illumio

WHITE PAPER

Migrating to the Public Cloud

Table of Contents

illumio

Use Case: Achieving Secure Cloud Migration with Illumio

Overview

Business Drivers

Organizations across industries are moving many of their applications to public clouds. With the cost of cloud computing and storage coming down significantly over the past few years, IT leaders see the value of not owning infrastructure and migrating at least some of their workloads to the public cloud.

While moving workloads to the cloud may provide business and cost advantages, ensuring application availability and security is critical to business continuity. Businesses still take a cautious approach to cloud migration because they lose the controls that they would otherwise have within their own data centers. Additionally, the differences in security considerations for workloads in the cloud have slowed adoption.

Moving workloads to the cloud lets enterprises:

- Take advantage of ondemand infrastructure
- Provide better availability, disaster recovery, and quality of service
- Scale applications without purchasing new hardware
- Reduce the cost of operations (people, processes, and facilities)
- Easily expand geographically with the ability to put a data center into operation, anywhere in the world, at any time
- Improve business agility

Current Challenges with Public Cloud Migration

- Security policies need to be reconfigured when applications migrate to public cloud.
- Virtual private clouds (VPCs) used for isolating applications and application tiers increase management complexity and introduce delays.
- Manual modifications to security policies for application changes across private data center and public cloud increase the risk of errors.
- Rudimentary security solutions offered by public cloud providers do not provide the enterprise-grade security controls available in private data centers.
- Public cloud providers offer different security solutions limiting application portability to other providers.
- Security solutions lack visibility to workload context and cannot adapt to application changes. Business is slowed down by security changes.

The Illumio Solution

- The Illumio Adaptive Security Platform® (ASP) enables security enforcement to follow workloads when they migrate to public cloud without any manual reconfigurations or customizations required for the cloud platform.
- Application security policies are consistent across public cloud and private data centers.
- Context-driven security policies are associated and enforced at the most granular level—from workload inception at scale to decommission.
- Continuous computation of application context ensures accurate enforcement of security policies regardless of workloads running in a private data center, public cloud, or hybrid cloud.
- Illumio ASP enables security to operate and scale at the same speed as the application infrastructure.

Current Approaches to Migrating to the Public Cloud

Limited Use of Public Cloud for Non-Critical Applications

Controls in the private data center do not work in the cloud. For instance, VLANs often provide security isolation for individual applications in a private data center, but it is difficult or impossible to replicate that VLAN infrastructure in a public cloud, where an organization does not own the infrastructure.

Without VLANs, it is difficult to replicate the firewall controls an organization has in its own data center—therefore security is often the biggest roadblock to cloud migration.

VLANs, zones, subnets, and even SDN are tied to the network. But having security controls tied to the network is probably the biggest roadblock to organizations achieving the agility that they can gain from migrating applications and workloads to public clouds.

Because controls are not in place, many organizations limit their use of the public cloud to:

- Ephemeral workloads for publicly facing apps (e.g., a consumer company runs an ad and drives users to a website that is hosted in a public cloud)
- Test and development environments

In both cases, the workloads placed into public cloud are not mission critical and do not carry sensitive data, therefore lack of stringent security controls, application separation, and workload isolation is acceptable.

Virtual Private Clouds

VPCs are effectively an extension of the organization. One way to think of a VPC is as a balloon that can be inflated within a public cloud environment. The balloon is created by placing a virtual gateway appliance (such as a virtual appliance firewall) in the public data center, and that appliance becomes the default gateway for all workloads that spin up inside the VPC. A VPN is then created from the edge of the enterprise's data center (from its firewall) into the virtual appliance firewall in the public cloud provider. Workloads can be spun up inside the VPC, but traffic generally comes back to the organization's private data center.

After basic connectivity has been established, traffic flows freely, however it is incumbent upon the organization to create rules on their gateway devices that reflect any security policies necessary for the virtual private cloud.

