illumio

WHITE PAPER

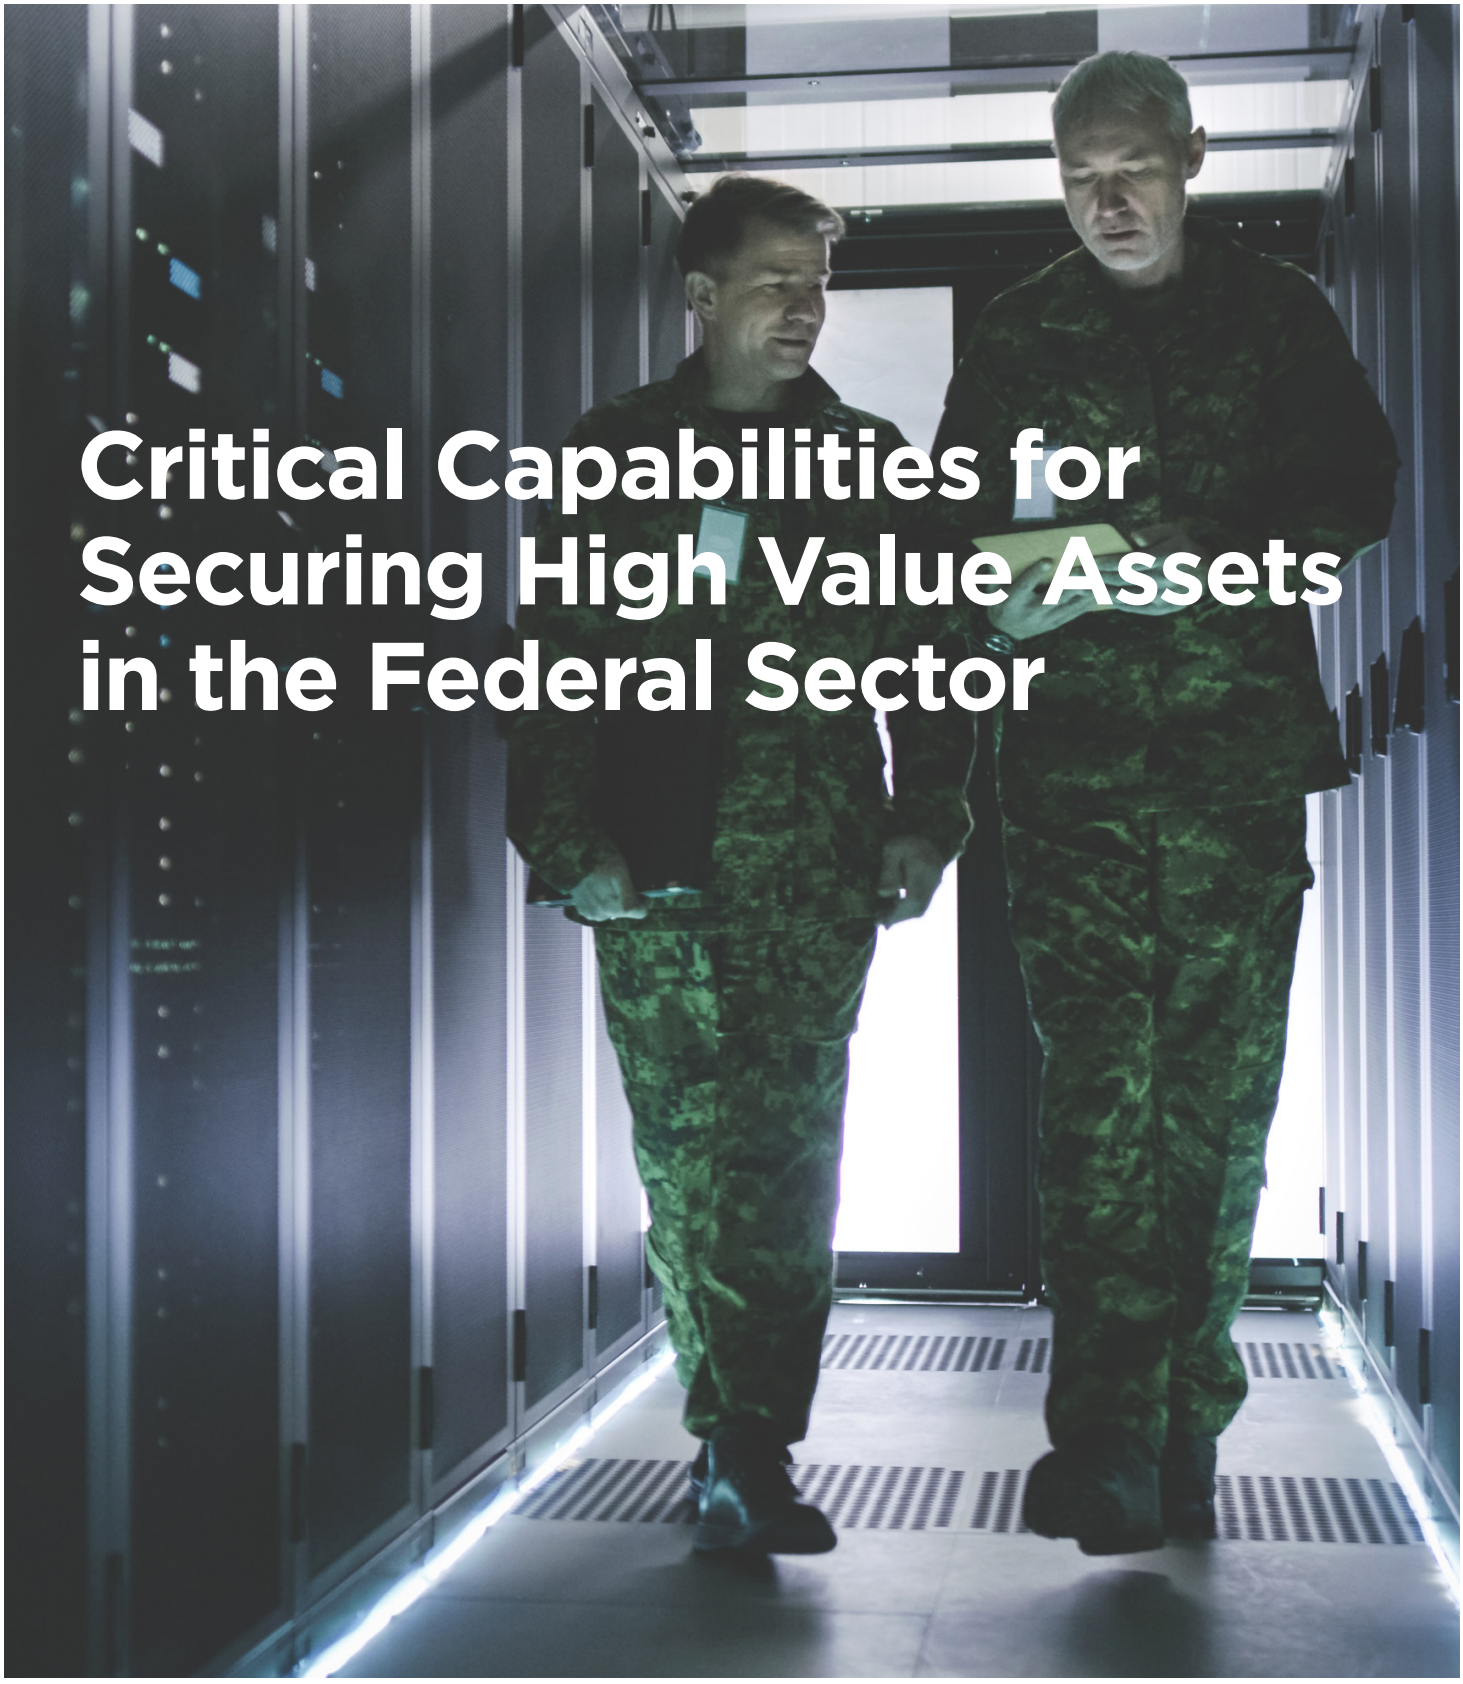

Critical Capabilities for Securing High Value Assets in the Federal Sector

Overview

Relying on perimeter network security alone is no longer sufficient protection against malicious attacks. The best way to achieve a cyber-resilient infrastructure is to assume a perimeter breach has already occurred, and contain the lateral movement of malicious actors inside the network – requiring every user, endpoint, and workload to verify and authenticate itself at every request for access or connections. This is the main principle of Zero Trust security. This white paper discusses the critical capabilities that enable federal agencies to employ Zero Trust principles in securing high value assets (HVAs).

This white paper covers:

1. The attributes of a federal HVA, the rationale behind the need to secure HVAs, and operational challenges federal agencies face in executing these security objectives.

2. An overview of the critical capabilities that a solution must have to effectively identify, secure, and manage HVAs at scale across heterogeneous compute platforms, endpoints, on-premises data centers, and multi-cloud and hybrid cloud environments.

3. Brief discussions on:

   - Using a Zero Trust strategy to operationalize an agency's HVA security objectives.
   - The critical capabilities for enabling HVA security.
   - The role of the mandatory FISMA CIO metrics in communicating the efficacy of the HVA security program.
   - A mapping of Illumio's capabilities to the HVA Process Framework.

4. Recommendations on how federal agencies can employ Illumio to operationalize Zero Trust and secure high value assets.

Defining high value assets

"A High Value Asset (HVA) is information or an information system that is so critical to an organization that the loss or corruption of this information or loss of access to the system would have serious impact to the organization's ability to perform its mission or conduct business. These assets, systems, and datasets may contain sensitive controls, instructions or data used in critical operations, or they may house unique collections of data. These sensitivities make HVAs of particular interest to criminal, politically-motivated, or state-sponsored actors for either direct exploitation of the data or to cause a loss of confidence by the public."[1]

