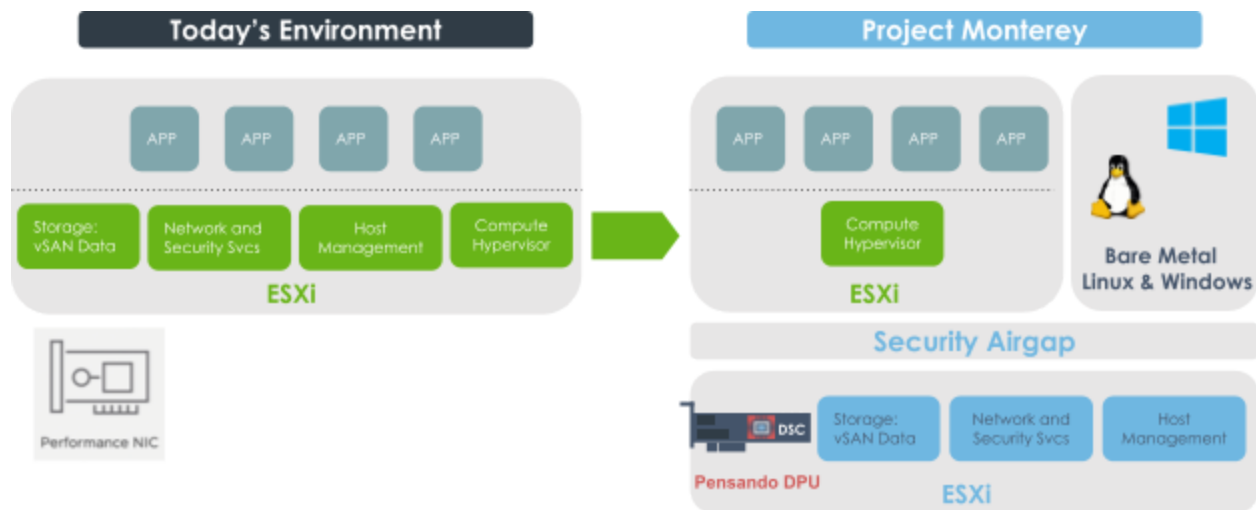


The Road to Monterey

Introduction

In September 2020, VMware announced *Project Monterey*, a significant shift in their hardware architecture focus. Taking a page from the hyperscalers, VMware is moving key infrastructure functions such as networking, security, and storage off of the server, where they contend with applications for expensive x86 server resources, and accelerating them in domain-specific hardware from Pensando and other vendors.

Infrastructure functions that formerly contended for CPU resources are now efficiently and securely hosted on domain-specific hardware.

