illumio

ONBOARDING GUIDE

The Zero Trust Segmentation Onboarding Guide

An essential guide to discovering how Illumio accelerates and simplifies the path to Zero Trust with Illumio CloudSecure.

Onboard With Confidence

Welcome to CloudSecure!

You've taken a step forward in your journey to achieving a Zero Trust Segmentation framework. We know this shift can be a complex one – it requires change and adopting a new security product into existing systems.

To help guide you through Illumio CloudSecure onboarding, we created this quick start guide to outline key steps and resources to consider as you get started on your hybrid- and multi-cloud segmentation journey. Ready to get started?

Week 0 | **Onboard CloudSecure**

Plan for your deployment and onboard your various cloud accounts. Onboarding CloudSecure is agentless and frictionless for your application or security teams.

Weeks 1 – 2 | **Discover real-time traffic flows**

Discover cloud inventory, topology, security controls and real-time traffic flows. Gain insights and awareness around traffic patterns in your cloud environment.

Weeks 3 – 4 | **Be ransomware-ready**

Utilize *Organizational Policies* for quick wins by providing "guardrail" policies to quickly allow or deny certain protocols across your cloud infrastructure, based on your implementation of label sets.

Weeks 5+ | **Protect your applications**

Utilize *Application Policies* to define, discover, and write policy sets towards workloads in your cloud, no matter what environment they are in.

illumio

Onboarding Guide Overview

Here's what you can to expect

Organized into weekly steps, our guide provides you with a structured roadmap to walk you through each stage as you onboard, discover, and write policy to protect your environment.

Our agentless CloudSecure onboarding happens seamlessly, without causing any disruption to your application or security teams – and it can take just minutes to onboard a single account, organization, or tenant.

As part of your journey, you'll have access to dedicated support resources and expertise through our Customer Success, Professional Services and Technical Account Management teams who are committed to your success. Our teams will ensure personalized guidance and proactive assistance tailored to your goals every step of the way.

Week 0

Onboard CloudSecure

- **Step 1** Planning Kickoff
- **Step 2** Meet Your Champions
- **Step 3** Cloud Security Review
- **Step 4** Cloud Account Targets
- **Step 5** Onboarding Cloud Accounts
- **Step 6** Flow Log Review

Weeks 1 – 2

Discover real-time traffic flows

- **Step 7** Cloud Map
- **Step 8** Inventory & Traffic Visibility
- **Step 9** Labels 101
- **Step 10** Tag to Label Mapping
- **Step 11** Running Reports

Weeks 3 – 4

Be ransomware-ready

- **Step 12** Policy Writing 101
- **Step 13** Organizational Security

Weeks 5+

Protect your applications

- **Step 14** Deployments & Applications
- **Step 15** Identifying Applications
- **Step 16** Writing Application Policy
- **Step 17** Impact Analysis
- **Step 18** Application Security

illumio

Week 0 | Onboard

Your cloud segmentation journey starts here.

Spending proper time preparing for your segmentation journey is critical – and significantly increases your probability of a successful implementation.

As part of your Illumio CloudSecure onboarding, we recommend mapping out your goals for Zero Trust Segmentation and having clearly defined expectations to ensure you are considering all your technical parameters as part of your detailed project plan. Additionally, we recommend getting buy-in early to keep the project on track by establishing a good communication cadence and regular check-ins with your champions at the project, management, and executive levels who will help communicate the implications of getting to enforcement.

Understanding exactly what is occurring during the onboarding phase and facilitating that understanding across teams is critical to obtaining and managing the necessary permissions related to connect CloudSecure with the related cloud accounts. We'll dive into this in Step 3. If you need additional assistance at any point during the onboarding process, contact your Illumio Account team or visit the Illumio Support Portal for 24/7 access to resources.

What to expect

Step 1 *Planning Kickoff*: Strong project planning is the first step to success – establish targets with cloud account owners, architects, and other stakeholders to demonstrate success quickly and wins to socialize internally.

Step 2 *Meet Your Champions*: Identify and align with your internal champions who will provide the mandate necessary to prioritize your Zero Trust Segmentation initiative for the cloud.

Step 3 *Cloud Security Review*: Review the necessary permissions need to successfully connect CloudSecure to your various cloud accounts. Determine who the appropriate team will be that can complete these initial tasks and gain any internal security approvals needed to proceed. This includes whether CloudSecure can be granted read-write for security enforcement points, or read-only access.

