illumio

WHITE PAPER

ELIMINATING FIREWALL RULE PROLIFERATION

CONTENTS

OVERVIEW 3
- Business drivers 3
- Current challenges with firewall rule proliferation 3
- The Illumio solution 4

CURRENT APPROACHES TO ELIMINATING FIREWALL RULE PROLIFERATION 5
- Manual security audits and change management 5
- Automated firewall rule management tools 5

SIX KEY CHALLENGES WITH EXISTING SOLUTIONS 5
- 1.Tool proliferation 5
- 2. Errors and delays 5
- 3. Increased security risk 6
- 4. Performance impact 6
- 5. Compliance issues 6
- 6. Managing firewalls from multiple vendors 6

THE ILLUMIO SOLUTION 7
- 1. Security that automatically adapts to changes 7
- 2. Security policies that are easy to write—and understand 8
- 3. Precise security policies with no redundancy 8
- 4. Simplified rule management 8
- 5. Simplified compliance management 9
- 6. Application workloads to be secured right from their inception 9
- 7. Simplified security audits 9
- 8. Integration with DevOps tools 9

USE CASE: USING ILLUMIO TO SECURE APPLICATIONS WHILE REDUCING FIREWALL RULES 10
- Writing security policies in natural language 10
- Generating Pairing Profiles 11
- Labels and workload identification 11
- Writing security policies based on labels 11
- The scope of security policies 12
- Security policies with Illumio vs. firewall rules 13

ABOUT ILLUMIO 14

OVERVIEW

BUSINESS DRIVERS

Organizations that maintain firewalls with thousands of rules aren't just facing technical challenges, they're also courting serious business risks.[1] Rule-base complexity contributes to configuration errors, conflicts, redundant rules, and security risks. The manual audits and verification processes to manage this complexity can slow down application rollouts and create unnecessary costs for business.

Enforcing security using network-based appliances relies on bringing traffic from the workloads to the enforcement point, where statically configured IP-based firewall policies are applied. Depending on how much motion, scale, and change there is in a data center, the underlying policies and rule bases can become very large and complex. Many of these rules and associated objects are also out-of-date, simply because the enterprise cannot keep up with the manual configurations required to match the changes in the underlying infrastructure. The sheer complexity can lead to mistakes, especially when it comes to multiple firewalls (possibly from different vendors) with complex rule sets. This complexity can result in:

- **Redundant rules** where one rule duplicates all or a portion of the access permitted or denied by an existing rule.
- **Shadowed rules** resulting from incorrect ordering of rules in a firewall rule base, which prevents the execution of one or more firewall rules.
- **Overly permissive rules** that allow more access than necessary to meet the stated business requirements.
- **Orphaned rules** that allow access to resources that have been decommissioned or migrated.

Over time the rule bases grow to be so complicated that redundant rules are not disabled or removed for fear of breaking an existing application. Rule explosion ultimately costs the business dearly due to resources spent managing and auditing obsolete rules, potential error, performance degradation, and security risks.

CURRENT CHALLENGES WITH FIREWALL RULE PROLIFERATION

- Large, intricate rule sets increase the risk of errors and out-of-date security rules; they also introduce delays in responding to policy change requests and troubleshooting.
- Firewall performance degrades due to the excessive tax of processing connection requests against large rule sets.
- Out-of-date rules and objects increase security risks by unintentionally opening up access to business-critical resources.

1 Michael J. Chapple, John D'Arcy, and Aaron Striegel, "An Analysis of Firewall Rulebase (Mis)Management Practices," ISSA Journal, February 2009.

- Managing complex rule sets across environments with a mixture of firewall platforms and vendors becomes progressively difficult and consumes IT time.
- Periodic audits to demonstrate compliance with security regulations and standards get increasingly complicated with large rule sets spread across multiple firewalls and can create unnecessary costs for business.
- The administrator must identify and remove redundant, old, unused, and shadowed rules to achieve optimized firewall performance.

THE ILLUMIO SOLUTION

- Simplifies security configuration with the enforcement of context-aware security policies without any modifications to the underlying infrastructure.
- Security adapts automatically to application or infrastructure changes using continuous computation of policies and workload context.
- Does not rely on perimeter-based enforcement for security. Avoids the need for internal firewalls to enforce security within the data center. Instead, fine-grained and precise security policies are directly applied to individual workloads to enforce permitted interactions.
- Simplifies security auditing with simple, natural-language policies expressed in the language of the application rather than network-centric parameters like IP addresses and security zones.
- Provides visibility to application traffic flows and enables validation of policy changes against existing traffic flows.
- Enforces security across public, private, and hybrid cloud, with policy consistency across environments and physical or virtual workloads.

