

Scalable, Flexible High Performance Computing in the Cloud

July 2020

Take advantage of the flexibility and scalability of the cloud to power modern, data-intensive and analytical applications with SUSE High Performance Computing and Microsoft Azure.

Introduction

SUSE and Microsoft—leaders in the HPC and hyperscale cloud worlds—have joined forces to make HPC in the cloud easier and faster to implement and capable of scaling to meet any enterprise challenge.

The world of high performance computing (HPC) is expanding rapidly. New workloads such as artificial intelligence and analytics have driven demand for the processing capabilities that were once limited to supercomputers. Today, however, enterprises no longer need specialized hardware. With HPC in the cloud, they can get the same capabilities with significantly reduced cost and effort.

The Growing Need for HPC

High performance computers have escaped the research laboratory. Organizations have begun to realize that they need HPC to match the volume of data and the complexity of analytics required to compete and create insight today. For that reason, parallel processing clusters have begun to appear in manufacturing, financial services, entertainment, retail and many other industries.

Organizations are using HPC to tackle massive and complex workloads, such as:

- **Simulations:** Manufacturing, engineering, life sciences, financial services and energy all use simulations to model and develop products and to understand complex scenarios such as fluid dynamics and loads on utility grids.
- **Artificial intelligence:** Artificial intelligence (AI) is already proving useful for image and voice recognition, as well as for a number of other uses across industries. Accurately training AI models requires large computational systems—hyperscale environments—often with graphical processing units (GPUs). AI learning models

have garnered most of the attention, but graph analytics will also play a crucial role with its unique ability to handle temporal and spatial relationships.

- **Machine learning:** Machine learning, a subset of AI, is an effective tool for fraud detection and other cybersecurity tasks, as well as for developing autonomous vehicles. It is in use or being explored by industries of all kinds.
- **Deep learning:** Organizations have begun to use deep learning for tasks such as disease identification and energy demand optimization. By combining AI and machine learning technologies with HPC cloud capabilities, organizations can turn enormous data sets into predictive analytics and insights.
- **Rendering:** Rendering images and 3D models in media production and engineering can be accomplished far faster with HPC than with traditional computing.
- **Visualization:** Similar to rendering, HPC systems can significantly speed up visualization of data and model, and can help scale up visualization systems to serve multiple users.
- **IoT:** Many organizations keep their IoT data in the cloud. Running HPC in the cloud is thus a logical way to analyze and draw value from that data, eliminating the expense of moving the data from the cloud to the data center.

Fortunately, for organizations putting these workloads in place, it has become easier than ever to achieve HPC.

Cloud HPC Goes Mainstream

Hyperion, a research firm in the HPC space, predicts that spending for HPC in the cloud will increase 83 percent by 2022.¹ There are multiple reasons for this growth. HPC in the cloud offers many of the same benefits as other workloads in the cloud, but often in more pronounced or dramatic ways:

Hyperion predicts that spending for HPC in the cloud will increase 83 percent by 2022.²

- **CapEx to OpEx.** The high cost of purpose-built, on-premises HPC systems makes the ability to trade capital expenditures for operating expenditures a powerful benefit of HPC in the cloud.
- **Agility and speed.** Traditional HPC systems suffer from long setup times, making the agility and speed benefits of the cloud highly valuable.
- **Scale.** The cloud's ability to reach massive scale economically has an enormous impact in a field where huge jobs are often the norm.

In addition, for organizations that store data in the cloud, it can be more efficient to also run the HPC-powered analysis in the cloud, thus avoiding the cost of moving all that data.

Given these benefits, it's not surprising that over 70 percent of HPC sites run some of their jobs in public clouds, up from only 13 percent in 2011.³ But, how do you use the cloud for HPC? How can you ensure the throughput, compute resources, workload management and more that are necessary to realize these benefits? Let's look at two ways of using cloud HPC and then look at how SUSE and Microsoft Azure can help.

Two Models for Cloud HPC

Organizations can choose to build their entire HPC system in the cloud or they can use cloud HPC capabilities to supplement their on-premises HPC systems, such as by bursting to the cloud.

ALL CLOUD

Running your entire HPC infrastructure in the cloud is ideal for when you need to set up a new HPC system

quickly. It also makes sense if you need HPC for limited occasions or if your need for it isn't consistent enough to justify capital expenditures. It might also be a good solution when networking limitations prohibit you from using the bursting model, which can demand high throughput.

