The bridge to possible

Cisco Hypershield

AI Scale | Cloud Native | Highly Distributed

Security for the AI-scale data center

AI is transformative, driving huge productivity gains. The engine of AI—the data center—is growing substantially because of the AI revolution and has fundamentally changed in two ways. One, the **infrastructure** is evolving to advanced computational systems like Graphic Processing Units (GPUs) and Data Processing Units (DPUs). Two, the **applications** that run on the data center have evolved from the traditional 3-tier architecture to using multiple microservices running on Virtual Machines (VMs) or containers.

To secure the AI-scale data center, we must reimagine security, as traditional appliances won't do the job.

Cisco Hypershield is the first truly distributed, AI-native security architecture that puts security wherever it needs to be: in every software component of every application running on the network, on every server, and in public or private cloud deployments.

It provides AI-powered management that automates security policy lifecycle and security infrastructure upgrades. At the same time, the system is designed to empower customers to determine the level of autonomy they are comfortable with - using test, record and report capabilities to earn trust. Imagine a network security solution that can write its own rules, test its own rules, deploy its own rules and lifecycle manage its own rules.
And it can even upgrade itself!

Hypershield deeply combines security and networking in a way only Cisco can, by taking the network security functions that used to come in a box and "melting them" into the network.

More a fabric than fence, Hypershield allows security enforcement to be placed everywhere it needs to be. It lets you embed security in VMs or Kubernetes clusters in public clouds. In the private cloud, security can be inserted in VMs and high-performance server DPUs. Our vision is to extend Hypershield to hardware accelerators running on networking devices such as switches, providing security beyond the data center to IoT/OT environments. Before long, a hospital will be able to secure its medical devices and other operational technology with Hypershield. Manufacturers will be able to do the same with the tech that sits on the factory floor.