.

CURRENT APPROACHES TO ELIMINATING FIREWALL RULE PROLIFERATION

MANUAL SECURITY AUDITS AND CHANGE MANAGEMENT

Maintaining a clean set of firewall rules is one of the most important firewall management functions, yet businesses continue to struggle with it.

Many security professionals advocate frequent reviews of firewall rule bases and other security controls to reduce the likelihood of errors. Accurate audits of the security configuration are time sensitive and require intimate knowledge of traffic flows and the underlying network topology to identify the relevant policy enforcement points. However, the sheer complexity of any given network can make this difficult. Enterprises often end up relying on rule "hit counts" to identify and disable unused rules. These half-baked approaches don't quite address the problem.

AUTOMATED FIREWALL RULE MANAGEMENT TOOLS

Some organizations may look to automation tools to help maintain their ACLs and firewall policies. These firewall rule management (FRM) tools employ firewall best practices and analyze the network environment to highlight gaps and weaknesses. They can help audit the rule base by comparing the rules to compliance guidelines. Some of these tools offer topology awareness, which can be leveraged to identify the firewalls that are affected by a proposed change.

Firewall analysis enabled by automation tools helps reduce operational costs related to firewall rules administration and compliance with regulatory mandates. However, these tools do not solve the fundamental problem of perimeter-based security and rule proliferation.

SIX KEY CHALLENGES WITH EXISTING SOLUTIONS

1. Tool proliferation

The current approaches for firewall rule-base management ultimately rely on static, network-based solutions that do not address the needs of today's highly scalable and dynamic data centers and clouds. Manual processes then must be put in place to configure the various network and security parameters to account for application changes including workload migrations, IP address changes, and decommissions. FRM tools that ultimately rely on these network-based solutions to enforce security cannot work around these limitations. Instead, they end up contributing to the proliferation of tools that busy administrators have to deal with.

2. Errors and delays

Adequately managing firewalls configured with thousands of rules places a considerable burden on organizations. Manually analyzing the effect of rule additions, changes, and deletions is not only tedious, it is also error prone. Moreover, as the rule base grows, it becomes increasingly difficult to identify errors once they occur.

3. Increased security risk

Having an unwieldy rule base is not just a technical nuisance, it also creates business risks. For instance, a firewall may unintentionally allow access to a business-critical resource that is reusing the IP address of an old decommissioned system referenced by an orphaned rule.

According to Gartner, "Through 2018, more than 95 percent of firewall breaches will be caused by firewall misconfigurations, not firewall flaws."[2] The implications of a security configuration error can therefore be severe, spanning service interruption to data breaches that cost organizations an average of $5 million.[3]

4. Performance impact

Beyond the management costs, there are also system costs due to complexity. When a firewall evaluates a connection, the entire rule base is parsed from top to bottom, comparing the connection characteristics to the rules until it finds a match. The larger the rule base, the more taxing it is for the firewall to evaluate new access attempts against the policy—leading to reduced performance.

5. Compliance issues

Most industries are subject to myriad security regulations and standards such as PCI-DSS, HIPAA, Basel II, NERC CIP, and SOX. Cluttered rule bases significantly complicate the auditing process, which often involves a review of each rule and its related business justifications. Owing to the size and dynamic nature of firewall rule bases, these audits can create unnecessary costs for the business and waste precious IT resources.

6. Managing firewalls from multiple vendors

Often, large organizations have deployed hundreds of firewalls from different vendors to control access, and the configurations of these devices are constantly changing. Even though firewalls from different vendors serve a similar purpose, there is no standardization in the way rules are specified. For instance, security zones are required for some firewalls, but are not used in others. This means that managing security can require a unique approach for each solution. A relatively straightforward change in business connectivity needs may require coordination across multiple firewalls and network devices, making it challenging to implement. Since most firewall platforms do not provide APIs, FRM tools end up relying on SSH access and command-line entry. As vendors make changes to commands with each new version, it becomes increasingly difficult to rely on third-party tools, which are always playing catch up.

