Solution Brief
Data Center | Hybrid Cloud

Get Fast Insights with a VMware Hybrid Cloud and Intel® Optane™ Technology

VMware Cloud Foundation, powered by Intel® technology, delivers high performance as well as simplified, user-friendly deployment and management of critical business workloads

Solution Benefits

- **Fast deployment.** Quickly launch database processing, analytics, AI, and machine-learning workloads on a platform that scales to future needs.
- **Easy management.** Migrate workloads between public clouds and private data centers without modifying applications or node configurations.
- **Excellent performance.** Whether running VMs or orchestrating containers, the latest Intel® technology provides the throughput and responsiveness critical for today's data-intensive workloads.

Executive Summary

The speed and ease with which companies are able to gain business insights can make the difference between having the competitive edge and simply trying to keep up. The digital landscape is changing rapidly, and traditional hosting applications and services can no longer deliver the necessary innovations at the speed businesses require. And with the cost of many new technologies, businesses are faced with difficult challenges.

Rapid software and hardware provisioning, deployment, and maintenance, as well as application development, testing, and delivery can help companies get ahead and stay ahead. VMware Cloud Foundation (VCF), powered by Intel® technology, offers a solution with a simplified hybrid cloud platform for easier management of virtual machines (VMs) and containers (see Figure 1). Now businesses can operate workloads, such as data analytics, artificial intelligence (AI), and machine learning (ML)—foundational applications for predicting trends—across private and public clouds with excellent performance and reliability without modifying node configuration or application refactoring. The platform is designed to scale as businesses grow and workload demands increase. Integrating Intel® Optane™ persistent memory (PMem) into VCF clusters improves transactions per minute (TPM) application performance, including data warehousing performance, by an average of 3.34x.[1]

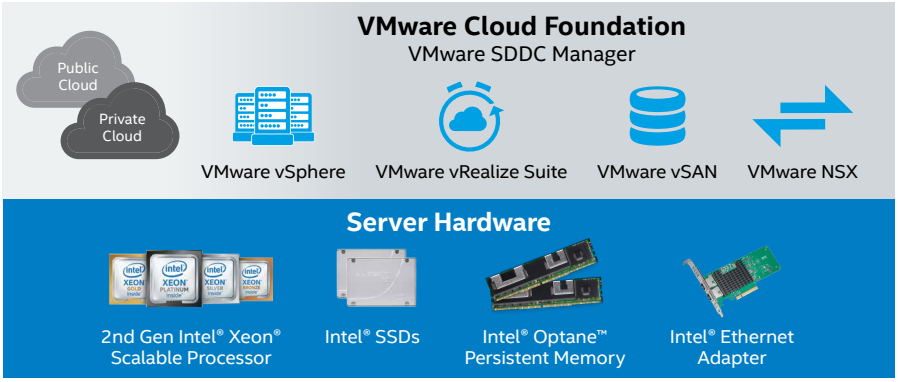

Figure 1. VMware Cloud Foundation, powered by Intel® technology, is a hybrid cloud solution that retains the benefits of on-premises infrastructure while providing the agility of a public cloud.

Author

Patryk Wolsza
vExpert, Cloud Solutions Architect
Intel Cloud & Enterprise Solutions

Business Challenge: Traditional Workload Management is Cumbersome

The digital landscape is constantly changing, and competing for market share requires rapidly identifying trends in enterprise data. But to do so, companies must accelerate workload deployment and simplify the overall management. The addition of end-user self-service solutions has further reduced expected time-to-market. Modern technology solutions must transcend the divide between hardware and software, delivering an end-to-end solution that encompasses hardware provisioning, deployment, and maintenance, as well as application development, testing, and delivery.

With hybrid cloud infrastructure, companies can retain the benefits, security, and control of a familiar on-premises infrastructure while gaining the flexibility and instant availability of the public cloud. VMware Cloud Foundation (VCF) with the latest Intel® technology provides the flexibility and agility companies need for everything from database processing and analytics to artificial intelligence (AI) and machine learning (ML), with scalability for future growth.