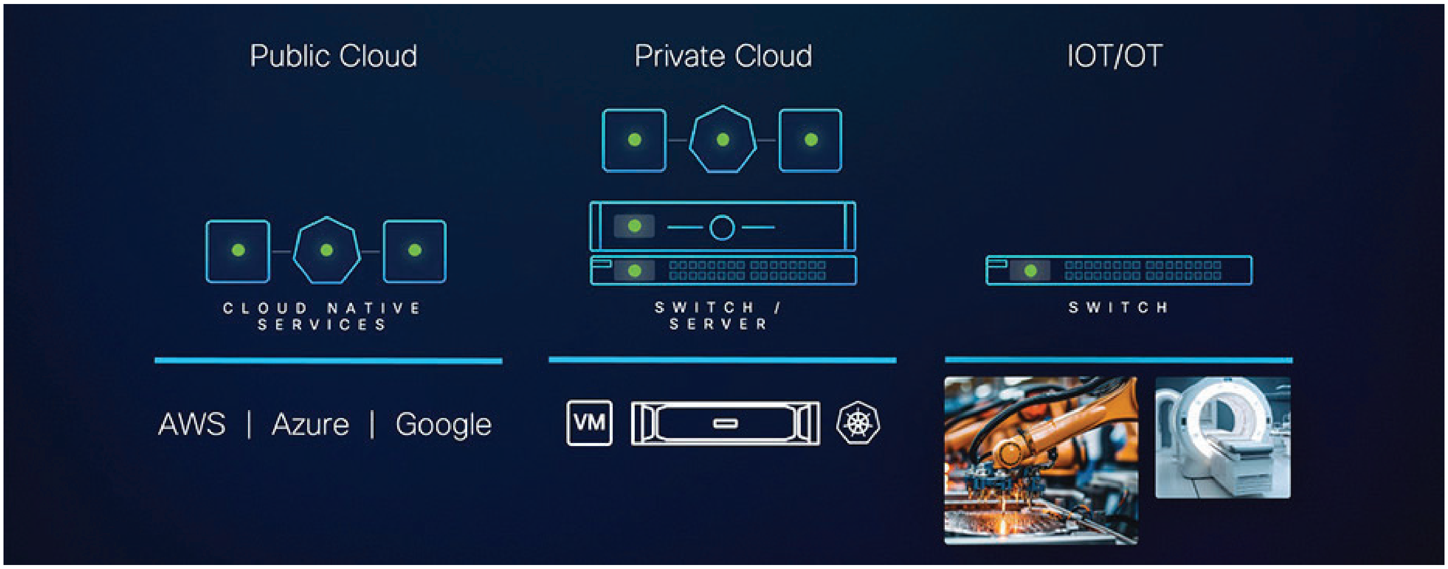

Figure 1. Hypershield puts security wherever you need it

Not the next generation of anything. The first generation of something new.

Cisco Hypershield reimagines security with unique capabilities that the solution is built around.

- **AI-native:** By designing Cisco Hypershield from the ground up to leverage the power of AI, it's orders-of-magnitude more autonomous than other security solutions. In fact, since Hypershield was built from the beginning around AI management, we think of it as AI-native, as opposed to an AI layer bolted on top of a traditional product.
- **eBPF enforcement:** Hypershield provides deep visibility and enforcement at the workload level using an open-source technology, eBPF. Co-created by Isovalent that was recently acquired by Cisco, eBPF is a software framework on modern operating systems that enables programs in user space to safely carry enforcement and monitoring actions via the kernel.
