

THE FIVE NEW RULES TO SECURE DYNAMIC DATA CENTER AND CLOUDS

CONTENTS

EXECUTIVE SUMMARY	3
IN THE RACE TO DYNAMIC DATA CENTERS AND CLOUDS, SECURITY HAS BEEN LEFT BEHIND	4
BUSINESS NEEDS ARE CHANGING AND SECURITY TEAMS HAVE THEIR BACKS TO THE WALL	4
THE MASLOW'S HAMMER PROBLEM—CHOKES POINT AND NETWORK SECURITY	4
THE FIVE NEW RULES TO SECURE DYNAMIC DATA CENTERS AND CLOUDS	5
Rule 1: Security cannot depend on infrastructure	5
Rule 2: Security is not effective without understanding application context	6
Rule 3: Security policies should be adaptable and human-readable	7
Rule 4: Visibility behind the firewall is necessary for security	7
Rule 5: Security should be enforced by automation	8
THE ILLUMIO ADAPTIVE SECURITY PLATFORM™	9
Context awareness	9
Dynamic policy computation for precise security	10
Policy abstraction and decoupling security from infrastructure	11
Multidimensional labeling—a new paradigm for flexible security policies	11
Fine-grained, human-readable security policies	12
Policy scope and scaling	13
Visualizing applications and their interactions behind the perimeter	13
Illumio Adaptive Security Platform architecture	14
CONCLUSION	15
ABOUT ILLUMIO	16

EXECUTIVE SUMMARY

In today's business environment, security is caught between a rock and hard place. High-profile data breaches, insider threats, and human errors have left IT security teams shouldering the weight of their enterprises' reputation and intellectual property. At the same time, business changes and the need for rapid development, infrastructure changes, and cost reductions mean security teams are scrambling to keep up in order to avoid being perceived as a bottleneck.

The current security model for protecting applications in data centers and clouds—one that relies on static network parameters and perimeter security—is broken. The network-centric security model is unable to keep up with the application changes and infrastructure evolution that are characteristic of dynamic data centers.

The rules for security have changed. The industry needs a fundamentally different approach to security—one that is adaptive, protects from the inside out, and decouples security from the underlying infrastructure.

IN THE RACE TO DYNAMIC DATA CENTERS AND CLOUDS, SECURITY HAS BEEN LEFT BEHIND

Data centers have undergone a massive transformation in the last few decades. Originally, enterprises owned infrastructure and protected it with tightly controlled physical and perimeter-based access controls. The traditional data center was based on this model and enterprise security strategies were built on the implicit trust that ownership provided. Security consisted of “guarding the fortress” with gateway security appliances at the boundary between trusted entities inside and untrusted entities outside.

This approach worked well when original assumptions about infrastructure, applications, and how those applications were accessed held true. However, the advent of server virtualization, cloud computing, and distributed, multitiered applications has resulted in a fundamentally new view of infrastructure and data centers. The security model, though, has remained largely unchanged, maintaining a focus on the perimeter and static, network-based policies. This has created a gap in the security capabilities and needs of dynamic data centers.

BUSINESS NEEDS ARE CHANGING AND SECURITY TEAMS HAVE THEIR BACKS TO THE WALL

Rapidly changing business needs require high availability of services, meeting compliance standards, the continuous delivery of applications, and leveraging affordable public-cloud infrastructure. The threat landscape has changed significantly—data breaches and advanced targeted attacks have relied on exploiting vulnerabilities inside the data center, behind perimeter defenses. And end-user demands brought on by the consumerization of IT have created new entry points and security concerns. All of these factors have combined to exert pressure on security and operations teams to be nimble in the face of change.