How CloudSecure connects with your cloud accounts:

For Amazon Web Services (AWS), CloudSecure requests the creation of an IAM role within the customer's AWS account. Then, CloudSecure assumes this role to perform actions in AWS such as reading resources and managing policies. This custom role is assigned the "SecurityAudit" standard AWS policy, as well as new policies created by CloudSecure to retrieve resource metadata and VPC flow logs and, optionally, write policies.

For Microsoft Azure, CloudSecure requests a new App Registration be created and assigned roles including 'Reader', 'Storage Blob Data Reader', and a custom role for managing network security groups.

TIP *A complete set of permissions requested by CloudSecure is available through our CloudSecure documentation. A list of CloudSecure public IP addresses is available as well if you wish to place a limit on access to objects like flow logs in storage accounts.*

Step 4

Cloud Account Targets: Identify the cloud accounts you wish to begin connecting to CloudSecure. Determine if you want to proceed at an individual account/subscription level, or a broader organizational/tenant perspective, based on the level of complexity within your environment.

CloudSecure supports onboarding both on an AWS Organization (root account) as well as individual AWS accounts. It does not yet support onboarding AWS Organizational Units (accounts under root can be onboarded without onboarding the root account as well). CloudSecure also supports onboarding both Azure subscriptions and tenants. Both these procedures involve running various pre-populated templates (an AWS Cloud Formation Stack, or a script run from Azure PowerShell) that will require the user to already be logged in to their cloud accounts.

Step 5

Onboarding Cloud Accounts: Onboard your cloud accounts with our agentless solution to begin collecting inventory, metadata, and traffic flows.

Stage 1 | Begin gathering inventory/metadata about your various cloud accounts, and optionally, grant access to manage security policies.

Stage 2 | Once CloudSecure discovers sources of flow logs, it will then trigger a second workflow to grant access to those storage accounts. This second step typically occurs within a few minutes of onboarding an account, as well as continuously if new sources of flow logs are discovered by CloudSecure.

TIP *If you are onboarding at a broader Organization/Tenant level, you will be provided an option to select which of the individual accounts/subscriptions you wish to sync data with CloudSecure.*

Step 6

Flow Log Review: Review where CloudSecure has discovered sources of flow logs and analyze whether this matches expectations, or if additional cloud accounts need to be onboarded as well.

TIP *From the 'Flow Log Access' button, you will see a list of all the flow log sources that CloudSecure has discovered to date, and whether access to those logs have been granted. Review these sources of flow logs to determine if this is adequate coverage for the resources being managed. At times, customers might send flow logs to accounts/subscriptions that are separate from where the resources reside. So, this is a good opportunity to review your onboarding strategy to see if other accounts/subscriptions need to be included.*

Weeks 1 – 2 | Discover

Discover real-time traffic and lay the foundation with labels.

Onboarding your accounts can be an eye-opening stage of your deployment as you begin to see our hierarchical view of all the resources CloudSecure discovered. Our Cloud Map lays out resources in a cloud-native structure, showing the various regions such as VPCs, VNets, and subnets. From one window, you'll gain visibility to your entire cloud landscape – and may even be surprised by a few things!

As part of this phase, we recommend exploring the various visualization tools to see the data collected with the *Inventory* and *Traffic* tabs. Expand this view and see how the resources are connected to other resources. Once you understand the data, we'll guide you through how to use our tag-to-label mapping tool to set the foundation for setting your tagging strategy.

Wrap up your Discover phase by running your first reports from the *Traffic* tab and leverage the results of the report to understand the 'risky' protocols and to begin conversations around ways to protect and secure your environment. We'll dive into reporting at a deeper level in the next "Be Ransomware Ready" phase so think of this as an introduction.

What to expect

Step 7 *Cloud Map*: Explore the CloudSecure UI and review the Cloud Map to get a hierarchical view of all the resources discovered and their traffic patterns.

The default Cloud Map view will highlight all regions with discovered resources. From this top-level view you can expand on each region to go further down the hierarchy where you'll see VPCs, subnets, and eventually individual resources.

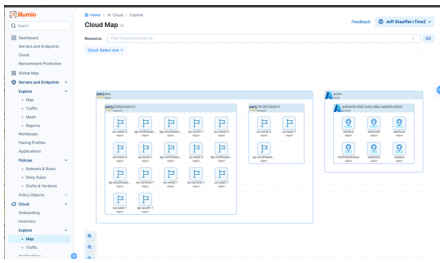

TIP *Clicking on a resource will open a side-panel that provides metadata about that object and associated objects. It also provides a Traffic table showing all flows from this object.*

TIP *Using the Filter panel is a great way to highlight objects wherever they reside, which will also overlay a set of animated traffic lines showing you what's going on in your environment. Try filtering on a region, or even a particular cloud to see how this works.*

Steps 8 – 11

Step 8 *Inventory & Traffic Visibility*: Explore other ways to visualize data collected by trying out the Inventory and Traffic sections, for a tabular view of resources and flows.

