illumio

WHITE PAPER

ENFORCING DATA RESIDENCY

CONTENTS

OVERVIEW

BUSINESS DRIVERS

Enterprises are grappling with a complex maze of laws, regulations, and policies that govern data protection and privacy. These requirements, which vary considerably around the globe, are designed to protect information about individuals, organizations, and national interests.

Initially, data protection controls and restrictions focused on organizational policies and procedures to safeguard data. However, in response to high-profile data breaches and foreign legislation that has provided authorities with access to enterprise data, several countries have enacted data residency laws that dictate where personally identifiable information (PII) is physically stored.

Meeting data residency requirements lets businesses:

- Prevent unauthorized—or undetected—access to data.
- Avoid financial and criminal penalties.
- Preserve access to services.
- Build customer trust.

Practical implementation of these laws requires data to be physically stored and accessed only from within the borders of a specific country. Businesses must be able control where data is stored and accessed to comply with data protection laws, regulations, and policies.

CURRENT CHALLENGES TO ENFORCING DATA RESIDENCY

- The store-anywhere and access-from-anywhere capabilities of data make it difficult to identify which data is being accessed and from where.
- Cloud computing and virtualization technologies are, by definition, distributed, which presents a conceptual challenge to meeting data residency requirements.
- The diversity of data protection requirements means businesses must maintain different policies for different countries, resulting in complex and difficult-to-manage security policies.
- The rapid evolution of data protection laws, regulations, and policies requires continuous updates to the security policies.
- Existing network-based security tools are manual and static, making it difficult and error prone to keep up with changes.

THE ILLUMIO SOLUTION

- The Illumination service provides graphical visualization and mapping of application workloads[1] and their every interaction to help identify non-compliance with data protection requirements.
- Security is computed using the dynamic context[2] of each individual workload and is not dependent on the infrastructure. Security policies remain in place and are consistent regardless of the underlying cloud computing or virtualization technology.
- Security policies are easy to write and update using natural language that reflects the data protection needs of the organization.
- Illumio Adaptive Security Platform (ASP)™ automatically adapts to infrastructure and application changes as they occur. Security policies are computed in real time and distributed to the workloads for immediate enforcement. Movements of workloads or changes to applications do not affect security.
- Changes to security policies are automatically distributed throughout the entire computing infrastructure to ensure that data protection requirements are met.

[1] A workload equates to a discrete OS instance. It can run on a physical device or VM, or as a cloud instance.

[2] 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, etc.). The context of workloads change as the application they are part of moves, changes, and scales up or down.

CURRENT APPROACHES TO ENFORCING DATA RESIDENCY

Organizations typically leverage one or more of the following technical solutions to comply with data protection laws ad regulations.

NETWORK-CENTRIC SECURITY SOLUTIONS

Many organizations use the physical separation between locations as a way to control access. This works well for restricting full access—if there is no network connection between two locations, then the locations are inaccessible to each other. However, most environments are not this simple. In most environments, access to some applications and data needs to be restricted, but other data and applications need to be accessible. Organizations typically deploy a chokepoint firewall at each location to control this access. All traffic from the workloads is then funneled through this enforcement point, where statically configured, IP-based firewall policies are applied to application flows crossing locations.

Organizations often also have specific policies that dictate which virtual LANs (VLANs) the application workloads need to be placed in. In these cases, the appropriate security policies are added or modified on a variety of network switches or perimeter security appliances in the traffic flow. This process is repeated for every new application being launched, migrated, or decommissioned as part of the application development life cycle. Many of these security changes require manual intervention from IT and ops teams, slowing down application rollouts.

Workloads are moving increasingly from dedicated physical servers to multiple virtual machines, which means an application can be instantiated anywhere in the data center, and can be migrated to a different data center or a public cloud infrastructure. Since network-based appliances don't have intelligence about the context of the workloads they are protecting, they tend to grow out of sync as these changes occur. This leads to a need for manual oversight, cumbersome policy audits, and approvals of responsible teams to update or maintain them.

SOFTWARE-DEFINED NETWORKING (SDN)

SDN-based solutions protect data by leveraging an overlay network built on top of existing network layer-2 and layer-3 technologies. However, this approach requires IT to monitor not only the physical network but also the SDN-based virtual overlays. This makes problems harder to troubleshoot, diagnose, and manage, including root cause analysis for compliance and forensics. SDN-based solutions also require complete control over the entire network infrastructure, which means they cannot be extended to a public or hybrid cloud environment. Finally, SDN technologies do not work for physical servers and often require vendor-specific virtual machines or networking hardware.