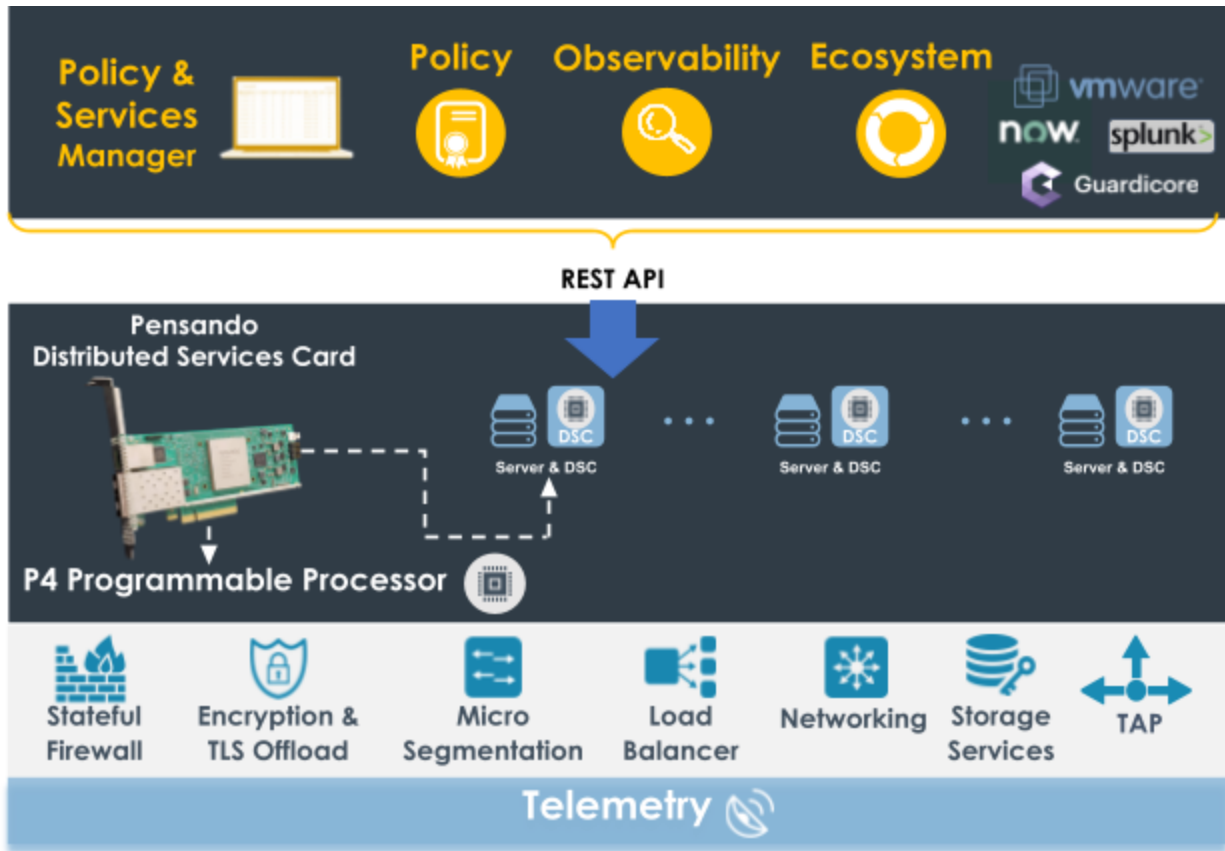


Over the last decade, the hyperscalers (AWS, Microsoft, Netflix, among others) have been identifying and addressing the limits of workloads deployed in a traditional leaf/spine architecture. These same limits are now becoming a reality for all levels of enterprise data centers.

Bringing Hyperscale to Enterprise Data Centers

In the past, the levels of integration and scale that are a necessity in the cloud world was not available to an enterprise (or carrier) data center—it depended on the investments and engineering staff of the hyperscalers. The **Pensando Distributed Services Platform** has democratized **cloud-level scale**, **security**, and **observability**, putting it within the reach of any organization.

The Pensando Distributed Services Platform is a comprehensive next-generation infrastructure for delivering software-defined services at the server edge.



Since function accelerators like the **Pensando Distributed Services Card (DSC)** are securely separated from the host operating system and hypervisor, they are protected from attacks based on malware or rogue root admin access. Regardless of what happens in the application layer, infrastructure services are isolated, protected, and unable to be compromised.

Not only is this the model VMware has chosen to support Project Monterey, it is the foundational vision of Pensando and the solutions we are bringing to our customers to scale their data centers, regardless of how they host their workloads.

The Road to Monterey

VMware, noting the lessons of the hyperscalers, is bringing their networking, security, and storage functions into the server edge. By optimizing the infrastructure services they deliver today on next-generation function accelerators, they no longer must rely on nor impact server CPU resources: the Pensando DSC can now host Project Monterey for VMware at scale with no impact on the server host. This allows any enterprise to scale their workloads on each server. This approach is not only valid for virtualized applications, but for the first time, bare-metal hosted services, too.

But now the focus moves from architecture to operations. While the need to pivot to infrastructure services at the server edge should now be apparent, the operational transition to supporting this architecture deserves an extended conversation. If the goal is to move functions formerly deployed in the ToR, appliance, or VMs in the data center, to the server edge, this new model for simplification and scale needs to be operationalized.