Using a VPC gives an organization the ability to deploy workloads on demand. This is very useful for test and development groups, which are not mission-critical and don't require compliance and controls. However if an organization wants to replicate the isolation, compliance and controls that they enjoy in their own data center, then deploying VPCs becomes very complex.

Here are some ways that organizations migrate workloads to a VPC:

- Basic application isolation: Gateway firewalls are configured to allow specific types of traffic flows from an outside (untrusted) network into a trusted area (zone or subnet). The firewall isolates the application from the outside world. One example of this is creating a VPC per application. This works well if an enterprise has few very simple applications. However as an enterprise deploys more applications, they require more VPCs, which become more and more unmanageable.

- Application grouping: If an enterprise groups applications by "classes" or types of applications, it can "tuck" similar groups of applications into the same VPC, which reduces VPC explosion. Application predictability can become problematic in this scenario because the virtual firewall appliance becomes the chokepoint. In an enterprise, the firewall is usually hardware based and provides highly predictable performance. However, if performance and scalability needs increase, the VPC may need to be split to address the scalability need. In addition, application grouping reduces VPC proliferation but does not provide application isolation. So the tradeoff is reduced complexity but reduced security.

Many cloud providers are trying to give organizations some of the controls that they enjoy inside their own data center, but these controls are not the equivalent and are generally locked to proprietary APIs that tie the organization to the cloud provider.

Four Challenges with Existing Solutions

1. Gateways limit application scale out and movement

Since the gateway is disconnected from the workload, it lacks context for and visibility into what is happening behind the perimeter. Application scale and motion will break gateway appliances. As workloads move and scale up or down, the firewall policies need to be updated. Moreover, firewall rules are written in terms of the network, but in the public cloud, organizations have limited control over their IP addresses and zones.

2. Policies are asymmetric across private data center and public cloud

Security policies that are tied to network chokepoints are not portable. With current solutions, security policies have to be rethought and rewritten when workloads move to the cloud. Going down this road creates asymmetric security policies (i.e., security policies enforced inside the private data center are different from the security policies and controls inside of the cloud providers.) Asymmetric policies are prone to human error, since policies inside the private data center have to be manually mapped into the cloud providers. Any changes in the security policy need to be remapped in multiple places.

3. Complexities with securing inter-application traffic

As soon as there is the need to separate different tiers of an application (e.g., separating database from web), another layer of gateways must be introduced into the public cloud. This requires a separate VPC since the enterprise cannot control traffic steering.

In this example, the gateway to the database tier must then be updated with all of the rules related to the web tier. In addition, the database VPC has more connections going into and out of it, which means the gateway firewall appliance once again becomes the chokepoint. If there are multiple applications that need to connect to the database, or if the database needs to synchronize with another database, the rules need to be put in place that allows that synchronization across VPCs.

The problem with all of these methods is that speed and scale in a private data center are impediments that get in the way in public cloud. Security becomes the biggest barrier to successful cloud migration.

4. Cloud security is not workload security

Public cloud vendors prioritize securing the underlying cloud fabric, whereas securing the workloads deployed onto that fabric are the responsibility of the customer. The cloud vendor secures the cloud, but securing the workloads is not their focus. The result is a weak set of workload security tools offered by the cloud vendor.

The Illumio Solution

The key to migrating workloads into the cloud is to eliminate the constraints imposed by networking without compromising security and controls.

The Illumio Adaptive Security Platform (ASP) solves this problem by attaching security and controls at the workload, eliminating the need for VLANs and zones, and allowing workloads to be migrated to the cloud without a chokepoint. This architecture provides application isolation without tying security to the network infrastructure.

Illumio ASP delivers enforcement at the workload via the Virtual Enforcement Node (VEN). The VEN is not in the data path. Rather, it enforces policy using the instruments that are in the operating system—specifically, IP Tables (for Linux operating systems) and Windows Filtering Platform (WFP) for Windows workloads.