The Maslow’s hammer problem—choke point and network security

“If all you have is a hammer, everything looks like a nail.”—Abraham Maslow

The industry has attempted to extend existing choke-point and network-based security tools to security tasks for which they were not originally designed. The virtualization of servers and networks has allowed enterprises to quickly spin up workloads and create multiple logical networks, but these virtual networks suffer from the same vulnerabilities as physical networks. Without visibility to inter-virtual machine (VM) traffic, security administrators can’t intercept traffic and determine whether it is safe. Perimeter firewalls have been extended by being moved into the data center to set up security zones and enforce access-control policies through segmentation and traffic flow management. However, none of these initiatives address the key issues facing dynamic data centers with distributed, heterogeneous applications where workloads (on VMs, bare-metal servers, or containers) can be spread across several distributed environments including multiple cloud providers.

At the root of the problem is the static, network-centric approach to securing infrastructure, which requires the use of rigid, predefined security policies and traffic steering based on IP addresses. Any changes to deployment strategies, infrastructure, or applications require careful orchestration across IT, security, and operations teams to implement manual changes to firewall rules and network parameters. Over the last few years, many new technologies, like next-generation firewalls and software-defined networking (SDN), have emerged, bringing with them their own layers of complexity and management. Many of these solutions are also vendor-specific (with dependencies on specific VMs or networking hardware) and do not work across data center and cloud environments. In the end, businesses are still held back when trying to securely migrate to the cloud, isolate environments, automatically scale applications, gain visibility to application interactions, and get a handle on complex firewall rule bases. The security considerations involved with these initiatives are bogging down IT security and operations teams.

THE FIVE NEW RULES TO SECURE DYNAMIC DATA CENTERS AND CLOUDS

There is an urgent need to rethink the way that security is architected and implemented. The industry needs a security model that can bridge the gap between dynamic business and IT environments and static, network-centric security approaches. Illumio has identified the five rules for this new model along with key questions for you to consider about your security strategy.

RULE 1: SECURITY CANNOT DEPEND ON INFRASTRUCTURE

Security that is tied to the infrastructure is doomed to fail. With the advent of virtualization and cloud computing, the traditional notions of IT infrastructure, its ownership, and how it is used, have changed. Organizations now have several choices for running their workloads, including VMs, physical servers, and containers. Further, these workloads can be hosted in public and private clouds, or a combination of the two.

The network-centric model squarely ties security to the infrastructure, requiring policies to be written in the language of the network. This results in inflexible policies, an explosion of network components, vendor lock-in, and a barrage of tools to manage security. At the same time, virtualization and orchestration tools have shrunk the amount of time it takes to launch new workloads, and businesses need applications to be delivered continuously. In today's fluid infrastructure models, characterized by speed and elasticity, security architecture and policies need to be completely independent of infrastructure. After all, an application is expected to produce the same results no matter where it runs and the application's code shouldn't change when infrastructure changes—why should security be any different?

KEY QUESTIONS TO CONSIDER:

- Can security policies remain unchanged if you modify subnets/VLANs?
- Can security policies work across any VM, bare-metal server, or cloud?
- Can you migrate to the cloud or between cloud providers with your security controls?
- How many subnets/VLANs, switches, and firewalls does your security architecture require?

RULE 2: SECURITY IS NOT EFFECTIVE WITHOUT UNDERSTANDING APPLICATION CONTEXT

The current way of defining security policies takes a “point in time” view of applications and the underlying infrastructure. Static security policies are configured based on source and destination IP addresses, with perimeter security appliances steering the traffic to the right location. But this approach has to change in order to keep up with dynamic data centers where applications are heterogeneous, computing is distributed, the perimeter has crumbled, and change is constant.

According to Gartner’s Neil McDonald,¹ “Security decisions that were largely black and white, and where policies were set statically in advance, become decisions with a multitude of shades of gray made dynamically at the time the request is made.”

But how are such security definitions created? The answer starts with considering the context and behavior of applications. Security must begin with a continuous understanding of the context of applications, so that well-informed and situation-aware security decisions can be made dynamically. For example, how should new workloads that are launched to scale an application behave? Will they receive identical security policies? What happens to an application if it is moved or replicated to a data center in a country where data residency rules apply? In each of these situations, understanding how the context of the application changes is key to applying the right security policies.