The Journey Is (Also) the Reward

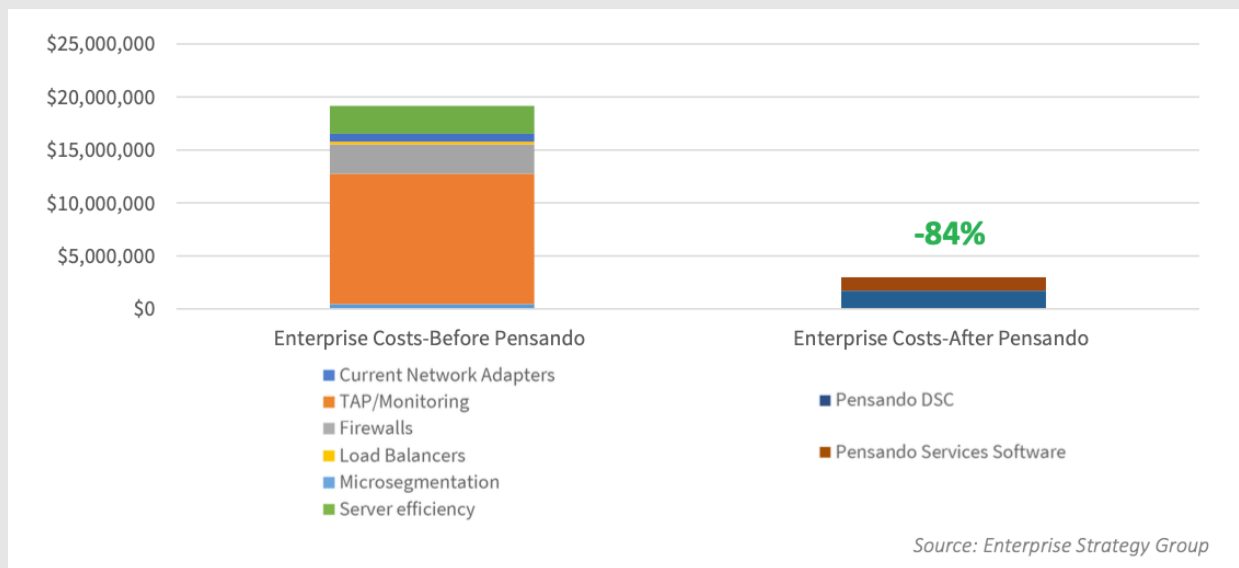
For the operations team, introducing any new solution must be an incremental approach, and the Pensando platform is designed to accommodate this. As new features are gradually introduced, policies can be developed, new insights can be integrated, and the handoff from server to network can be defined.

Pensando is uniquely positioned to provide a phased transition—the *Road to Monterey*—not as an abrupt shift that suddenly breaks existing management models before any new benefits can be realized, but in a way that customers can gain the practices and skills required for this new architecture without disrupting operations. Pensando is unique because it has built a robust platform to support data center/cloud deployments, based on a P4 programmable processor, a full software stack for that function accelerator, and centralized policy, services, and life cycle management to abstract and simplify all that the platform can now deliver.

Three-year TCO Benefits Provided by Pensando in the Enterprise Data Center: The Enterprise Strategy Group (ESG) reviewed the Pensando platform and uncovered economic benefit for enterprises, carriers, and cloud service providers.

ESG's analysis found that Pensando's scale-out software-defined services approach enabled organizations to centralize management, simplify administration, and optimize performance, confirming that the Pensando platform allows customers to consolidate network monitoring, eliminate appliances, reduce network downtime, increase server utilization efficiency, and improve security.

ESG's modeled scenarios demonstrate significant savings for both traditional enterprises and cloud service providers, and conversations with real-world customers confirm that. ESG found that an enterprise with 2,000 servers could save **84% over three years** by consolidating network monitoring, east-west firewalls, load balancers, and micro-segmentation.



Pain Point: Complexity

As applications grow in number, scale, and sophistication, the resulting networking challenges require an ever-expanding range of discrete network appliances:

- Top of Rack (ToR) switches to provide networking, telemetry, and TAP (Traffic Access Point)/mirroring support
- Firewalls (either physical or virtual) to support security between workloads
- Load balancers (either physical or virtual) to support scale between workloads
- Encryption/decryption devices (sometimes physical, often virtual) to protect data in transit
- Storage networks running across Ethernet or on their own dedicated networks

While each of these components makes sense on its own to address a specific problem, the sheer weight of all of them is a networking design/scale headache, with complex dependencies on application workloads: Do I need to add another firewall for every 10 new workloads, or is it 12? Are my load balancers in the right locations to optimize flows between new applications? Can I encrypt my traffic, or will that reduce my application scale on each server—driving up complexity and cost? How do I ensure that the traffic between two workloads can be properly inspected?

Modeling how many of these virtual/physical appliances need to be added to the network becomes yet another element of complexity.

Solution: The hyperscalers took a different approach: by migrating all of these functions to run in software on each server (software-defined networking, or SDN), network services scaled as every new server was added.

Scaling the physical network itself was now simplified to a matter of focusing on connectivity and flow isolation.

Pain Point: Overhead

The next challenge emerged as East/West bandwidth continued to increase. When data center servers were interconnected at 10G, VM/hypervisor deployed Infrastructure services burden on server resources was not significant. But as networks migrate to 25G, this burden has grown to consume 30% of the server core workload, impacting application performance and scale.