- **Self-qualifying updates:** Hypershield was designed to be self-upgrading and updating. Because of the distributed architecture, the eBPF agents that send in the telemetry also act as enforcement points, using a patent-pending design that brings the continuous update CI/CD model of the cloud to premises-based systems, whether at the network, workload, file or process level.

On top of these platform capabilities, there are **modules** that deliver specific outcomes like segmentation and protecting against vulnerability exploits that we'll cover in the next section. Hypershield is a **subscription-based software product** that sits on top of the hardware.

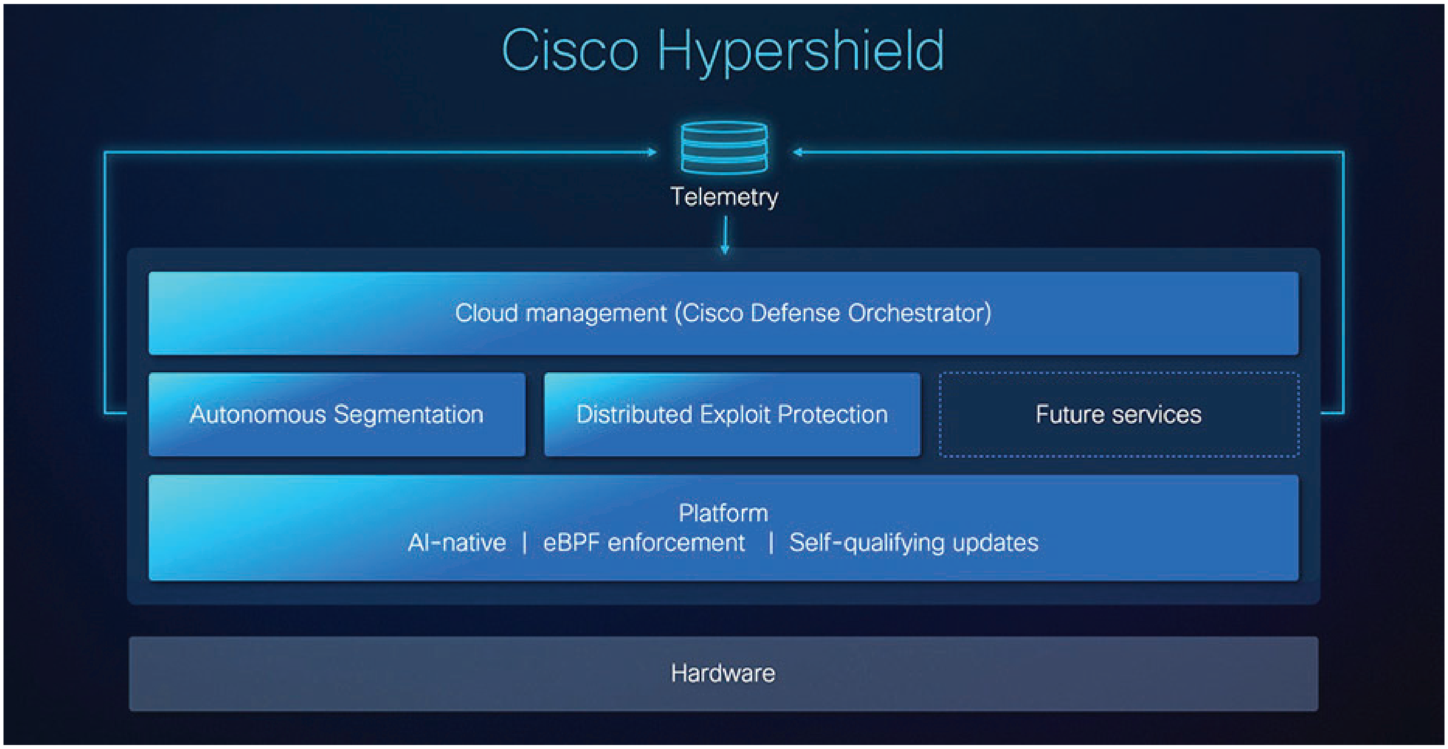
Cisco Hypershield
Telemetry
Cloud management (Cisco Defense Orchestrator)
Autonomous Segmentation
Distributed Exploit Protection
Future services
Platform
AI-native | eBPF enforcement | Self-qualifying updates
Hardware