KEY QUESTIONS TO CONSIDER:

- How do you adjust security policies when applications change or new applications come up?
- How are application environments (e.g., dev, test, staging, and production) separated?
- Can workloads move from one zone or location to another without affecting security?
- How do you isolate applications or application tiers for threat containment or compliance?

¹ Gartner, “[The Future of Information Security Is Context Aware and Adaptive](#)” (ID: G00200385), published May 14, 2010, refreshed February 6, 2014.

RULE 3: SECURITY POLICIES SHOULD BE ADAPTABLE AND BASED ON NATURAL LANGUAGE

Humans address complexity by encapsulating it with simpler constructs and providing logical abstraction layers. Abstraction models have allowed us to better understand complex computing systems and to develop tools to work faster. For example, networking teams have benefited from the abstraction provided by subnets, VLANs, and virtual switches—they don't interact with the physical network every time. Application developers have benefited from object-oriented programming and scripting—they do not develop applications in machine language. But security administrators have not had a policy-abstraction model that allows security rules to be written in natural language—instead they depend on static IP addresses and ports. As a result, security policies are not adaptive or dynamic when IT infrastructure or applications change. This means that if an IP address changes, it will break the policies. Often, manual processes are used to revise or rewrite security policies. But these manual processes introduce the possibility of errors, potential security breaches, and a slowing down of business priorities.

The [recent data breach at Healthcare.gov](#) is an example of human error. A development server was supposed to remain segmented and unavailable externally, but someone inadvertently connected it to the Internet, which allowed the introduction of malware that could have been used in denial-of-service (DoS) attacks. Incidents like this heighten the caution that security teams already feel; they are then slowed down further due to the lack of a policy-abstraction model and the volume of manual processes. Security teams are under pressure from application developers who are ready to go with newly developed apps, and from infrastructure/networking teams who can launch new VMs or migrate to public clouds with ease.

When security policies are specified with the attributes of the application they are intended to protect, it abstracts them from static policy constructs forced by network-based parameters like IP addresses, ports, subnets, VLANs, and security zones.

KEY QUESTIONS TO CONSIDER:

- How easy is it to apply security policies when workloads are added, moved, or changed?
- How long does it take to make the necessary security changes when applications change?
- How many firewall rules do you have and how often do they change?
- Do you understand what each of your firewall rules is protecting?

RULE 4: VISIBILITY BEHIND THE FIREWALL IS NECESSARY FOR SECURITY

Gateway security solutions are like the doormen to a building—they know what goes in through the perimeter but not what happens inside. However, a large part of securing enterprise applications is knowing exactly what happens between workloads inside the data center or cloud perimeter. At an even more granular level, it is important to know both

sides of the communication path. For example, if workload A needs to communicate with workload B, then B should also be set up to accept communications from A. How else can one know of an insider threat or a masquerading intruder that has compromised a workload and is able to quietly initiate communications with other workloads, gather data or intelligence, and eventually report back to a malicious agent?

Workloads, the services that they use or provide to other workloads, and the communications that occur between them must be fully visible to security and operations teams for them to make well-informed and effective security decisions. Having this visibility is key to preventing the lateral spread of attacks inside the data center. IT, DevOps, and security teams also need this information when making decisions on application scaling, decommission, and data center or cloud migration.

KEY QUESTIONS TO CONSIDER:

- Can you visually map all the workloads that belong to an application?
- Can you visualize the traffic flows that occur between workloads running anywhere?
- Can you validate your policy changes against existing applications before they are enforced?
- Can you lock down individual workloads to specific inbound/outbound communications?

