get"

ILLUMIO_SA_EMAIL="illumio-onboarding@cs-prod-01.iam.gserviceaccount.com"

GREEN='\033[32m'
RESET='\033[0m'
RED='\033[31m'

# Resource tracking
created_resources__service_accounts=""
created_resources__roles=""
RESOURCE_TYPE="$RESOURCE_TYPE_PROJECT" # using default resource type as
object, based on if organization id is passed in arguments, it changes to
"$RESOURCE_TYPE_ORGANIZATION"

# Function to add a resource to the tracking list - these lists will be
used to clean up resources in case of an error
add_resource() {
    local resource_type=$1
    local resource_value=$2
    local var_name="created_resources__${resource_type}"
    if [[ -z "${!var_name}" ]]; then
        eval "$var_name=\"$resource_value\""
    else
        eval "$var_name=\"${!var_name},$resource_value\""
    fi
}

# Cleanup function for errors (triggered by trap)
cleanup() {
    if [ "$1" == "0" ]; then
        return
    fi
    echo -e "${RED}Cleaning up resources...${RESET}"

    # Delete service accounts
    IFS=',' read -ra SAS <<< "${created_resources__service_accounts}"
    for sa in "${SAS[@]}"; do
```

```
        echo -e "${RED}Deleting service account: $sa${RESET}"
        if ! gcloud iam service-accounts delete "$sa" --quiet --
project="$PROJECT_ID"; then
            echo -e "${RED}Failed to delete service account: $sa${RESET}"
        fi
    done

    # Delete roles
    IFS=',' read -ra ROLES <<< "${created_resources__roles}"
    for role in "${ROLES[@]}"; do
        echo -e "${RED}Deleting IAM role: $role${RESET}"
        if [ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
            if ! gcloud iam roles delete "$role" --
organization="$ORGANIZATION_ID" --quiet; then
                echo -e "${RED}Failed to delete IAM role: $role${RESET}"
            fi
        else
            if ! gcloud iam roles delete "$role" --project="$PROJECT_ID" --
quiet; then
                echo -e "${RED}Failed to delete IAM role: $role${RESET}"
            fi
        fi
    done
    echo -e "${GREEN}DONE${RESET}"
}

# Error handling setup
set -e
trap 'cleanup $?' EXIT

print_usage() {
    echo "Usage: $0 --project-id PROJECT_ID --auth-key AUTH_KEY --auth-
secret AUTH_SECRET --tenant-id CS_TENANT_ID  [--organization-id
ORGANIZATION_ID] [--read-write] [--role-name ROLE_NAME] [--sa-name SA_NAME]
[--sa-display-name SA_DISPLAY_NAME] [--illumio-sa-email ILLUMIO_SA_EMAIL]
[--post-url POST_URL]"
    echo "  --project-id PROJECT_ID             Specify the GCP project ID
(mandatory)"
    echo "  --post-url POST_URL                 Specify the POST URL
(optional)"
    echo "  --role-name ROLE_NAME               Specify the role name
(optional)"
    echo "  --sa-name SA_NAME                   Specify the service account
name (optional)"
    echo "  --sa-display-name SA_DISPLAY_NAME   Specify the service account
display name (optional)"
    echo "  --illumio-sa-email ILLUMIO_SA_EMAIL Specify the Illumio service
account email (optional)"
    echo "  --auth-key AUTH_KEY                 Specify the authentication
key for basic auth to Illumio endpoint (mandatory)"
    echo "  --auth-secret AUTH_SECRET           Specify the authentication
secret for basic auth Illumio endpoint (mandatory)"
    echo "  --tenant-id CS_TENANT_ID            Specify the tenant ID for
HTTP requests (mandatory)"
    echo "  --organization-id ORGANIZATION_ID   Specify the GCP
```

```
organization ID (optional)"
    echo "  --folder-id FOLDER_ID               Specify the GCP folder ID
(optional)"
    echo "  --read-write                        Enable read-write mode
(optional)"
    exit 1
}


# Parse arguments
READ_WRITE_MODE=false
while [[ $# -gt 0 ]]; do
    case $1 in
        --project-id) PROJECT_ID="$2"; shift 2 ;;
        --post-url) POST_URL="$2"; shift 2 ;;
        --role-name) ROLE_NAME="$2"; shift 2 ;;
        --sa-name) SA_NAME="$2"; shift 2 ;;
        --sa-display-name) SA_DISPLAY_NAME="$2"; shift 2 ;;
        --illumio-sa-email) ILLUMIO_SA_EMAIL="$2"; shift 2 ;;
        --organization-id) ORGANIZATION_ID="$2"; shift 2 ;;
        --folder-id) FOLDER_ID="$2"; shift 2 ;;
        --read-write) READ_WRITE_MODE=true; shift ;;
        --auth-key) AUTH_KEY="$2"; shift 2 ;;
        --auth-secret) AUTH_SECRET="$2"; shift 2 ;;
        --tenant-id) CS_TENANT_ID="$2"; shift 2 ;;
        *) echo "Unknown option: $1"; print_usage ;;
    esac
done


# Validate mandatory parameters
if [ -z "$PROJECT_ID" ] || [ -z "$AUTH_KEY" ] || [ -z "$AUTH_SECRET" ] ||
[ -z "$CS_TENANT_ID" ]; then
#    echo "$PROJECT_ID $AUTH_KEY $AUTH_SECRET $CS_TENANT_ID"
    echo "Error: Project id, auth key, secret and tenant id are mandatory."
    print_usage
fi

# Set default values if not provided
ROLE_NAME="${ROLE_NAME:-$DEFAULT_ROLE_NAME}"
SA_NAME="${SA_NAME:-$DEFAULT_SA_NAME}"
SA_DISPLAY_NAME="${SA_DISPLAY_NAME:-$DEFAULT_SA_DISPLAY_NAME}"
POST_URL="${POST_URL:-$DEFAULT_POST_URL}"

# Validate PROJECT_ID
if ! [[ "$PROJECT_ID" =~ ^[a-zA-Z][a-zA-Z0-9_-]*$ ]]; then
    echo -e "${RED}Error: Invalid PROJECT_ID. It must start with a letter
and can only contain alphanumeric characters, hyphens, or underscores.$
{RESET}"
    exit 1
fi

# Validate POST_URL
if ! [[ "$POST_URL" =~ ^https?:// ]]; then
    echo -e "${RED}Error: Invalid POST_URL. It must start with http:// or
```