Figure 2. Cisco Hypershield architecture

Centralized security policy

With the rise of a highly distributed IT environment came the challenge of management and consistent enforcement of policies across multiple domains. The scale of policies within an enterprise is staggering, with hundreds of thousands of rules that go in but never come out, written by folks who might no longer be within the organization. This level of complexity has been difficult to manage – until now. With AI-powered capabilities, we now have an opportunity to help security administrators implement strong, consistent and dynamic policies at scale.

Hypershield is AI-powered and uniquely architected to implement a truly intent-based policy model that is centralized and easy to manage. No matter the form factor or location of the enforcement point, the policy being enforced is organized at a central location by Hypershield's management console. When a new policy is created or an old one is updated, it is "compiled" and intelligently placed on the appropriate enforcement points. Security administrators always have an overview of the deployed policies, no matter the degree of distribution in the enforcement points. Policies can follow workloads as they move, from on premises to the native public cloud.

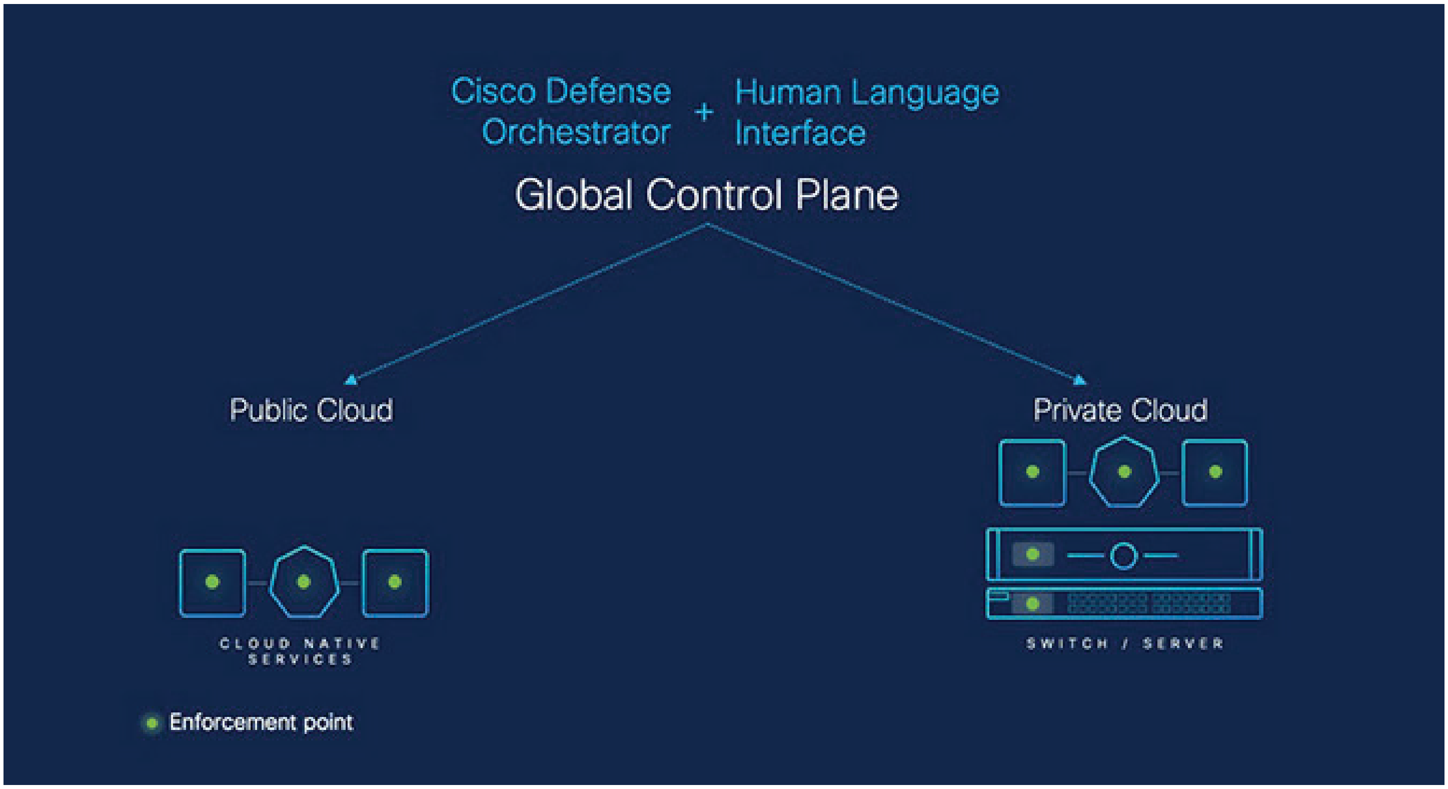

Figure 3. Centralized management for distributed enforcement

Solving real customer challenges

Cisco Hypershield can solve the critical problems organizations are facing today.

1. Autonomous Segmentation

Companies have been using segmentation for decades to specify which workloads and applications can access which parts of the network to prevent lateral movement in case of an attack. But segmentation is hard. Customers tell us it can take 40 days or more to define segmentation rules for a single application. That's way too slow. Current segmentation tools lack a deep understanding of the application and try to baseline application behavior based on time only without taking into account app-specific events, which isn't effective.

Take the example of a factory app that does sheet metal delivery. Traditional tools will observe the app for a certain period of time like 90 days and formulate segmentation policies based on the behavior observed in that period. But what if the factory runs out of sheet metal on day 91 and this event triggers communications from the app to the various systems to place a new order? This legitimate behavior will come across as random when it's not. That's why you need a continuous and intimate understanding of the application.

With the AI-native Hypershield, we look beyond the network flows that other products focus on. The full scope of observed behaviors is informed by what's happening across all the environments it's protecting; what threat intelligence teaches it about behaviors that should never happen; the latest attack vectors, techniques, and vulnerabilities; what the system has learned and observed based on best practices that model how the customer modifies recommended policies; and what the customer does when they step in when under attack.

The system begins with understanding what's in the environment and who is talking to who. Based on that, macro guardrails are created to capture governance requirements and keep the business safe. As Hypershield learns more based on the attributes mentioned above, it tightens those policies further. This is a continuous and dynamic process, so if the application changes or moves, the segmentation policies will relax, and as the system learns about the new behavior, we can tighten it back again.

The result is higher-confidence, data-backed recommendations based not on what might have happened in the past but on what is happening now. The recommendations are automatically tested against live traffic and presented to the user with results on effectiveness and any performance impact of the policies. Once the user is satisfied and accepts the recommendation, only then it is deployed. So the system is autonomous but it earns trust and is continuously learning.

