

VIRTUALIZING HIGH-PERFORMANCE COMPUTING (HPC) ENVIRONMENTS

Reference Architecture

September 2018

Index

1. AUDIENCE	3
2. INTRODUCTION	3
3. WHAT IS HPC?	3
4. MAJOR TYPES OF HPC WORKLOADS	4
4.1 Message Passing Interface (MPI) Workloads.....	4
4.2 Throughput Workloads.....	4
5. BENEFITS OF VIRTUALIZING HPC	5
5.1 Benefits of Virtual Machines.....	5
5.2 Security.....	6
5.3 Resilience and Redundancy.....	6
5.4 Performance.....	6
6. DESIGN	7
6.1 Traditional HPC.....	7
6.2 Basic Virtualized HPC (vHPC) Architecture.....	8
6.3 Secure HPC Cloud Architecture.....	9
6.4 Management Cluster.....	11
6.5 Compute Clusters.....	12
6.6 Hardware Accelerators.....	13
7. SAMPLE ARCHITECTURE	16
7.1 Scenario A: MPI Workloads.....	16
7.2 Scenario B: Throughput Workloads.....	17
8. CONCLUSIONS	19
9. AUTHORS	19

This reference architecture for virtualizing high-performance computing (HPC) environments describes the infrastructure and configuration of an HPC deployment based on VMware technologies. In addition, it provides information about the components of traditional and virtualized HPC (vHPC) environments.

1. AUDIENCE

This document is intended for virtualization architects, IT infrastructure administrators, and HPC systems administrators who intend to design, deploy, and maintain vHPC workloads. These IT professionals are generally new to either HPC or virtualization environments and typically collaborate to deploy vHPC environments.

2. INTRODUCTION

While virtualization has proven to provide enterprises with cost-effective, scalable, and reliable IT computing, the approach to modern HPC has not evolved. Rather, it remains bound by the use of physical resources with its inherent lack of flexibility.

This white paper identifies the ways in which virtualization accelerates the delivery of HPC environments and provides a system design that demonstrates how virtualization and HPC technologies work together to deliver a secure, elastic, fully managed, self-service, virtual HPC environment.

Before diving into the reference architecture for HPC, it is worth reviewing what HPC is, the primary workload types, and the specific ways in which virtualization significantly improves HPC operational efficiency.

3. WHAT IS HPC?

HPC is the use of multiple computers and parallel-processing techniques to solve complex computational problems. HPC systems have the ability to deliver sustained performance through the concurrent use of computing resources, and they are typically used for solving advanced scientific and engineering problems and performing research activities through computer modeling, simulation, and analysis.

HPC is continuously evolving to meet increasing demands for processing capabilities. It brings together several technologies, including computer architecture, algorithms, programs and electronics, and system software under a single canopy to solve advanced problems effectively and quickly. Multidisciplinary areas that can benefit from HPC include:

- Aerospace
- Biosciences
- Energy
- Electronic Design
- Environment and Weather
- Finance
- Geographic Information
- Media and Film

4. MAJOR TYPES OF HPC WORKLOADS

This reference architecture targets two major types of HPC workloads:

4.1 Message Passing Interface (MPI) Workloads

MPI workloads consist of multiple processes running simultaneously across a group of compute nodes that need to communicate with each other, often with extremely high frequency, making their performance sensitive to interconnect latency and bandwidth.

Figure 1 illustrates the coupling characteristic of MPI workloads: An MPI job is decomposed into a number of small tasks, and these tasks communicate via MPI and are mapped to available processor cores. For example, weather forecast modeling is computationally intensive with the demand for computing power increasing at higher resolutions. In order to run the models in a feasible amount of time, they are decomposed into multiple pieces, with the calculation results from each piece affecting all adjacent pieces. This requires continual message-passing between nodes at extremely low latency. Typically, these models are run on a cluster of machines using the MPI standard to communicate during job execution. Other examples of MPI workloads include molecular dynamics, computational fluid dynamics, oil and gas reservoir simulation, jet engine design, and emerging distributed machine learning workloads.

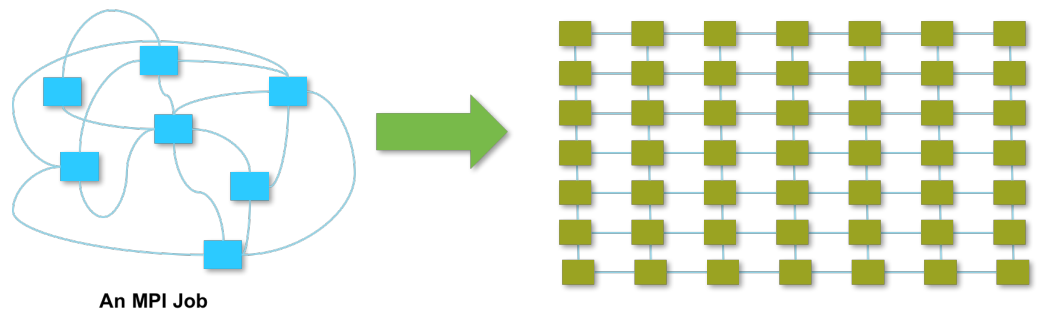
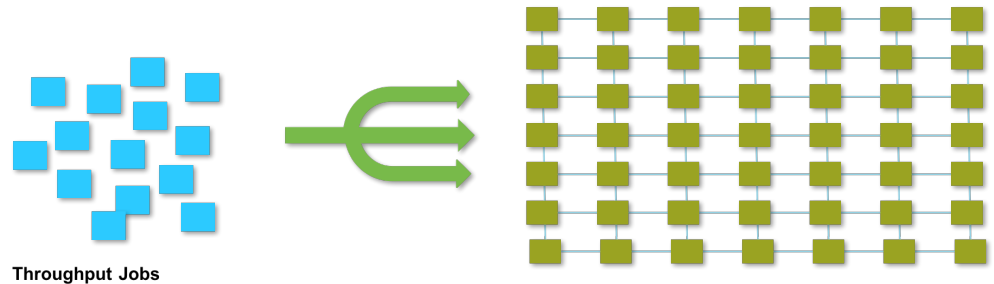


Figure 1. MPI Workloads Run on Clusters and Are Characterized by Their Intensive Communication Between Processes.

4.2 Throughput Workloads

With throughput workloads, a task is divided into many small jobs that run simultaneously but, unlike MPI workloads, there is no communication between jobs, as illustrated in Figure 2. Digital movie-rendering is a typical example of a throughput workload. A movie will be divided into individual frames and distributed across the cluster so that each frame can be rendered simultaneously on a different processor core and then assembled at the end to produce a complete movie. Other examples of throughput workloads include Monte Carlo simulations in financial risk-analysis, electronic design-automation, and genomics analysis, where each program runs in a long-time scale or features hundreds, thousands, or even millions of executions with varying inputs.