Per OMB guidelines:

"An agency may designate Federal information or a Federal information system as an HVA when it relates to one or more of the following categories:

- Informational Value – The information or information system that processes, stores, or transmits the information is of high value to the Government or its adversaries.
- Mission Essential – The agency that owns the information or information system cannot accomplish its Primary Mission Essential Functions (PMEF), as approved in accordance with Presidential Policy Directive 40 (PPD-40) National Continuity Policy, within expected timelines without the information or information system.
- Federal Civilian Enterprise Essential (FCEE) – The information or information system serves a critical function in maintaining the security and resilience of the Federal civilian enterprise."[2]

Key milestones in federal requirements for securing HVAs

OMB, NIST, and DHS are striving to segment and secure every HVA in the federal government and guidelines have evolved to reflect updated priorities and new technology challenges. The key dates are:

- May 2018: DHS issues a Binding Operational Directive (18-02) that requires agencies to protect HVAs through proper access controls, configuration management, vulnerability scanning, and by segmenting them from other network traffic.
- December 2018: OMB M-19-03 reclassifies HVAs from a single definition into three categories that provide agencies more flexibility in designating the protections needed. It also updates the required approaches for reporting, assessing, and remediating those assets against the threat of a cyberattack. This memorandum rescinds M-16-04, Cybersecurity Strategy and Implementation Plan (CSIP) for the Federal Civilian Government, and M-17-09, Management of Federal High Value Assets
- November 2019: DHS releases "micro-segmentation" as a recommended capability under the Continuous Diagnostics Mitigation Program (CDM).
- November 2019: OMB M-20-04 updates FISMA reporting deadlines and guidance, which will be based on NIST security baselines and OMB M-19003 and DHS BOD 18-02 requirements.