Solution: Hyperscalers began looking into how to reduce the impact of SDN impact on the server CPU environment, and explored offloading them to network coprocessors (function accelerators¹/SmartNICs).

As companies like AWS, with their "Nitro" accelerator initiative (based on their 2015 acquisition of hardware company Annapurna Labs) moved virtualized networking, security, telemetry, and storage services from their x86 cores to dedicated hardware accelerators, they restored server scalability and at the same time gained faster, more scalable, and more secure cloud capabilities.

It also allowed them to decouple server workload innovation and network infrastructure innovation cycles, with the added security of isolating network services from x86 software, removing a significant threat vector.

Where to Begin

The path to optimizing at the server edge is already available: Organizations can leverage functions today to build the muscle memory and operational processes for this new world. The earliest services that can be leveraged have been historically running in the Top of Rack (ToR) switches in the data center, and now are available at the server edge.

¹ Function accelerators: Next-generation domain-specific processors with capabilities beyond classic SmartNIC technologies—sometimes also known as data processing units or function offload co-processors

Metrics and Telemetry/IPFIX

The Pensando DSC is in an ideal position to gather analytics on both the PCIe and Ethernet sides of the server, providing a unique level of insight that is not possible in traditional ToR-based models, at the same time addressing complexity and overhead by eliminating monitoring appliances and the need for “tromboning” This new extension of telemetry is key to answering the first and most fundamental question in application troubleshooting: Is it a server/application issue, or a network issue?

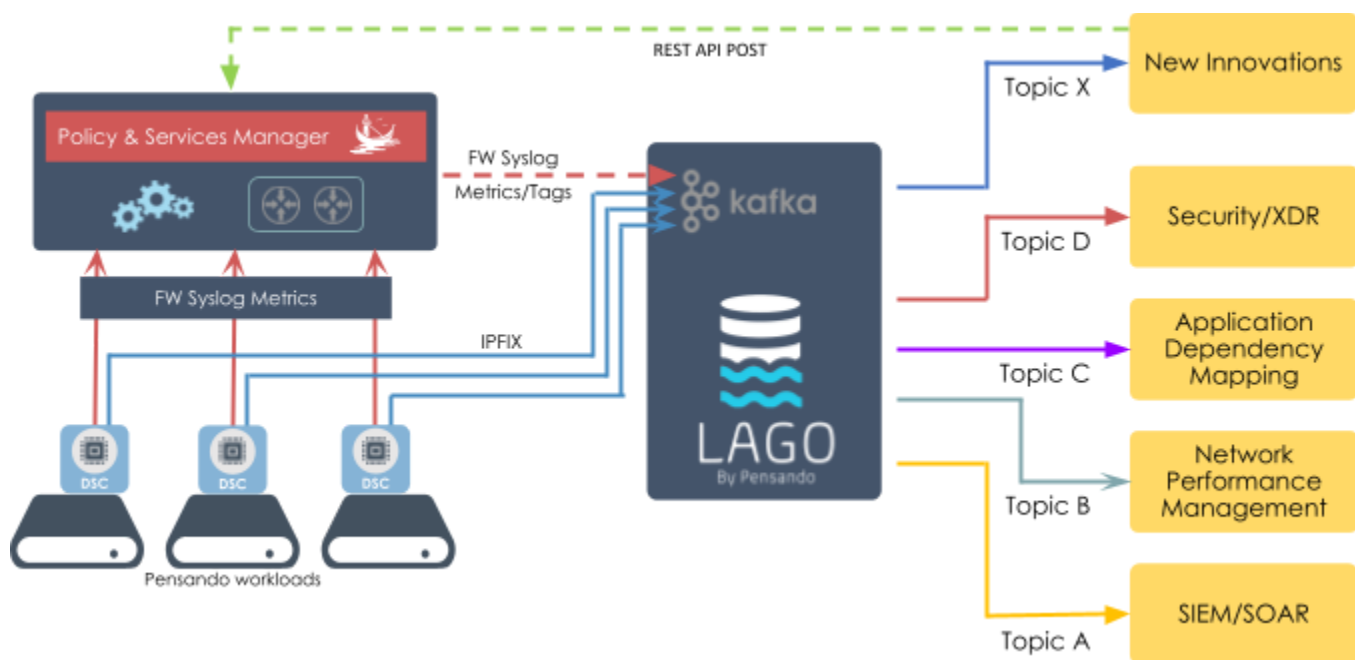
Traditionally, determining the answer would require installing an agent or sensor on a workload that could identify TCP errors such as TCP resets or retransmissions, or analyzing a TCP dump or full packet capture. TCP error detection is a native service and is applied to every flow that passes through the Pensando DSC.

