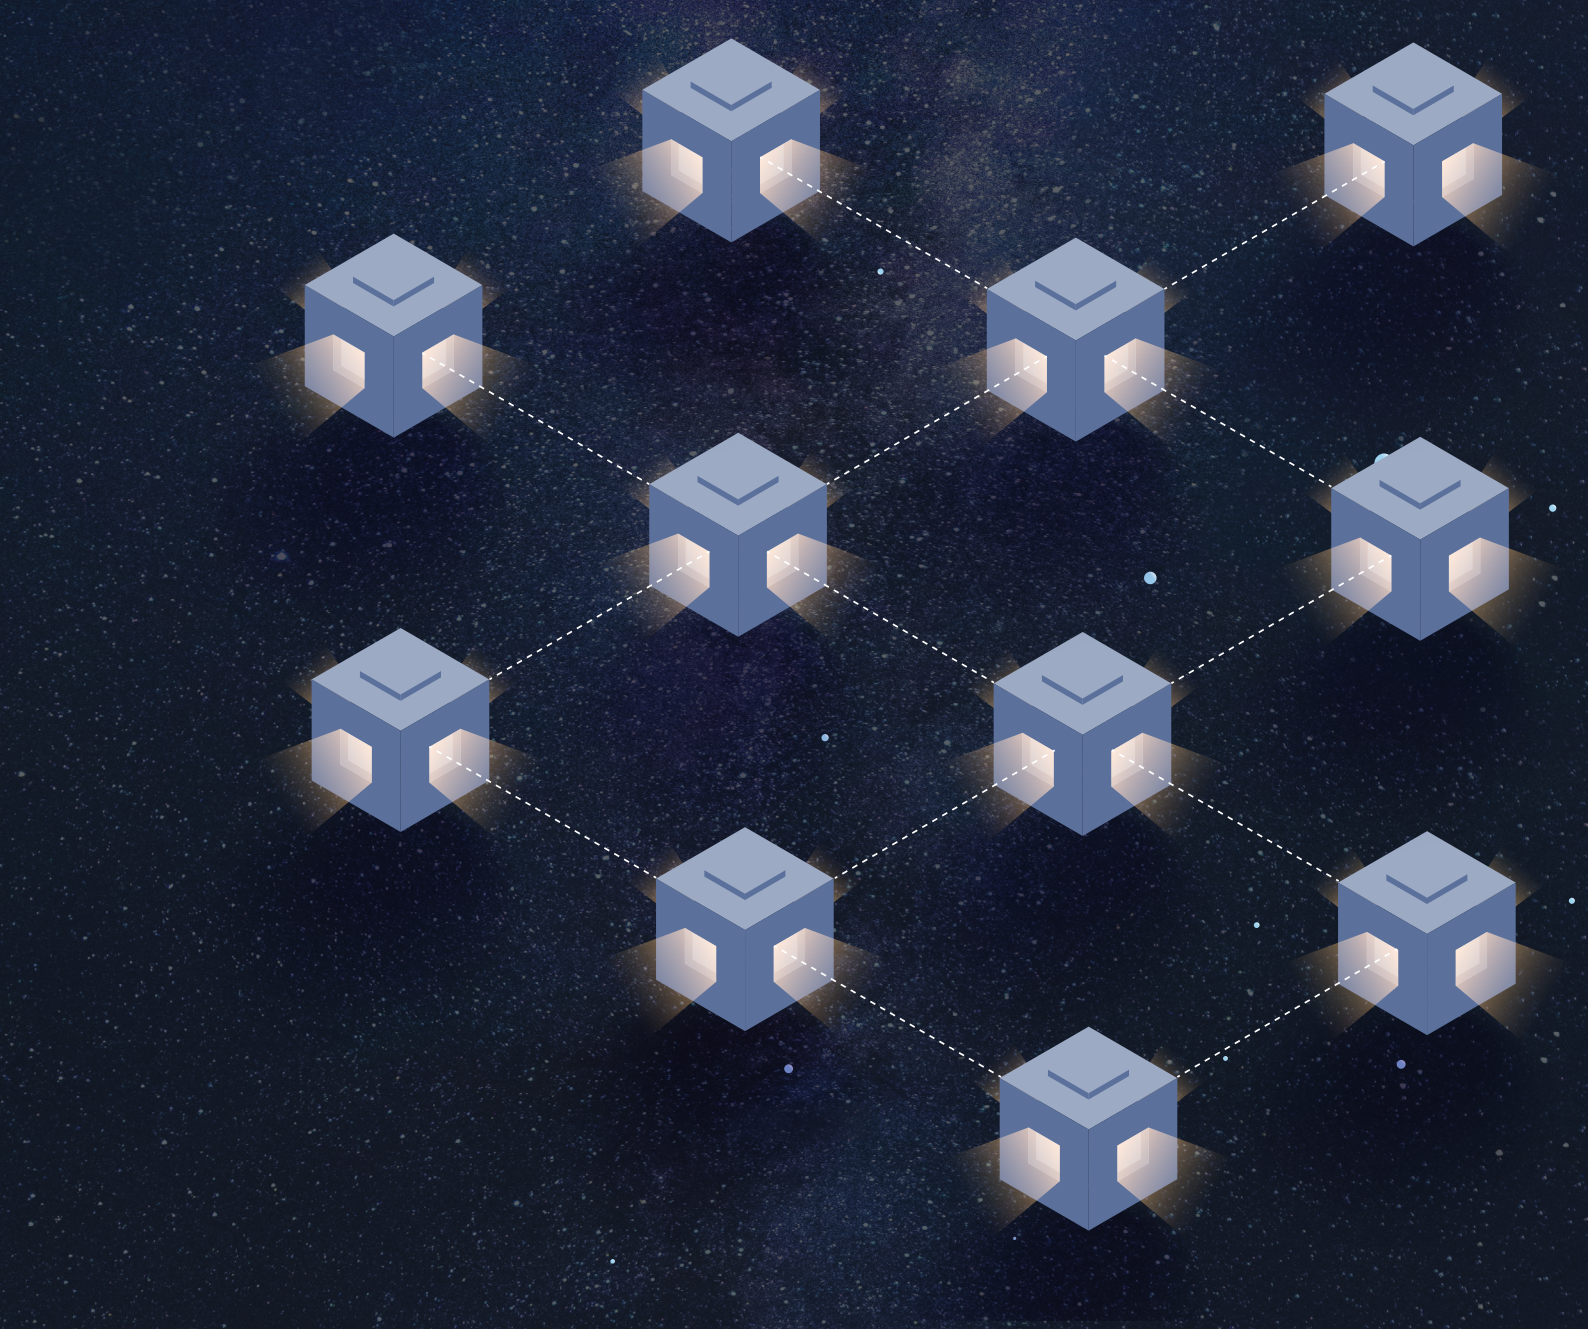

Core Concepts Guide

Implementing Zero Trust Principles for Continuous Security Monitoring

HYPATIA ★ VETS

Decoding Core Concepts

Achieving Zero Trust and Continuous Compliance
New Jargon for the New Normal

The composition of the modern data center, the granularity of workloads and the speed of development is changing rapidly. As cloud computing and artificial intelligence continue to pervade daily life, a new attack surface has emerged. New concepts of risk and mitigation strategies are defining conversations in cybersecurity communities of practice. The "new normal" in cloud-native cybersecurity has implications for continuous monitoring methods and requirements. This guide is a starting point and not intended to be a comprehensive terminology guide or a substitute for well-established lexicons such as the NISTIR 7298 Glossary of Key Information Security Terms.

Foundation of Zero Trust with Five Pillars

What is Zero Trust?

Zero Trust employs the principles of “never trust, always verify” and treats everything inside and outside of a network as suspect. It shifts the focus from securing the **network perimeter** to securing **data**, **assets**, and ators within a networked environment or ecosystem.

In the government setting, Zero Trust requires all users and devices, whether in or outside a federal agency’s network, to be authenticated, authorized, and continuously validated for security configuration and posture before being granted access or keeping access to cloud systems, applications, and data.

The President and the Office of Management and Budget (OMB) have issued important guidance on Zero Trust implementation for government agencies in the *Executive Order on Cybersecurity* and the *Memo on Zero Trust Strategy*. To secure the government’s digital infrastructure, agencies are adopting Zero Trust frameworks like those outlined by the Cybersecurity and Infrastructure Security Agency (CISA), Department of Defense (DoD), and National Institute of Standards and Technology (NIST).