```
https://.${RESET}"
    exit 1
fi

# Validate ROLE_NAME
if [ ${#ROLE_NAME} -gt 64 ]; then
    echo -e "${RED}Error: ROLE_NAME must not exceed 64 characters.${RESET}"
    exit 1
fi

# Validate SA_NAME
if [ ${#SA_NAME} -gt 30 ]; then
    echo -e "${RED}Error: SA_NAME must not exceed 30 characters.${RESET}"
    exit 1
fi

# Validate SA_DISPLAY_NAME
if [ ${#SA_DISPLAY_NAME} -gt 100 ]; then
    echo -e "${RED}Error: SA_DISPLAY_NAME must not exceed 100 characters.$
{RESET}"
    exit 1
fi

# Validate ILLUMIO_SA_EMAIL
if ! [[ "$ILLUMIO_SA_EMAIL" =~ ^[a-zA-Z0-9._%+-]+@[a-zA-Z0-9_-]+
\.iam\.gserviceaccount\.com$ ]]; then
    echo -e "${RED}Error: Invalid ILLUMIO_SA_EMAIL. It must be a valid
email address.${RESET}"
    exit 1
fi

# Validate PROJECT_ID
if ! [[ "$PROJECT_ID" =~ ^[a-zA-Z][a-zA-Z0-9_-]*$ ]]; then
    echo -e "${RED}Error: Invalid PROJECT_ID '$PROJECT_ID'. It must start
with a letter and can only contain alphanumeric characters, hyphens, or
underscores.${RESET}"
    exit 1
fi

# Determine resource type based on folder ID or organization ID presence
if [ -n "$FOLDER_ID" ] && [ -n "$ORGANIZATION_ID" ]; then
    RESOURCE_TYPE="$RESOURCE_TYPE_FOLDER"
    RESOURCE_ID="$FOLDER_ID"
    echo "Operating at folder level: $FOLDER_ID"

    # Validate folder ID format (typically numeric)
    if ! [[ "$FOLDER_ID" =~ ^[0-9]+$ ]]; then
        echo -e "${RED}Error: Invalid FOLDER_ID. It must be numeric.$
{RESET}"
        exit 1
    fi

    # Validate organization ID format (typically numeric)
    # This is required for folder onboarding to create custom IAM role
    if ! [[ "$ORGANIZATION_ID" =~ ^[0-9]+$ ]]; then
```

```
        echo -e "${RED}Error: Invalid ORGANIZATION_ID. It must be numeric.
Organization id is required to create custom IAM role for folder onboarding$
{RESET}"
        exit 1
    fi

elif [ -n "$ORGANIZATION_ID" ]; then
    RESOURCE_TYPE="$RESOURCE_TYPE_ORGANIZATION"
    RESOURCE_ID="$ORGANIZATION_ID"
    echo "Operating at organization level: $ORGANIZATION_ID"

    # Validate organization ID format (typically numeric)
    if ! [[ "$ORGANIZATION_ID" =~ ^[0-9]+$ ]]; then
        echo -e "${RED}Error: Invalid ORGANIZATION_ID. It must be numeric.$
{RESET}"
        exit 1
    fi
else
    RESOURCE_TYPE="$RESOURCE_TYPE_PROJECT"
    RESOURCE_ID="$PROJECT_ID"
    echo "Operating at project level: $PROJECT_ID"
fi


# READ_WRITE_MODE
if $READ_WRITE_MODE; then
    echo "Read-write mode enabled."
else
    echo "Read-only mode."
fi

# Function to set the active project
set_project() {
    local project_id=$1
    echo "Setting active project to $project_id..."
    gcloud config set project "$project_id"
    echo -e "${GREEN}DONE${RESET}"
}

# Function to create service account
create_service_account() {
    local project_id=$1
    local sa_name=$2
    local sa_display_name=$3

    echo "Creating service account..."
    gcloud iam service-accounts create "$sa_name" \
        --display-name="$sa_display_name" \
        --project="$project_id" \
        --quiet
    #creating service account email using the service account name and
project id
    local sa_email="${sa_name}@${project_id}.iam.gserviceaccount.com"
    echo "$sa_email"
}
```

```
# Function to create IAM role
create_iam_role() {
    local resource_type=$1
    local resource_id=$2
    local role_name=$3
    local permissions=$4
    local organization_id=$5

    echo "Creating custom IAM role..."
    if [ "$resource_type" == "$RESOURCE_TYPE_PROJECT" ]; then
        gcloud iam roles create "$role_name" \
            --project="$resource_id" \
            --title="$role_name" \
            --description="Custom role for listing and getting storage and
VPCs" \
            --permissions="$permissions" \
            --stage="GA" \
            --quiet > /dev/null
    elif [ "$resource_type" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
         gcloud iam roles create "$role_name" \
             --organization="$resource_id" \
             --title="$role_name" \
             --description="Custom role for organization-level permissions"
\
             --permissions="$permissions" \
             --stage="GA" \
             --quiet > /dev/null
    elif [ "$resource_type" == "$RESOURCE_TYPE_FOLDER" ]; then
#      iam role could only be created at project/organization level, for
folder onboarding creating it at organization level
       gcloud iam roles create "$role_name" \
           --organization="$organization_id" \
           --title="$role_name" \
           --description="Custom role for folder-level permissions" \
           --permissions="$permissions" \
           --stage="GA" \
           --quiet > /dev/null
    fi
    add_resource "roles" "$role_name"
    echo -e "${GREEN}DONE${RESET}"
}

# Function to bind IAM role to service account
bind_role_to_service_account() {
    local resource_type=$1
    local resource_id=$2
    local sa_email=$3
    local role_name=$4
    local organization_id=$5

    echo "Binding custom IAM role to service account..."
    if [ "$resource_type" == "$RESOURCE_TYPE_PROJECT" ]; then
        gcloud projects add-iam-policy-binding "$resource_id" \
            --member="serviceAccount:$sa_email" \
            --role="projects/$resource_id/roles/$role_name" \
            --condition=None \
            --quiet > /dev/null
    elif [ "$resource_type" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
         gcloud organizations add-iam-policy-binding "$resource_id" \
             --member="serviceAccount:$sa_email" \
             --role="organizations/$resource_id/roles/$role_name" \
             --condition=None \
             --quiet > /dev/null
    elif [ "$resource_type" == "$RESOURCE_TYPE_FOLDER" ]; then
         gcloud resource-manager folders add-iam-policy-binding
"$resource_id" \
             --member="serviceAccount:$sa_email" \
             --role="organizations/$organization_id/roles/$role_name" \
             --condition=None \
             --quiet > /dev/null
    fi
    echo -e "${GREEN}DONE${RESET}"
}