The Pensando platform makes telemetry an effective and pervasive tool: it’s easy to gather telemetry at each server and stream it all in real time—without impacting the server’s performance or degrading network performance. And since a DSC is located in each server, telemetry capability automatically scales as the application footprint expands, no longer limited by appliance oversubscription models.

The DSC uses standard flow-based protocols such as IPFIX to stream all of this data to existing collectors in each data center, adding scale to proven, refined networking capabilities and bringing them to the server edge.

Having scalable access to telemetry makes it possible to feed machine learning engines across security, ADM, networking, and DDOS integrations. The Pensando platform is ideally situated to leverage existing 3rd-party data center vendors, as the telemetry “single source of truth”, and augment with new ones whenever they become available. Pensando has designed a Kafka-based solution, Lago, which aggregates the multiple DSC telemetry flows and implements simple pub/sub integration for 3rd party AI/ML engines, making it easy to provide the right feed to the right partner solution to support their value within the data center. Additionally, any necessary actions determined by ML partners can be immediately pushed back to the PSM via its open REST API to update network policies. This is true data center automation based on ML insights addressing the concerns of networking, security, or even storage insights.

Pensando feeds high-quality network data to existing and future elements of the AI/ML ecosystem.



ERSPAN/Smart-Capture/TAP

Similarly, granular packet capture functions are now instantly available at the server edge. If telemetry is the first line of operational insight, TAPs are the next critical step to determine what is happening between two workloads and what needs to be done to fix the issue. While the ability to TAP or deliver ERSPAN traffic has been available in switched networks for a while now, leveraging it has been challenging.

In the early days of network troubleshooting, there was almost always a one-to-one relationship between a server and workload, so it was trivial to understand what server a workload was running on, what ToR it was attached to. If we ever needed telemetry or needed to TAP the flow between two workloads, it was easy to determine where in the network to request it. That is not the case in today's virtualized world.

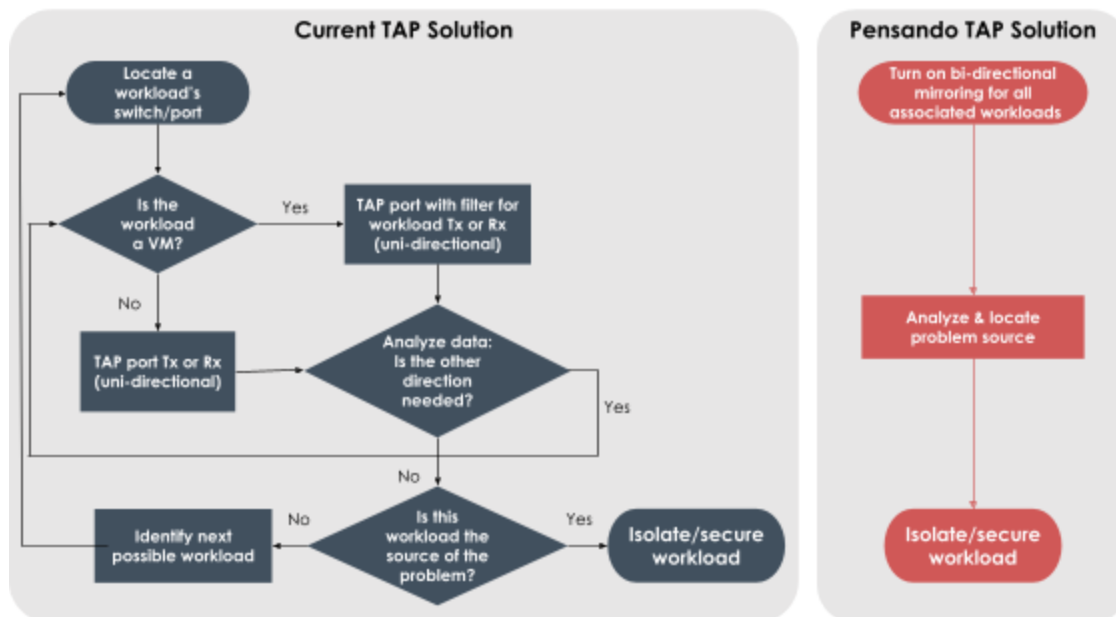
Today, to TAP the traffic between two workloads, we need to:

- Determine which server is currently hosting that workload—VMs and Pods can move from server to server based on design and scale demands
- Determine which ToR that server is attached to and whether that ToR can TAP traffic. If it does, we then TAP the flow (or trunk) unidirectionally
- Perform the same process in reverse to TAP the return traffic

Customers report that prior to implementing the Pensando platform, the time from an application alert to telemetry investigation to implementing a bi-directional TAP was typically 45 minutes to 2 hours.