Flexible Hybrid Cloud Platform for a Variety of Modern Workloads

Many important enterprise workloads fall into two broad areas: data warehousing and AI/ML. Most of these applications, like Microsoft SQL Server, Oracle Database, and DataRobot, can run as VMs next to other workloads in Kubernetes-based containers on the same platform. Containers are treated as first-level citizens. VCF supports both approaches (see Figure 2), so that enterprises can deploy a workload in the way that makes most business sense.

VCF enables businesses to provision infrastructure from a pool of virtualized resources on premises, at the edge, or in the cloud. Modern apps can take advantage of built-in automation and orchestration to simplify management of VM- or container-based workloads. And when the business grows, VCF deployments can scale as necessary.

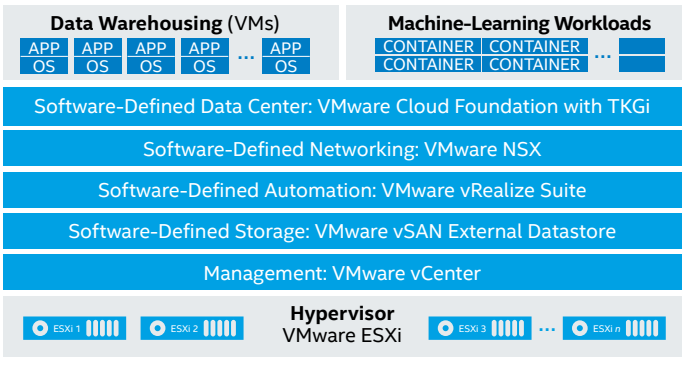

Figure 2. VMware Cloud Foundation can run a wide variety of data analytics, AI, and machine-learning workloads.

Solution Value: Scalable Processing Now and in the Future

To demonstrate the value of deploying modern workloads on VCF using the latest Intel technologies, Intel recently conducted performance tests for both data warehousing and machine-learning workloads. The tests results showed that software optimized for Intel® hardware, plus innovations such as Intel® Optane™ persistent memory (PMem) can boost application performance. For detailed information about testing and results, as well as a bill of materials, see the full Reference Architecture.

Data Warehousing Benefits from Intel® Optane™ PMem

Testing showed that Microsoft SQL Server performance improves when Intel Optane PMem in App Direct Mode is added. The test compared the number of transactions per minute (TPM) achieved by a "Base cluster" without Intel Optane PMem to the TPM achieved by a "Plus cluster" with Intel Optane PMem and a high-performance Intel® Xeon® Platinum processor. The Plus cluster provided an average of 3.34x improvement in TPM and scaled linearly with the addition of VMs (see Figure 3).[2]

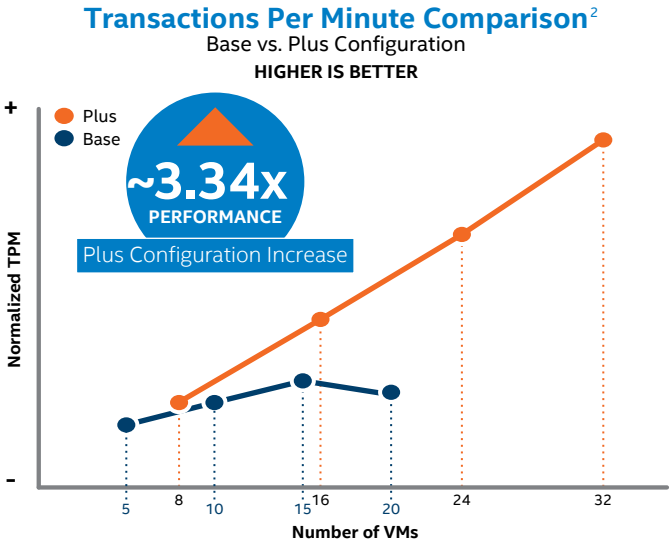

Figure 3. The Plus cluster achieved an average of 3.34x more TPM compared to the Base cluster and is highly scalable.