Throughput Jobs

Figure 2. Throughput Workloads Run Independently in Parallel on Clusters with No Communication Between Processes.

5. BENEFITS OF VIRTUALIZING HPC

HPC workloads have traditionally been run only on bare-metal hardware due to the performance penalty incurred by virtualization having added a layer of software between the operating system and hardware. However, advances in both hypervisor and hardware virtualization support in x86 microprocessors have yielded dramatic increases in performance for these highly computation-intensive workloads. Combined with the numerous operational benefits that virtualization offers, virtualization of HPC environments is becoming more common. In this section, we discuss the particular benefits of virtualization that enhance HPC environments and productivity.

5.1 Benefits of Virtual Machines

The fundamental element of virtualization is the virtual machine (VM)—a software abstraction that supports running an operating system and its applications in an environment whose resource configuration may be different from that of the underlying hardware. The benefits of VMs in an HPC environment include:

- **Heterogeneity:** By using VMs, different resource configurations, operating systems, and HPC software can be flexibly mixed on the same physical hardware. In addition, with a self-provisioning model, IT departments can deliver various environments with decreased time-to-solution for researchers, scientists, and engineers per each user's requirements.
- **Increased control and research reproducibility:** Infrastructure and HPC administrators can dynamically resize, pause, take snapshots, back up, replicate to other virtual environments, or simply wipe and redeploy VMs based on their role-based permissions. Since configurations and files are encapsulated within each VM, the VMs can be archived and rerun for research purposes, such as compliance.
- **Improved resource-prioritization and balancing:** Compute resources for VMs can be prioritized individually or in a pool. It's also possible to migrate running VMs and their encapsulated workloads across the cluster for load-balancing. This migration increases overall cluster efficiency compared with a bare-metal approach.
- **Fault-isolation:** By running jobs in an isolated VM environment, each job is protected from potential faults caused by jobs running in different VMs.

5.2 Security

Security rules and policies can be defined and applied based on environment, workflow, VM, physical server, and operator, including:

- Actions controlled via user permissions and logged for audit reporting. For example, root access privileges are only granted as needed and based on the specified VMs, preventing compromise of other HPC workflows.
- Isolated workflows where sensitive data cannot be shared with other HPC environments, workflows, or users running on the same underlying hardware.

5.3 Resilience and Redundancy

HPC VMs provide fault-resilience, dynamic recovery, and other capabilities not available in traditional HPC environments. Specifically, HPC VMs enable:

- Hardware maintenance without impacting operational HPC workflows or serviceability.
- Automatic restart on another physical servers within the cluster following a server failure.
- Live migration to another physical host when resources of a given host are at capacity.

5.4 Performance

Performance is paramount in HPC and technical computing environments. Throughput workloads generally run at close to full speed in a virtualized environment—with less than 5% performance degradation compared to native and just 1-2% in many cases. This has been verified through tests of various applications across multiple disciplines, including life sciences, electronic design automation, and financial risk analysis.

Furthermore, the study [Virtualizing HPC Throughput Computing Environments](#) demonstrates that with CPU overcommitment and the concept of creating multitenant virtual clusters on VMware vSphere®, performance of high-throughput workloads in a virtual environment can sometimes exceed the performance of bare-metal environments.

The performance degradations of MPI workloads are often higher than throughput workloads due to latency requirements combined with intensive communication among processes. The performance study [Virtualized HPC Performance with VMware vSphere 6.5 on a Dell PowerEdge C6320 Cluster](#) has shown that performance degradations for a range of common MPI applications can be kept under 10%, with the highest scale testing (using 32 nodes with 640 cores) showing larger slowdowns in some cases.

For either throughput or MPI applications that leverage compute accelerators, such as HPC workloads using General-Purpose Graphics Processing Units (GPGPU), the ability to map the physical PCIe resources directly to VMs (VMware Direct Path I/O) delivers virtualized performance that is near that of bare metal.

6. DESIGN

Designing vHPC environments requires understanding how the architecture will compare to that of traditional bare-metal HPC environments. In this section, we first introduce the architecture of a traditional HPC environment and then present how it can be converted into a virtual environment through the use of VMware software.

6.1 Traditional HPC

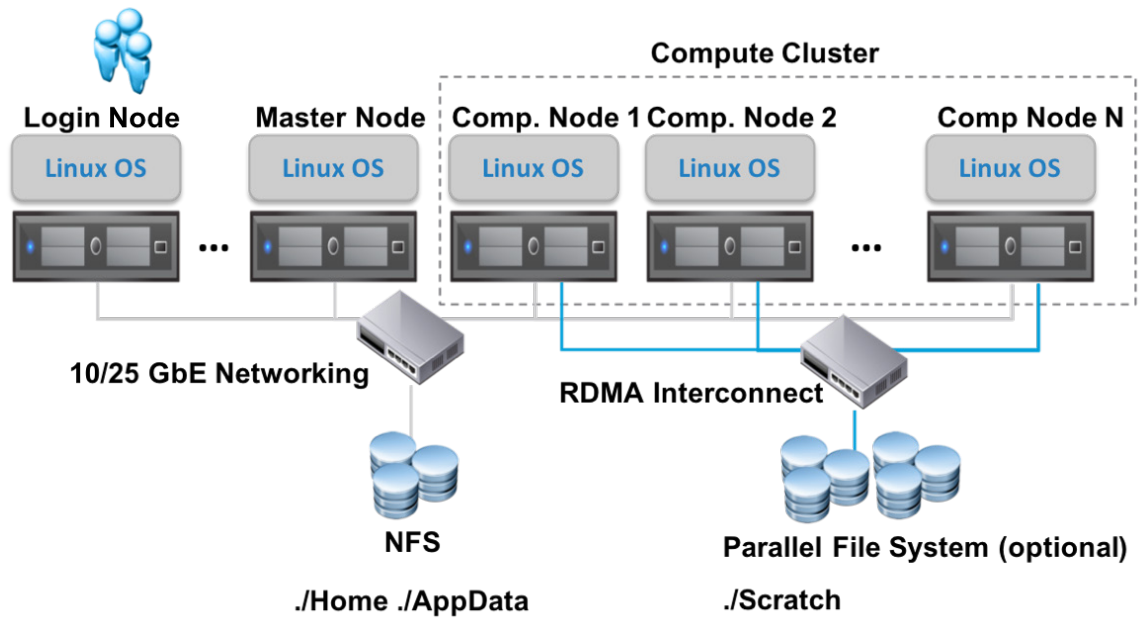


Figure 3. Traditional Bare-metal HPC Cluster

As illustrated in Figure 3, the typical configuration of a traditional bare-metal HPC cluster includes the following components:

Management Node(s):

- Login node manages user logins
- Master node for scheduling jobs
- Other services, such as DNS and gateway
- For small clusters, login, job scheduler, and other services can run on same node, however, it is necessary to separate them when the number of nodes grows larger.
- All management nodes typically run a Linux OS.