# Function to assign predefined role to service account
assign_predefined_role() {
    local resource_type=$1
    local resource_id=$2
    local sa_email=$3
    local role=$4

    echo "Assigning predefined role $role to service account..."
    if [ "$resource_type" == "$RESOURCE_TYPE_PROJECT" ]; then
        gcloud projects add-iam-policy-binding "$resource_id" \
            --member="serviceAccount:$sa_email" \
            --role="$role" \
            --condition=None \
            --quiet > /dev/null
    elif  [ "$resource_type" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
        gcloud organizations add-iam-policy-binding "$resource_id" \
                --member="serviceAccount:$sa_email" \
                --role="$role"\
                --condition=None \
                --quiet > /dev/null
    elif [ "$resource_type" == "$RESOURCE_TYPE_FOLDER" ]; then
        gcloud resource-manager folders add-iam-policy-binding
"$resource_id" \
            --member="serviceAccount:$sa_email" \
            --role="$role" \
            --condition=None \
            --quiet > /dev/null
    fi
    echo -e "${GREEN}DONE${RESET}"
}

# Function to enable impersonation permissions
enable_impersonation_permissions() {
    local sa_email=$1
    local project_id=$2
    local illumio_sa_email=$3

    echo "Configuring impersonation permissions..."
    gcloud iam service-accounts add-iam-policy-binding "$sa_email" \
        --member="serviceAccount:$illumio_sa_email" \
        --role="roles/iam.serviceAccountTokenCreator" \
        --project="$project_id" \
        --condition=None \
        --quiet > /dev/null
    echo -e "${GREEN}DONE${RESET}"
}

send_data_to_endpoint() {
    local sa_email=$1
    local resource_id=$2
    local project_id=$3
    local post_url=$4
    local tenant_id=$5
    local auth_key=$6  # Basic Authentication key
    local auth_secret=$7  # Basic Authentication secret
    local event=$8
    local resource_type=$9

    if [[ "$post_url" !=  *"proxy"* ]] || [[ "$post_url" ==
*"sunnyvale"* ]]; then
        post_url="$post_url/api/v1/integrations/cloud_credentials"
    fi

    # Determine onboarding_type based on resource_type
    local integration_type=""
    if [ "$resource_type" == "project" ]; then
        integration_type="GcpProject"
    elif [ "$resource_type" == "folder" ]; then
        integration_type="GcpFolder"
    elif [ "$resource_type" == "organization" ]; then
        integration_type="GcpOrganization"
    else
        echo "Error: Invalid resource_type provided."
        exit 1
    fi

    echo "Sending resource ID, service account email, and project ID to
$post_url"

    # Perform the HTTP POST request with Basic Authentication and custom
headers
    response=$(curl -s -w "\n%{http_code}" -X POST \
         -H "Content-Type: application/json" \
         -H "Authorization: Basic $(echo -n "$auth_key:$auth_secret" |
base64 | tr -d '\n')" \
         -H "X-Tenant-Id: $tenant_id" \
         -d
"{\"sa_email\":\"$sa_email\",\"gcp_resource_id\":\"$resource_id\",\"project_
csp_id\":\"$project_id\",\"type\":\"$event\",\"integration_type\":\"$integra
tion_type\"}" \
         "$post_url")

    # Extract the response body and status code
    body=$(echo "$response" | sed '$d')  # Remove last line (status code)
    status_code=$(echo "$response" | tail -n1)  # Extract last line (status
code)

    # Check the HTTP status code
    if [ "$status_code" -ge 200 ] && [ "$status_code" -lt 300 ]; then
        echo "POST request successful. Response:"
        echo "$body"
    else
        echo "Error: POST request failed with status code $status_code.
Response:"
        echo "$body"
        exit 1  # Trigger cleanup via ERR trap if there's an error.
    fi
}

# Function to check if gcloud CLI is available
check_gcloud_availability() {
    if ! command -v gcloud &> /dev/null; then
        echo -e "${RED}Error: gcloud CLI is not available. Please install
and configure it before running this script.${RESET}"
        exit 1
    fi
    echo "gcloud CLI is available. Proceeding with the script."
}

# Function to enable specific APIs for a project
enable_apis_for_project() {
    local project_id=$1
    echo "Enabling APIs (iamcredentials.googleapis.com,
cloudresourcemanager.googleapis.com) for project $project_id..."
    if ! gcloud services enable iamcredentials.googleapis.com
cloudresourcemanager.googleapis.com --project="$project_id" --quiet; then
        echo -e "${RED}Failed to enable required APIs for project
$project_id.${RESET}"
        exit 1
    fi
    echo -e "${GREEN}DONE${RESET}"
}

# Main execution sequence
main() {
    check_gcloud_availability
    if [ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_PROJECT" ]; then
        set_project "$PROJECT_ID"
    fi
    # TODO: Check user permissions before proceeding

    SA_EMAIL=$(create_service_account "$PROJECT_ID" "$SA_NAME"
"$SA_DISPLAY_NAME" | tail -n 1)
    # doing add_resource here for service account because we are using a
pipe (| tail -n 1) when calling the function, which creates a subshell.
```

```
    # Variables modified in subshells do not propagate back to the parent
shell, hence the added service account was not visible in cleanup function
    add_resource "service_accounts" "$SA_EMAIL"

    # Prompt user for consent to enable APIs for Illumio-supported resources
    echo "Do you allow Illumio to enable APIs for supported resources at
the time of ingestion? (y/N)"
    read -r consent
    if [[ "$consent" == "y" || "$consent" == "Y" ]]; then
        echo "User consented to enable APIs for Illumio-supported
resources."
        create_iam_role "$RESOURCE_TYPE" "$RESOURCE_ID"
"$DEFAULT_API_ENABLE_ROLE_NAME" "$API_ENABLE_PERMISSION" "$ORGANIZATION_ID"
        bind_role_to_service_account "$RESOURCE_TYPE" "$RESOURCE_ID"
"$SA_EMAIL" "$DEFAULT_API_ENABLE_ROLE_NAME" "$ORGANIZATION_ID"
    else
        echo "User did not consent to enable APIs for Illumio-supported
resources."
    fi