TIP *From the Traffic tab, utilize the filtering capabilities to isolate IP Addresses, protocols, and flow status (i.e., allowed vs. denied).*

TIP *From the Inventory tab, you can see a list of all the resources discovered from your cloud accounts. Check out the 'Resource Graph' option to get an expanded view of how resources are connected to other resources. For example, it's not an easy task to figure out what all virtual machines are using a particular Security Group natively in your cloud provider's portal. CloudSecure has performed all the necessary tasks to understand these associations and their inter-connectedness. You can quickly visualize this, for example, by selecting an AWS Security Group and viewing what object types have associations to this group. Then, you can expand to view all the interfaces and instances that also share this Security Group.*

Step 9 *Labels 101*: You may have already noticed an array of tags coming from your cloud environments. It's time to begin discussing your internal processes for identifying workloads and their purpose. Learn how to use our tag-to-label mapping tool to consolidate your various tagging schemes into one set of CloudSecure label types that can be used to begin authoring policy.

CloudSecure uses a multidimensional label system where each label has its own independent usage in the policy model. You can create your own custom label types such as 'Role', 'Location', and 'Operating System'. There are two fundamental label types in CloudSecure – Application and Deployment, which will be addressed in a later step.

TIP *If your organization is already using tagging strategies in managing your cloud, you can easily incorporate these schemas with the CloudSecure Tag-To-Label-Mapping tool:*

- *Map the tag key to a CloudSecure label type and have the tag value flow through to populate the CloudSecure label value.*
- *Perform a one-to-one mapping, or map multiple tags to a single label type. Perhaps you have multiple tagging schemas defined in various part of your organization, or have made acquisitions, as examples of many-to-one mappings.*

As resources come and go in your cloud environment, this label mapping exercise happens automatically and will appropriately adjust how resources are labeled in CloudSecure.

Step 10 *Tag to Label Mapping*: CloudSecure uses labels as the mechanism to apply policy to your cloud resources, thus, it's important to have labels associated with a significant number of your resources. Create some Tag-to-Label mapping rulesets to begin generating CloudSecure label types. You can create rules that just focus on certain cloud accounts as part of your matching criteria as well.

TIP *Once you have some labeling constructs in place, you can use this additional metadata in Cloud Map queries. For example, you could search on a 'Role' label of 'database' and view all resources across your cloud infrastructure that have traffic flows to/from any database, across your multi-cloud environment.*

Step 11 *Running Reports*: Based on traffic data gathered by CloudSecure run a report from the Traffic tab to determine what protocols are running in your environment that are considered 'risky' from the perspective of threat targets. Begin conversations around what might be simple ways to begin protecting your environment.

STAY TUNED as we expand on this important area of various reports available to download from CloudSecure! You can begin by viewing an exhaustive list of what CloudSecure considers to be 'Risky Ports' that deserve an early focus, such as RDP, SMB and RPC protocols. You can view this by top source/destination, number of flows, etc. to better understand their usage in your environment.

Weeks 3 – 4 | Be ready

Create and provision policies, simply.

To fully maximize the value of Illumio, establishing a strong groundwork is crucial. As you continue your journey, our primary aim is to empower you to proactively protect your environment against potential cyber threats. Critical to this is guiding you through creating and implementing policies with CloudSecure.

The focus of this phase and the next is to provide you with an overview of our policy framework, highlighting key aspects of CloudSecure, such as policy attributes and the nuances between organization and application policies, as well as how to write both varieties effectively.

We realize that policy writing is an iterative process, and our team is ready here to help. If you need additional information around these topics and want to go deeper, we recommend our documentation:

- CloudSecure Policy Model
- CloudSecure Policy Attributes
- Organization Policy versus Application Policy

Additionally, for any queries or clarifications, we recommend tapping into the expertise of our community by engaging with our team on the on Illumio Community.

What to expect

Step 12 *Policy Writing 101*: Use information gathered in previous weeks – from reviewing traffic flows, risky ports, and labeling strategies to produce a plan as to how to begin implementing security policy.

Review how CloudSecure takes your **policy intent** and converts into security rulesets.

There are two primary paths available to you at this point:

- **Organizational wide policy**: CloudSecure provides a simple way to apply an organization-wide policy, based on labels you've already defined. This policy can include allow or deny rules. Perhaps you have certain corporate mandates in place around security policy. This is one quick way to obtain compliance with blocking certain protocols that are no longer in use anywhere in your cloud deployment.
- **Individual, application-based approach**

STAY TUNED as we'll dive into the individual, application-based path in the next few steps.