Compute Nodes:

- Compute nodes are dedicated for computational tasks.
- All compute nodes run the same Linux OS as the management nodes and are managed by a job scheduler running on the master node.

Networking and Interconnect:

- Management traffic typically communicates with 10/25 Gb/s Ethernet among management nodes and compute nodes.
- If there are needs for fast inter-node communication, this is achieved through remote direct memory access (RDMA) capable interconnects, such as InfiniBand or RoCE (RDMA over Converged Ethernet) host-channel adapters (HCA) and switches.

Storage:

- Network File System (NFS) is often used for home directories and project space mounted across all nodes. It can also often provide a space for research groups sharing data.
- Parallel file systems such as Lustre or IBM Spectrum Scale can be used for HPC workloads that require access to large files, simultaneous access from multiple compute nodes, and massive amounts of data. The implementation of parallel file systems makes it easy to scale in terms of capability and performance. Such file systems take advantage of RDMA transfers with large bandwidth and reduced CPU usage. The parallel file system is usually used as scratch space and intended for work that requires optimized I/O. Examples include workload setup, pre-processing, running, and post-processing.

6.2 Basic Virtualized HPC (vHPC) Architecture

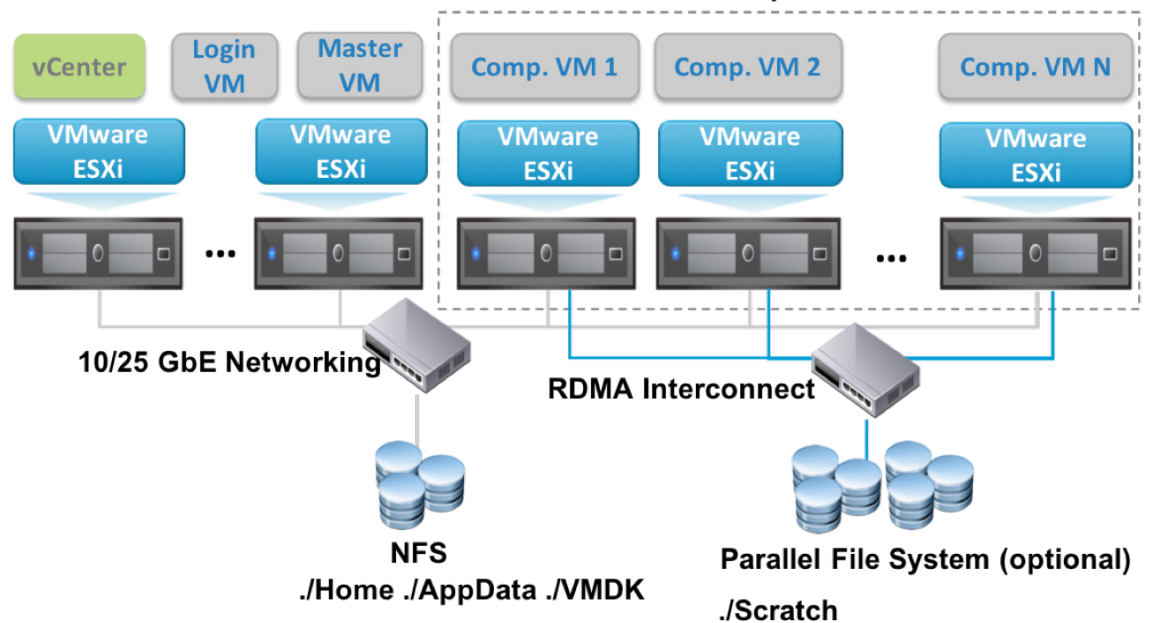


Figure 4. vHPC Cluster with Single VM per Node (Basic Architecture)

As illustrated in Figure 4, based on the architecture of traditional HPC, the management nodes and compute nodes are easily virtualized with VMware vSphere, which includes two core components: ESXi™ and vCenter Server®. VMware ESXi is a Type-1 hypervisor installed directly on bare-metal servers, providing a layer of

abstraction for VMs to run while mapping host resources such as CPU, memory, storage, and network to each VM. vCenter Server provides centralized management of the hosts and VMs, coordinating resources for the entire cluster.

The management login node, master node for job scheduling, DNS, and gateway run as VMs on one or more management nodes within the same cluster. Each VM will run Linux as the “guest OS,” which will be installed on virtual machine disks (VMDKs) that are stored securely on a shared datastore. Microsoft Windows or other operating systems may be used as well.

6.3 Secure HPC Cloud Architecture

Using the vHPC cluster as a foundation, other VMware products and technologies can be leveraged to deploy a secure private-cloud environment for HPC, as shown in Figure 5.