[1] cisa.gov/sites/default/files/publications/CISAInsights-Cyber-SecureHighValueAssets_S508C.pdf
[2] whitehouse.gov/wp-content/uploads/2019/12/M-19-03.pdf

Why securing HVAs matters

Attacks on federal agencies by cybercriminals and nation states are more frequent and more aggressive. This is evidenced by the December 2020 SolarWinds hack, which is the largest and most significant cyber event in American history. According to the 2019 FISMA Report to Congress, there were 31,107 security incidents reported in 2018 – 10,169 (27%) were of unknown attack vectors, and 579 involved multiple attack vectors.[3]

Historically, security investments have largely been focused on prevention and detection. With the adversary always staying one step ahead, NIST asserts they are either already inside the network or will be in the near future. Dwell time describes the period of time between malware executing within an environment and when it is detected. According to FireEye's M-Trends Report 2020, the median dwell time for federal agencies in 2019 was 56 days, down 15 days from the prior year's median of 71 days.[4] While these improvements are encouraging, it is worthwhile to keep in mind that bad actors only need a few days to wreak havoc on critical federal systems.

Operational challenges

In order to effectively secure HVAs, agencies need to be aware of and have a plan for addressing the following operational challenges:

- Have visibility into the application, its behavior, its relationships to other systems, connections with distributed endpoints, and its vulnerability exposure across complex heterogeneous and geographically-distributed networking and compute environments.
- Monitor and detect changes in the agency's data centers to develop, maintain, and regularly update the HVA inventory list.
- Tightly couple firewall change management with IT change, SLDC, and DevOps so that security parity is maintained continuously.
- Avoid limitations of relying on network devices and third-party firewalls to secure East-West traffic in complex, dynamic, heterogeneous data centers.
- Be able to proactively prevent unauthorized peer-to-peer connections across endpoints, especially in remote work scenarios.
- Ensure that remote VDI users do not have broad access to the agency's data center and cloud resources once it is able to VPN into the network.

Understanding application dependencies and traffic matters

Knowing an agency's most important assets is foundational to Zero Trust. This will help baseline security capabilities, design and make architecture decisions, manage the scope of Zero Trust implementations, and plan a Zero Trust roadmap. Whether migrating legacy applications to the cloud or segmenting HVAs, cloud and security teams must keep systems operational while achieving these objectives. Most organizations have little idea how their applications actually work; for example, what roles, applications, and application tiers are allowed to communicate with a critical application. This poses obstacles and challenges for cloud migration and security teams to meet segmentation requirements.

[3] whitehouse.gov/wp-content/uploads/2019/08/FISMA-2019-Report-FINAL-to-post.pdf
[4] fireeye.com/current-threats/annual-threat-report/

Figure 1 is an example of a data center with 15 applications across 80 workloads with thousands of connections (red and green lines). Trying to solve this problem for thousands of workloads across on-premises data centers and public clouds with a full portfolio of HVAs is what the typical federal agency is up against.

FIGURE 1

Applying a Zero Trust strategy to secure HVAs

Federal agencies and command centers are adopting Zero Trust as a strategy for operationalizing the various standards and operational directives for securing HVAs. NIST Special Publication (SP) 800-207 offers a framework for building out an agency's Zero Trust architecture, recommending the following posture: "Assume adversary will compromise system."[5] This mindset shifts the emphasis from detection and prevention to containment and remediation (often called cyber resiliency). The goal is to limit bad actors' ability to traverse the internal network and reach HVAs by adopting a Zero Trust architecture. Micro-segmentation, which focuses on controlling East-West connections to only legitimate traffic, is a foundational component of this containment approach and has become standard good hygiene for any security-conscious enterprise.

Zero Trust is a journey, and effectively securing HVAs is one of its desired outcomes. At a high level, an agency's Zero Trust strategy should include:

- Visibility of HVAs, including the workloads, devices, endpoints, and legitimate connections.
- A baseline of the agency's existing Zero Trust capabilities for securing HVAs across key pillars: data, applications/workloads including cloud, networks, device, and users. This also involves understanding how its current and desired Zero Trust capabilities map to the NIST SP 800-207 Zero Trust Architecture Framework.
- A solid understanding of the agency's current and desired state of its networking, compute, cloud, and edge infrastructure.

[5] "Building Cyber Resilient Systems, A National Security and Economic Imperative" Dr. Ron Ross, May 2018

This analysis will help the agency design and execute the 800-207 compliant Zero Trust architecture that tightly couples with its networking and data center strategy.

Critical capabilities for securing HVAs

Zero Trust is not a one and done proposition. Once an agency has designed its Zero Trust strategy for securing HVAs, one of the key elements that go into its Zero Trust architecture roadmap is a list of critical capabilities that would enable it to effectively operationalize and execute ongoing Zero Trust operations. The following critical capabilities enable Zero Trust and help mitigate the operational challenges that federal agencies face during projects to secure HVAs:

1. Visibility and application insight across endpoints, data centers, clouds, and containers

2. Label-based policy model

3. OS compute infrastructure agnostic, host-based micro-segmentation

4. Automation and orchestration with IT operations and DevOps

5. Continuous monitoring and enhanced vulnerability management

1. Visibility and application insight across endpoints, data centers, clouds, and containers

You cannot protect what you cannot see. The first step to securing HVAs is to gain insight into how applications are connected and communicating. A real-time application dependency map enables you to quickly understand the relationships and dependencies between applications and their individual components. This leads to a greater understanding of application groupings, their classification, and upstream and downstream relationships. It ultimately provides the ability to block unnecessary connections without breaking applications. For example, security teams can use this information to assess any control gaps that result from migrating applications to the public cloud.

The figures that follow illustrate how to use the map to visualize application groupings, connections, and policy state.

Figure 2 is a snapshot of the Core Services Application Group in a federal data center with segmentation implemented. It shows that there are 7 applications that are consuming services from Core Services and 6 applications that are providing services to Core Services. Expanding the consuming apps view shows the list of the 7 consuming applications.