DATA ENCRYPTION AND TOKENIZATION

The desire to leverage cloud technologies while meeting data residency laws has given rise to a new class of solutions that use encryption and tokenization technologies. Encryption uses algorithms to transform plain text information into non-readable ciphertext. A key is required to decrypt the information and return it to plain text. Companies developing encryption technologies to meet data residency requirements argue that when data is encrypted, PII is completely unidentifiable and therefore not subject to data privacy regulations. However, the validity of this approach is questionable. Although encrypted, the data is still not stored within the borders of the originating country and therefore is not under its legal jurisdiction. The holder of the data in the foreign country is subject to that country's legal jurisdiction and could be compelled to provide the keys necessary to decrypt the data.

Tokenization uses data substitution to replace the real value with a token. A token is a set of random characters unrelated to the original data and there is no "key" that can decipher the token and turn it back into real data. Instead, a table mapping the token to the original data is used to decipher the data. The mapping table can reside in the originating country while the tokens can reside outside of the country. Because the tokens do not contain the original information (they are only a link), compliance with data residency laws is maintained. However, this approach requires custom application development to integrate with a specific vendor's tokenization solution; it is not broadly applicable to the large number of commercial software applications in use within the organization.

FIVE KEY CHALLENGES WITH EXISTING SOLUTIONS

1. Lack of visibility to data flows

Networking enables data to be stored and accessed from anywhere; its sole purpose is to enable connectivity. To meet data protection requirements, security policies must be designed and implemented to restrict this connectivity. The challenge is in determining the effectiveness of these policies. While existing solutions can use IP connectivity to provide part of the picture, they cannot show data flows at an application level. To assess the effectiveness of data protection policies, organizations must be able visualize—at the application level—where data is stored and who can access it.

2. Lack of endpoint context

While network devices such as switches, routers, and firewalls can detail access information at the IP level, they lack the context of the individual workloads that comprise an application. They also lack an understanding of the interactions between the workloads that comprise an application. For example, a firewall would not know that a single application is delivered by three workloads—a web server, a processing server, and a database server—and that those servers should only be accessible by users in a specific location. Without this information, it is difficult for an organization to get a complete and accurate picture of where data is stored and control who can access it.

3. Difficulty translating data protection needs into firewall rules

Data protection needs are expressed in natural language, such as: “Only employees in Germany are allowed to access the German CRM application.” Firewall rules, however, are expressed using a technical language that requires a deep understanding of the underlying network. For example:

- Source: IP addresses in Germany
- Destination: German CRM web server IP address
- Service: TCP/443
- Direction: Inbound
- Action: Permit

With existing solutions, data protection requirements must be translated into complex firewall rules. The translation process is manual and tedious, and it introduces a high likelihood for errors. Once the rules are translated, it is difficult to verify that they actually meet the data protection requirements. Then, any future changes to the requirements require additional translation and updates to an already complex list of firewall rules.

4. Firewall rule explosion and out-of-sync security policies

As data protection laws and regulations evolve, so too must the associated security policies and network firewall rules. Because data protection policies must be manually translated into firewall rules, it is difficult and time consuming to put new ones into effect. Over time, this manual implementation results in a complex tangle of firewall rules that are risky to change since it is difficult to understand their complex interactions. The result is a list of out-of-date rules that organizations are reluctant to reconcile out of fear of breaking existing applications. The firewall rules organically grow out of sync with the data protection requirements, exposing the organization to potential data breaches and non-compliance with data privacy laws and regulations.

5. Delays due to manual processes

Translating security policies into the technical language of firewall rules requires careful review to ensure that the rules are accurate. The risk of errors and misconfigurations leading to potential security holes has led many businesses to put extensive validation and approval processes in place. Manual configuration and custom scripts are then used to implement the rules in the network infrastructure. Manually adding or modifying firewall policies requires careful coordination between the application, the network, and security teams, which slows down the implementation—and an organization’s ability to respond to data protection needs.

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.