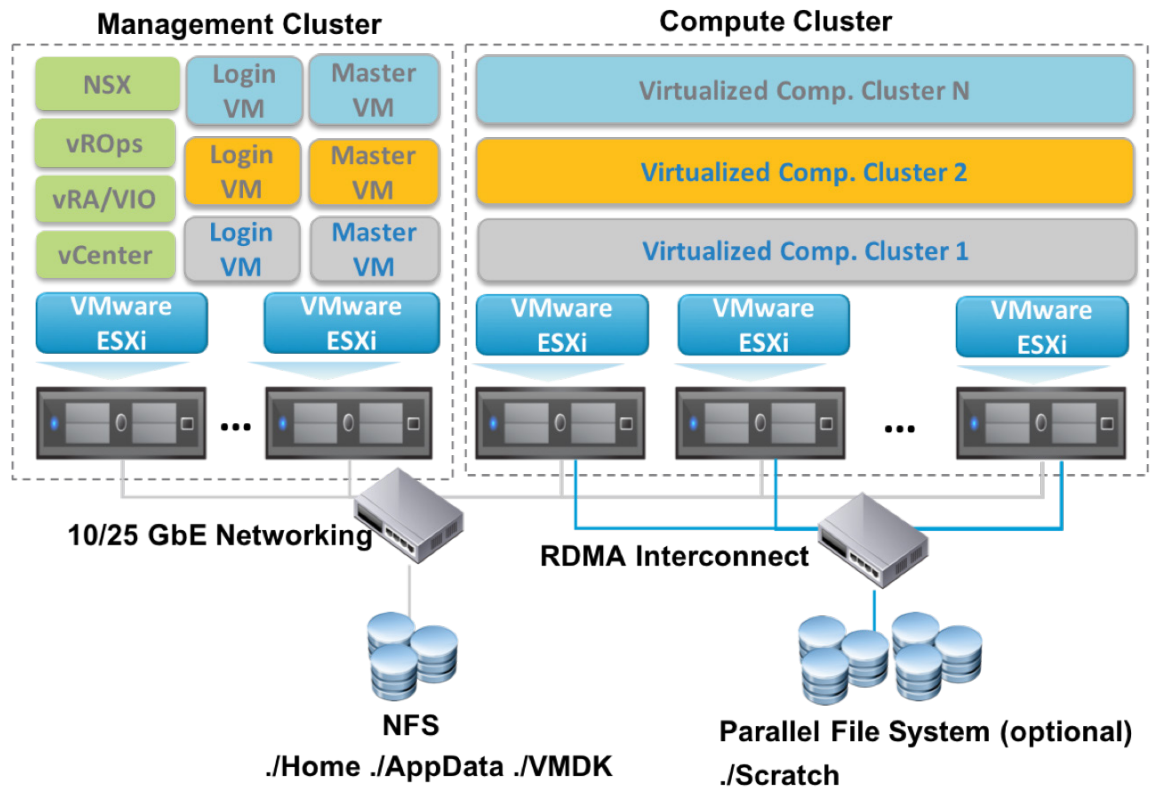


Figure 5. Secure Private HPC Cloud Architecture

The [additional VMware software](#) for operations that can be leveraged includes:

- **VMware vSAN™.** A high-performance hyper-converged storage solution that aggregates SSDs across servers to provide a shared datastore for VMDKs in the management cluster. This provides a simple, cost-effective, and highly available solution that is easy to manage and won't impact the performance of storage used in the compute clusters.

- **VMware NSX®:** A network virtualization platform that delivers networking and security entirely in software, abstracted from the underlying physical infrastructure. It provides switching, routing, load-balancing and firewall to the management and compute VMs through optimized networking and security policies that are distributed throughout the cluster. Network and security policies are handled at the virtual switch and apply only to the Ethernet network, not to the RDMA interconnects, which should be configured separately.
- **VMware vRealize® Automation™ (vRA):** Accelerates the deployment and management of applications and compute services, empowering IT and researchers to automate the deployment of vHPC workloads that can be requested and delivered on-demand through a self-service portal. The concept of a “blueprint” is used to describe the deployment and configuration of VMs, virtual networks, storage, and security properties of an application as it’s provisioned to a cluster.
- **VMware vRealize® Operations™ (vROps):** A robust operations-management platform that delivers performance optimization through efficient capacity-management, proactive planning and intelligent remediation. It continuously monitors the health, risk, and efficiency of the vHPC clusters through real-time predictive analytics to identify and prevent problems before they occur.
- **VMware Integrated OpenStack (VIO):** A VMware-supported OpenStack distribution that makes it easy for IT to run an enterprise-grade OpenStack cloud on top of VMware virtualization technologies, boosting productivity by providing simple, standard, and vendor-neutral OpenStack API access to VMware Infrastructure.

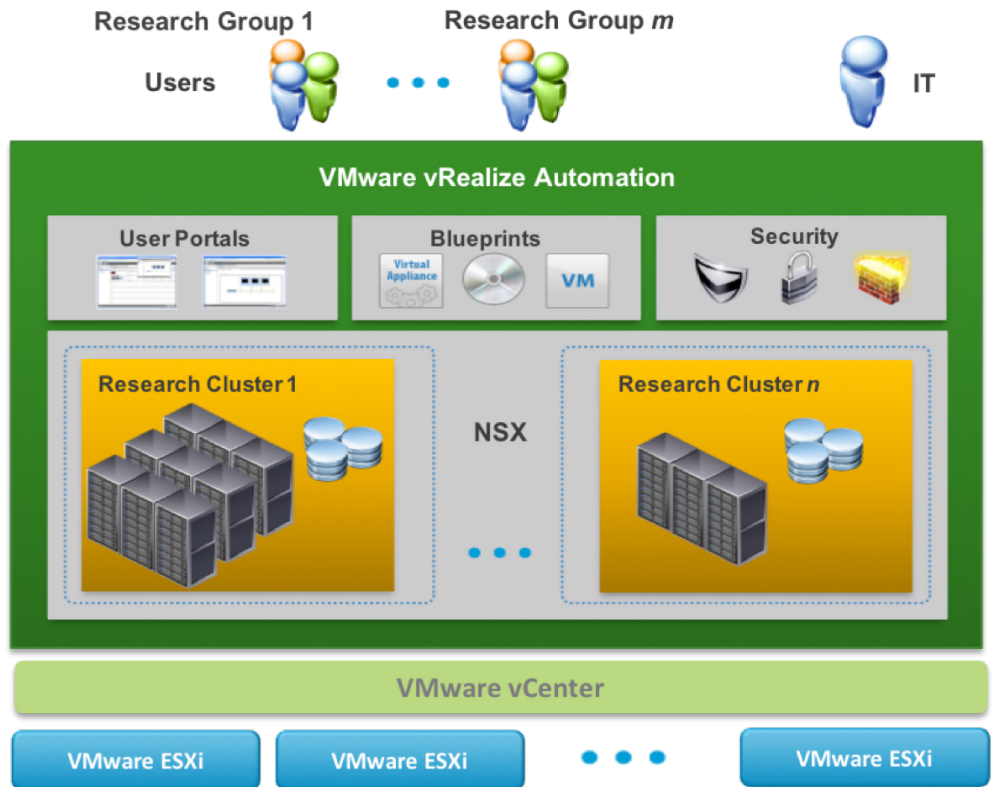


Figure 6. Secure Private HPC Cloud Architecture with vRealize Automation and NSX

Figure 6 presents a secure private-cloud architecture with vRealize Automation and NSX. This architectural approach takes the basic vHPC architecture as shown in Figure 4 and wraps it in a private-cloud infrastructure. This adds self-provisioning, which allows individual departments or researchers and engineers to instantiate the resources they need for their project without needing to wait for the IT department to create the resource for them. When an end-user instantiates a virtual HPC cluster, it is done using a blueprint which specifies the required machine attributes, the number of machines, and the software that should be included in the VM, including operating system and middleware, allowing full customization to their requirements. By the same token, the blueprint approach also allows the central IT group to enforce corporate IT requirements—security and data-protection policies, for example—by including those in the blueprint as well.