In a server edge model, the complexity of manually determining a workload's flow from server to ToR goes away. The Pensando Policy and Services Manager (PSM), which manages all deployed DSCs, handles the underlying details: the operator selects the two workloads in question, and a bi-directional TAP/mirror is configured on the appropriate DSC, no matter where it is in the data center. A TAP/mirror can be started manually by an operator in just a few minutes, or even automated via thresholds in seconds.

The Pensando platform transforms TAP from a complex process to an always-available service.



As with telemetry, dynamic TAP/mirror brings immediate value to the data center's productivity, uptime, and operational simplicity. Perhaps even more significantly, the benefits realized by having these services readily available will drive operating model changes over time. Historically a NOC has not had access to data at the server level, being restricted to the ToR--but they do with the Pensando platform and will continue to do so with

Project Monterey. Long before Project Monterey shifts the full set of network infrastructure functions to the Server Edge, operations can immediately begin to increase their visibility into server issues by leveraging edge functionality, which will also simplify the migration from the ToR to the Server Edge.

One example is the ability to implement TAP-as-a-service, where an application owner can create and manage their own TAP sessions without having to make a request to the networking team--which still retains full visibility into the process. Since TAPs are centrally controlled via the PSM, and the DSC has direct visibility to the workload; the Pensando platform simplifies determining where and how to TAP. The application owner controls their troubleshooting in the same way they can in the cloud--improving MTTR and automating what was formerly a series of manual tasks for the NOC.

Moving More Services to the Edge: Encryption of Data in Flight

Other functions outside of the more Network Operations-centric services just discussed are also migrating to the server edge, such as East/West encryption.

Encrypting traffic within the data center has been a challenge for years. While the technology to support it has been available, it's been difficult to deliver it at scale and make it easy to manage. There have been several models brought to market:

- Deploy a device in the network attached to the ToR, similar to a firewall or load balancer, to encrypt server traffic crossing the data center backbone. This appliance-based solution quickly becomes untenable in the face of VM/Pod mobility.
- Deploy IPsec on a Performance NIC between servers. This introduces several challenges:
 - Visibility and control of per-workload traffic;
 - Provisioning on a server-by-server basis, with no centralized controller available;
 - As with appliances, VM/Pod mobility breaks tunnels that are set up statically.
- Deploy IPsec as an agent running in software. This approach demands more CPU and memory resources from each server in order to scale. It also makes it impossible to gather any telemetry in the network.
- Make encryption a requirement of the application. As with the agent-based model, this can slow overall application performance and reduce scale on a per-server basis. It also shifts the burden to each application, increasing deployment complexity.