A **Zero Trust** approach to the cloud removes trust assumptions when users, devices and applications connect: each asset must individually validate the others trying to communicate with it before consenting to share data. In order to implement a zero trust model, organizations need to integrate security into the workloads themselves, and have the security 'travel' with the instances and data as they migrate between/among cloud environments.

A closely related concept is the **secure access service edge (SASE). Zero-trust browsing, or remote browser isolation (RBI)**, is an approach that enables organizations to open up web access while improving threat-prevention and security. RBI "walls off" malware from the endpoint, regardless of the level of trust and organization places on the site, by executing active web content in a remote, isolated container in the cloud. For example, a media stream representing a website is sent to the endpoint's browser, providing a seamless user experience; whether the user browses to a malicious site on his own or reaches one by clicking a URL embedded in a phishing email or a malicious PDF document, no web content is ever executed directly on the device.

Adaptive Microsegmentation facilitates application security and provides a mechanism for implementing zero trust model, starting with the premise of denying all traffic and permitting only communications that are whitelisted, either explicitly, or through a machine learning mechanism that evaluates behaviors within the environment and dynamically adapts policies to changes.

Microsegmentation, also referred to as Zero Trust or identity-based segmentation, delivers on segmentation requirements without the need to re-architect. Security teams can isolate workloads in a network in order to limit the effect of malicious lateral movement. Microsegmentation controls can be assimilated into three categories:

- **Agent-based** solutions use a software agent on the workload and enforce granular isolation to individual hosts and containers. Agent-based solutions may leverage the built-in host-based firewall or derive isolation abilities based on workload identity or attributes.
- **Network-based** segmentation controls rely on the network infrastructure. This style leverages physical and virtual devices, such as load-balancers, switches, software-defined networks (SDN), and overlay networks to enforce policy.
- **Native cloud controls** leverage capabilities embedded in the cloud service provider (e.g., Amazon security group, Azure firewall, or Google Cloud firewall).

Community Immunity is a cybersecurity model based on increasingly popular practices of community integration and threat intelligence sharing. Advanced CWPP offerings share threat intelligence across their community of users, helping to identify interenterprise patterns that are not visible in a single organization alone. By sharing telemetry and analysis, there is value in broader "community immunity." By obfuscating the telemetry that is shared, CWPP vendors balance the enterprise need for privacy with the community need for protection. The implementation of community immunity is a modern form of **integrated risk management,** providing a software-defined approach that enables an organization to share risk information and risk analysis and to synchronize independent yet complementary risk management strategies.

Console and Integrations. The consoles provided by Cloud Workload Protection Platforms (CWPP) are only effective when they support the logical grouping of workload type and policy application. As agencies select vendors, they will want to ensure their console provides the ability to define agency-defined logical naming and tagging to workloads, and that it provides native integration with the APIs of Azure, AWS, Google and others. The console should be able to import and understand the tagging of the underlying cloud platform for the simplification of policy formation. Agencies have options: they may choose inherited consoles or they will need to commit console design expertise to develop a configuration driven entirely by scripts and CI/CD pipeline tools.

Continuous Adaptive Risk and Trust Assessment (CARTA) is a framework Gartner introduced in 2019 in response to broad adoption of cloud-native technologies. CARTA proposes a new strategic approach for information security based on the emerging class of vulnerabilities resulting from **Cloud Infrastructure and Platform Services (CIPS).** The increasing adoption of cloud and platform services (IaaS and PaaS), combined with a lack of cloud skills (including security), have left enterprise information and workloads exposed. Compounding the issue is a lack of comprehensive visibility into programmatic cloud infrastructure – this means that incorrect and noncompliant configurations go undetected for extended periods of time.

Cybersecurity posture is a relative measure of the overall security strength of an enterprise across policies, processes and controls for proactive and reactive protection of digital assets. While perfect protection is not possible, with a continuous monitoring strategy that optimizes adaptive tools designed for cloud-native environments, agencies can risk-optimize its security posture to reduce the likelihood of a successful attack, and, when an attacker gains a toehold, reduce the ability to cause damage. **Cloud Security Posture Management (CSPM)** vendors protect workloads from the outside by assessing secure and compliant configuration of cloud platform's control plane.