FIGURE 2

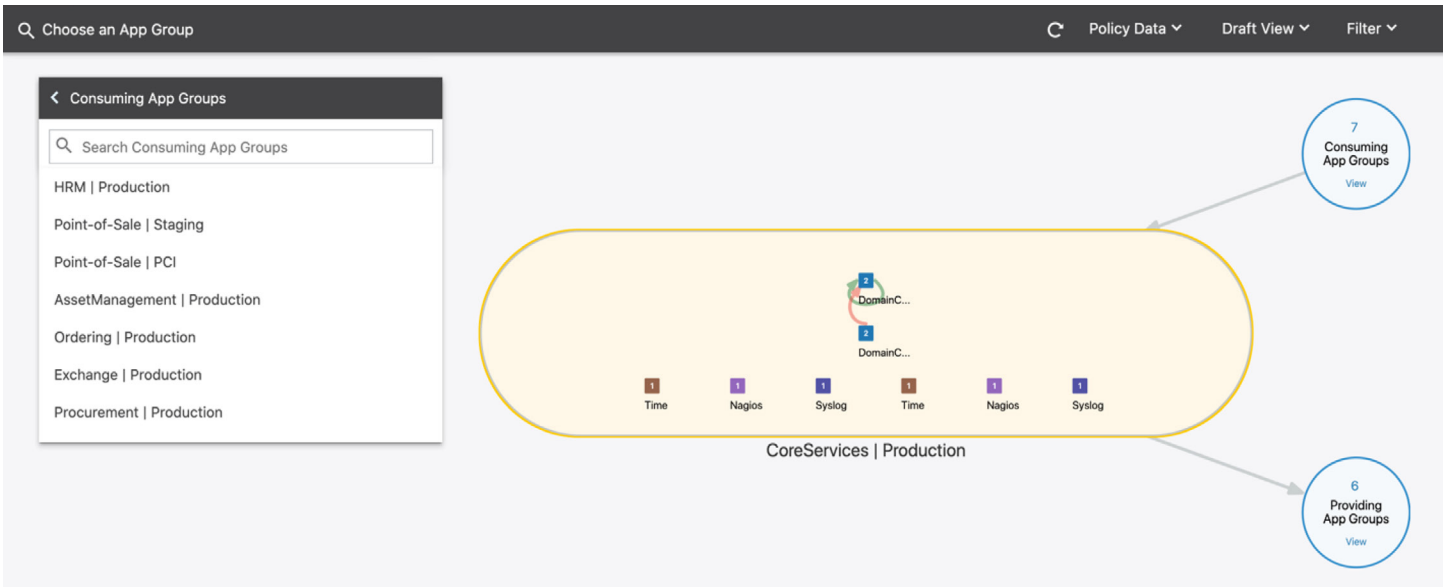

By selecting “Exchange | Production” in the Consuming Apps Group View and then expanding the group roles view, the map will show that traffic was detected between the exchange-edge-transport servers in “Exchange | Production” and the domain controllers in “CoreServices | Production.” The red lines indicate that there is no policy governing traffic between these servers.

FIGURE 3

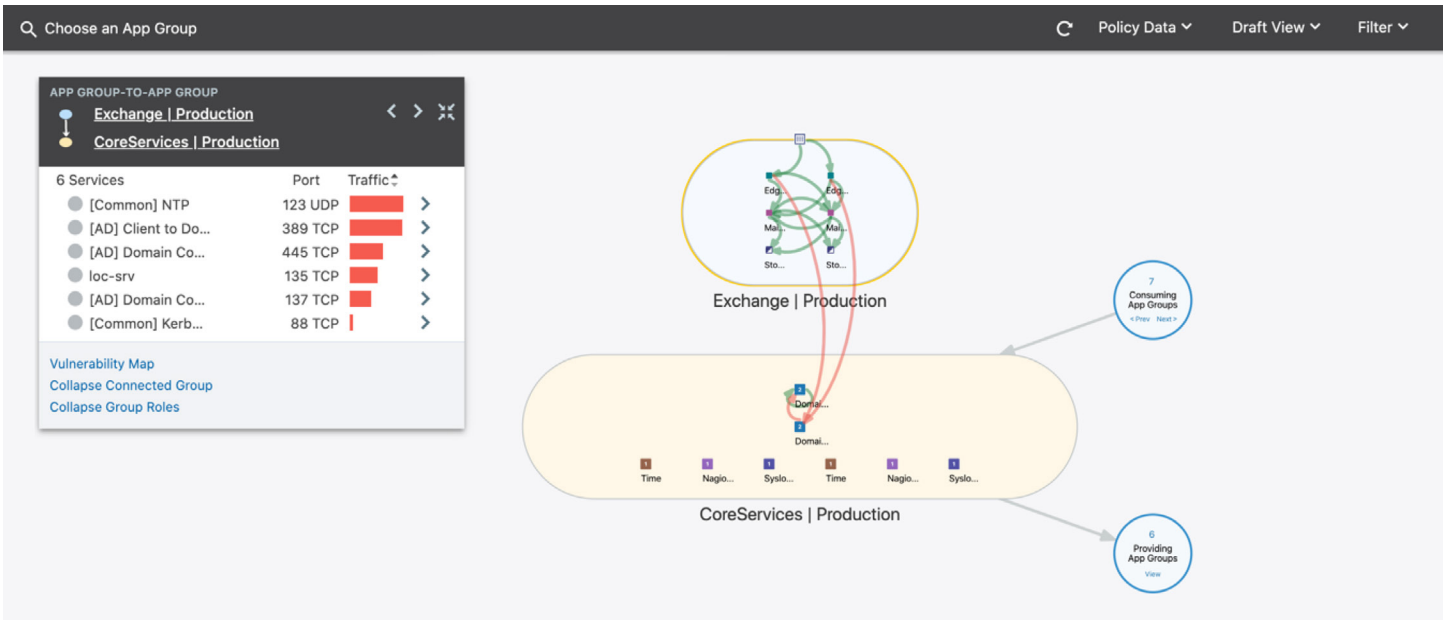

illumio

By selecting “HRM | Production” in the Consuming Apps Group View and then expanding the group roles view, the map will show that segmentation rules are enforced on traffic between the 2 web servers and 2 database servers in the “HRM | Production” and the Time and Syslog servers in “CoreServices | Production.”