Policy is computed using the centralized Policy Compute Engine (PCE), which receives contextual information about workloads as telemetry from all of the VENs. The PCE takes that live telemetry data to create an application dependency map. This map visually displays traffic flows between applications and workloads. It baselines an application's connectivity so that security teams can right-size the segmentation policies for their applications, as well as detect anomalous behavior.

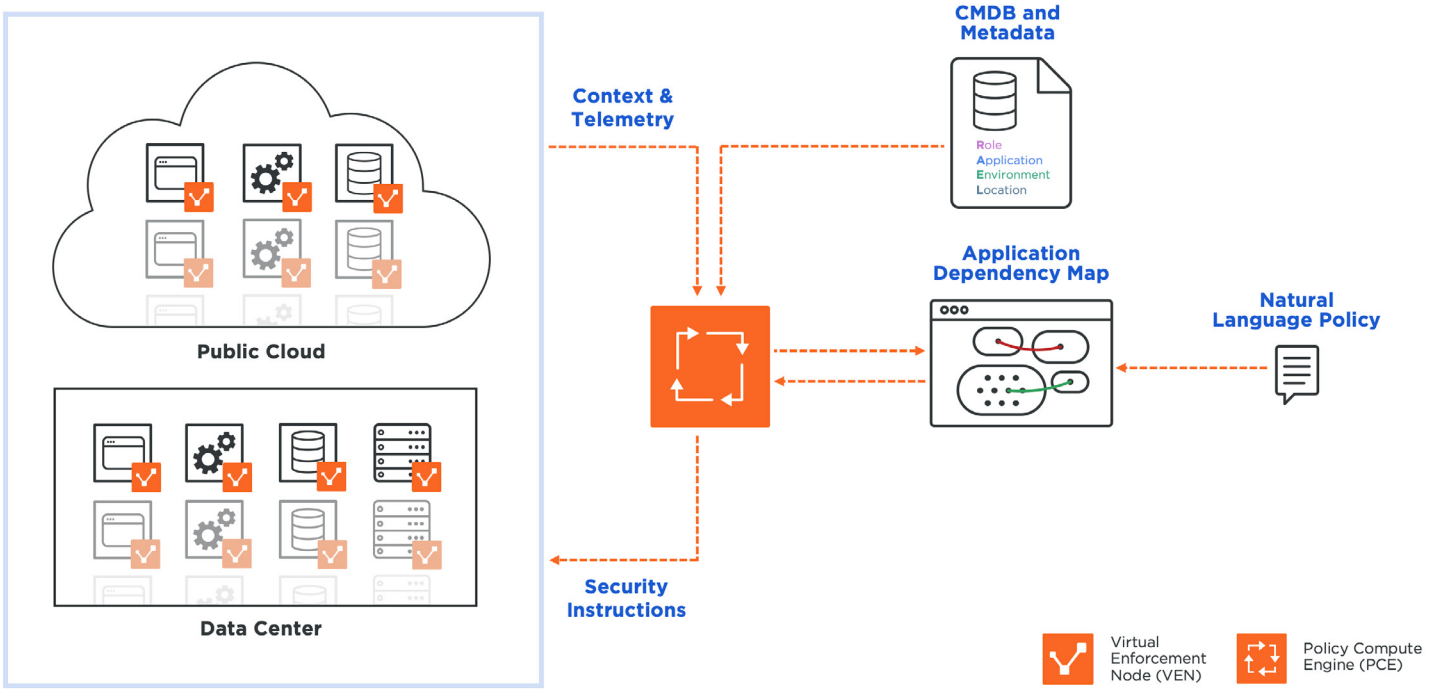

With Illumio ASP, when an organization wants to move an application into a public cloud, it can use the exact same security policy that was in place inside of its own private data center. This gives an enterprise a uniform approach to security, and in every location, the workloads and the application the workloads comprise are secure.

Here are six ways Illumio ASP helps enterprises with their cloud migration.

1. Granular security that moves with workloads

Illumio ASP provides enforcement at the individual workload rather than using the network or an artificial gateway inside of a cloud service provider. This allows enterprises to create a protection profile around a workload, or set of workloads, inside their existing data center or in a public data center—or they can split tiers of an application between their own data center and a cloud provider. The approach gives the enterprise complete autonomy in choosing where and how their applications will be deployed without having to worry about the security controls.