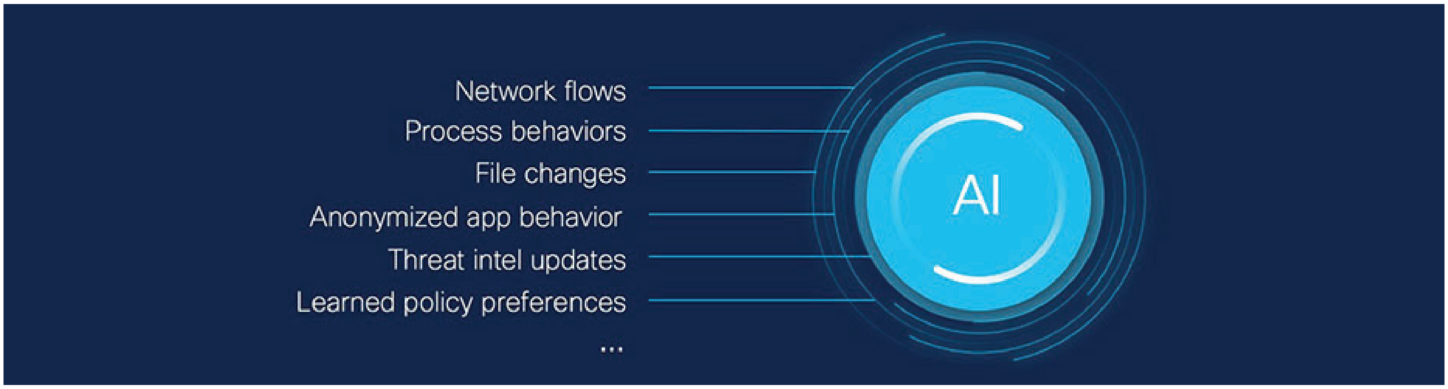

Figure 4. Comprehensive inputs for segmentation policy creation

2. Distributed Exploit Protection

Attackers have gotten very good at either stealing a credential or compromising the services of an application so that they can move through legitimate application pathways. You need to understand the services that make up an application and its vulnerabilities. But even if the vulnerabilities are identified, patching them is hard; it takes time (weeks or months) and, in some cases, is not possible. Whereas attackers begin exploits within a few hours of a vulnerability being published.

Hypershield draws a picture of your entire inventory and integrates with your existing vulnerability tools. However, with around 500-1000 CVEs being published every week, it is difficult to prioritize and fix those vulnerabilities. This is where Hypershield's AI capabilities and deep understanding of the application help prioritize the most critical vulnerabilities specific to the organization's environment based on three key questions:

- Is the vulnerable code module running in memory?
- Is the vulnerability theoretical or being exploited in the wild?
- Is this vulnerability affecting a high-value asset?

While the application team takes the time to qualify the patch, Hypershield applies a surgical compensating control in the path of the application to prevent an exploit. Once the patch is applied, the compensating control gets removed automatically.