AI/Analytics: Deep-Learning Inference

In recent tests, Intel benchmarked two image classification topologies, comparing the standard version of TensorFlow versus the TensorFlow distribution from the Deep Learning Reference Stack with pretrained models from Model Zoo for Intel® Architecture. The test consisted of two experiments, one with a fat VM and another with a Tanzu Kubernetes Grid Integrated (TKGi) instance. The results demonstrate the benefits of software optimized for Intel architecture for both topologies we tested.

For the fat VM test, using the Deep Learning Reference Stack version of TensorFlow improved the performance of the Inception v3 topology for the Base configuration by 2.4x, and improved the throughput of the Plus configuration by 3x (see Figure 4).[3]

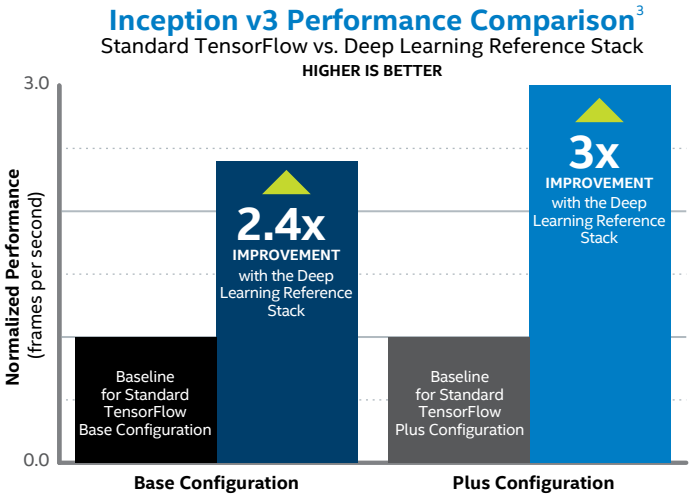

Figure 4. Software optimized for Intel® architecture accrues significant performance benefits.

For the Kubernetes cluster provisioned by TKGi, we measured throughput scaling for the Deep Learning Reference Stack container by running the benchmark on one to six workers in parallel for the deep-learning workloads. We observed improvement in throughput with additional jobs running in parallel, demonstrating the effectiveness of ESXi scheduler (see Figure 5).[4]

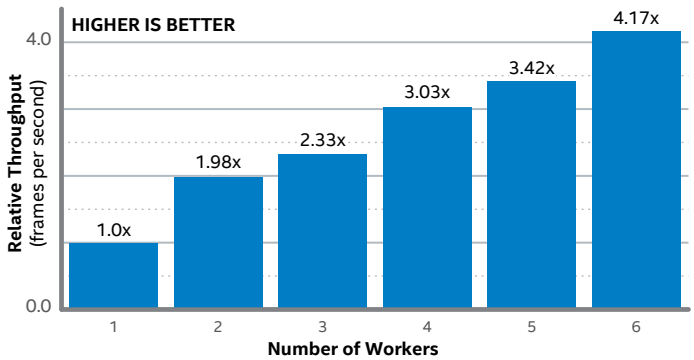

Figure 5. Optimized software, combined with high-performance hardware, results in excellent scalability.

Solution Architecture: High Performance, User-Friendly, Scalable Processing

User-friendly VCF for hybrid cloud computing uses Intel technology to help achieve high performance and scalability (see Figure 6 for the solution architecture). The solution includes the following technologies:

Intel® Technology

- **Intel Optane PMem** helps improve server performance by supplying large amounts of persistent storage with low-latency access. It provides capacity that is unavailable with traditional DRAM modules, and operating modes can be configured using the platform BIOS or memory management tool.
- **2nd Generation Intel Xeon Scalable processors** are workload optimized for analytics, AI, and in-memory database workloads to process ever-increasing amounts of data.
- **Intel® Optane™ SSDs** in the cache layer and large-capacity NVMe-based 3D NAND SSDs in the capacity layer help VMware vSAN optimize performance.
- **Intel® Volume Management Device (Intel® VMD) technology for NVMe drivers** enables serviceability of NVMe-based SSDs by supporting hot swap replacement from the PCIe bus without shutting down.
- **Intel® Ethernet 700 Series** accelerates performance for VMware hybrid cloud platform to help meet enterprise-level requirements for data resiliency, service reliability, and ease-of-provisioning.