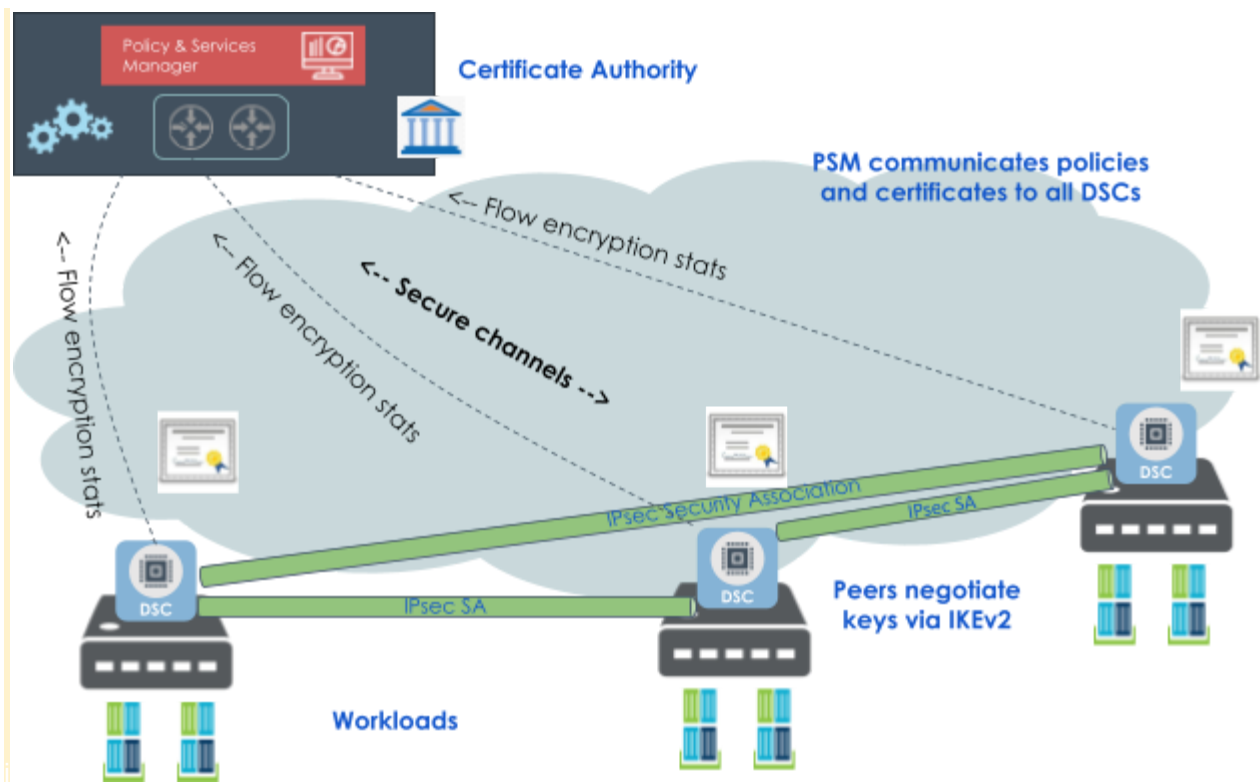
*The DSC's scalable architecture allows IPsec to leverage **the power of AND**: telemetry, TAP, and encryption together, running at line rate at the server edge, providing extreme flexibility, functionality and scale.*

The DSC model of encrypting East/West traffic solves the three most significant issues found with encryption:

1. **Simplicity:** Encryption policy is centrally managed on a workload-by-workload basis: the PSM presents an abstracted network/workload configuration, allowing the operator to easily select which workload-to-workload flows require encryption and which don't. This both simplifies setup and makes configurations far more maintainable. Since the PSM is integrated with services such as vCenter, policies are set to automatically follow any workload migrations.

2. **Performance:** IPsec scales at line rate with no impact to the application and with no increased requirements on the x86 side of the server. Decisions about which workloads require encryption can be made with no concerns about application performance impact.
3. **Observability:** A traditional encryption model intrinsically breaks telemetry services such as IPFIX or ERPSAN, but since those functions are now delivered on the DSC, telemetry can be interposed ahead of encryption.

Pensando platform encryption is centrally managed, including fully automated handshaking and key management for all DSCs, requiring no crypto awareness or overhead on the compute platform. Flow encryption statistics are reported back to the PSM.



There are many reasons to encrypt flows in the data center, including privacy concerns, regulatory needs, and intellectual property protection. Add to those an emerging paradigm with a need for encryption: Software-Defined Storage.

The continuing decomposition and distribution of storage architectures has resulted in models for accessing data across multiple servers over standard Ethernet interfaces. Typically, stored data is encrypted at rest, but traffic between a storage "source" server and "target" server has been in cleartext. The DSC's ability to easily encrypt all of this traffic in transit now allows all data to be protected both in motion and at rest.

With the server edge model, the DSC can provide telemetry for all of this storage traffic, to quickly identify any performance issues and proactively address them before they impact service. This visibility can be integrated with a NOC's SIEM tools, such as Splunk and Elastic, and integrated into SDS vendors' controllers, further simplifying any remediation tasks.

The Operational Roadmap

The above examples are just the start of the network infrastructure services that can and will be delivered at the server edge. Once Monterey is realized, the NSX network, security, and VSAN storage stacks will be

distributed to the Pensando Platform. The convergence of infrastructure services to the Server Edge will further blur the lines between Server Ops, Network Ops, and Security Ops.