2. Context-aware security enforcement

Illumio ASP is fully context aware. It understands the context of each workload (the intrinsic properties and relationships to other workloads), which ensures that the PCE gets the right security answer every time. A flexible, multi-dimensional labeling mechanism is used to define a workload based on its role (database, web server, mail server etc.), the application that it serves (Payroll, Sales etc.), the environment it runs in (dev, test, production, etc.), and its location (US, Atlanta, Rack #3, etc.). All dimensions can have an infinite depth: as they need more labels within a dimension, they can simply be added.

The Illumio PCE maps the labeled workloads with the policies, specified in natural language to dynamically compute workload specific rules using the telemetry provided by the individual VENs. This allows security policies to be resilient to changes to applications or the underlying network infrastructure. Illumio provisions the most accurate security policies for every workload—regardless of it being in a private data center, or in any service provider network.

Because the security is attached to each workload, it eliminates the need to rely on gateway devices to enforce workload and application security. If any legacy segmentation techniques are in place, Illumio ASP can still be used with full functionality. There is no need to rearchitect an existing VPC in a cloud service provider, or change an enterprise's VLAN, zones, or even SDN architecture.

3. Operates at the speed and scale of application infrastructure

Illumio ASP allows an enterprise to operate at speed, meaning it doesn't have to configure and deploy a gateway device that may have non-deterministic scaling characteristics. Illumio ASP offers predictable performance since it is instrumenting capabilities that are already in the kernel. Without the context provided by the VEN, it would be difficult to manually instrument thousands of individual iptables rules. Illumio ASP can scale to hundreds of thousands of workloads.

4. Security with no dependency on IT infrastructure

By enforcing security on the workload, Illumio ASP completely decouples security from the underlying network infrastructure. An enterprise does not need to rearchitect or change their existing segmentation technology or network topology. They can simply integrate Illumio ASP and then allow their separation technology to provide connectivity and forwarding fabric—the capabilities those solutions were designed to offer. This enables enterprises to secure applications running on bare-metal servers and VMs across private data centers and public cloud infrastructures including AWS, Azure, Rackspace, Google Compute, Oracle Cloud, and IBM Cloud.

5. Segmentation for containers and serverless compute resources in cloud

In addition to virtual machine workloads in the cloud, Illumio supports segmentation for containers and serverless compute resources. In addition to the VEN agent used for bare-metal and VM compute resources, Illumio offers a containerized instance of the VEN agent, to enable segmentation capabilities around constructs used in Kubernetes and OpenShift: namespaces, pods, and services.

Illumio also supports a rich set of APIs and offers suggested scripted API-driven solutions to enable the configuration of AWS Security Groups for tools such as RDS using a community-supported model via Illumio Labs: https://labs.illumio.com.

6. Application visualization and policy validation

Illumio ASP's application dependency map, Illumination, gives you live visibility of applications, workloads, and flows to help understand application behaviors within a cloud provider, and across multiple cloud providers and legacy data centers. Illumination is the basis for building and testing segmentation policies. The map enables you to visually model policies and provide immediate feedback about any flows that may be blocked when moving into enforcement mode. This capability provides assurance that segmentation will not break an application even as workloads and applications are migrated. It also helps to validate that segmentation policy is being applied consistently before, during, and after migration.

Use Case: Achieving Secure Cloud Migration with Illumio

To understand how to use Illumio ASP to migrate workloads to the cloud, consider a three-tier customer processing application with the following security constraints:

- Apache service provided by the web tier is open to the Internet.
- Tomcat service provided by the processing tier is used by the web tier.
- PostgreSQL service provided by the database tier is used only by the processing tier.