RULE 5: SECURITY SHOULD BE ENFORCED BY AUTOMATION

The movement of application workloads across infrastructure in dynamic data centers is akin to water finding its level—once the security, cost, and access models have been considered and worked out, economic decisions should move computing to where it makes the most sense. But businesses have come to expect that security simply will not evolve at the same speed as applications or the underlying infrastructure in today's dynamic data center. Security teams have a long history of dealing with manual changes, including careful change-management processes and lengthy firewall rule management strategies.

The need to secure the enterprise is often at odds with the need to drive business growth, leading to zero-sum tradeoffs between IT security and business agility. This is an unfortunate consequence of the way security has (or hasn't) evolved.

The rise of DevOps automation and orchestration practices is driving faster application life cycles, predictable release times, and better quality at many enterprises. But security needs a fundamentally new approach to take advantage of the automation processes brought about by the DevOps practices. Security automation and integration with DevOps tools, along with context awareness, infrastructure independence, and flexible policies, are essential components of an adaptive security architecture. This integration helps avoid the unnecessary zero-sum tradeoffs that businesses have been forced to make.

KEY QUESTIONS TO CONSIDER:

- Are lines of business, or application developers, held back by security considerations?
- Is security perceived as a bottleneck?
- Can you automate your current security process with DevOps tools?
- Is your security solution tied to your networking/virtualization vendor?

THE ILLUMIO ADAPTIVE SECURITY PLATFORM™

The new model for security requires that it can work when:

- It is no longer possible to know in advance where workloads might run.
- IT may not own or control the infrastructure.
- Clear perimeter demarcations are not possible.
- Applications can grow and shrink elastically.

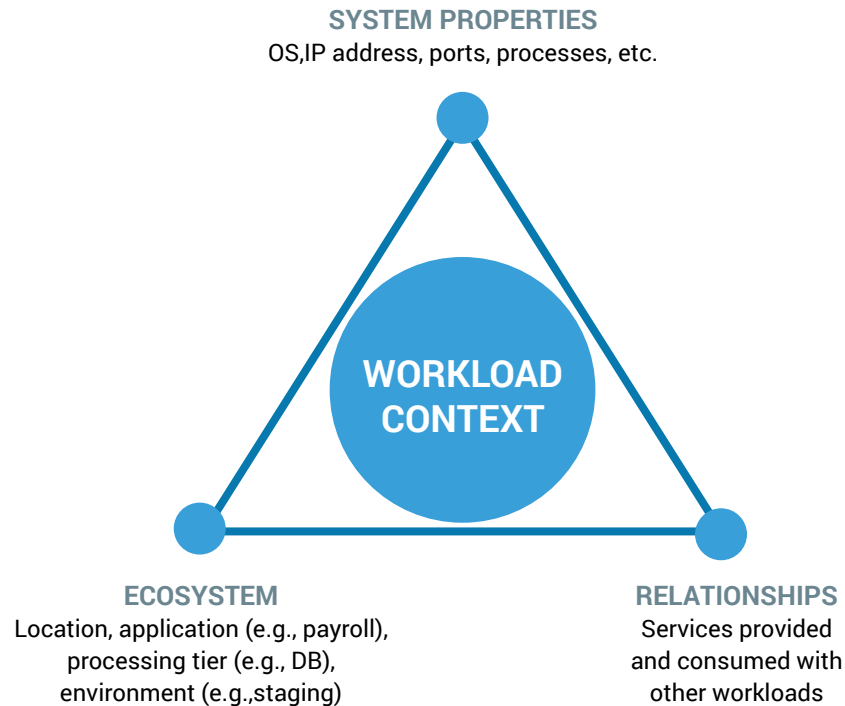
Illumio has developed a pathbreaking security platform that addresses these issues and frees security from being tied to infrastructure. The Illumio Adaptive Security Platform (ASP) allows enterprises to secure applications across private, public, or hybrid cloud environments on any VM (vSphere, Hyper-V, Xen, etc.) or physical server. It provides visibility and enforcement services to dynamically keep pace with the motion, change, and automation of the underlying IT infrastructure and applications.