2 Greg Young, "One brand of firewall is a best practice for most enterprises," Gartner Inc., November 2012.
3 Eric Chabrow, "Regulations' Impact on Data Breach Costs," bankinfosecurity.com, June 11, 2013.

THE ILLUMIO SOLUTION

Illumio Adaptive Security Platform (ASP)™ secures enterprise applications in data centers and private, public, or hybrid cloud environments by decoupling security from the underlying infrastructure. The platform continuously computes security for enterprise applications by using the dynamic context of individual workloads running on virtual machines or physical servers.

The **Illumination** service provides visualization of applications and workload interactions with a graphical view of application topology to help inform security and policy decisions. Since traffic flows are correlated to the configured policies, Illumination can also be used to verify enforced security policies. This is particularly useful when all workload interactions of an existing application may not be fully known. In such cases, Illumination can be used in a workflow mode to visualize the workloads that comprise an application, label them, and then build the rules based on permitted flows.

The **Policy Compute Engine (PCE)** is a centralized controller than manages all of the state and policies of the computing environment it visualizes and protects. It examines the relationships among workloads, computes the rules required to protect each workload, and distributes those rules out to the **Virtual Enforcement Nodes (VENs)** on the workloads.

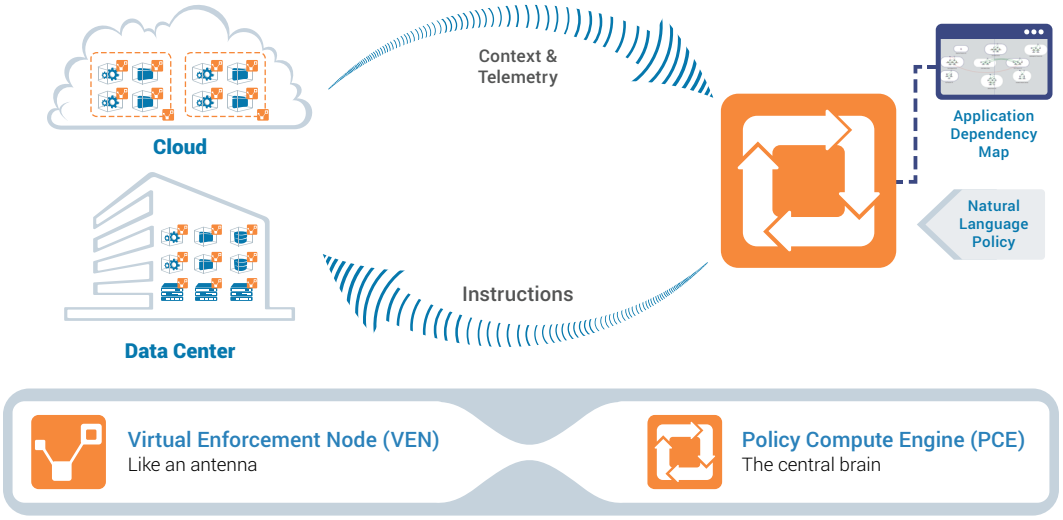

The Illumio solution addresses the challenges listed above by enabling:

1. Security that automatically adapts to changes

In the Illumio model, security policies are attached to individual workloads—not the infrastructure—to create the most accurate enforcement point. New workloads are automatically assigned the correct security policies as soon as they are instantiated. These policies follow the workloads as they migrate within or between data centers, public clouds, or virtual private clouds, and are automatically removed when the workloads are decommissioned.

Decoupling security from the underlying infrastructure also simplifies network configuration, eliminating the need to implement VLANs, security zones, and rules across multiple networks and firewalls. With this approach, enterprises can secure applications running on bare-metal servers, VMs, and Linux containers across private data centers and public cloud infrastructures including AWS, Microsoft Azure, Rackspace, and Google Compute. Illumio ASP orchestrates the enforcement of policies, ensuring that the workloads are always secured with the most up-to-date rule set, without any redundant rules.

2. Security policies that are easy to write—and understand

Security policies are easy to write and update using natural language that reflects the data protection needs of the organization. Illumio ASP uses a flexible, multidimensional labeling mechanism to define a workload based on its role (e.g., database, web server, mail server), the application that it serves (e.g., CRM, e-commerce, payroll), the business environment in which it runs (e.g., development, test, production), and its location (e.g., Germany, Australia, United States). Using dynamic application attributes, rather than static network parameters like IP addresses, allows organizations to express the relationships between workloads in human-readable policies. For example: "Only employees in Germany are allowed to access the German CRM application." Illumio ASP does the hard work of translating the human-readable policy into the technical rules needed by the workloads to enforce the security policy.