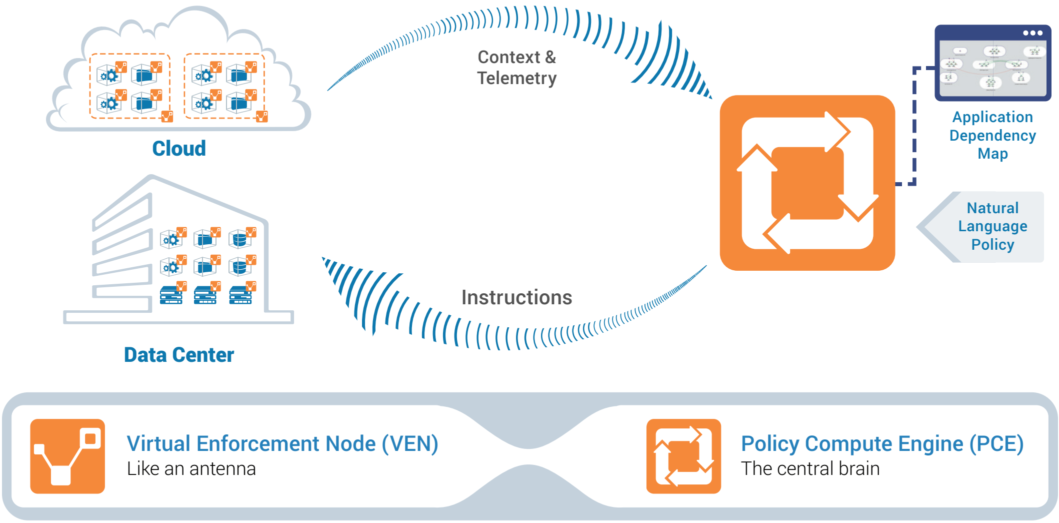

The Illumio solution addresses the challenges listed above by:

1. Validating data protection compliance using Illumination

The Illumination service provides visualization of application workloads and their interactions to identify non-compliance with data protection requirements. Illumio ASP has full visibility into workloads and their context, enabling it to dynamically compute a graph of their relationships. Illumination displays this interactive graph and provides powerful insight into the services provided by workloads and the traffic flows between them. Using this information, organizations can identify where data is stored and determine who is accessing it. Appropriate rules can then be created to block unauthorized access.

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

Security is associated with the context of each individual workload—not the infrastructure. Illumio ASP is fully context aware; it understands the intrinsic properties of each workload and the relationships between workloads. Using this information, the centralized PCE dynamically computes security policies. By enforcing security on each workload, using its context and policies abstracted from the network, Illumio completely decouples security from the infrastructure. This enables organizations to easily secure applications regardless of the underlying cloud computing or virtualization technology.

3. Enabling easy-to-write and understand policy language

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 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.

4. Making security policies easy to update

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. Automatically adapting to changes

Illumio ASP automatically adapts to infrastructure and application changes as they occur. Illumio removes the dependency on the underlying physical or virtual network by performing security policy management on the workload. The workload remains under management throughout its life cycle and is dynamically updated whenever the infrastructure changes. Similarly, changes to the application are also automatically reflected in the rules applied to the workloads. No manual intervention is required; changes are computed in real time and distributed to the workloads for immediate enforcement.

USE CASE: MEETING DATA RESIDENCY REQUIREMENTS WITH ILLUMIO

To illustrate how Illumio ASP can be used to meet data residency requirements, consider a multinational company—ABC Corp.—with a customer relationship management (CRM) application. To meet Germany's data residency requirements, ABC Corp. must ensure that all data about German customers resides within Germany, and is accessible only to employees in Germany.

IDENTIFY WORKLOADS WITH MULTI-DIMENSIONAL LABELS

ABC Corp. uses Illumio ASP's flexible labels to define the context—the role, application, environment, and location—for the workloads that make up its CRM applications in the United States and Germany. In this example, ABC Corp.'s CRM application (and its workloads) in the US and Germany are distinguished by their location values.

ABC Corp's multi-tier app and its labels

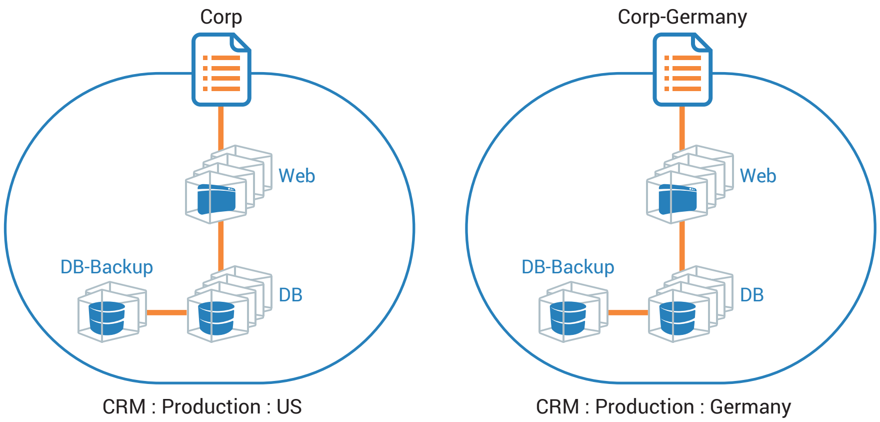

	ROLE	APPLICATION	ENVIRONMENT	LOCATION
CRM : Production : US				
Web workloads	Web	CRM	Production	US
Database workloads	DB	CRM	Production	US
Database-Backup	DB-Backup	CRM	Production	US
CRM : Production : Germany				
Web workloads	Web	CRM	Production	US
Database workloads	DB	CRM	Production	US
Database-Backup	DB-Backup	CRM	Production	US