FIGURE 4

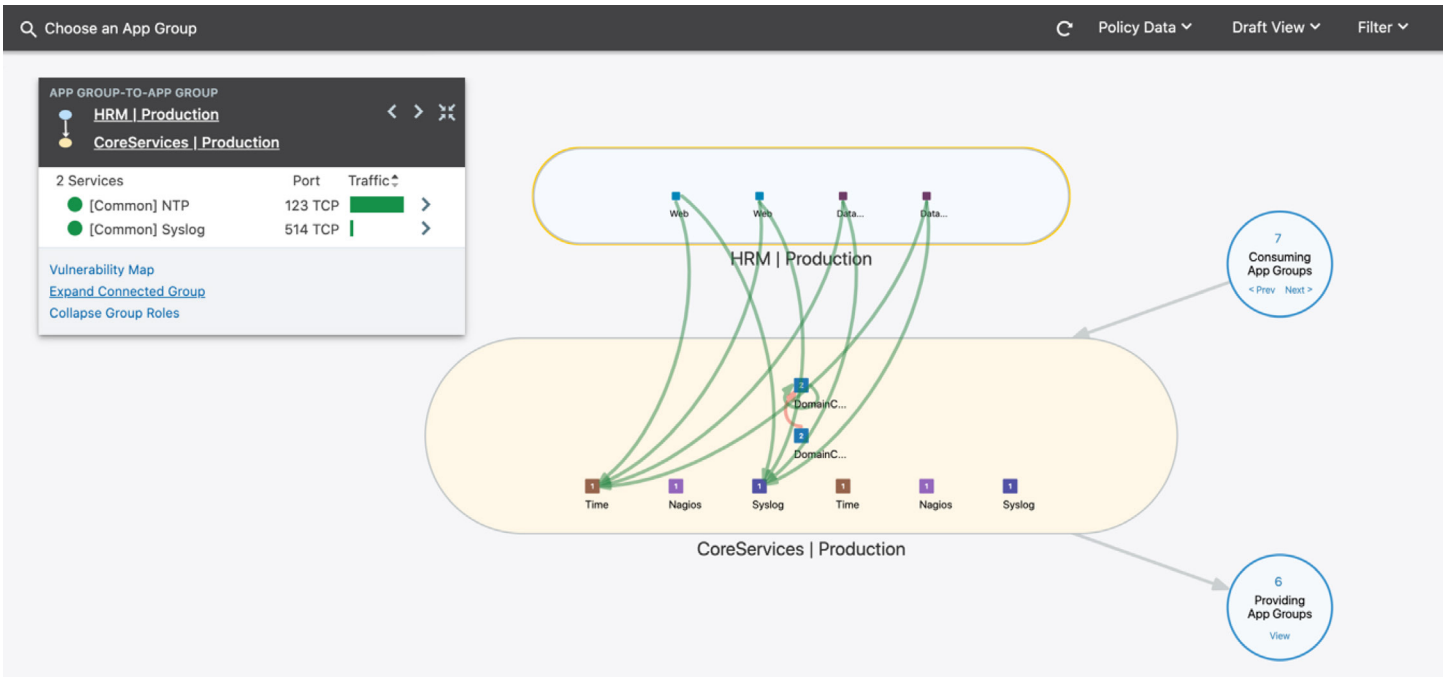

The map and application groupings will also help agencies map HVAs to the categories defined in OMB M-19-03 – Informational Value, Mission Essential, or Federal Civilian Enterprise Essential – and develop and apply the applicable security protocols.[6]

As agencies transition to remote work operations, visibility and control into the inbound peer-to-peer connections of agency-issued and sanctioned endpoints is also critical to securing HVAs. Visibility will help detect changes in the attack surface in the remote work settings, then proactively contain traffic by default.