3. Precise security policies with no redundancy

Perimeter firewalls apply broad IP-based policies between groups of workloads, with no visibility or security control over the east-west traffic within the group. By contrast, Illumio enforces a fine-grained approach that attaches security policies to the individual workloads. This ensures that any traffic to and from the workload is secured, independent of the underlying infrastructure. Explicit rules ensure that every workload is only able to access resources that are necessary for the application's legitimate purpose. By compartmentalizing the application workloads, the surface area of attack is reduced—preventing the lateral spread of attacks even if a workload is compromised.

4. Simplified rule management

The Illumination service provides application-specific flow visualization correlated to the configured policies, making it possible to identify unused rules that can be eliminated. Any policy changes can be evaluated against existing application flows even before they are enforced. Comprehensive alerting options provide visibility into traffic connections that could be dropped if the security policies were enforced, eliminating concerns about rule changes breaking applications.

5. Simplified compliance management

The traditional approach to specifying security policies can result in thousands of rules spread across multiple network elements, making it time-consuming and complicated to conduct an audit. By contrast, Illumio ASP enforces an accurate set of whitelisted rules based on the context of the workload. Human-readable policies make it easier to review each rule and its related business justification. Finally, a versioning system keeps an archive of all the security configuration modifications, allowing operators to step back in time to audit the exact contents of the change, identify who made the modification, and compare the differences between versions.

With Illumio ASP, applications can be secured with the most up-to-date, fine-grained policies enforced on the workloads, while eliminating the proliferation of rules across perimeter firewalls. This enables the perimeter firewall to go back to doing what it is meant for: guarding the walls of the data center.

6. Application workloads to be secured right from their inception

Illumio ASP's workload pairing capability automatically secures newly instantiated workloads of the customer processing application without the need for writing additional security policies. The VEN can be built into the image (AMI in AWS, VM templates in VMware, etc.) or ISO file for instantiating new workloads of the customer processing application. Newly instantiated workloads can are then automatically configured with the labels and associated security policies as soon as they are paired with the PCE. This gives an enterprise the ability to spin up the application anywhere and ensure that the security policy is in place—regardless of location or cloud provider.

7. Simplified security audits

Security auditing is greatly simplified since security is continuously computed and enforced with a minimal set of rules that are always up-to-date. Illumio ASP provides easy visualization of the rules and lists out all the enforced security rules in the context of the selected workload. Rules can be correlated to the active application flows using the Illumination service. Any user-initiated changes to the rule sets of the customer processing application are captured as part of the audit logs, which can be used to analyze and verify system changes.

8. Integration with DevOps tools

Illumio ASP simplifies security and allows security to match the speed of application development with DevOps integration. If an organization uses orchestration to instantiate workloads, obtaining the correct pairing key can be part of the orchestration procedure for a workload. For instance, if using Chef, the Chef recipe can use Illumio RESTful API to get the right pairing key for a workload. This ensures that wherever a workload is brought online, the right pairing profile is attached to the workload.

USE CASE: USING ILLUMIO TO SECURE APPLICATIONS WHILE REDUCING FIREWALL RULES

Illumio ASP computes and enforces fine-grained security at each workload. This automatically reduces rules on perimeter-based firewalls where IP-based rules must be configured for every north-south interaction to workloads behind the chokepoint. To illustrate this concept, we will consider ABC Corp., an enterprise with a three-tier customer data processing application.

ABC Corp's customer data processing application is composed of a web tier that runs Apache, an application processing tier running Tomcat, and a database tier running PostgreSQL. Similar instances of this application are running across different infrastructure in Azure, AWS, and a private data center in San Jose, Calif.

All instances of the customer processing application have the following security constraints:

- Apache service hosted on the web tier is open to the internet.
- PostgreSQL service hosted on the database tier is accessible only from the processing tier.
- The instances in the respective locations should not be allowed to communicate with each other.

ABC Corp. is using Illumio ASP to secure these instances with a single rule set, without the need for configuring additional firewall rules, installing any additional security appliances or modifying the network infrastructure.