6.4 Management Cluster

The management cluster runs the VMs that manage the vHPC environment. As shown in Figure 7, these include vSphere and vSphere-integrated components such as vSAN, NSX, vCenter Server, vRealize Operations, vRealize Automation, and HPC administrative components such as the master VM for workload scheduling and login VM. The management cluster provides high availability for these critical services. Permissions for the management cluster limit access only to administrators, protecting VMs running management, monitoring, and infrastructure services from unauthorized access.

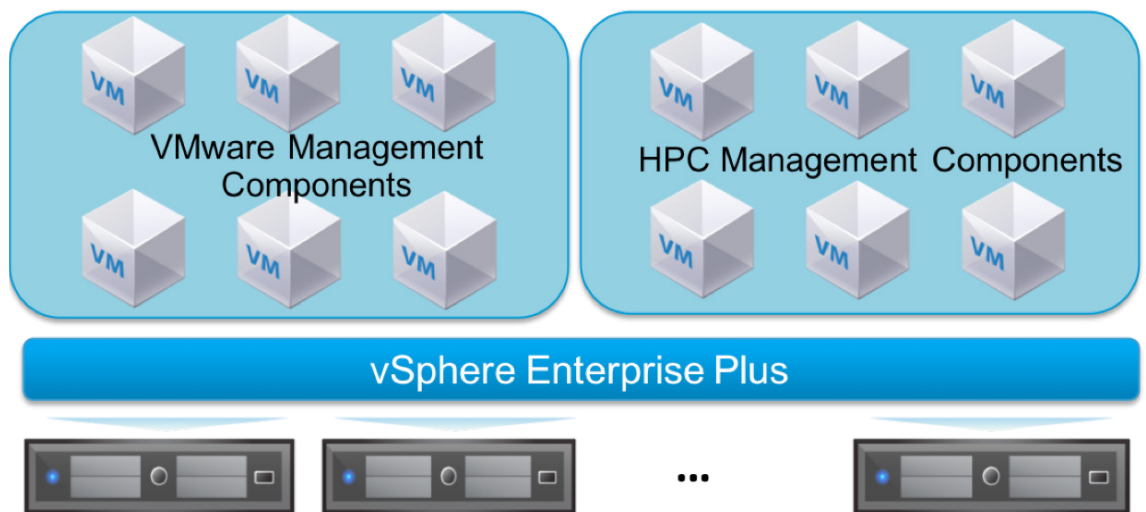


Figure 7. vHPC Management Cluster with VMware vSphere® Enterprise Plus Edition™

The management cluster, based on [VMware Validated Designs](#), recommends a minimum 4-node vSAN cluster in order to tolerate failure of a node even with one node removed for maintenance. The new cluster should be sized based on projected management workload, including additional headroom for growth. Capacity analysis should be performed on an existing management cluster and adjusted to ensure that there is enough capacity to add HPC management components.

Due to the critical nature of workloads with many single points of failure, it is recommended that this cluster be licensed with VMware vSphere Enterprise Plus Edition for high availability and other advanced features. VMware vSphere Enterprise Plus Edition provides VMware vSphere High Availability (HA), Distributed Resource Scheduler (DRS), and other advanced capabilities that help reduce downtime for these critical workloads.

6.5 Compute Clusters

The compute clusters run the HPC workloads for different scientific and engineering groups. As shown in Figure 8, VMware's [vSphere Scale-Out](#) license supports HPC workloads at a cost-effective price point and can be leveraged for these compute clusters.

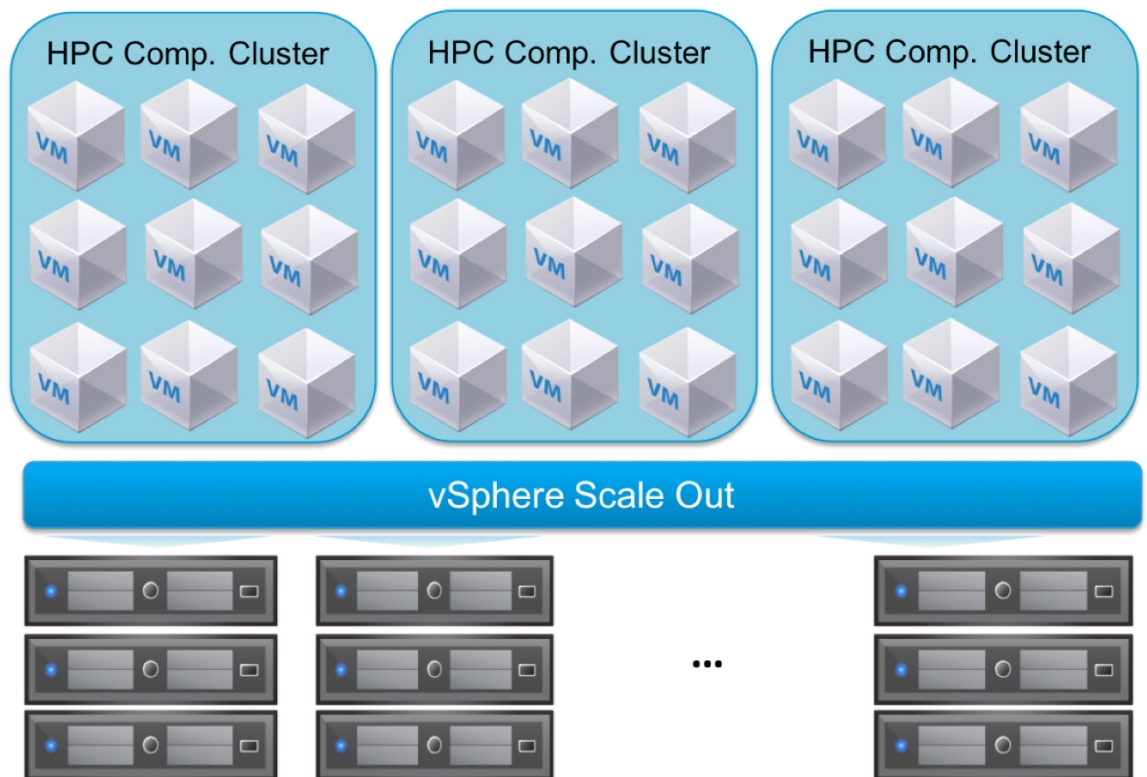


Figure 8. vHPC Compute Clusters with VMware vSphere Scale-Out

6.6 Hardware Accelerators

Compute accelerators are specialized hardware designed for highly parallel mathematical operations. The use of compute accelerators for handling massive, parallelized computation across many cores is a strong trend in HPC and machine learning. By offloading the intensive computations from CPU to accelerators, each accelerated server may replace dozens of commodity CPU servers.

Typical hardware compute accelerators include GPUs, Intel Xeon Phi, and field-programmable gate arrays (FPGA). NVIDIA, AMD, Intel, and Xilinx are examples of vendors that provide such compute accelerators. GPUs have been the most common accelerators in recent years.