[6] whitehouse.gov/wp-content/uploads/2018/12/M-19-03.pdf

FIGURE 5

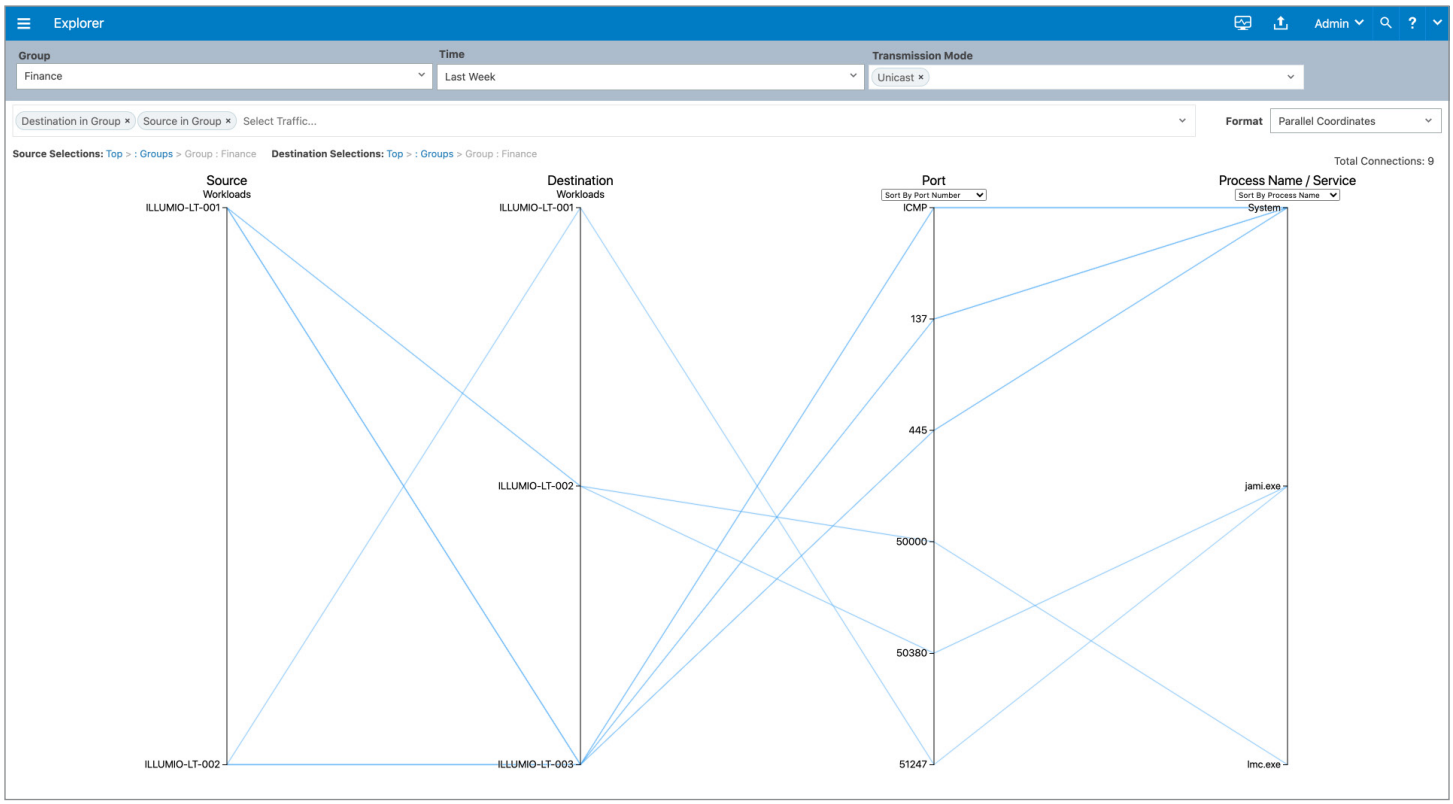

2. Label-based policy model

Historically, security, IT service management, and IT operations tools relied on a combination of IP addresses, networking metadata, and configuration attributes to uniquely identify application components in the data center. Security tools used IP addresses to define security policies and automate provisioning and enforcement of firewall rules. Admins had to pay special attention to maintaining the correct priority order of firewall rules as part of the firewall change management process so it doesn't break applications. This approach is complex and difficult to scale, especially when application stacks are increasingly distributed across heterogeneous platforms and multiple clouds and the agency does not have a strict governance model for creating tags. Containers, microservices, and DBaaS (all of which use ephemeral and dynamic IP addresses) compound this management complexity.

A solution that doesn't rely on IP addresses, but instead uses metadata from existing systems of record like a CMDB to identify, provision, and enforce policies, combined with an easy-to-understand labeling and governance model to describe components, makes it easier to create and manage security policies at scale. The security operator no longer has to keep track of all the IP connections and the priority order of security rules. With this approach, two workloads that have the same labels will have the same security policies applied to them. It also makes it easy to illustrate and understand the application groupings and interdependencies visually.

3. Host-based micro-segmentation architecture

Micro-segmentation is not intended to replace your perimeter firewalls or any other infrastructure you may already have in place. However, maintaining a flat network makes it easy for bad actors to traverse the network and eventually reach your HVAs. Most agencies' environments are complex and typically include multiple geographically-distributed data centers, shared goverment cloud/public cloud resources, application stacks running on different compute resources, and have a combination of physical networks and SDN. It could also include active users across global regions. Using networking/SDN and data center firewalls to control East-West traffic in this environment is complex, unwieldy, and adds more chokepoints to manage. You do not want to add more firewalls and re-architect your networking architecture to achieve more fine-grained isolation of East-West traffic.

Moving security closer to the host and utilizing the existing stateful firewalls you already have in your operating systems offers several benefits:

- By programming the native stateful firewalls in each host, you do not have to buy more networking equipment, SDN, and hardware firewalls and appliances. You get more value out of the existing investments across your infrastructure.
- You will also get real-time visibility and control of the East-West traffic across bare-metal servers, virtual machines, containerized hosts, load balancers, and public clouds from a single console instead of cobbling together two or more solutions.
- You streamline visibility and control between multi-cloud and on-prem to public cloud connections. Since federal IT modernization initiatives involve the increased use of public cloud computing resources, the ability to provide visibility and fine-grained segmentation across heterogeneous environments, including multi-cloud and hybrid-cloud scenarios, is highly critical.
- You gain the ability to create, test, and enforce security policy without the risk of breaking applications.
- To secure your endpoints, programming the native stateful firewall in each laptop host to restrict inbound traffic to only the "allowlist" is easier to manage than GPOs and NAC.
- To secure your endpoints, you will also want a tool that takes advantage of and integrate with your existing investments in Endpoint Protection Platform (EPP) solutions like CrowdStrike.

4. Automation and orchestration with IT operations and DevOps

Federal initiatives and security standards like CDM, the HVA Process Framework, NIST CSF, and Zero Trust include language that requires you to maintain your segmentation posture continuously. A solution that integrates with IT operations ensures that can maintain your security posture even while the environment and IP connections change. This solution should also be tightly coupled with DevOps so that you have visibility into server-to-container and container-to-container traffic, for example, and ensure that Zero Trust segmentation is provisioned at the "birth" of a server or container component.

5. Continuous monitoring and enhanced vulnerability management

In order to continuously maintain your segmentation posture, you will want a solution that not only monitors for changes in the environment and IP connections, but also reports on multiple failed attempts to connect to blocked traffic. Integration with SIEM tools allows security staff to analyze these activities in a broader context and helps them assess if these actions are indicators of a malicious nature.