To help streamline security management, CloudSecure handles resources associations within security policies:

- CloudSecure maintains knowledge about the resources associated with the labels you reference in your security policy.
- CloudSecure effectively manages the association of security enforcement points to these resources.

This detail is abstracted from the user, freeing you to concentrate on authoring policy intent and entrusting the platform to determine the necessary actions required to enforce policies effectively across your cloud environment.

Step 13 *Organizational Security*: Provision your first Organization-level policy.

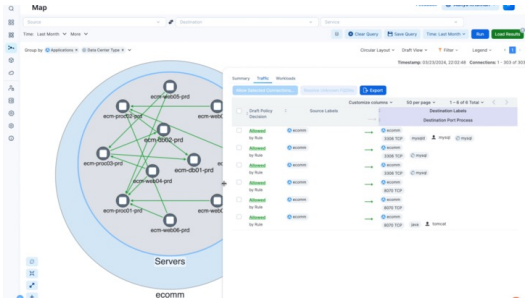

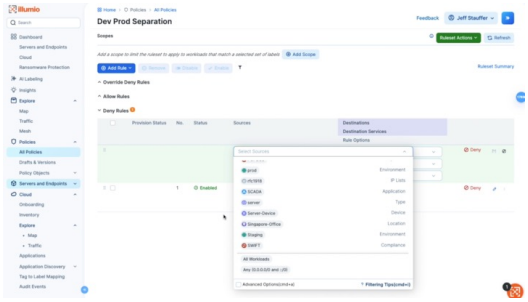

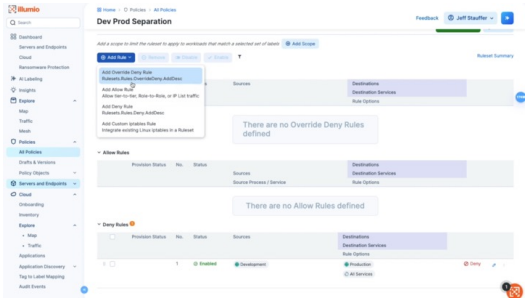

Start with a simple policy, perhaps one that blocks protocols that are no longer in use in your environment – RDP, NetBios, SMB are some common ones that are considered 'risky' ports.

NOTE *Keep in mind that security groups rule counts can be a limited resource especially in AWS environments where the default rule count is 60. So, make sure to use granular policy here, with perhaps incorporating "Any (0.0.0.0/0)" as a source address as to reduce the permutations of src/dst addressing.*

Weeks 5+ | Protect

Take the next step in your journey toward Zero Trust.

CloudSecure operates on the foundational understanding that your cloud infrastructure comprise various applications deployed across diverse environments. For instance, you might have a payment app that runs both production and development instances or an HR application that has production, development, and staging environments.

With CloudSecure, we offer tools to identify these applications, pinpoint their locations, and manage security policies at granular levels.

As you progress through the final stage, we'll walk you through leveraging Application Policies to define, discover, and write policy sets for workloads in your cloud, regardless of environment.

What to expect

Step 14 *Deployments & Applications*: Explore options for how CloudSecure can assist with identifying various Applications as well as where they're deployed within your cloud environment. Review how CloudSecure creates 'Deployments' and 'Applications'. Create a gameplan for managing Applications.

DEPLOYMENTS Based on how you define where you roll out your various applications, CloudSecure will map these using our Deployment label. Create deployments such as Production, Staging, Engineering, QA-Team-East, or using your organizational structure. Deployments can be defined based on:

- Cloud accounts
- Regions
- Virtual networks / VPCs / subnets
- Cloud tags

NOTE *Defining deployments is optional. CloudSecure will not apply a 'Deployment label if a found resource does not match any of the Deployment definitions you've created.*

APPLICATIONS CloudSecure provides the ability to manage policy for individual Applications. Similar to the Deployment definition workflow, you can also create Application definitions based on similar metadata such as cloud accounts, VPCs/VNets, or cloud tags. Once you've defined the characteristics, CloudSecure will scan all discovered resources to identify objects that match your criteria. It then generates an 'Application' label that is associated with all matching resources.

APPLICATION DISCOVERY RULES We recognize that you may have numerous Applications running in your environment and that it's not practical to create these definitions individually. We provide an automated mapping option called *Discovery Rules* that can take advantage of standard tagging schemes you may already have in use to automatically generate multiple 'Application' labels based on this pattern. For example, let's say you have three applications developed and their resources are marked with a cloud tag of "App:HR", "App:Payroll", and "App:Finance." You can create one rule that simply maps any value of a tag called "App" to a new Application label, thus creating three Application labels in one easy step: HR, Payroll, and Finance.