CONTEXT AWARENESS

At its core, Illumio ASP starts out with the most atomic unit of computing in data centers: the workload. Applications today are distributed, multitiered (i.e., they contain separate tiers for databases, web servers, application processing, etc.), and heterogeneous (i.e., they run across several workloads). Individual workloads then represent the most granular and accurate enforcement point for securing computing resources and communications inside and outside applications. The context of a workload can change dynamically and describes the workload completely at any given time.

A workload's context includes its system properties (OS, IP address, ports, running processes, etc.), its relationships and dependencies to other workloads within the application and beyond, and its ecosystem (location, application details, life cycle environment—test, staging, production, etc.). Illumio ASP works by attaching fine-grained security at the level of individual workloads to be continuously aware of their context. Illumio ASP uses this context to create adaptive security.

Figure 1: Workload context




DYNAMIC POLICY COMPUTATION FOR PRECISE SECURITY

Illumio ASP “listens” to each workload to collect all of its contextual information. It then creates a graph of dependencies between workloads to “understand” multitiered applications that are spread across the data center (see “Figure 4: Illumination—visualizing workloads and their interactions” below). Finally, Illumio ASP “acts” to compute and enforce security by combining context information with configured security policies.

The process of listening to and computing security based on context is continuous, allowing the platform to apply accurate policies when changes occur. Real-time policy computation lets Illumio ASP (see the “Illumio Adaptive Security Platform architecture” section below) rely on the power of automation based on algorithms instead of requiring manual modifications to adjust and enforce security rules for each workload. This eliminates the need for manual reconciliation of security policies, reduces human errors, and enables the solution to scale massively.

Figure 2: Illumio—dynamic policy computation



Illumio ASP’s fine-grained policy enforcement capability locks down the workload to inbound or outbound communications that are explicitly allowed with other workloads. The Illumio model obviates the need to implicitly trust workloads based on where they reside in the network—a workload isn’t trustworthy simply because it is behind the perimeter. This approach helps mitigate insider threats, isolates potentially compromised workloads, and prevents the lateral spread of attacks.

POLICY ABSTRACTION AND DECOUPLING SECURITY FROM INFRASTRUCTURE

Humans process and remember words more easily than numbers. People have learned to describe websites, and even phone numbers, using words and mnemonics for ease of use. Vanity phone numbers allow us to dial 1-800-COMCAST instead of 1-800-266-2278 and the Domain Name System (DNS) allows us to type <https://www.google.com> instead of <https://74.125.25.99>. However, this type of simplification is not available when describing security policies for the data center; security policies are still written using IP addresses and port numbers. This has resulted in static, predetermined security policies that cannot adapt to changes in the underlying infrastructure. It leads to complicated firewall rules that proliferate over time, as well as manual processes, errors, and inflexible configurations.

MULTIDIMENSIONAL LABELING—A NEW PARADIGM FOR FLEXIBLE SECURITY POLICIES

Illumio has developed a new syntax to specify security policies using a set of multidimensional labels for describing application parameters. The process to abstract security policies from network parameters starts with the workload. The Illumio labeling process allows a workload to be characterized based on its unique persona. This means a workload can be described based on its role (database, web server, mail server, etc.), application type (HR, sales, finance, etc.), environment (development, test, production, etc.), or location (Denver, Europe, rack #3, etc.).

For example, one workload might be the database for an Online-Store application, running in the production environment, and located in the US data center. Another workload could be the web server for the same application, in the same environment and location.

Table 1: Labeling of workloads with Illumio

WORKLOAD	ROLE	APPLICATION	ENVIRONMENT	LOCATION
Online-Store DB	Database	Online-Store	Production	US
Online-Store Web server	Web server	Online-Store	Production	US

With this labeling process, security policies can be written based on what workloads do, where they belong, and with whom they are allowed to communicate, rather than being constrained by static network constructs.