You also want to ensure that your vulnerability management program does not prioritize patching based only on the CVSS scores that third-party vulnerability scanning tools provide. These tools look at each host alone and do not contextualize the risk across the application environment. You need a solution that overlays vulnerability information on top of the application dependency map so that your security analysts understand the exploitability of vulnerable ports. This way, they can identify a bad actor's potential lateral attack pathways. This solution enables IT Ops and SecOps to see that the low to medium severity on a vulnerable and highly connected workload is a higher risk than a high severity workload with no live connections. Then they can prioritize accordingly. Ultimately, your solution should enable IT Ops to use vulnerability-enhanced segmentation as a compensating control when a patch is not immediately feasible.

Mapping Illumio controls to key agency obligations to improve the management and security of HVAs

The Office of the Federal Chief Information Officer publishes and maintains the HVA Process Framework, which defines key tasks, stakeholders, and requirements to help agencies execute a standard process for managing and securing HVAs. Figure 6 is a visual overview of the Agency HVA Process Framework, including the specific actions that comprise the process.

FIGURE 6

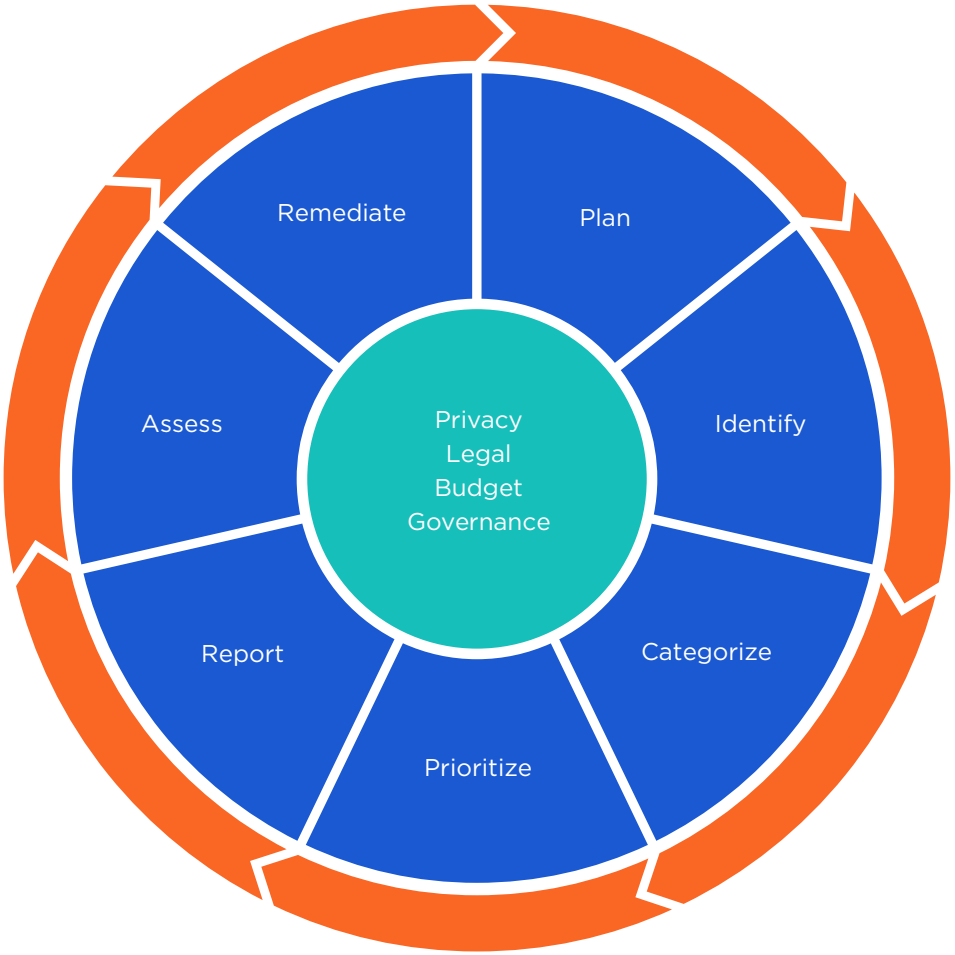

Source: policy.cio.gov/hva/process/

Illumio's host-based micro-segmentation software helps federal agencies effectively fulfill the following requirements.

HVA Process	Description of HVA Requirements	How Illumio Can Help
Plan	Agencies must develop, maintain, and regularly update their HVA inventory lists, at least annually.	Illumio's real-time application dependency map helps create and validate your HVA inventory list. It monitors and detects for changes in workloads and IP connections in real time, and identifies mislabeled and orphaned servers. The Illumio map will also show the endpoints that are connecting with the HVAs in an agency's data center and cloud. A traffic report displays the inbound peer-to-peer traffic to agency-issued laptops. This information is valuable for understanding the impact of remote work in your attack surface.
	Incorporate HVA activities into broader agency IT and information security and privacy management planning activities, including change management and Information Security Continuous Monitoring (ISCM) Strategy.	Illumio integrates with third-party ITSM, CMDB, IT Ops, and SIEMs via REST APIs and OOTB plug-ins so that data and processes from Illumio can be incorporated into change management and continuous security monitoring processes.
Identify, Categorize & Prioritize	Identify, categorize, and prioritize HVAs to ensure that information systems performing or enabling mission essential functions have been considered as potential HVAs.	Illumio's architecture and label-based policy model allow you to identify, validate, and categorize HVAs.
	Take a Federal enterprise-wide perspective of the risks posed by their HVAs and of their mission responsibilities to both identify their most critical functions, information, and data and to use that information to categorize information systems as critical mission enablers or mission essential functions.	Illumio's real-time application dependency map shows the essential connections across servers and applications. The traffic report on inbound peer-to-peer connections to agency-issued, remote laptops will improve visibility into your remote attack surface.