ILLUMIO ASP MAKES IT EASY TO WRITE SECURITY POLICIES WITHOUT NETWORK DEPENDENCIES

Once ABC Corp. has labeled its workloads, it can write security policies to capture the interactions that are explicitly allowed (whitelisted policies) between the workloads. By definition, all other interactions will be denied.

Both the Germany and US location would have the following two rules:

RULE 1: The MySQL service running on the database servers will be accessible from the web servers.

RULE 2: The MySQL service running on the database servers will be accessible from the database-backup servers over an encrypted connection.
Note: This encryption of data in motion is provided by the SecureConnect capability in Illumio ASP and is represented by the icon shown in the policy tables below.

The third rule, defining the access granted to the web tier, would be different in each location. For the United States:

RULE 3: The Apache service running on the web servers will be accessible from any company machine, which is represented using an IP list.

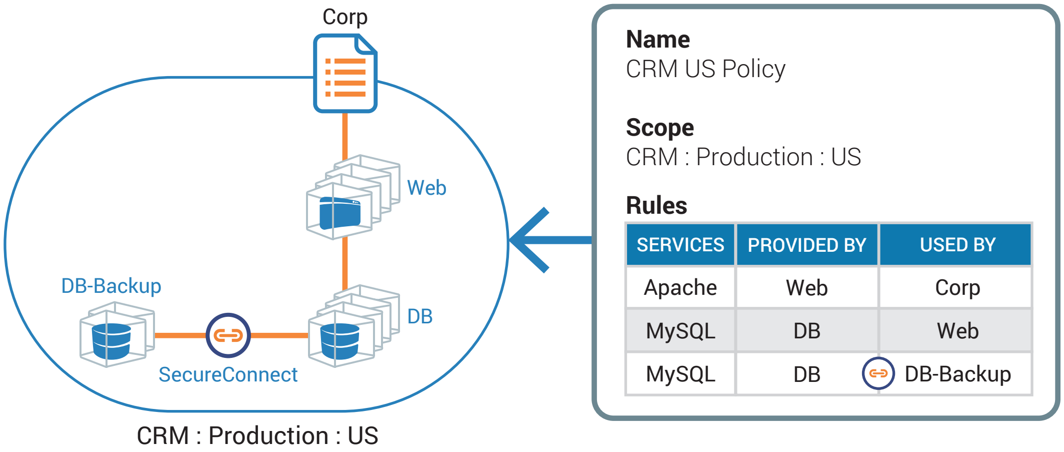

Name
CRM US Policy

Scope
CRM : Production : US

Rules

SERVICES	PROVIDED BY	USED BY
Apache	Web	Corp
MySQL	DB	Web
MySQL	DB	🔗 DB-Backup

For Germany:

RULE 3: The Apache service running on the web servers will be accessible only from Germany, which is represented using an IP list.

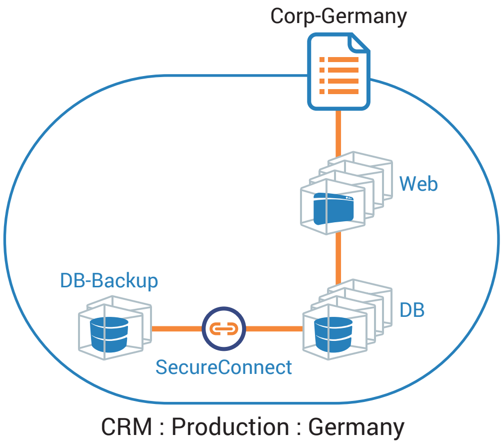

Name
CRM Germany Policy

Scope
CRM : Production : Germany

Rules

SERVICES	PROVIDED BY	USED BY
Apache	Web	Corp
MySQL	DB	Web
MySQL	DB	🔗 DB-Backup

Note: Data in motion between the primary database and the backup database in both locations is encrypted with the Illumio ASP's SecureConnect—a policy-driven, single-click capability to set up IPsec connectivity.