    # Convert the comma-separated list into an array and loop through it
    IFS=',' read -ra ROLES_ARRAY <<< "$PREDEFINED_ROLES"
    for role in "${ROLES_ARRAY[@]}"; do
        assign_predefined_role "$RESOURCE_TYPE" "$RESOURCE_ID" "$SA_EMAIL"
"$role"
    done
    enable_impersonation_permissions "$SA_EMAIL" "$PROJECT_ID"
"$ILLUMIO_SA_EMAIL"

    # required to list projects under the organization or folder
    if [ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_ORGANIZATION" ] ||
[ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_FOLDER" ]; then
      assign_predefined_role "$RESOURCE_TYPE" "$RESOURCE_ID" "$SA_EMAIL"
"roles/browser"
    fi

    # Check if READ_WRITE_MODE is true to create and bind IAM role with
WRITE_PERMISSIONS
    if $READ_WRITE_MODE; then
        echo "Read-write mode is enabled. Creating IAM role with write
permissions..."
        create_iam_role "$RESOURCE_TYPE" "$RESOURCE_ID"
"$DEFAULT_WRITE_ROLE_NAME" "$WRITE_PERMISSIONS" "$ORGANIZATION_ID"
        bind_role_to_service_account "$RESOURCE_TYPE" "$RESOURCE_ID"
"$SA_EMAIL" "$DEFAULT_WRITE_ROLE_NAME" "$ORGANIZATION_ID"
    fi

    # Enable APIs for the project to allow impersonation and
cloudresourcemanager to read child resources
    enable_apis_for_project "$PROJECT_ID"

    send_data_to_endpoint "$SA_EMAIL" "$RESOURCE_ID" "$PROJECT_ID"
"$POST_URL" "$CS_TENANT_ID" "$AUTH_KEY" "$AUTH_SECRET" "GCPRole"
"$RESOURCE_TYPE"
}
```

```
main
echo -e "${GREEN}Script executed successfully.${RESET}"#!/bin/bash

#Constants
RESOURCE_TYPE_PROJECT="project"
RESOURCE_TYPE_ORGANIZATION="organization"
RESOURCE_TYPE_FOLDER="folder"
DEFAULT_ROLE_NAME="illumio_role_$(date +%s)"
DEFAULT_SA_NAME="illumio-sa-$(date +%s)"
DEFAULT_SA_DISPLAY_NAME="Illumio Service Account"
DEFAULT_POST_URL="https://cloud.illum.io"
PREDEFINED_ROLES="roles/iam.securityReviewer,roles/compute.viewer,roles/
cloudasset.viewer"
DEFAULT_WRITE_ROLE_NAME="illumio_write_role_$(date +%s)"
WRITE_PERMISSIONS="compute.firewalls.create,compute.firewalls.delete,compute
.firewalls.get,compute.firewalls.update,compute.networks.updatePolicy"
DEFAULT_API_ENABLE_ROLE_NAME="illumio_api_enable_role_$(date +%s)"
API_ENABLE_PERMISSION="serviceusage.services.enable,serviceusage.services.li
st,serviceusage.services.get"

# Resource tracking
created_resources__service_accounts=""
created_resources__roles=""
RESOURCE_TYPE="$RESOURCE_TYPE_PROJECT" # using default resource type as
object, based on if organization id is passed in arguments, it changes to
"$RESOURCE_TYPE_ORGANIZATION"

# Function to add a resource to the tracking list - these lists will be
used to clean up resources in case of an error
add_resource() {
    local resource_type=$1
    local resource_value=$2
    local var_name="created_resources__${resource_type}"
    if [[ -z "${!var_name}" ]]; then
        eval "$var_name=\"$resource_value\""
    else
        eval "$var_name=\"${!var_name},$resource_value\""
    fi
}

# Cleanup function for errors (triggered by trap)
cleanup() {
    if [ "$1" == "0" ]; then
      return
    fi
    echo "Cleaning up resources..."

    # Delete service accounts
    IFS=',' read -ra SAS <<< "${created_resources__service_accounts}"
    for sa in "${SAS[@]}"; do
        echo "Deleting service account: $sa"
        gcloud iam service-accounts delete "$sa" --quiet --
project="$PROJECT_ID" || true
    done
    # Delete roles
    IFS=',' read -ra ROLES <<< "${created_resources__roles}"
    for role in "${ROLES[@]}"; do
        echo "Deleting IAM role: $role"
        if [ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
            gcloud iam roles delete "$role" --
organization="$ORGANIZATION_ID" --quiet || true
        else
            gcloud iam roles delete "$role" --project="$PROJECT_ID" --quiet
|| true
        fi
    done
}

# Error handling setup
set -e
trap 'cleanup $?' EXIT

ILLUMIO_SA_EMAIL="illumio-onboarding@cs-prod-01.iam.gserviceaccount.com"

print_usage() {
    echo "Usage: $0 --project-id PROJECT_ID --auth-key AUTH_KEY --auth-
secret AUTH_SECRET --tenant-id CS_TENANT_ID  [--organization-id
ORGANIZATION_ID] [--read-write] [--role-name ROLE_NAME] [--sa-name SA_NAME]
[--sa-display-name SA_DISPLAY_NAME] [--illumio-sa-email ILLUMIO_SA_EMAIL]
[--post-url POST_URL]"
    echo "  --project-id PROJECT_ID             Specify the GCP project ID
(mandatory)"
    echo "  --post-url POST_URL                 Specify the POST URL
(optional)"
    echo "  --role-name ROLE_NAME               Specify the role name
(optional)"
    echo "  --sa-name SA_NAME                   Specify the service account
name (optional)"
    echo "  --sa-display-name SA_DISPLAY_NAME   Specify the service account
display name (optional)"
    echo "  --illumio-sa-email ILLUMIO_SA_EMAIL Specify the Illumio service
account email (optional)"
    echo "  --auth-key AUTH_KEY                 Specify the authentication
key for basic auth to Illumio endpoint (mandatory)"
    echo "  --auth-secret AUTH_SECRET           Specify the authentication
secret for basic auth Illumio endpoint (mandatory)"
    echo "  --tenant-id CS_TENANT_ID            Specify the tenant ID for
HTTP requests (mandatory)"
    echo "  --organization-id ORGANIZATION_ID   Specify the GCP
organization ID (optional)"
    echo "  --folder-id FOLDER_ID               Specify the GCP folder ID
(optional)"
    echo "  --read-write                        Enable read-write mode
(optional)"
    exit 1
}
```