WRITING SECURITY POLICIES IN NATURAL LANGUAGE

Illumio ASP uses a flexible, multidimensional labeling mechanism to characterize a workload, based on its role (database, web server, mail server, etc.), the application that it serves (Payroll, Ordering, etc.), the environment it runs in (dev, test, production, etc.) and its location (US, Atlanta, Rack #3, etc.). Administrators can create a library of labels that are unique to their environment and use them to write security policies in natural language format. The labels are used to specify accurate policies using a service provider/consumer syntax for east-west and north-south interactions between workloads within and across applications, wherever they are running. The Illumio PCE maps the labels and configured rules to dynamically compute workload specific rules using real-time telemetry provided by the individual VENs. These labels can be automatically assigned as part of pairing the workload (a process used to bring them under management) using Pairing Profiles.

GENERATING PAIRING PROFILES

The Pairing Profile is a configuration template that specifies labels that are to be applied to newly instantiated workloads. The Pairing Profile can also be used to generate unique pairing keys that are used by newly instantiated workloads to pair themselves to the PCE. When the new workloads are paired they acquire the labels and the associated security policies within the scope of their labels.

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

ABC Corp.'s has labeled the workloads in its customer processing application as follows:

	ROLE	APPLICATION	ENVIRONMENT	LOCATION
Web tier	Web	Customer Processing	Production	AWS or Azure or San Jose
Processing tier	Processing	Customer Processing	Production	AWS or Azure or San Jose
Database tier	Database	Customer Processing	Production	AWS or Azure or San Jose

Depending on the location (or cloud provider) that the application instances are running, the appropriate location labels are assigned to the workloads.

WRITING SECURITY POLICIES BASED ON LABELS

Once the workloads have been labeled, ABC Corp. can write security policies to capture the explicitly allowed interactions (whitelisted policies) between the workloads. Interactions not captured are simply denied. These rules can be configured through the GUI or leveraging the REST API of the PCE.

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

Rule 1: Only the Apache service running on the web servers will be accessible from anywhere.

Rule 2: The Tomcat service running on the processing servers will be accessible from the web servers.

Rule 3: The PostgreSQL service running on the database servers will only be accessible from the processing servers.