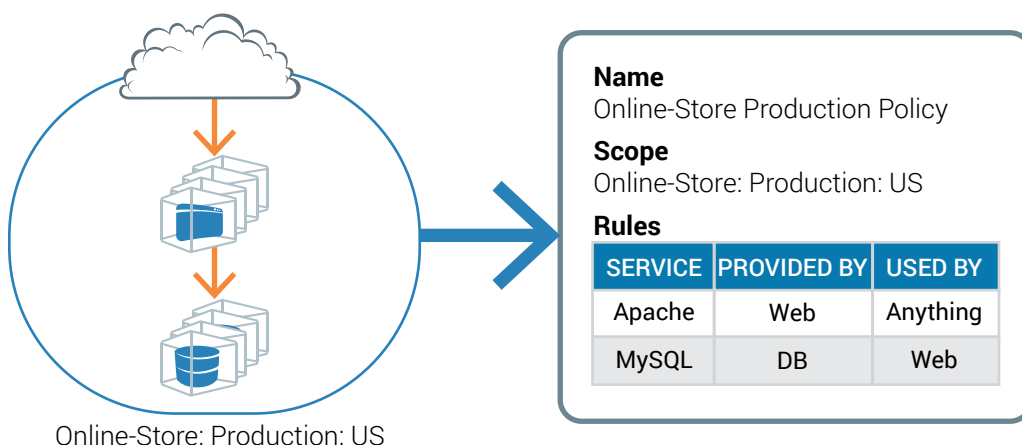
FINE-GRAINED, HUMAN-READABLE SECURITY POLICIES

Once workloads are labeled, administrators can define simple policies in natural language. For the example above, the enterprise may want to define the following enforcement rules:

1. The web servers of the Online-Store application are accessible from anywhere.
2. The database servers of the Online-Store application are only accessible from the web server.

The policy descriptions (see “Figure 3: Security policies with Illumio” below) specify only the explicit communications that are allowed. The policy example above prevents any workloads other than the web servers of the Online-Store application in the production environment in the United States from initiating any connections into or out of the database workload.

Figure 3: Security policies with Illumio



POLICY SCOPE AND SCALING

The scope of policies defines the extent to which a set of policies apply. In the above example, the policy scope has been narrowed down to all workloads in the Online-Store application, in the production environment, in the United States. However, the scope could be expanded by simply specifying that rules apply to Online-Store applications in any environment and any location. This allows security policies to be reused and expanded to encompass the desired workloads.

Policies can also scale horizontally and vertically across workloads. Once workloads have been labeled and initial policies have been set, Illumio ASP provides an innovative label- and policy-inheritance mechanism that creates pairing profiles to eliminate the need to relabel or recode policies for new related workloads. For example, new web-tier workloads that may be launched to auto scale the application share the same persona as existing web-tier workloads. Workloads are automatically protected from their inception all the way to their decommissioning.

By using application parameters rather than network parameters, security policies can be defined without knowing the details of the underlying infrastructure. For the first time, security can keep pace with dynamic changes in infrastructure and applications.