CSPM offerings provide a strategic complement to **Cloud Workload Protection Platforms (CWPP),** which are designed to protect workloads from attack. CWPP vendors provide workload-centric security offerings (network segmentation, application control, behavioral monitoring, host-based intrusion prevention, anti-malware protection) that target the unique protection requirements of server workloads in modern cloud-native environments and multicloud data center architectures.

Limiting the Blast Radius is a concept that applies to the scope of cyber attack damage – that is, containing it to the smallest possible surface area. In this construct, attackers are prevented from leveraging one compromised asset to access another because you deploy mechanisms that control asset-to-asset communication (see adaptive microsegmentation concept) and evaluate the applications running, assessing what these applications are trying to do. This is increasingly important with artificial intelligence applications.

Containerization. As cloud computing infrastructure continues to improve, so does application architecture. Containerization offers a lightweight alternative to traditional, resource-intensive virtualization. Vendors in this category deploy and/or manage container solutions for cloud environments and facilitate all-sale Dev Ops for software factories and artificial intelligence applications. Container usage for production deployments in enterprises is still constrained by concerns regarding security, monitoring, data management and networking; however, container security tools are maturing.

Models of Virtuous Deception include security protection capabilities that creates fake vulnerabilities, which may share “honey data” or “honey tokens” to lure individual or coordinated attackers. If an attacker or attacker group tries to access or use these fake resources, this is a strong indicator that an attack is in progress, because a legitimate workload should not see or try to access these resources. A closely related concept is a botnet: a collection of computers or applications compromised by malicious code and controlled across a network. A Bot Master or Bot Herder is the controller of a botnet that, from a remote location, provides direction to the compromised computers or applications in the botnet.

Host-based segmentation. This model enables segmentation across highly distributed, heterogeneous environments with both consistency and granularity, decoupling security segmentation from the network. The idea is to support many disparate environments by orchestrating policy through centralized controls. This model allows organizations to run workloads in any environment and protect applications wherever they run. Microsegmentation is a related but distinct concept: By creating very granular segments within an IT infrastructure, an organization effectively limits the size of their network’s attack surface by breaking it into a lot of small pieces. If a particular segment gets compromised, the other segments are “walled-off” and protected. Next-generation firewalls (NGFWs) are key segmentation-enabling tools that require complementary techniques beyond NGFW, such as software-defined Perimeters (SDPs), cloud-access service brokers (CASBs), encryption and proxies to roll out a microsegmentation approach that protects their data wherever it resides.

Cloud Native Monitoring provides organizations a console with tools that provide the ability to monitor and analyze the performance of cloud infrastructure in real time. The increasing complexity of cloud environments makes addressing cloud complications in a timely manner difficult. Importantly, as organizations adopt running their applications and data on third-party infrastructure, they are often surprised to learn that they no longer have visibility into, or control over, their entire network. To address this market gap, vendors are developing solutions that monitors cloud infrastructure and application performance.

Pervasive Secure Sockets Layer/Transport Layer Security (SSL/TLS) is replacing "man-in-the-middle" traffic decryption approaches by performing decryption and inspection at the host workload, where the session is terminated. This is strategy is valuable for inspecting traffic that moves laterally ("east/west") from service to service in microservices-based architectures.

Running without runtime protection agents. With containers and serverless architectures, the workloads should be scanned in development, meeting foundational requirements for application control. With prescanning, runtime protection needs (i.e. segmentation, behavioral monitoring) may be delivered outside of the workload. This strategy is increasingly used with the adoption of immutable infrastructure and container/serverless PaaS life spans, which are measured in minutes, not hours or days.

Contact Information

Innovation Hub
(Palo Alto, CA)

Integration Center
(Round Rock, TX)

Headquarters &
Training Center
(Washington, DC)

Digital Lab
(Fredericksburg,VA)

Contact Us for Zero Trust Architecture and Implementation Planning Services

Website: www.hypatia-vets.com

Email: meg@hypatia-vets.com

Office: 571-346-7612

HQ Address: 1818 Library Street, Suite 500, Reston, VA 20190

Hypatia UEI: HF9JYGZE2N15 (WOSB)

HypatiaVETS UEI: YYAJPKWUFM39 (WOSB + SDVOSB)

We provide vendor management and 24-hour pricing across best-in-class cybersecurity, cloud and infrastructure OEMs. We've been serving the Department of Defense and the Fortune 500 since 2016.