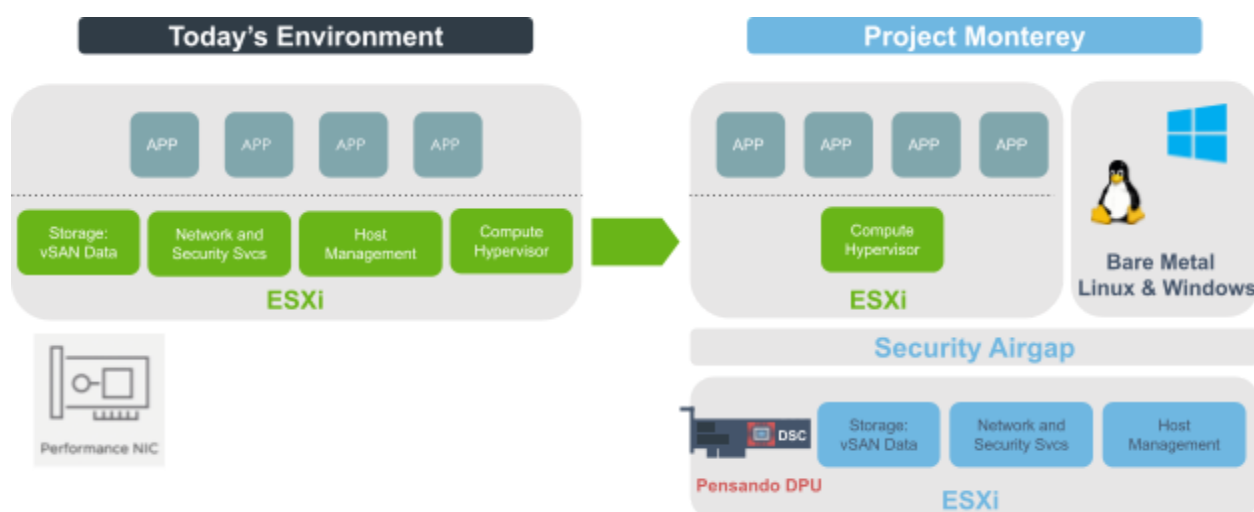


# Pensando and VMware Project Monterey

## Introduction

In September 2020, VMware announced *Project Monterey*, a rethinking of their overall software architecture to take advantage of one of the most significant shifts in how data centers are built: distributed services at the server edge. Taking a page from the hyperscalers' playbook, VMware is moving key infrastructure functions such as networking, security, and storage off of the x86 complex, where they contend with application workloads for critical server resources, and accelerating them in domain-specific hardware.



*Project Monterey: Infrastructure functions that formerly contended for application server resources are instead efficiently and securely hosted on domain-specific hardware.*

As part of Project Monterey, Pensando is collaborating with VMware to integrate their technologies with the **Pensando Distributed Services Card (DSC)**, featuring its groundbreaking fully-programmable data processing unit (DPU). Pensando has a leadership position in providing cloud-like agility, security and operational simplicity with a distributed services platform, resulting in its partnership with VMware in its next phase of redefining hybrid cloud services.

The Pensando DSC is ideally suited for Project Monterey, given its high-performance, fully programmable architecture, market success through supported OEM vendors and established relationships with system providers such as Dell and HPE.

## What is Project Monterey?

Project Monterey is a re-architecture of VMware Cloud Foundation to support the rapidly evolving requirements of next-gen applications. As application resource demands continue to grow, Project Monterey takes advantage of accelerators like the Pensando DSC to offload the main server, freeing up precious CPU cycles, and at the same time bringing key services such as networking and security closer to where applications

actually run and operate on data. Project Monterey is looking at the infrastructure design with a fresh perspective and builds the foundation for next-generation infrastructure services and management with improved efficiency, performance, security, and TCO.

## Bringing Hyperscale to Enterprise Data Centers

The **Pensando Distributed Services Platform** has democratized **cloud-level scale, security, and observability**, putting it within the reach of any organization. Since DPUs like the Pensando 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 the Pensando product family and the solutions we are bringing to our customers to scale their data centers, regardless of how they host their workloads.

## Project Monterey Goals

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 on DPUs, they no longer depend 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 approach allows support for not only traditional virtualized or container based applications, but for the first time, those same network, security, storage, etc. services can be extended to bare-metal instances too.

## Why Pensando?

The Pensando Distributed Services Platform is a natural choice for Project Monterey, offering unmatched capabilities and programmability.

**Not all DPUs are created equal:** The Pensando DSC is a generation ahead of other DPUs, offering leading-edge capabilities with orders of magnitude better performance at scale, latency and jitter, and a low power profile (20W @ 25G, ~30W at 100G). It is capable of processing streams at rates of up to tens of millions packets per second, with a scale of millions of routes and ACLs, while maintaining latency and jitter in the range of microseconds.

**Fully Programmable:** The Pensando DSC is a custom ASIC optimized for the P4 programming language, providing both high performance and extreme flexibility. Unlike competing architectures, this allows partners such as VMware to accelerate their iterative development process, leveraging the Pensando P4 match processing units (MPUs) to deliver new functionality with faster time to market, without the performance overhead of a sea of Arm cores or power-hungry FPGAs.

"Customers around the globe rely on VMware Cloud Foundation to deploy and manage modern applications across multiple environments. Together with Pensando, we are building the next generation of more secure and agile infrastructure, leveraging the new breed of accelerators exemplified by the Pensando Distributed Services Platform, to support the evolving requirements of these applications. The work between VMware and Pensando on Project Monterey will help mutual customers benefit substantially from increased performance, enhanced security and a consistent operating model."

—Krish Prasad, senior vice president and general manager,  
Cloud Platform Business Unit, VMware

As one example, changes to VMware protocols can be rapidly designed and delivered via software updates that run at ASIC performance levels, without costly investments in development and delivery of new chip ASIC implementations.

**Full Programmability = Performance AND Scale:** The Pensando design team's experience with large cloud providers has allowed us to develop and validate the performance benefits of the P4 architecture: we are seeing **4X to 5X** performance improvements when compared to Arm-based platforms. (Although the Pensando DPU does include Arm cores, they are only used to run management plane functions; data streams are handled exclusively within the P4 MPUs, yielding optimized performance without sacrificing flexibility.)

**Full Programmability = Pensando-Unique Capabilities:** Pensando is working very closely with VMware to leverage the full capabilities of the hardware and software stack. P4 programmability makes it possible to deliver significant value-added functionality; for example, the ability to use vMotion with an accelerated datapath. In legacy environments, this is an "OR" proposition: vMotion or VM acceleration. The combination of the Pensando DPU and VMware brings the *Power of AND*: delivering both at the same time.

A fully programmable environment ensures that this is just the start: additional unique capabilities can be added without compromising existing hardware and deployment investments. The Pensando platform is designed to support a broad range of differentiated infrastructure services, such as **security, storage, network, and real time telemetry**—and **enable them simultaneously** (e.g. in-flight and at-rest encryption with visibility at line rate, telemetry, TAP network, and NVMe over fabrics).