Using Labels in Security Policies

Illumio ASP allows administrators to create a library of labels unique to their environment. These labels are used to describe the role, application, environment, and location for every workload and can be automatically assigned as part of pairing the workload (i.e., bringing them under management). If an organization does not understand how an application works, it can use Illumination to find the workloads that comprise an application, label those workloads, and build the rules.

Generating Pairing Profiles

The Pairing Profile is a configuration template that specifies the labels to be applied to newly instantiated workloads. The Pairing Profile can also be used to generate the unique pairing keys used to identify newly instantiated workloads to the PCE. When the new workloads are paired, they then acquire the labels and associated security policies.

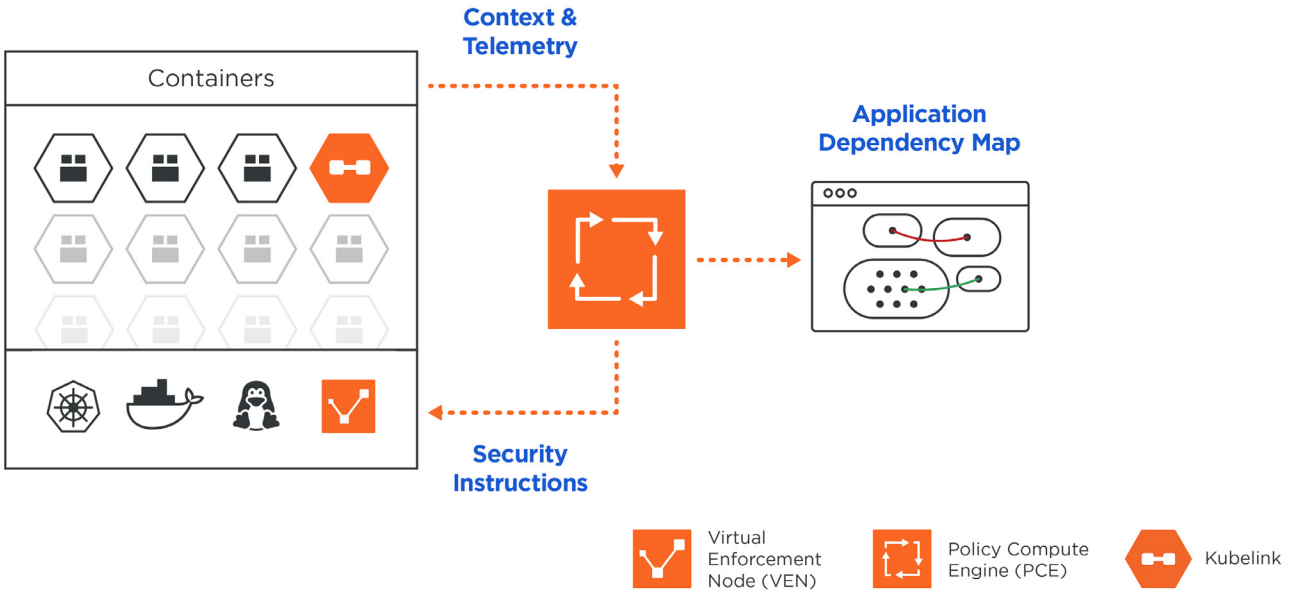

The following image shows the Pairing Profile for the customer processing web tier. This Pairing Profile will be used to generate the pairing key and script.