```
# Parse arguments
READ_WRITE_MODE=false
while [[ $# -gt 0 ]]; do
    case $1 in
        --project-id) PROJECT_ID="$2"; shift 2 ;;
        --post-url) POST_URL="$2"; shift 2 ;;
        --role-name) ROLE_NAME="$2"; shift 2 ;;
        --sa-name) SA_NAME="$2"; shift 2 ;;
        --sa-display-name) SA_DISPLAY_NAME="$2"; shift 2 ;;
        --illumio-sa-email) ILLUMIO_SA_EMAIL="$2"; shift 2 ;;
        --organization-id) ORGANIZATION_ID="$2"; shift 2 ;;
        --folder-id) FOLDER_ID="$2"; shift 2 ;;
        --read-write) READ_WRITE_MODE=true; shift ;;
        --auth-key) AUTH_KEY="$2"; shift 2 ;;
        --auth-secret) AUTH_SECRET="$2"; shift 2 ;;
        --tenant-id) CS_TENANT_ID="$2"; shift 2 ;;
        *) echo "Unknown option: $1"; print_usage ;;
    esac
done


# Validate mandatory parameters
if [ -z "$PROJECT_ID" ] || [ -z "$AUTH_KEY" ] || [ -z "$AUTH_SECRET" ] ||
[ -z "$CS_TENANT_ID" ]; then
#    echo "$PROJECT_ID $AUTH_KEY $AUTH_SECRET $CS_TENANT_ID"
    echo "Error: Project id, auth key, secret and tenant id are mandatory."
    print_usage
fi

# Set default values if not provided
ROLE_NAME="${ROLE_NAME:-$DEFAULT_ROLE_NAME}"
SA_NAME="${SA_NAME:-$DEFAULT_SA_NAME}"
SA_DISPLAY_NAME="${SA_DISPLAY_NAME:-$DEFAULT_SA_DISPLAY_NAME}"
POST_URL="${POST_URL:-$DEFAULT_POST_URL}"

# Validate PROJECT_ID
if ! [[ "$PROJECT_ID" =~ ^[a-zA-Z][a-zA-Z0-9_-]*$ ]]; then
    echo "Error: Invalid PROJECT_ID. It must start with a letter and can
only contain alphanumeric characters, hyphens, or underscores."
    exit 1
fi

# Validate POST_URL
if ! [[ "$POST_URL" =~ ^https?:// ]]; then
    echo "Error: Invalid POST_URL. It must start with http:// or https://."
    exit 1
fi

# Validate ROLE_NAME
if [ ${#ROLE_NAME} -gt 64 ]; then
    echo "Error: ROLE_NAME must not exceed 64 characters."
    exit 1
fi

# Validate SA_NAME
```

```
if [ ${#SA_NAME} -gt 30 ]; then
    echo "Error: SA_NAME must not exceed 30 characters."
    exit 1
fi

# Validate SA_DISPLAY_NAME
if [ ${#SA_DISPLAY_NAME} -gt 100 ]; then
    echo "Error: SA_DISPLAY_NAME must not exceed 100 characters."
    exit 1
fi

# Validate ILLUMIO_SA_EMAIL
if ! [[ "$ILLUMIO_SA_EMAIL" =~ ^[a-zA-Z0-9._%+-]+@[a-zA-Z0-9_-]+
\.iam\.gserviceaccount\.com$ ]]; then
    echo "Error: Invalid ILLUMIO_SA_EMAIL. It must be a valid email
address."
    exit 1
fi

# Validate PROJECT_ID
if [ ${#PROJECT_ID} -lt 6 ] || [ ${#PROJECT_ID} -gt 30 ]; then
    echo "Error: PROJECT_ID must be between 6 to 30 characters."
    exit 1
fi

# Determine resource type based on folder ID or organization ID presence
if [ -n "$FOLDER_ID" ] && [ -n "$ORGANIZATION_ID" ]; then
    RESOURCE_TYPE="$RESOURCE_TYPE_FOLDER"
    RESOURCE_ID="$FOLDER_ID"
    echo "Operating at folder level: $FOLDER_ID"

    # Validate folder ID format (typically numeric)
    if ! [[ "$FOLDER_ID" =~ ^[0-9]+$ ]]; then
        echo "Error: Invalid FOLDER_ID. It must be numeric."
        exit 1
    fi

    # Validate organization ID format (typically numeric)
    # This is required for folder onboarding to create custom IAM role
    if ! [[ "$ORGANIZATION_ID" =~ ^[0-9]+$ ]]; then
        echo "Error: Invalid ORGANIZATION_ID. It must be numeric.
Organization id is required to create custom IAM role for folder onboarding"
        exit 1
    fi

elif [ -n "$ORGANIZATION_ID" ]; then
    RESOURCE_TYPE="$RESOURCE_TYPE_ORGANIZATION"
    RESOURCE_ID="$ORGANIZATION_ID"
    echo "Operating at organization level: $ORGANIZATION_ID"

    # Validate organization ID format (typically numeric)
    if ! [[ "$ORGANIZATION_ID" =~ ^[0-9]+$ ]]; then
        echo "Error: Invalid ORGANIZATION_ID. It must be numeric."
        exit 1
    fi
```

```
else
    RESOURCE_TYPE="$RESOURCE_TYPE_PROJECT"
    RESOURCE_ID="$PROJECT_ID"
    echo "Operating at project level: $PROJECT_ID"
fi


# READ_WRITE_MODE
if $READ_WRITE_MODE; then
    echo "Read-write mode enabled."
else
    echo "Read-only mode."
fi

# Function to set the active project
set_project() {
    local project_id=$1
    echo "Setting active project to $project_id..."
    gcloud config set project "$project_id"
}

# Function to create service account
create_service_account() {
    local project_id=$1
    local sa_name=$2
    local sa_display_name=$3

    echo "Creating service account..."
    gcloud iam service-accounts create "$sa_name" \
        --display-name="$sa_display_name" \
        --project="$project_id" \
        --quiet
    #creating service account email using the service account name and
project id
    local sa_email="${sa_name}@${project_id}.iam.gserviceaccount.com"
    echo "$sa_email"
}