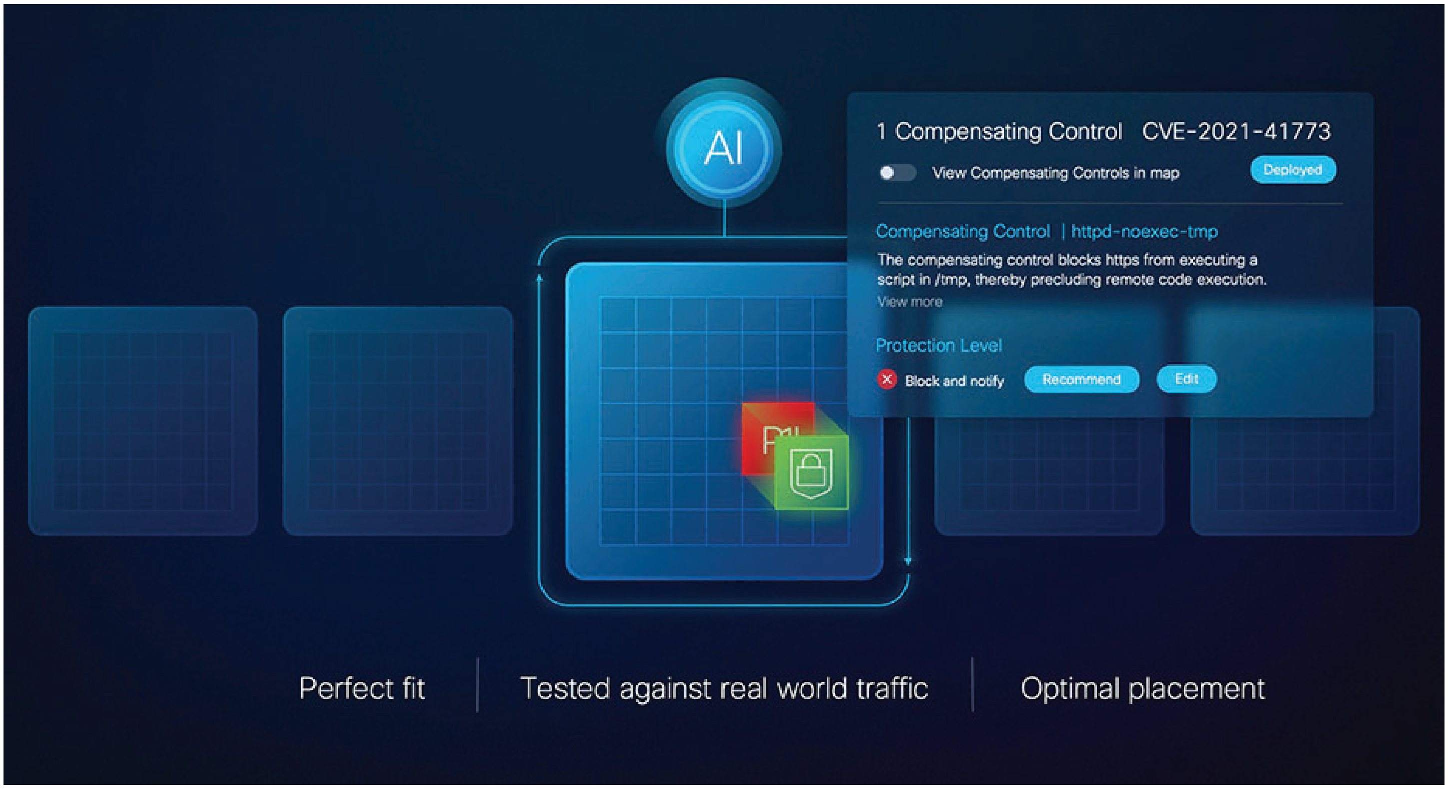

1 Compensating Control CVE-2021-41773

View Compensating Controls in map

Compensating Control | httpd-noexec-tmp

The compensating control blocks https from executing a script in /tmp, thereby precluding remote code execution.

Protection Level

Block and notify

Figure 5. Hypershield applies compensating controls to protect against vulnerability exploits

3. Self-qualifying updates

The nature of security controls is such that they tend to get outdated quickly. Sometimes, this happens because a new software update has been released. Other times, new applications and business processes force a change in security policy. Traditionally, neither scenario has been accommodated well by enforcement points – both acts can be disruptive to the IT infrastructure and present a business risk that few security administrators want to undertake. A mechanism that makes software and policy updates normal and non-disruptive is called for.

Cisco Hypershield has precisely such a mechanism, called the **dual dataplane**. It supports two data paths: a primary dataplane and a shadow dataplane. Live, real-world traffic is replicated between the primary and the shadow dataplane - a digital twin running in every enforcement point within your environment, not in a lab or a simulation. Software updates are first applied to the shadow dataplane, and when fully vetted and accepted by the user, the roles of the primary and shadow dataplanes are switched. Similarly, new security policies can be applied first to the shadow dataplane, and when everything looks good, the shadow becomes the primary.