Typical areas that use accelerators include life sciences, image processing, quantum chemistry, and machine learning. Particularly, the emergence of deep learning and the enhancement of accelerators have driven adoption of machine learning applications across a broad spectrum of industries and use cases, such as facial recognition, medical diagnosis, robotics, automobile safety, and text and speech recognition.

In general, most compute accelerators can be configured in DirectPath I/O (passthrough) mode on VMware vSphere¹, which allows a guest OS to directly access the device, essentially bypassing the hypervisor. Because of the shortened access path, applications accessing accelerators in this way can achieve levels of performance very close to that of bare-metal systems. With DirectPath I/O, one or multiple accelerators can be configured into a single VM. Each accelerator is dedicated to a VM and there is no accelerator sharing among the VMs. For example, Figure 9 shows configuring one or three GPUs in DirectPath I/O mode.

¹ While VMware supports the VM Direct Path I/O feature, the official support statement for a particular card and/or card+system configuration would have to be made by the system vendor.

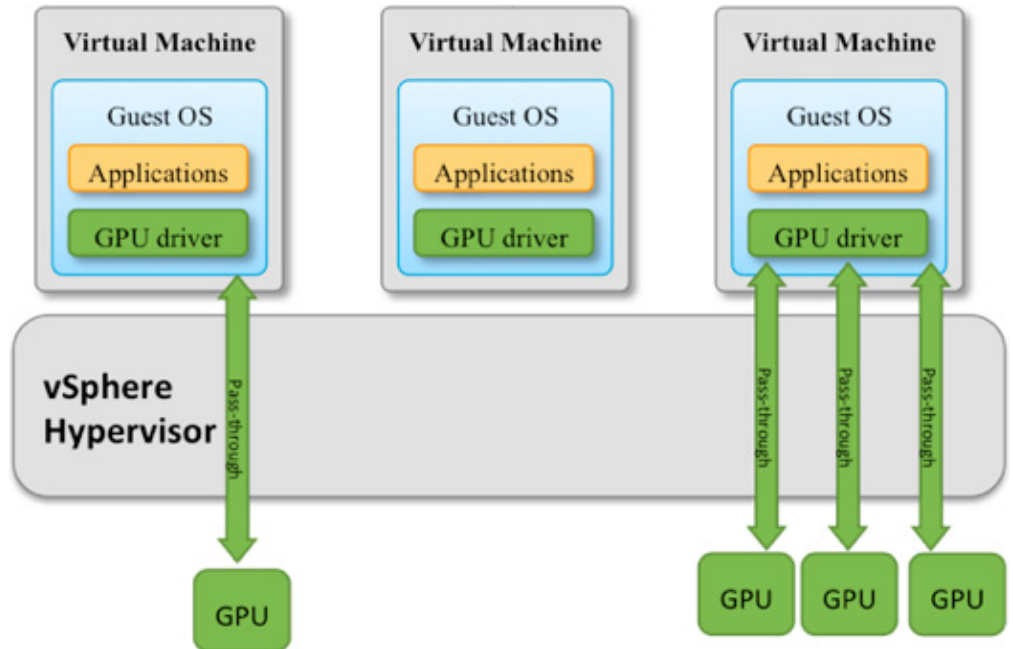


Figure 9. Configuring GPU in VMware DirectPath I/O (Passthrough) Mode

In addition to DirectPath I/O, in the VMware virtual environment there are multiple ways of sharing GPUs. Using NVIDIA GRID vGPU technology in VMware vSphere, sharing an NVIDIA GPU card across multiple VMs can be enabled by creating multiple logical virtual GPU (vGPU) devices, each of which can be assigned to a VM (Figure 10). vSphere 6.7 also provides the capability to execute suspend/resume for vGPU enablement. Another GPU-sharing technology is FlexDirect, which is a GPU virtualization solution provided by Bitfusion, a VMware partner. It allows machine learning and other applications running inside a VM to access one or multiple GPUs installed in remote hosts. It also supports multiple VMs sharing a single physical GPU.

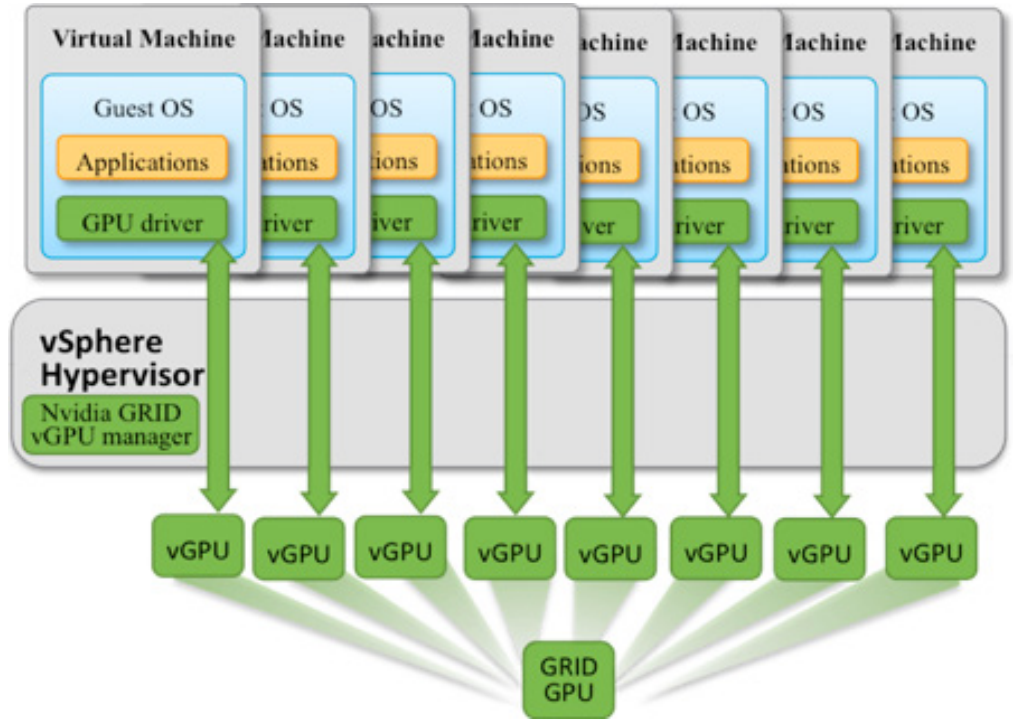


Figure 10. NVIDIA GRID vGPU

7. SAMPLE ARCHITECTURE

The following sections demonstrate the ways in which various common HPC scenarios can be virtualized.