Cloud Infrastructure

- **VMware Cloud Foundation** offers a simplified path to hybrid cloud computing through an integrated software platform for both private and public cloud environments. It includes a complete set of software-defined services for compute, storage, network, and security, as well as application-focused cloud management capabilities.
- **VMware Software-Defined Data Center (SDDC) Manager** manages VCF system bring-up, as well as creates and manages workload domains and performance lifecycles. It also manages logical and physical resources.
- **VMware vSphere** extends virtualization to storage and network services, and automated, policy-based provisioning and management.
- **VMware NSX-T Data Center** (formerly NSX-T) is a network virtualization platform that enables software-defined virtualization by working like a network hypervisor. It reproduces Layer 2-7 network services: routing, switching, access control, firewalls, quality of service (QoS), and Dynamic Host Configuration Protocol (DHCP) in software.
- **VMware Cloud** (on AWS) is a hybrid cloud solution that allows easy extension, migration, and modernization of applications. It also protects applications in the public cloud.
- **TKGi** is a container solution with advanced networking, a private container registry, and life cycle management. It simplifies deployment and operation of Kubernetes clusters so you can run and manage containers at scale on private and public clouds.

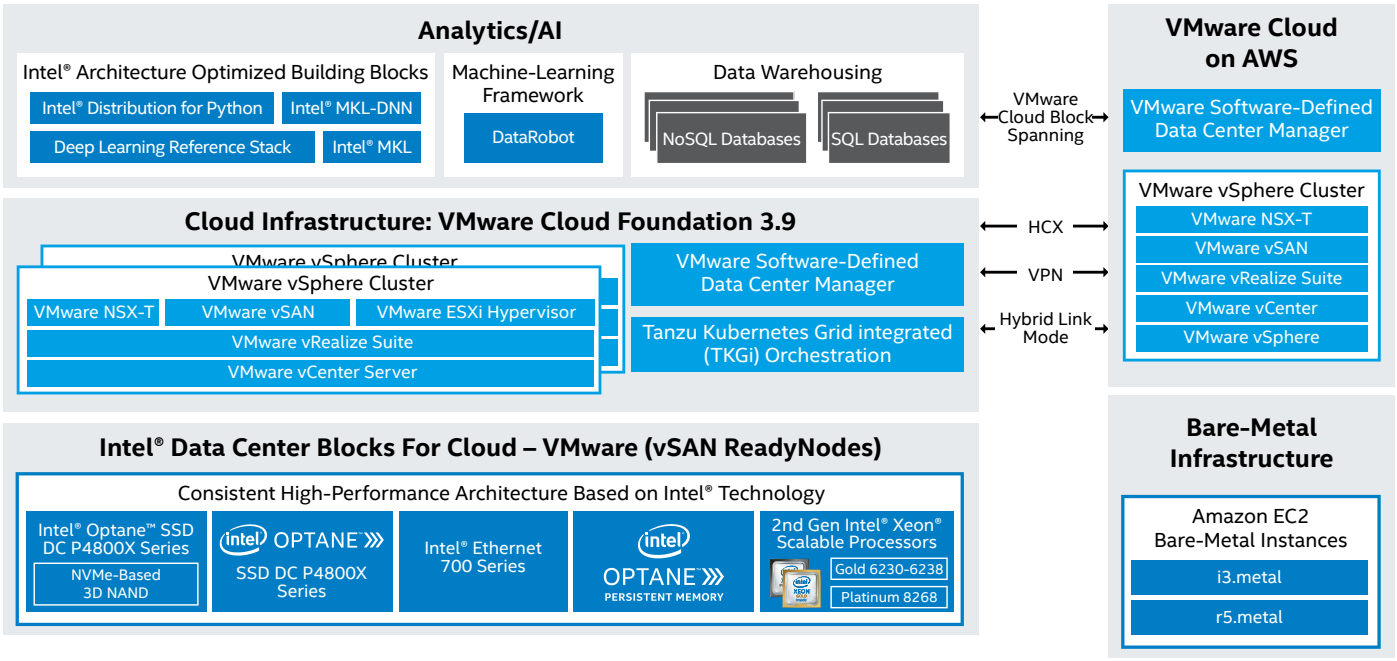

Figure 6. VMware Cloud Foundation with Intel® technology offers an end-to-end solution to meet today's needs, as well as future growth.