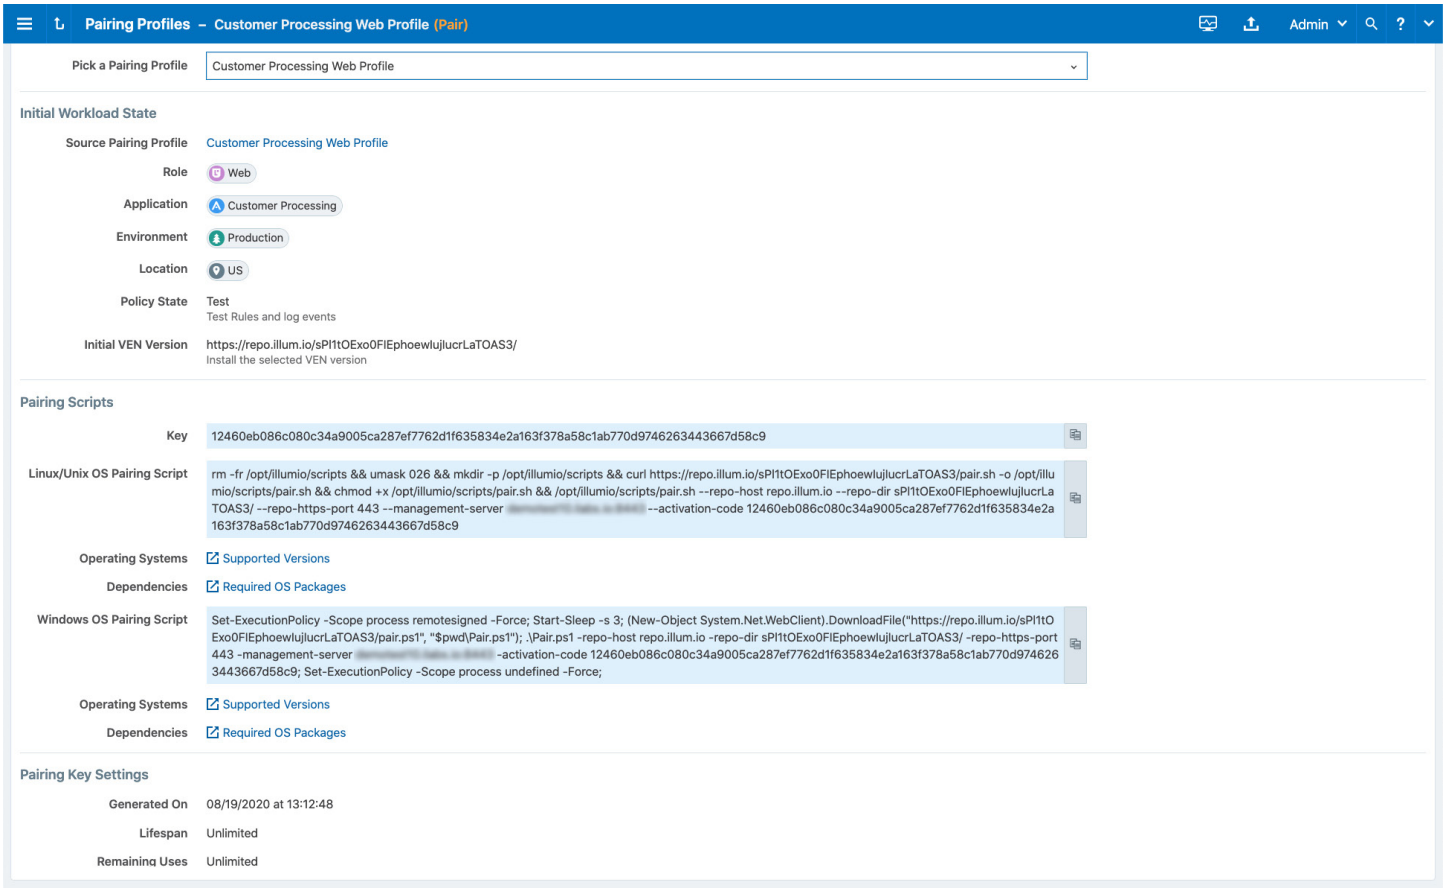

If specific Pairing Profiles have not yet been created, then individual workloads can be paired using a default profile and manually relabeled later.

Labels and Workload Identification

The three tiers of the sample customer processing application are labeled as follows:

	Role	Application	Environment	Location
Web Workloads	Web	Customer Processing	Production	US
Processing Tier	Processing	Customer Processing	Production	US
Database Tier	Database	Customer Processing	Production	US

Writing Natural-Language Security Policies Based on Labels

Once the workloads have been labeled, security policies can be written to capture the explicitly allowed interactions (whitelisted policies) between the workloads. Interactions that are not captured are simply denied. Illumio enables the use of natural language to write security. Rules within Illumio ASP are written the way a developer is likely to think about the application. For instance, in the case of the customer processing application,anyone from the Internet can access the web tier because that is where signups happen. The web tier uses the processing tier and the processing tier uses the database tier.