7.1 Scenario A: MPI Workloads

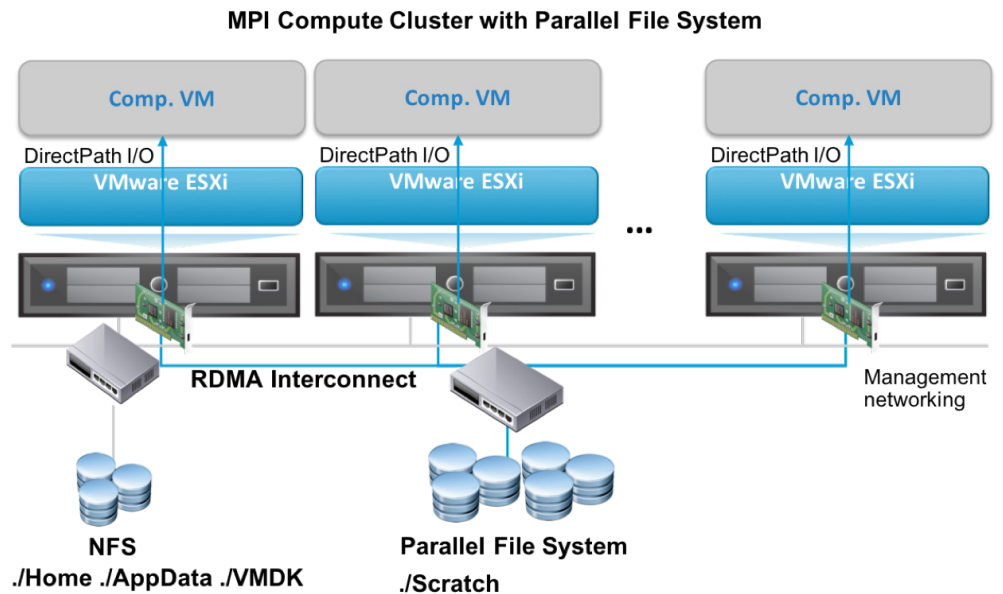


Figure 11. Virtualized Compute Cluster Architecture for MPI Workloads with Access to Parallel File System via RDMA Interconnect (DirectPath I/O)

As illustrated in Figure 11, the sample architecture for a vHPC environment that runs MPI workloads consists of:

Hardware

- Multiple compute nodes: The number of nodes determined by the workload computational needs.
- Management nodes: A minimum of four nodes is recommended for enterprise-class redundancy.
- Existing network storage for VMDK placement and long-term application data storage
- Parallel file system for application scratch data (optional)
- High-speed interconnects (such as 100 Gb/s Ethernet or RDMA) for achieving low-latency and high-bandwidth for HPC application message exchanges or accessing the parallel file system
- Ethernet cards with 10/25 Gb/s connectivity speed for management
- GPUs or other accelerators for application acceleration needs (optional)

Software

- VMware vSphere
- VMware vSAN
- VMware NSX
- VMware vRealize Suite/VMware Integrated OpenStack: It's possible to leverage existing VMware management clusters that support enterprise workloads to also manage the vHPC environment.
- HPC management and operations solutions, such as HPC batch scheduler
- Lustre for parallel file system

VM sizing

- One compute VM per host because MPI workloads are CPU-heavy and can make use of all cores. CPU and memory overcommit would greatly impact performance and are not recommended. For further details, please see “VM-Sizing, Placement, and CPU/Memory Reservation” in Running High-Performance Computing Workloads on VMware vSphere Best Practices Guide.
- The size of other management VMs should be determined by the considerations of management requirements.

7.2 Scenario B: Throughput Workloads

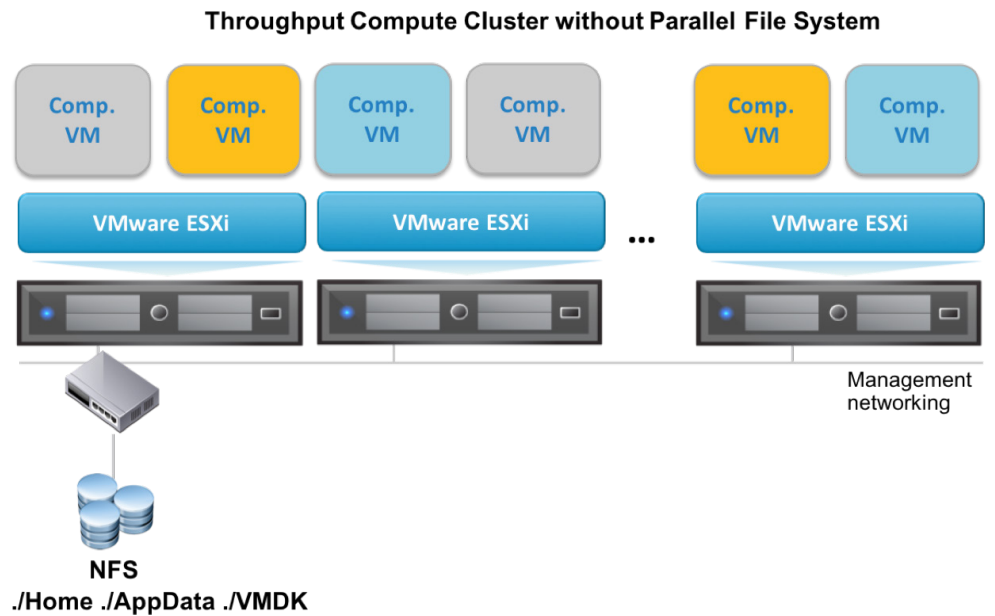


Figure 12. Virtualized Compute Cluster Architecture for Throughput Workloads Without Parallel File System