Conclusion

Companies are increasingly pressed to gain faster insights from their data to remain competitive. Traditional hosting applications and services cannot deliver the necessary innovations at the speed businesses require today and as they grow in the future. More than ever, businesses are turning to modern applications, such as analytics, AI, and ML to overcome these difficult challenges. And modern applications require dynamic cloud-native environments that offer ease of use, scalability, rapid hardware provisioning, deployment, and maintenance, as well as application development, testing, and delivery.

VCF powered by Intel architecture provides a simplified hybrid cloud platform for easier management of VMs and containers that also allows AI/ML workloads to coexist with legacy workloads on the same platform. Businesses can now operate workloads across private and public clouds with excellent performance and reliability (without modifying node configurations) in one scalable solution.

Find the solution that is right for your organization. Visit intel.com/optane or contact your Intel representative.

Learn More

You may also find the following resources useful:

- Building a VMware Hybrid Cloud Platform Reference Architecture
- Intel® Data Center Blocks for Cloud – vSAN ReadyNodes
- Intel® Optane™ persistent memory home page
- Intel Optane Persistent Memory – Virtualized Performance
- Deep Learning Reference Stack
- 2nd Generation Intel® Xeon® Scalable processors
- Intel® Select Solutions for VMware vSAN ReadyNodes
- Intel® Optane™ Solid State Drives
- Intel® Solid State Drives Data Center Family
- Intel® Framework Optimizations
- VMware vSAN
- VMware Cloud Foundation
- VMware Cloud on AWS

Solution Provided By:

[1] Testing by Intel, June 15-23, 2020. Each cluster (Base, Plus, and the Management cluster that were used for workload generation) consists of four machines. All machines are identical within the given cluster. On top of those clusters the VMs are placed: five VMs per node for Base and eight VMs per node for Plus. We had four scenarios for each environment (Base and Plus). The first scenario used only one node in the cluster, the second scenario used two, and so on, up until all four nodes were used. Each scenario was repeated three times; the results given here are the average values for these three runs for each scenario.

Base 4-node cluster configuration: Intel® Xeon® 6248 processor (2.5 GHz, 20 cores); 384 GB DRAM (12x 32 GB 2933 MHz); Intel® Hyper-Threading Technology ON; Intel® Turbo Boost Technology ON; Storage: VMware vSAN, disk group: 1x Intel® Optane™ SSD DC P4800X 375 GB and 2x Intel® SSD DC P4510 2 TB, two disk groups per node; BIOS = 02.01.0010; microcode = 0x0500002c; OS = VMware ESXi 6.7.015160138; 1x Intel® Ethernet Adapter XXV710-DA2.

Base VM configuration: 8 vCPUs; 44 GB vMemory; vDisks: 100 GB – OS drive (LSI controller), 4x 18 GB – Data DB (1 SCSI controller), 4 files per disk, 2x 200 MB – Temp DB (1 SCSI controller), 200 MB – Temp log (same SCSI controller as User transaction log), 10 GB – User transaction log (1 SCSI controller); vNetwork = VMXNET3; OS = Windows Server 2019 Datacenter 17763.rs5.180914-1434; Database = MSSQL Server 2019 15.0.2070.41.