VISUALIZING APPLICATIONS AND THEIR INTERACTIONS BEHIND THE PERIMETER

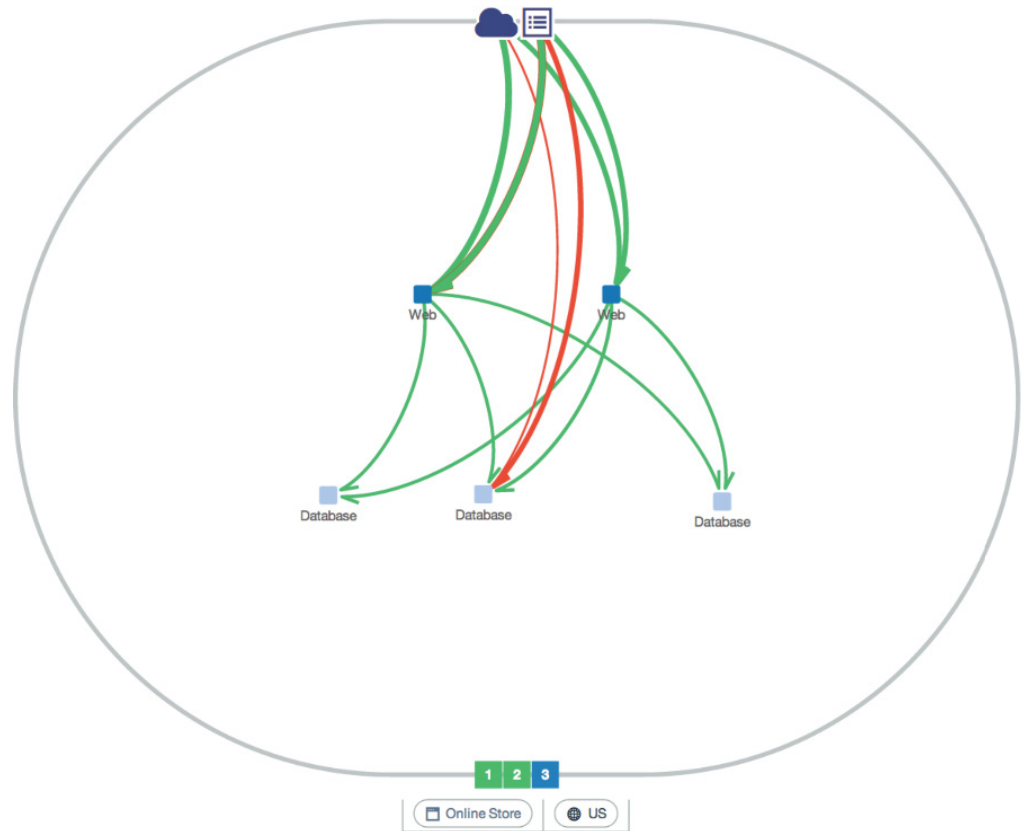
Perhaps one of most difficult challenges faced by security and operations teams is not having visibility into the communications between workloads behind the security perimeter. This lack of visibility means that administrators cannot tell whether a particular interaction between two workloads is authorized or necessary. This makes it hard to pin down insider threats or malicious reconnaissance through compromised systems.

Illumio ASP provides a service called Illumination, an industry-first innovation that sheds light on internal traffic flows between workloads using an interactive, graphical layout that displays the topology of the workloads that make up an application. Illumio ASP provides two modes of operation.

3. Illumination mode allows security administrators to view the current state of affairs of the workloads and their interactions. Administrators can use this mode to visualize workloads and traffic flows, perform policy analyses, assess security gaps and potential impacts to application flows, and even discover unused workloads.
4. Enforcement mode lets administrators write security policies using natural-language terms to describe desired communications between workloads. Once the policies are enforced, workloads are locked down to interact only with the specifically allowed paths.

Illumination allows security teams to make well-informed policy decisions and apply security policies selectively. It offers a stepping-stone to enforce accurate policies on applications. After enforcement is applied, Illumination gives visibility into blocked connection attempts that might point to a compromised system.