# Function to create IAM role
create_iam_role() {
    local resource_type=$1
    local resource_id=$2
    local role_name=$3
    local permissions=$4
    local organization_id=$5

    echo "Creating custom IAM role..."
    if [ "$resource_type" == "$RESOURCE_TYPE_PROJECT" ]; then
        gcloud iam roles create "$role_name" \
            --project="$resource_id" \
            --title="$role_name" \
            --description="Custom role for listing and getting storage and
VPCs" \
            --permissions="$permissions" \
            --stage="GA"
    elif [ "$resource_type" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
         gcloud iam roles create "$role_name" \
             --organization="$resource_id" \
             --title="$role_name" \
             --description="Custom role for organization-level permissions"
\
             --permissions="$permissions" \
             --stage="GA"
    elif [ "$resource_type" == "$RESOURCE_TYPE_FOLDER" ]; then
#      iam role could only be created at project/organization level, for
folder onboarding creating it at organization level
       gcloud iam roles create "$role_name" \
           --organization="$organization_id" \
           --title="$role_name" \
           --description="Custom role for folder-level permissions" \
           --permissions="$permissions" \
           --stage="GA"
    fi
    add_resource "roles" "$role_name"
}

# Function to bind IAM role to service account
bind_role_to_service_account() {
    local resource_type=$1
    local resource_id=$2
    local sa_email=$3
    local role_name=$4
    local organization_id=$5

    echo "Binding custom IAM role to service account..."
    if [ "$resource_type" == "$RESOURCE_TYPE_PROJECT" ]; then
        gcloud projects add-iam-policy-binding "$resource_id" \
            --member="serviceAccount:$sa_email" \
            --role="projects/$resource_id/roles/$role_name" \
            --condition=None \
            --quiet
    elif [ "$resource_type" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
         gcloud organizations add-iam-policy-binding "$resource_id" \
             --member="serviceAccount:$sa_email" \
             --role="organizations/$resource_id/roles/$role_name" \
             --condition=None \
             --quiet
    elif [ "$resource_type" == "$RESOURCE_TYPE_FOLDER" ]; then
         gcloud resource-manager folders add-iam-policy-binding
"$resource_id" \
             --member="serviceAccount:$sa_email" \
             --role="organizations/$organization_id/roles/$role_name" \
             --condition=None \
             --quiet
    fi
}

# Function to assign predefined role to service account
assign_predefined_role() {
    local resource_type=$1
    local resource_id=$2
    local sa_email=$3
    local role=$4

    echo "Assigning predefined role $role to service account..."
    if [ "$resource_type" == "$RESOURCE_TYPE_PROJECT" ]; then
        gcloud projects add-iam-policy-binding "$resource_id" \
            --member="serviceAccount:$sa_email" \
            --role="$role" \
            --condition=None \
            --quiet
    elif  [ "$resource_type" == "$RESOURCE_TYPE_ORGANIZATION" ]; then
        gcloud organizations add-iam-policy-binding "$resource_id" \
                --member="serviceAccount:$sa_email" \
                --role="$role"\
                --condition=None \
                --quiet
    elif [ "$resource_type" == "$RESOURCE_TYPE_FOLDER" ]; then
        gcloud resource-manager folders add-iam-policy-binding
"$resource_id" \
            --member="serviceAccount:$sa_email" \
            --role="$role" \
            --condition=None \
            --quiet
    fi
}

# Function to enable impersonation permissions
enable_impersonation_permissions() {
    local sa_email=$1
    local project_id=$2
    local illumio_sa_email=$3

    echo "Configuring impersonation permissions..."
    gcloud iam service-accounts add-iam-policy-binding "$sa_email" \
        --member="serviceAccount:$illumio_sa_email" \
        --role="roles/iam.serviceAccountTokenCreator" \
        --project="$project_id" \
        --condition=None \
        --quiet
}

send_data_to_endpoint() {
    local sa_email=$1
    local resource_id=$2
    local project_id=$3
    local post_url=$4
    local tenant_id=$5
    local auth_key=$6  # Basic Authentication key
    local auth_secret=$7  # Basic Authentication secret
    local event=$8
    local resource_type=$9

    if [[ "$post_url" !=  *"proxy"* ]] || [[ "$post_url" ==
*"sunnyvale"* ]]; then
```

```
        post_url="$post_url/api/v1/integrations/cloud_credentials"
    fi

    # Determine onboarding_type based on resource_type
    local integration_type=""
    if [ "$resource_type" == "project" ]; then
        integration_type="GcpProject"
    elif [ "$resource_type" == "folder" ]; then
        integration_type="GcpFolder"
    elif [ "$resource_type" == "organization" ]; then
        integration_type="GcpOrganization"
    else
        echo "Error: Invalid resource_type provided."
        exit 1
    fi

    echo "Sending resource ID, service account email, and project ID to 
$post_url"

    # Perform the HTTP POST request with Basic Authentication and custom 
headers
    response=$(curl -s -w "\n%{http_code}" -X POST \
         -H "Content-Type: application/json" \
         -H "Authorization: Basic $(echo -n "$auth_key:$auth_secret" | 
base64 | tr -d '\n')" \
         -H "X-Tenant-Id: $tenant_id" \
         -d 
"{\"sa_email\":\"$sa_email\",\"gcp_resource_id\":\"$resource_id\",\"project_
csp_id\":\"$project_id\",\"type\":\"$event\",\"integration_type\":\"$integra
tion_type\"}" \
         "$post_url")

    # Extract the response body and status code
    body=$(echo "$response" | sed '$d')  # Remove last line (status code)
    status_code=$(echo "$response" | tail -n1)  # Extract last line (status 
code)

    # Check the HTTP status code
    if [ "$status_code" -ge 200 ] && [ "$status_code" -lt 300 ]; then
        echo "POST request successful. Response:"
        echo "$body"
    else
        echo "Error: POST request failed with status code $status_code. 
Response:"
        echo "$body"
        exit 1  # Trigger cleanup via ERR trap if there's an error.
    fi
}

# Function to check if gcloud CLI is available
check_gcloud_availability() {
    if ! command -v gcloud &> /dev/null; then
        echo "Error: gcloud CLI is not available. Please install and 
configure it before running this script."
        exit 1
    fi
    echo "gcloud CLI is available. Proceeding with the script."
}


# Main execution sequence
main() {
    check_gcloud_availability
    if [ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_PROJECT" ]; then
        set_project "$PROJECT_ID"
    fi
    # TODO: Check user permissions before proceeding

    SA_EMAIL=$(create_service_account "$PROJECT_ID" "$SA_NAME"
"$SA_DISPLAY_NAME" | tail -n 1)
    # doing add_resource here for service account because we are using a
pipe (| tail -n 1) when calling the function, which creates a subshell.
    # Variables modified in subshells do not propagate back to the parent
shell, hence the added service account was not visible in cleanup function
    add_resource "service_accounts" "$SA_EMAIL"