Plus 4-node cluster configuration: Intel® Xeon® 8268 processor (2.9 GHz, 24 cores); 384 GB DRAM (12x 32 GB 2933 MHz) plus 1,536 GB Intel® Optane™ persistent memory (12x 128 GB); Intel Hyper-Threading Technology ON; Intel Turbo Boost Technology ON; Storage: VMware vSAN, disk group: 1x Intel Optane SSD DC P4800X 375 GB and 2x Intel SSD DC P4510 2 TB, two disk groups per node; BIOS = 02.01.0010; microcode = 0x0500002c; OS = VMware ESXi 6.7.0 15160138; 1x Intel Ethernet Adapter XXV710-DA2; PMem used a 2-2-2 configuration in App Direct Mode; PMem firmware version = 01.02.00.5417.

Plus VM configuration: 8 vCPUs; 44 GB vMemory; vDisks: 100 GB – OS drive (LSI controller), 4x 18 GB Data DB (NVDIMM), 4 files per disk, 2x 200 MB Temp DB (NVDIMM non-DAX volume), 200 MB Temp log (NVDIMM non-DAX volume), 10 GB User Transaction log (NVDIMM non-DAX volume); vNetwork = VMXNET3; OS = Windows Server 2019 Datacenter 17763.rs5.180914-1434; Database = MSSQL Server 2019 15.0.2070.41.

[2] See endnote 1.

[3] Testing by Intel, June 15-23, 2020. Workload: Inception V3 topology at fp32 precision. Default TensorFlow tensorflow/tensorflow:1.15.0-py3; optimized TensorFlow: clearlinux/stacks-dlrs-mkl:v0.5.0.

Base 4-node cluster configuration: Intel® Xeon® 6248 processor (2.5 GHz, 20 cores); 384 GB DRAM (12x 32 GB 2933 MHz); Intel® Hyper-Threading Technology ON; Intel® Turbo Boost Technology ON; Storage: VMware vSAN, disk group: 1x Intel® Optane™ SSD DC P4800X 375 GB and 2x Intel® SSD DC P4510 2 TB, two disk groups per node; BIOS = 02.01.0010; microcode = 0x0500002c; OS = VMware ESXi 6.7.015160138; 1x Intel® Ethernet Adapter XXV710-DA2.

Base Fat VM configuration: 80 vCPUs; OS = CentOS Linux release 8.1.1911; kernel = 4.18.0-147.5.1.el8_1.x86_64; other SW: VMware vSphere 6.7, VMware Cloud Foundation 3.9, VMware ESXi hypervisor 6.7 update 3.

Normalized Base Results: 2.4X more throughput using the Deep Learning Reference Stack version of TensorFlow, compared to the default TensorFlow performance.

Plus 4-node cluster configuration: Intel Xeon 8268 processor (2.9 GHz, 24 cores); 384 GB DRAM (12x 32 GB 2933 MHz) plus 1,536 GB Intel Optane persistent memory (12x 128 GB); Intel Hyper-Threading Technology ON; Intel Turbo Boost Technology ON; Storage: VMware vSAN, disk group: 1x Intel Optane SSD DC P4800X 375 GB and 2x Intel SSD DC P4510 2 TB, two disk groups per node; BIOS = 02.01.0010; microcode = 0x0500002c; OS = VMware ESXi 6.7.0 15160138; 1x Intel Ethernet Adapter XXV710-DA2; PMem used a 2-2-2 configuration in App Direct Mode; PMem firmware version = 01.02.00.5417.

Plus Fat VM configuration: 96 vCPUs; otherwise, identical to the Base Fat VM configuration.

Normalized Plus Results: 3X more throughput using the Deep Learning Reference Stack version of TensorFlow, compared to the default TensorFlow performance.

[4] Testing by Intel, June 15-23, 2020. Workload: Inception V3 topology at fp32 precision. 6 worker VMs with 16 vCPU each. See endnote 3 for physical infrastructure configuration. Default TensorFlow: tensorflow/tensorflow:1.15.0-py3; optimized TensorFlow: clearlinux/stacks-dlrs-mkl:v0.5.0. **Worker Node configuration:** OS = Ubuntu 16.04.6 LTS; kernel = 4.15.0-55-generic; other SW: VMware vSphere 6.7, VMware Cloud Foundation 3.9, VMware ESXi hypervisor 6.7 update 3.