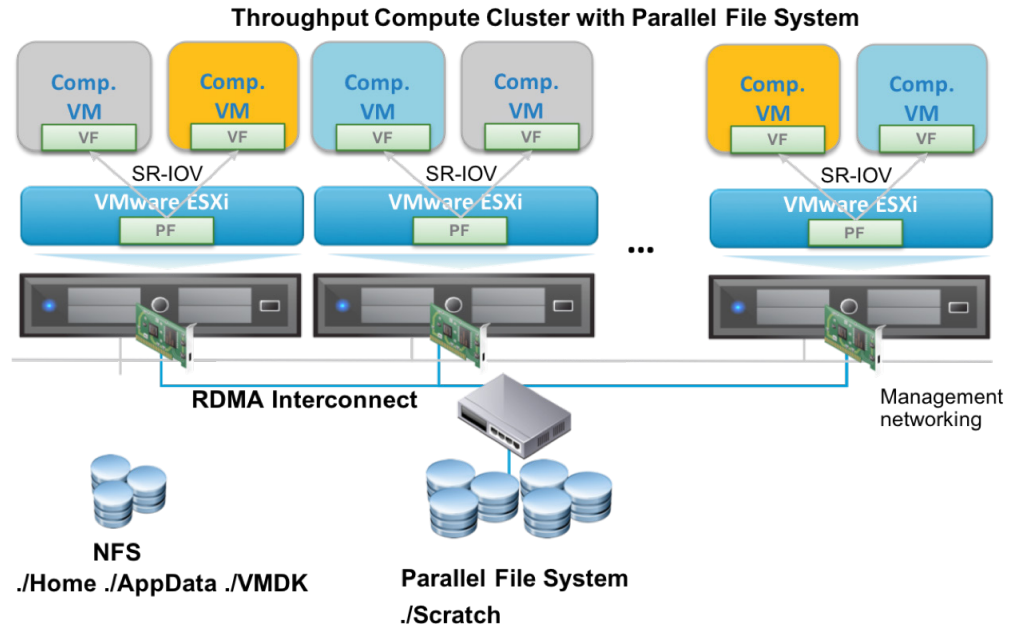


Figure 13. Virtualized Compute Cluster Architecture for Throughput Workloads with Access to Parallel File System via RDMA Interconnect (SR-IOV)

As illustrated in Figure 12 and Figure 13, the sample architectures for a vHPC environment running throughput workloads with or without a parallel file system consists of:

Hardware

- Multiple compute nodes: The number of nodes determined by the workload computational needs.
- Management nodes: A minimum of four nodes is recommended for enterprise-class redundancy.
- Existing network storage for VMDK placement and long-term application data storage
- Parallel file system nodes for application scratch data (optional)
- Throughput workloads don't require low-latency and high-bandwidth for HPC application message exchanges, however, they can still leverage parallel file systems. If there is a need for a parallel file system, high-speed interconnects (such as 100 Gb/s Ethernet or RDMA) are preferred. DirectPath I/O can only pass one InfiniBand device to a single VM but most throughput cases involve running multiple VMs per host. In that case, using SR-IOV (single root I/O virtualization) technology can enable a single physical PCI device to be shared by multiple VMs. As shown in Figure 13, SR-IOV functionality of the network adapter enables virtualization and exposes the PCIe physical function (PF) into multiple Virtual Functions (VF).
- Ethernet cards with 10/25 Gb/s connectivity speed for management
- GPUs for application acceleration needs (optional)

Software

- VMware vSphere
- VMware vSAN
- VMware NSX
- VMware vRealize Suite/VMware Integrated OpenStack: It's possible to leverage existing VMware management clusters that support enterprise workloads to also manage the vHPC environment.
- HPC management and operations solutions, such as HPC batch scheduler
- Lustre for parallel file system (optional)

VM sizing

- Multiple compute VMs per host. It's recommended to use different sizes of VMs conforming with NUMA boundaries. CPU oversubscription can be leveraged to achieve higher overall cluster throughput than bare-metal environments.
- The size of other management VMs should be determined by the considerations of management requirements.

8. CONCLUSIONS

Virtualization offers tremendous benefits for HPC environments. With a simple understanding of HPC components and technologies, any enterprise organization can leverage HPC best practices and guidelines to adopt a VMware virtualized infrastructure. Virtualization provides the building blocks for individualized or hybrid clusters designed to host important HPC applications. Meanwhile, scale-out licensing for vSphere and HPC allows administrators to virtualize compute nodes at a low cost.

Through the addition of optional VMware management solutions, HPC further benefits from increased security, self-service, multi-workflow environments and remote-console accessibility. As the broader HPC community looks toward the cloud, HPC virtualization is a must to protect and tune IT infrastructure for the future.

9. AUTHORS

Justin King has been involved with the IT industry for over 20 years, holding various roles and responsibilities from administration to architecting solutions. Since joining VMware in 2009, Justin has supported sales teams as a sales engineer, evangelized VMware technologies, and currently instills confidence by designing and testing end-to-end data analytic solutions from VMware and Dell Technologies.

Mohan Potheri is VCDX#98 and has more than 20 years in IT infrastructure, with in-depth experience on VMware virtualization. He currently focuses on evangelization of High-Performance Computing (HPC) and big data virtualization on vSphere. He also has extensive experience with business-critical applications such as SAP, Oracle, SQL, and Java across UNIX, Linux, and Windows environments. Mohan is an expert on SAP virtualization and has been a speaker at multiple VMworld and PEX events. Prior to VMware, Mohan worked at many large enterprises where he has engineered fully vHPC solutions. He has planned, designed, implemented, and managed robust, highly available, DR-compliant virtual environments in UNIX and x86 environments.

Na Zhang is senior member of technical staff working on High-Performance Computing (HPC) within VMware's Office of the CTO. She has broad experience in HPC and vHPC. She has been working on various vHPC topics, including performance tuning for throughput workloads, MPI workloads, and financial services workloads in virtual environment, design, and implementation of vHPC tools, accelerator solutions, and integration of HPC middleware with VMware products. She received her Ph.D. in Applied Mathematics from Stony Brook University in 2015. Her research primarily focused on design and analysis of parallel algorithms for large- and multi-scale simulations running on world-class supercomputers. She has served on the Technical Program Committee for more than 10 international HPC conferences and workshops, including SC (The International Conference for High-Performance Computing, Networking, Storage, and Analysis, vHPC, HPC&S), vHPC (Workshop on Virtualization in High-Performance Cloud Computing), and HPC&S (The 2018 International Conference on High Performance Computing & Simulation).

Shawn Kelly is a Staff Engineer in San Diego. He works directly with customers to develop innovative solutions in multiple disciplines, including HPC and Research Computing. He earned Master of Science degrees in both Software Engineering and Information Technology Management from the Naval Postgraduate School in 2010, with a focus on cloud computing, and subsequently served at Marine Corps Systems Command developing solutions for the Marine Corps.

Josh Simons is the Chief Technologist for HPC. He currently leads an effort within VMware's Office of the CTO to bring the value of virtualization to HPC. He has over 20 years of experience in High-Performance Computing (HPC). Previously, he was a Distinguished Engineer at Sun Microsystems with broad responsibilities for HPC direction and strategy. He joined Sun in 1996 from Thinking Machines Corporation, a pioneering company in the area of Massively Parallel Processors (MPPs), where he held a variety of technical positions. Josh has worked on developer tools for distributed parallel computing, including language and compiler design, scalable parallel debugger design and development, and MPI. He has also worked in the areas of 3D graphics, image processing, and real-time device control. Josh has an undergraduate degree in Engineering from Harvard College and a Masters in Computer Science from Harvard University. He has served as a member of the OpenMP ARB Board of Directors since 2002.