Steps 15 – 16

Step 15 *Identifying Applications*: Identify what Applications are good candidates to begin authoring policy.

TIP *Start small but think big! Start with a well-defined application that you can utilize to try out the Application Definition process. Is this application contained to a particular virtual network? A cloud account? Or are there well-documented tags in place to identify resources?*

Once you can communicate to CloudSecure what defines the 'boundaries' of this application, CloudSecure will scan your cloud resources to determine which ones meet the application criteria and categorize them by whatever deployment definitions you've already created. This way, you can begin to view inventory, traffic maps, and policy for each of the various environments where you have this application deployed.

You can also utilize the 'Discovery Rules' option if you have many applications that adhere to a particular naming convention: Based on the general rule you define, CloudSecure will create multiple Application definitions based on the pattern you've provided. For example, if all of your applications use a simple cloud tag of "App:<application-name>", CloudSecure can create one application definition for each value found in that tag.

Now that you have an Application defined, you can begin to utilize the same visualization tools that were available at a global level, but now with a very granular view of just this Application. Check out traffic tables, inventory, and a Cloud Map diagram all focused just on this Application!

Step 16 *Writing Application Policy*: Define your first Application within CloudSecure and author security policy to be implemented.

There are a few important concepts to point out as you begin to write your Application Policy:

- *Policy Scope*: As you view individual application policy pages, you'll notice that there is a mechanism to control the scope of where these rules are applied. You can select individual deployments that you defined, or you can select the "All Deployments" option. Selecting 'All Deployments' will create a copy of any rules written and apply to each of the specific deployments available.
- *Restrictions*: You'll notice that the destination application field is restricted to the application that you are currently editing. The intention here is that you can only control what resource is accessing your application. Accessing some other application should be written from that application's policy page.
- *Intra- and inter-Application Policy*: If you're writing rules for intra-application policy, the rulesets will get applied as outbound rules to the source workload, and inbound rules to the destination workload. If you're writing inter-application rules, these will only apply inbound rules to the destination workloads. Consequently, matching rules would also need written from the source application policy page to complete the process.

If you have policies that you want applied across all the various deployments, make sure to utilize the "All Deployments" view within the application, so that a copy of your rulesets get pushed into each application environment.

TIP *If you have policies that you want applied across all the various deployments, make sure to utilize the "All Deployments" view within the application, so that a copy of your rulesets get pushed into each application environment.*

Steps 17 – 18

Step 17 *Impact Analysis*: Look at the 'Impact Analysis' to help analyze how your policy rulesets will affect existing security points.

NOTE *Before you provision a policy, you may wish to gauge what its impact will be. Within the policy authoring page there is a "Show Impact" button that will display for you a list of what security enforcement points will be affected by this change. This view will display your draft rules along with any other rules already provisioned from other CloudSecure sources so that you can properly determine what this rule change will look like on any individual device. Only CloudSecure-written rules will be displayed at this time.*

Step 18 *Application Security*: Provision your first Application Policy. This is a big step towards a Zero Trust Segmentation journey!

TIP *Review your procedures, determine how processes might be improved, and begin the journey to expand your segmentation coverage across your multi-cloud environment.*

Applying security policy is an iterative process. As you wrap up the onboarding guide, consider which additional applications and workloads could benefit from protection to further advance in your Zero Trust journey.

We're here to help.

Get support on identifying issues, making software fixes and more.

Our Technical Support team is available online and offline 24/7 for your Illumio solution.

Visit our Support Portal >

- Open a new support case
- Email support@lllumio.com

Direct hotline support (Toll Free)

- Australia: 1 800 995 469
- France: 08 00 94 00 75
- New Zealand: 0800 369 529
- Singapore: 800 321 1271
- United Kingdom: 0800 069 8795
- United States: 1-888-631-6354

Direct hotline support (Toll)

- United States: +1-408-831-6354
- International: +1-408-831-6354

About Illumio

Illumio, the Zero Trust Segmentation company, stops breaches and ransomware from spreading across the hybrid attack surface. The Illumio ZTS Platform visualizes all traffic flows between workloads, devices, and the internet, automatically sets granular segmentation policies to control communications, and isolates high-value assets and compromised systems proactively or in response to active attacks. Illumio protects organizations of all sizes, from Fortune 100 to small business, by stopping breaches and ransomware in minutes, saving millions of dollars in application downtime, and accelerating cloud and digital transformation projects.