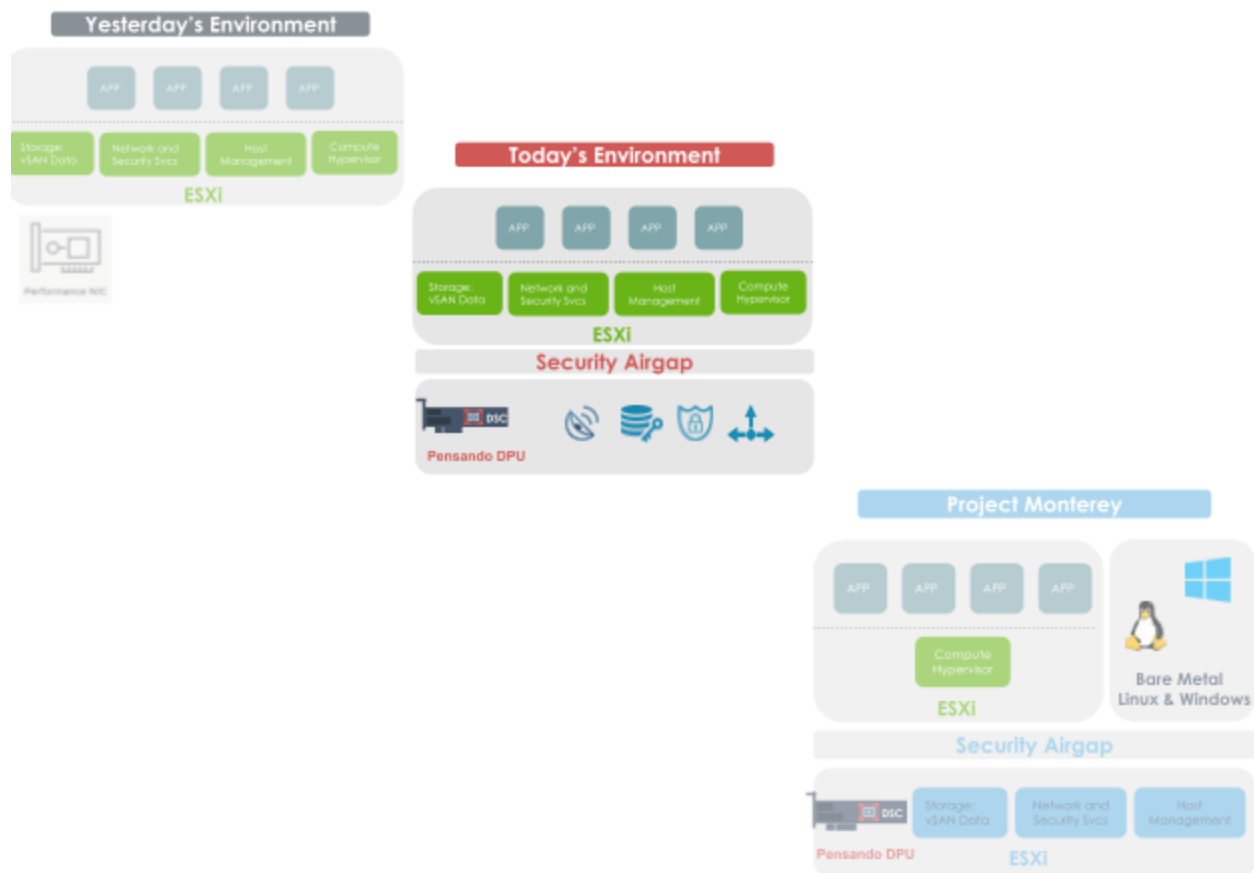
You can already begin building the skills that will be required in this converged world: start distributing services today to virtualized and bare-metal workloads and reap the benefits of a distributed services platform.

By shifting focus to highly-scalable Ethernet fabrics connecting workloads, without designing services based on appliances, you can reduce the complexity of service insertion, lack of East/West visibility at the application layer, and appliance sprawl in the data center. Each server now brings all the necessary network, visibility, and security infrastructure services with it, just as we see with today's hyperscalers.

Changes—even good ones—come with challenges. Refining operational processes, learning new diagnostic capabilities, and refining how workloads are deployed and scaled will be an evolving process. Waiting to make a large shift, such as Project Monterey, puts the customer in the risky position of migrating all of these functions at once, and will extend the amount of time it takes for them to make the final transition.

The Pensando platform gives you the ability to start this transition today, implement the services described above (with more to follow!), and learn how this architectural revolution can further simplify your processes. Start now, learn lessons, refine your models, and then incorporate new services—without disruption. The Road to Monterey will be just as rewarding as the actual arrival of Project Monterey.

The Pensando Distributed Services Platform delivers advanced services today, while paving the way for even further capabilities in the future.



About Pensando

Founded in 2017, Pensando Systems is pioneering distributed computing designed for the New Edge, powering software-defined cloud, compute, networking, storage and security services to transform existing architectures into the secure, ultra-fast environments demanded by next generation applications. For more information, please visit pensando.io.