    # Prompt user for consent to enable APIs for Illumio-supported resources
    echo "Do you want to allow Illumio to enable APIs for supported
resources? (y/N)"
    read -r consent
    if [[ "$consent" == "y" || "$consent" == "Y" ]]; then
        echo "User consented to enable APIs for Illumio-supported
resources."
        create_iam_role "$RESOURCE_TYPE" "$RESOURCE_ID"
"$DEFAULT_API_ENABLE_ROLE_NAME" "$API_ENABLE_PERMISSION" "$ORGANIZATION_ID"
        bind_role_to_service_account "$RESOURCE_TYPE" "$RESOURCE_ID"
"$SA_EMAIL" "$DEFAULT_API_ENABLE_ROLE_NAME" "$ORGANIZATION_ID"
    else
        echo "User did not consent to enable APIs for Illumio-supported
resources."
    fi

    # Convert the comma-separated list into an array and loop through it
    IFS=',' read -ra ROLES_ARRAY <<< "$PREDEFINED_ROLES"
    for role in "${ROLES_ARRAY[@]}"; do
        assign_predefined_role "$RESOURCE_TYPE" "$RESOURCE_ID" "$SA_EMAIL"
"$role"
    done
    enable_impersonation_permissions "$SA_EMAIL" "$PROJECT_ID"
"$ILLUMIO_SA_EMAIL"

    # required to list projects under the organization or folder
    if [ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_ORGANIZATION" ] ||
[ "$RESOURCE_TYPE" == "$RESOURCE_TYPE_FOLDER" ]; then
      assign_predefined_role "$RESOURCE_TYPE" "$RESOURCE_ID" "$SA_EMAIL"
"roles/browser"
    fi

    # Check if READ_WRITE_MODE is true to create and bind IAM role with
WRITE_PERMISSIONS
```

```
    if $READ_WRITE_MODE; then
        echo "Read-write mode is enabled. Creating IAM role with write
permissions..."
        create_iam_role "$RESOURCE_TYPE" "$RESOURCE_ID"
"$DEFAULT_WRITE_ROLE_NAME" "$WRITE_PERMISSIONS" "$ORGANIZATION_ID"
        bind_role_to_service_account "$RESOURCE_TYPE" "$RESOURCE_ID"
"$SA_EMAIL" "$DEFAULT_WRITE_ROLE_NAME" "$ORGANIZATION_ID"
    fi

    send_data_to_endpoint "$SA_EMAIL" "$RESOURCE_ID" "$PROJECT_ID"
"$POST_URL" "$CS_TENANT_ID" "$AUTH_KEY" "$AUTH_SECRET" "GCPRole"
"$RESOURCE_TYPE"
}