HVA Process	Description of HVA Requirements	How Illumio Can Help
	Protect that collection of HVAs according to the handling directions.	Illumio makes it easy for stakeholders to design and create policies for controlling allowed connections and traffic between servers that are in-scope as HVAs. It continuously monitors traffic within HVA servers and between HVAs and non-HVA systems. Based on policies, it can block and/or report on failed attempts to connect to HVA systems. For endpoints, locking down inbound traffic to agency-issued laptops will obstruct lateral movement via a compromised home network.
Report	All Federal agencies are responsible for keeping their internal HVA lists up-to-date. All CFO Act agencies are required to report all of their HVAs, including the prioritized top 10 list, to DHS on an annual basis. DHS will coordinate with OMB and other interagency partners to ensure appropriate oversight and governance across the Federal Government.	Illumio's historical database maintains an inventory of your HVAs. Lightweight agents monitor and detect for any changes in the HVA systems, including new servers and new IP connections, and also detects mislabeled servers. Authorized users can run a natural language query on Illumio's historical database using the "Explorer" capability or use the application dependency map to identify the top 10 HVAs visually. See section on FISMA CIO metrics below.

HVA Process	Description of HVA Requirements	How Illumio Can Help
	Implement and validate security controls including: Routine vulnerability scanning and remediation, increased monitoring and analysis of relevant audit logs, network segmentation, appropriate boundary protections, maintenance of 100% automated asset visibility and control.	Illumio enables agencies to design, build, and operationalize a Zero Trust architecture that scales massively. Illumio's default-deny model blocks unauthorized connections. It continuously monitors traffic and connections between servers, so that blocked and quarantined traffic and failed attempts to connect can be viewed via Explorer for further analysis. You can also export to tools like Splunk and QRadar for additional analysis and investigation. Illumio integrates with third-party SIEM and vulnerability management tools via rich APIs or OOTB plug-ins, and can also ingest third-party vulnerability scan data from tools like Qualys, Tenable, and Rapid7.
	Agencies must identify the connections between HVAs and other systems, including other HVAs and non-HVAs, to understand critical dependencies.	Illumio's real-time application dependency map, combined with a label-based policy model, presents the relationships and dependencies in easy-to-understand business context. Users can see the application groups, application tiers, and geolocation of connections across HVAs instead of just viewing a group of servers with IP addresses.
Remediate	The agency must complete its remediation plan expeditiously and should treat it as a priority. The remediation plan must include actions, milestones, and timelines for remediating the weaknesses or deficiencies identified in the assessment's findings.	Illumio helps agencies build and operationalize a Zero Trust architecture. You can use Illumio to identify gaps (such as overly-broad segments) and then ringfence the HVAs. You can also use process-level segmentation as a compensating control.

Note: For a detailed view of the HVA Process and its requirements, see: policy.cio.gov/hva/process

FISMA CIO metrics for HVA security

FISMA requires agency heads to report on the adequacy and effectiveness of their enterprise's information security policies, procedures, and practices. The 2020 CIO FISMA metrics focus on each agencies' progress on strengthening federal cybersecurity. The metrics are organized around the NIST Cybersecurity Framework, NIST SP 800-37 and 800-39, as well as standards that specifically address HVAs, namely OMB Memo M-19-03 and DHS BOD 18-02. FISMA metrics have defined target levels that define the minimum thresholds. Agencies use the metrics as indicators of a security program's efficacy and to identify critical gaps and update roadmaps.

The critical capabilities mentioned in the previous section can be used to prepare and validate the information in the FISMA CIO metrics report. Here are a few examples:

- Identify – Validate and update the HVA asset inventory and enhance the accuracy of the asset count using real-time application dependency mapping and reporting from a historical traffic database.
- Protect – Increase the count of HVA systems that are protected by micro-segmentation and data at rest encryption. Agencies can also improve protection scores by using vulnerability map-enhanced micro-segmentation as a compensating control for workloads with vulnerabilities.
- Detect – Improve network defense scores by using micro-segmentation as a preventive tool that, by default, denies connections from unauthorized workloads and devices.
- Respond – Enhance respond scores by increasing the count of HVA systems that are continuously monitoring new and changes to IP connections and automatically blocking traffic that is not in the "allowlist."

Conclusion

Securing HVAs is an immediate and imperative priority for federal agencies and commands. An agency's mission can be compromised when internal networks are flat, leaving HVAs unprotected from lateral movement attacks. Implementing host-based micro-segmentation to operationalize Zero Trust in your networks, workloads, and applications will increase your cyber resiliency, stop lateral movement, and improve your FISMA score.

Illumio helps agencies identify, validate, and prioritize their HVAs, as well as their interdependencies with other systems. Agile host-based micro-segmentation is a proven strategy to monitor distributed environments, shut down the attack surface, and reduce the exposure of your most critical assets. Illumio further helps agencies operationalize these objectives while avoiding the management complexity, costs, and risks that stem from using networking and hardware firewalls to control connections across HVA assets in dynamic and multi-cloud environments.

Illumio certifications include:

Common Criteria

FIPS VALIDATED 140-2

AICPA SOC

GSA Section508.gov GSA Government-wide IT Accessibility Program

[7] cisa.gov/sites/default/files/publications/FY_2020_IG_FISMA_Metrics.pdf

Illumio is a cybersecurity software company enabling end-to-end Zero Trust in Defensive Cyberspace Operations. The company helps agencies, commands, and organizations achieve Zero Trust and prevent attacker lateral movement by protecting high value assets, critical applications, and workloads through real-time application dependency mapping, coupled with host-based micro-segmentation. Illumio is FIPS 140-2 validated and NIAP Common Criteria Protection Profile Certified. Illumio can be placed in multi-vendor hardware environments, using existing infrastructure to improve agencies' cybersecurity postures and effectively accomplish their missions.

Gartner peerinsights™

See what customers have to say about Illumio.
gartner.com/reviews/market/cloud-workload-protection-platforms/vendor/illumio