DEFINE THE EXTENT OF SECURITY POLICIES BY SPECIFYING THE SCOPE

The scope identifies the set of workloads on which the security rules are applied.

ABC Corp. has applied its rules across the workloads of the CRM application based on their location. Since there are no security policies configured to explicitly permit traffic flows between Germany and the rest of the company, data about German customers remains in Germany and is only accessible by machines in Germany.

If ABC Corp. needs to enforce data residency in a different country, the labels and security policy are easily updated to reflect the new data protection requirement. For example, a new "Australia" location label would be created along with a new IP List for Australia. ABC Corp. would then create a rule set to implement the new policy.

VISUALIZING APPLICATION FLOWS WITH ILLUMINATION

ABC Corp. uses Illumination to visualize the application flows specific to the CRM application. It then uses these traffic flows to verify the enforced security policies.

Any traffic flows that do not match the configured rules will be considered unauthorized attempts—and therefore blocked. They will show up in red in the Illumination graphs.

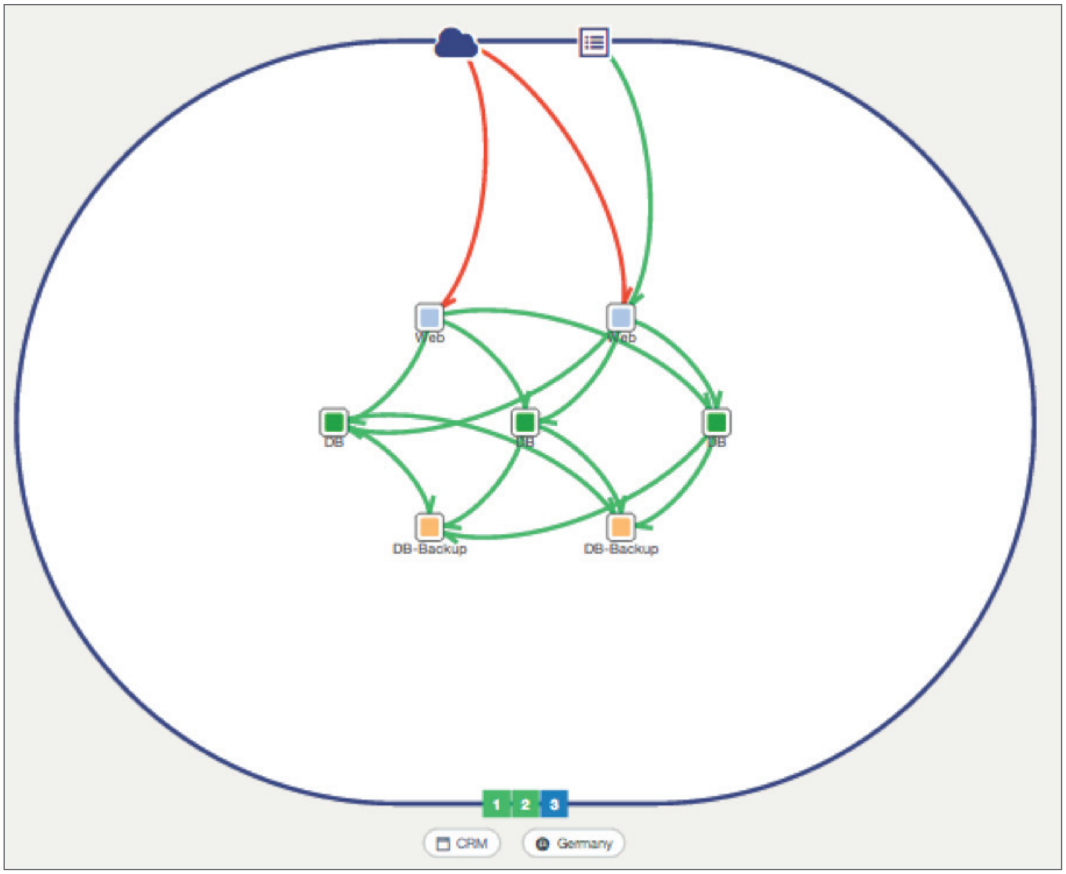

ABOUT ILLUMIO

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

Copyright © 2018 Illumio, Inc. All rights reserved. This document is protected by U.S. and international copyright and intellectual property laws. Illumio's products and services are protected by one or more U.S. and international patents listed at https://www.illumio.com/patents. Illumio® is a trademark or registered trademark of Illumio, Inc. or its affiliates in the U.S. and other countries. To review a list of Illumio's trademarks, go to https://www.illumio.com/trademarks. Third-party trademarks mentioned in this document are the property of their respective owners.