The figure below shows the ruleset that describes the relationships between the workloads of the customer processing application.

- Only the Apache service running on the web servers will be accessible from anywhere.
- The Tomcat service running on the processing servers will be accessible from the web servers.
- The PostgreSQL service running on the database servers will only be accessible from the processing servers.

CUSTOMER PROCESSING: PRODUCTIONS: US

Scope
Customer Processing: Productions: US

Rules

Service	Provided By	Used By
Apache	Web	Anything
Tomcat	Processing	Web
MySQL	Database	Processing

The Scope of Security Policies

The scope identifies the set of workloads to which security rules apply. In the above example, the rules are applied across all the workloads of the customer processing application running as part of the production environment in the United States.

At this point, there is an "application container" in place. If a new instance of the customer processing application is brought up in (or migrated to) a public cloud infrastructure (in this case, AWS), the only change required would be to modify the scope of the ruleset to include the new location as indicated below:

- Only the Apache service running on the web servers will be accessible from anywhere.
- The Tomcat service running on the processing servers will be accessible from the web servers.
- The PostgreSQL service running on the database servers will only be accessible from the processing servers.

CUSTOMER PROCESSING: PRODUCTIONS: AWS

Scope
Customer Processing: Productions: US
Customer Processing: Production: AWS

Rules

Service	Provided By	Used By
Apache	Web	Anything
Tomcat	Processing	Web
MySQL	Database	Processing

Securing Applications Launched Using an Image

An image (e.g., AMI in AWS or VM templates in VMware) or ISO with the VEN "baked in" is used for instantiating new workloads of the customer processing application. The workloads will be automatically configured with the labels and associated security policies as soon as they connect (i.e., pair) with the PCE.

The following procedure is used to prepare the workload before the creation of an image:

- Create a pairing profile for the workload from the web UI or using the RESTful API.
- Copy and edit the generated pairing script to replace all occurrences of "pair.sh" with "prepare.sh".
- Execute this script on the workload to install and program the VEN to use the pairing key.
- Generate an image from this workload.

Any time a new workload is spawned from this prepared image, it will check in with the PCE for any labels that were ascribed to it during the prepare process. This gives an enterprise the ability to spin up the application anywhere and ensure that the correct security policy is in place—regardless of location or cloud provider.

Integration with DevOps Tools

For organizations using orchestration to spin up workloads, then getting the proper pairing key can be part of the recipe for a workload. For instance, if using Puppet, the recipe can call the Illumio PCE Rest API to get the right pairing key for a workload. This ensures that wherever a workload is spun up, the right pairing profile is literally attached to the workload.

Conclusion

While the benefits of moving to the cloud are many, getting there can be a daunting task. Ensuring application availability and security is critical. Secure and efficient migration requires knowing (and seeing) how your applications work and communicate, and ensuring security moves with your applications and workloads before, during, and after migration.

Illumio enables you to implement workload segmentation without dependency on the underlying cloud fabric security tools or networking constructs. You have a choice to automate workload segmentation directly at the host and containerized construct level, as well as integrate with automation tools provided by your cloud vendor. And a real-time map to visualize all workload connections and flows. The end result is segmentation that consistently follows the workload across network segments to the cloud, between clouds, and across hybrid cloud architectures.

Illumio enables organizations to realize a future without high-profile breaches by preventing the lateral movement of attackers across any organization. Founded on the principle of least privilege in 2013, Illumio provides visibility and segmentation for endpoints, data centers or clouds. The world's leading organizations, including Morgan Stanley, BNP Paribas, Salesforce, and Oracle NetSuite, trust Illumio to reduce cyber risk. For more information, visit www.illumio.com/what-we-do