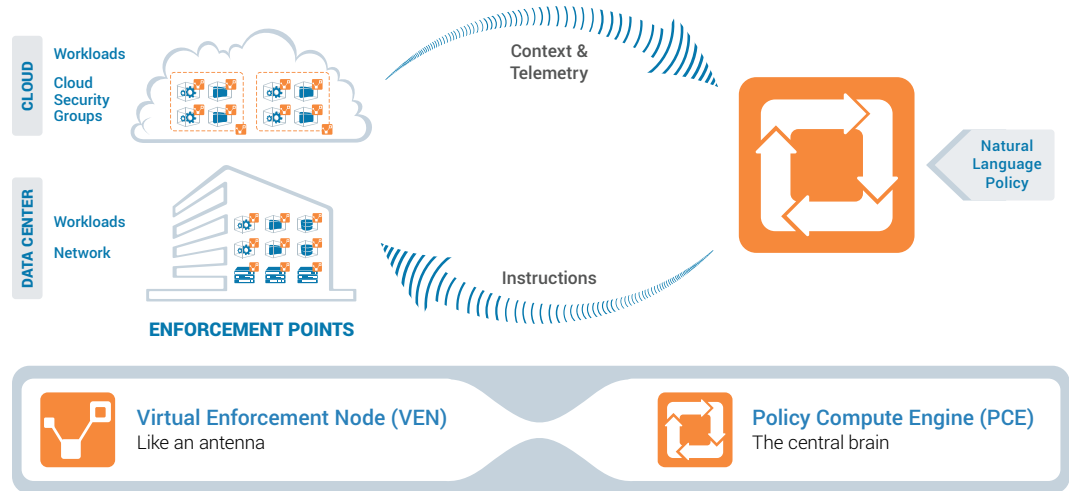
Figure 4: Illumination—visualization of applications and their communications



ILLUMIO ADAPTIVE SECURITY PLATFORM ARCHITECTURE

Illumio ASP was built to adapt security to the needs of modern data centers and clouds. It does this by completely removing dependencies on the underlying infrastructure when specifying and enforcing security. The platform is architected as a distributed and asynchronous system with Virtual Enforcement Nodes (VENs) attached to individual workloads (running on any VM, physical server, or private, public, or hybrid cloud) and a centralized Policy Compute Engine (PCE). The PCE and VENs work together to continuously monitor and apply security without the need for any intermediate choke points. The VENs listen continuously for the context of their workload and relay the information to the PCE, which computes enforcement rules by combining the workload's context with configured security policies. The enforcement policies are then sent to the VEN, which modifies the appropriate iptables or Windows Filtering Platform parameters on the workload. Since the policies are specified using a zero-trust model, where all communications to and from the workload are explicitly specified, any fraudulent transactions are automatically prevented. If the context of workloads changes due to movements, scale outs, or application changes, the VENs report the new context to the PCE, where precise policy changes are computed and pushed back out to the VENs. In the final step of this automated process, the VENs immediately apply the changes to the workloads.

Figure 5: Illumio ASP architecture



CONCLUSION

We built Illumio with the vision that security needs a quantum leap in thinking to protect computing resources in a fast-changing model of infrastructure and applications. For security to be a strategic decision with a long-term vision, and yet be reactive to the block-and-tackle nature of everyday computing decisions, the industry needs an adaptive architecture that empowers security, IT infrastructure, and operations teams to do their jobs efficiently. We created Illumio ASP to fulfill this vision and address the needs of enterprises with dynamic and advanced data center and cloud deployments.

ABOUT ILLUMIO

Illumio, the leader in micro-segmentation, prevents the spread of cyber threats inside data centers and cloud environments. Enterprises such as Morgan Stanley, BNP Paribas, Salesforce, Workday, and Oracle NetSuite use Illumio to reduce cyber risk and achieve regulatory compliance. Illumio's Adaptive Security Platform™ uniquely protects critical information with real-time application dependency mapping and micro-segmentation that works in any data center, public cloud, or across hybrid deployments on bare-metal, virtualization, and containers. For more information about Illumio, visit www.illumio.com/what-we-do or follow [@Illumio](https://twitter.com/Illumio).

- [Engage with Illumio on Twitter](#)
- [Follow Illumio on LinkedIn](#)
- [Like Illumio on Facebook](#)
- [Join Illumio on G+](#)
- [Subscribe to the Illumio YouTube Channel](#)

CONTACT US

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

Illumio Adaptive Security Platform and Illumio ASP are trademarks of Illumio, Inc. All rights reserved.