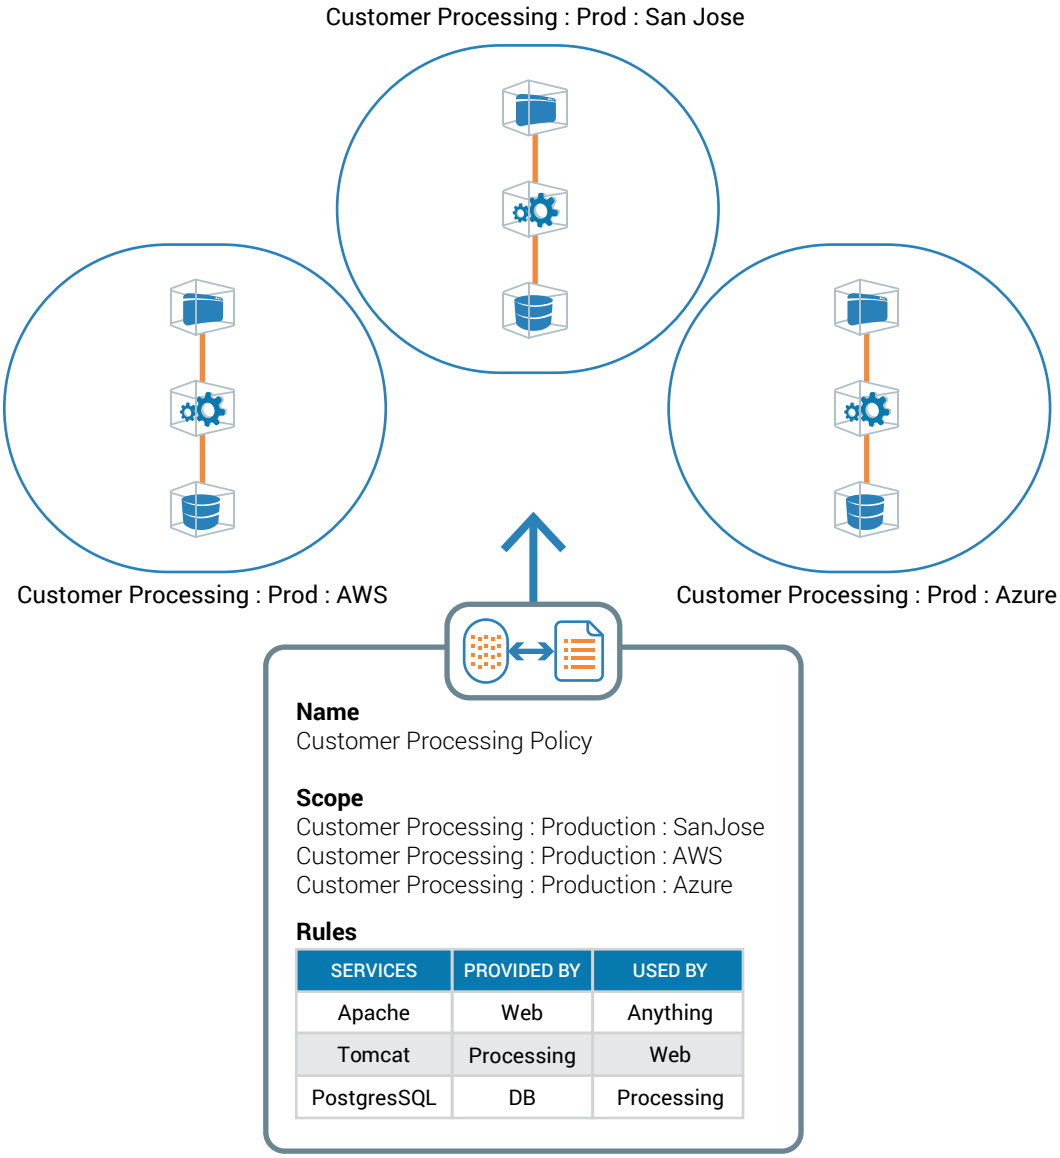

THE SCOPE OF SECURITY POLICIES

The scope identifies the set of workloads to which the security rules apply. In the above example, the scope "Customer Processing : Prod : San Jose" addresses the set of workloads belonging to the customer processing application in the production environment in the San Jose location. Administrators do not need to rewrite the policies to duplicate the rules across all the locations. By including the scope specific to the individual locations in the application policy the rules can be shared across all the workloads irrespective of the location they reside in.

Since ABC Corp. hasn't configured any security policies that explicitly permit traffic flows between the application instances, any communication between the workloads of San Jose, AWS, and Azure will be blocked.

Consistent and accurate security is enforced regardless of the number of workloads instantiated as part of the application. If any of the workloads are migrated or decommissioned, these changes are automatically detected by the PCE. Related security policies will be adjusted instantly—without any manual reconfigurations.

SECURITY POLICIES WITH ILLUMIO VS. FIREWALL RULES

What would it look like if ABC Corp. used a traditional firewall to secure its customer processing application, instead of using Illumio ASP? Assuming a workload per tier, three security rules would need to be configured on the firewalls at each location (San Jose, Azure, AWS). With a redundant high-availability pair, this would amount to a total of 18 rules [3 (rules) * 3 (locations) * 2 (firewalls in HA pair)]. The number of rules would increase dramatically with the need to scale out the application tiers. For example, a scale out of 10 web workloads would increase the number of firewall rules that address the web tier workloads by a multiple of 10. By contrast, Illumio requires a single policy with three rules-across all locations and scale-out scenarios.

ABOUT ILLUMIO

Illumio, the leader in micro-segmentation, prevents the spread of breaches inside data center and cloud environments. Enterprises such as Morgan Stanley, BNP Paribas, Salesforce, and Oracle NetSuite use Illumio to reduce cyber risk and achieve regulatory compliance. The Illumio Adaptive Security Platform® uniquely protects critical information with real-time application dependency and vulnerability mapping coupled with micro-segmentation that works across any data center, public cloud, or hybrid cloud deployment on bare-metal, virtual machines, and containers. For more information about Illumio, visit www.illumio.com/what-we-do or follow @Illumio.

- Engage with Illumio on Twitter
- Follow Illumio on LinkedIn
- Like Illumio on Facebook
- Subscribe to the Illumio YouTube Channel

CONTACT US

For more information about Illumio ASP and how it can be used to achieve environmental separation, email us at illuminate@illumio.com or call 855-426-3983 to speak to an Illumio representative.

Illumio, Inc. 920 De Guigne Drive, Sunnyvale, CA 94085 Tel (669) 800-5000 www.illumio.com