main
echo "Script executed successfully."
```

GCP flow log support

Illumio Segmentation for the Cloud supports VPC and firewall flow logs.

See Set up flow logs in your CSP environment.

OCI Flow Log Access IP Addresses

Illumio Segmentation for the Cloud uses TCP port 443 to access your flow logs, so open that port for the IP addresses listed in this section.

Illumio Control Plane (For all OCI Regions)

The Illumio Segmentation for the Cloud control and data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound/outbound allowed list:

- 35.167.22.34
- 52.88.124.247
- 52.88.88.252

Illumio Data Plane for all OCI Regions

The Illumio Segmentation for the Cloud data plane uses the following public IP addresses to reach customer networks, so add them to your firewall inbound allowed list for OCI.

- 13.57.69.111
- 52.8.11.104
- 52.8.120.46

Illumio IP addresses accessed by the Kubernetes Cloud Operator

The Kubernetes Cloud Operator deployed into each Kubernetes cluster uses TCP port 443 to connect to Illumio Segmentation for the Cloud. The operator uses this port to report Kubernetes resources and flow logs, and to retrieve configuration. You must allow access to that port for the IP addresses listed for the control plane and the Illumio Regions where each cluster is onboarded. For an overview of Agentless Containers, see Agentless Containers overview [148]. For the Illumio Cloud Operator code, which is open source under Apache License 2.0, see GitHub.

Illumio control plane (for all Kubernetes clusters)

All customers must permit the following public IP addresses to successfully onboard clusters. These IPs are required for the Kubernetes Cloud Operator to authenticate and communicate with Illumio Segmentation for the Cloud. Ensure that you add these IPs to your firewall's outbound allow list.

- 35.80.225.104
- 100.20.246.114
- 52.42.243.65

Illumio AWS US West 2 Region (data plane)

The Kubernetes Cloud Operators onboarded into this Illumio Region access the following public IP addresses to report Kubernetes resources and flow logs, and to retrieve the configuration. Add them to your firewall outbound allowed list for each cluster onboarded into this Illumio Region.

- k8sclustersync.aws.us-west-2.prod.cloud.illum.io
- 35.82.131.82
- 52.89.200.143
- 54.214.36.211

Illumio AWS AP Southeast 2 Region (data plane)

The Kubernetes Cloud Operators onboarded into this Illumio Region access the following public IP addresses to report Kubernetes resources and flow logs, and to retrieve the configuration. Add them to your firewall outbound allowed list for each cluster onboarded into this Illumio Region.

- k8sclustersync.aws.ap-southeast-2.prod.cloud.illum.io
- 54.79.89.106
- 3.24.74.41
- 13.211.119.109

Illumio AWS US West 1 Region (data plane)

The Kubernetes Cloud Operators onboarded into this Illumio Region access the following public IP addresses to report Kubernetes resources and flow logs, and to retrieve the configuration. Add them to your firewall outbound allowed list for each cluster onboarded into this Illumio Region.

- k8sclustersync.aws.us-west-1.prod.cloud.illum.io
- 54.153.101.43
- 52.52.76.163

Illumio AWS EU West 2 Region (data plane)

The Kubernetes Cloud Operators onboarded into this Illumio Region access the following public IP addresses to report Kubernetes resources and flow logs, and to retrieve the configuration. Add them to your firewall outbound allowed list for each cluster onboarded into this Illumio Region.

- k8sclustersync.aws.eu-west-2.prod.cloud.illum.io
- 13.43.35.249
- 52.56.199.135
- 35.177.86.66

Illumio Azure US East 2 Region (data plane)

The Kubernetes Cloud Operators onboarded into this Illumio Region access the following public IP addresses to report Kubernetes resources and flow logs, and to retrieve the configuration. Add them to your firewall outbound allowed list for each cluster onboarded into this Illumio Region.

- k8sclustersync.azure.eastus.prod.cloud.illum.io
- 172.190.182.192

Illumio Azure Germany West Central Region (data plane)

The Kubernetes Cloud Operators onboarded into this Illumio Region access the following public IP addresses to report Kubernetes resources and flow logs, and to retrieve the configuration. Add them to your firewall outbound allowed list for each cluster onboarded into this Illumio Region.

- k8sclustersync.azure.gwc.prod.cloud.illum.io
- 9.141.21.191

Illumio Azure West US 2 Region (data plane)

The Kubernetes Cloud Operators onboarded into this Illumio Region access the following public IP addresses to report Kubernetes resources and flow logs, and to retrieve the configuration. Add them to your firewall outbound allowed list for each cluster onboarded into this Illumio Region.

- k8sclustersync.azure.westus2.prod.cloud.illum.io
- 52.175.211.226

Google APIs enabled during GCP onboarding

The following Google APIs are enabled during GCP onboarding so that Illumio Segmentation for the Cloud can fetch inventory:

- compute.googleapis.com
- batch.googleapis.com
- storage.googleapis.com
- container.googleapis.com
- cloudresourcemanager.googleapis.com
- logging.googleapis.com
- pubsub.googleapis.com
- sqladmin.googleapis.com
- networkconnectivity.googleapis.com
- servicenetworking.googleapis.com
- cloudasset.googleapis.com
- iamcredentials.googleapis.com

Troubleshooting

The content in this category explains how you troubleshoot Illumio Segmentation for the Cloud issues.

Troubleshoot Azure traffic flow log shrinkage or lack of traffic

If the size of the traffic flow log has shrunk or you are not seeing traffic, it may indicate a problem with either your permissions or the PT1H.json file.

Symptoms

- You no longer see traffic on the Traffic page
- The traffic flow log has shrunk. For example, if the current log is 100MB but the previous log size was 700MB, traffic data may be missing.

Cause

If you don't see traffic in the traffic logs, possible causes include:

- Incorrect permissions were granted to the Illumio application
- The provided storage account does not contain logs

Resolution

To solve this issue, follow these step-by-step instructions.

1. If the file exists, ensure that the correct permissions are granted to the Illumio application.
2. Verify that the PT1H.json file exists for the storage account 'flowlogstoragedaita' and is located in your container path.

   For example: insights-logs-flowlogflowevent/flowLogResourceID=+ <FLOWLOG-SOURCE> / y=2023/m=09/d=02/h=03/m=00/macAddress=<address>/PT1H.json"
3. If the PT1H.json file is missing, see the Microsoft website.
4. If this doesn't resolve the issue, contact your Illumio Support representative.

Troubleshoot system-generated Azure Firewall messages

If you see error messages on the Illumio Segmentation for the Cloud Events page System Events tab related to Azure Firewalls, you may need to adjust settings according to the table.

List of system-generated Azure Firewall messages

Use the following table to assist you with troubleshooting system-generated error messages.

Message	Cause	Resolution
Enforcement for <azure-firewall-csp id> failed	Error sent from Azure Cloud when Illumio Segmentation for the Cloud tries to enforce rules	Examine the 'details' field to see the error that Azure returned. 1. RESPONSE Not Found - This error means that either a resource group or firewall policy has been deleted, so Illumio Segmentation for the Cloud fails to enforce rules. No action is required. 2. RESPONSE Forbidden - This error means that Illumio Segmentation for the Cloud lacks permissions to update the Azure Firewall Policy rules. Check your permissions in the Azure portal to make sure that Illumio Segmentation for the Cloud has write permissions in addition to read permissions.
CloudSecure authored rule <rule-id> from <azure-firewall-csp id> was modified. Correct rule will be enforced again.	An administrator might have updated a rule enforced by Illumio Segmentation for the Cloud	Ignore the message if changes were made by mistake. This is because Illumio Segmentation for the Cloud automatically fixes the policy. See Tamper Protection [272]. If changes were intended, then review and change the policy in Illumio Segmentation for the Cloud.
CloudSecure authored rule <rule-id> from <azure-firewall-csp id> was removed. Correct rule will be enforced again.	An administrator might have deleted a rule enforced by Illumio Segmentation for the Cloud	Ignore the message if changes were made by mistake. This is because Illumio Segmentation for the Cloud automatically fixes the policy. See Tamper Protection [272]. If changes were intended, then review and change the policy in Illumio Segmentation for the Cloud.
CloudSecure authored Rule Collection Group <rcg-id> from <azure-firewall-csp id> was modified. Correct rule collection group will be enforced again.	An administrator might have modified a Rule Collection Group created by Illumio Segmentation for the Cloud	Ignore the message if changes were made by mistake. This is because Illumio Segmentation for the Cloud automatically fixes the policy. See Tamper Protection [272]. If changes were intended, then review and change the policy in Illumio Segmentation for the Cloud.
CloudSecure authored Rule Collection Group <rcg-id> from <azure-firewall-csp id> was removed. Correct rule collection group will be enforced again.	An administrator might have deleted a Rule Collection Group created byIllumio Segmentation for the Cloud	Ignore the message if changes were made by mistake. This is because Illumio Segmentation for the Cloud automatically fixes the policy. See Tamper Protection [272]. If changes were intended, then review and change the policy in Illumio Segmentation for the Cloud.
CloudSecure authored Rule Collection <rcg-id> from <azure-firewall-csp id> was modified. Correct rule collection will be enforced again.	An administrator might have modified a Rule Collection created by Illumio Segmentation for the Cloud	Ignore the message if changes were made by mistake. This is because Illumio Segmentation for the Cloud automatically fixes the policy. See Tamper Protection [272]. If changes were intended, then review and change the policy in Illumio Segmentation for the Cloud.
CloudSecure authored Rule Collection <rc-id> from <azure-firewall-csp id> was removed. Correct rule collection will be enforced again.	An administrator might have deleted a Rule Collection created by Illumio Segmentation for the Cloud	Ignore the message if changes were made by mistake. This is because Illumio Segmentation for the Cloud automatically fixes the policy. See Tamper Protection [272]. If changes were intended, then review and change the policy in Illumio Segmentation for the Cloud.

Legal Notice

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

Resources

- Legal information
- Trademarks statements
- Patent statements
- License statements

Contact Information

- Contact Illumio
- Contact Illumio Legal
- Contact Illumio Documentation