The dual dataplane concept enables security administrators to implement software upgrades and policy updates on enforcement points without fear of business disruption or impact on performance. This critical functionality helps deliver the outcomes of Autonomous Segmentation and Distributed Exploit Protection discussed above by testing and deploying policy updates as well as compensating controls.

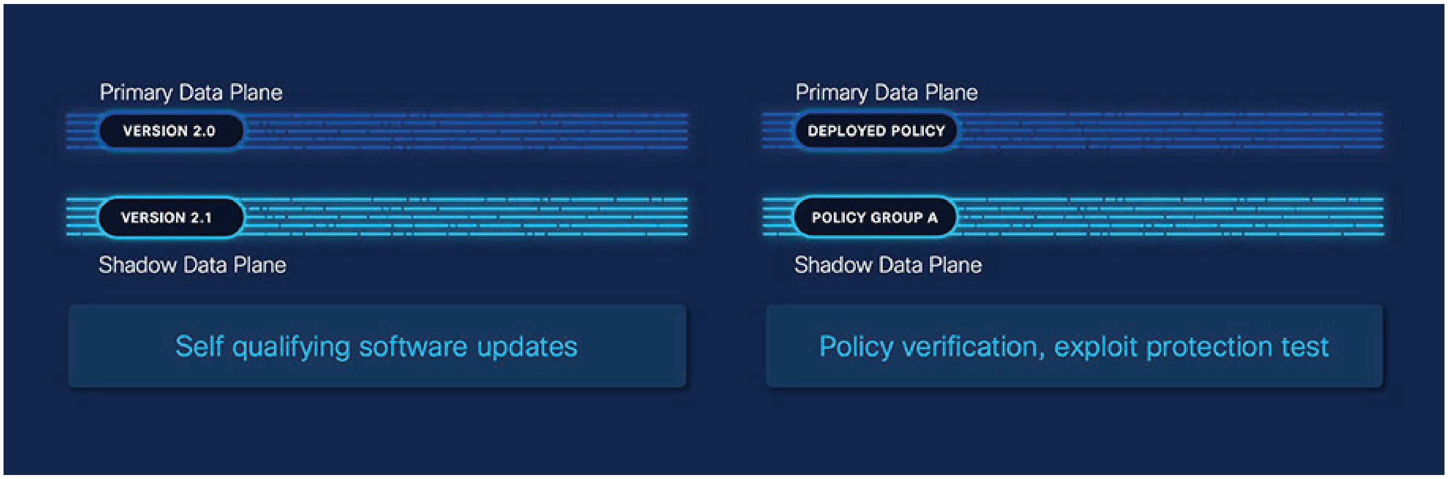

Figure 6. Cisco Hypershield dual dataplane

Conclusion

With Cisco Hypershield, we have fundamentally reimagined security for the AI-scale data center. The AI-powered capabilities, the deep visibility and enforcement right down to the kernel level and the self-qualifying updates with the dual dataplane, make Cisco Hypershield a powerful network security solution that writes its own rules, tests its own rules, deploys its own rules and lifecycle manages its own rules. And it can even upgrade itself. At the same time, you can set the dial for autonomy within the AI too, increasing it as the system earns your trust with its ability to test, record and report everything. This remarkable, almost magical capability is only possible because it was purpose built with AI management, another example of being AI native.

Additional resources

- Solution web page
- Video: Unveiling a New Era of AI-native Security - Cisco Hypershield launch
- Blog: Cisco Hypershield: Security Reimagined - Hyper-Distributed Security for the AI-Scale Data Center by Jeetu Patel, EVP and GM, Cisco Security and Collaboration Business Group
- Blog: A New Era of Distributed, AI-Native Security by Tom Gillis, SVP Cisco Security Business Group
- Blog: Reimagining Security by Craig Connors, VP and CTO Cisco Security Business Group
- Blog: Our Vision to Combat Unknown Vulnerabilities by Craig Connors
- To stay up-to-date on product availability, demos and other news, please sign up here