Because the head node, compute nodes and storage are all in the cloud, finding the right cloud vendor is paramount.

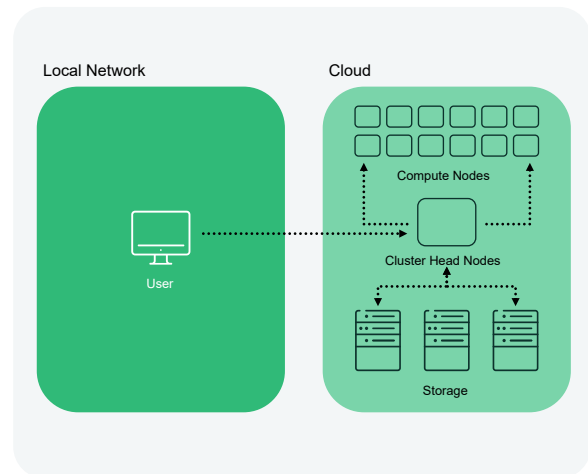


Figure 1. In the all-cloud option, the entire HPC cluster is in the cloud. Users connect to the system from outside.

BURSTING

Using the cloud to support an on-premises system is ideal if your organization already has an HPC cluster running. Even if you don't yet have an on-premises system, it can make sense to invest in one if your use will be consistent and steady over a long period. By adding cloud bursting to this on-premises system, you can keep the on-premises system small and cost-effective, without sacrificing the ability to run large jobs.

1 Hyperion Research Update: Research Highlights In HPC, HPDA-AI, Cloud Computing, Quantum Computing, and Innovation Award Winners, Hyperion Research, 2019
2 Ibid

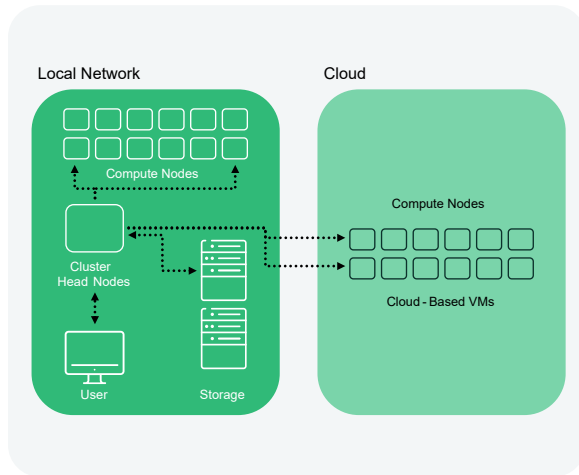


Figure 2. With cloud bursting, you run a job locally and then burst to the cloud when you need additional capacity.

Regardless of the model you adopt, it is crucial to find the right partners to help you build your cloud infrastructure.

SUSE and Microsoft Azure

SUSE Linux is the leading operating system for HPC. It runs on 21 of the top 50 supercomputers and 37 percent of the top 100 systems. In all, over half of the paid Linux systems in the HPC TOP500 use SUSE Linux Enterprise.⁴

SUSE is also the first supported enterprise Linux HPC distribution available on Microsoft Azure. Azure offers the most productive HPC infrastructure of any cloud provider. In fact, in a recent study by independent researcher Exabyte⁵, Azure came out as the best performing—ahead of other cloud competitors and even most traditional on-premises clusters.

With SUSE Linux Enterprise High Performance Computing on Microsoft Azure, you get the capabilities you need to succeed with HPC in the cloud.

AUTOMATED SCALING

SUSE Linux Enterprise High Performance Computing comes with SUSE-supported versions of popular workload management tools such as Slurm and Ganglia, which can reduce your workload when scaling and managing clusters. When you implement the SUSE solution on Azure in an auto-scaling model, you get similar benefits. Its built-in horizontal autoscaling means your operators don't need to

continually monitor performance and make decisions about adding or removing resources.

FASTER OPERATIONS WITH HIGH-SPEED DATA TRANSPORT

An HPC cluster is only as good as the interconnects and networking that support it. Azure supports Linux remote direct memory access (RDMA) technology, which allows Linux virtual machines in Azure to directly communicate through the InfiniBand network.

This enables faster communication between virtual machines at low latency for parallel workloads and applications that take advantage of the Message Passing Interface (MPI) stack.

Meanwhile, ExpressRoute helps you maintain a fast, private connection to the cloud. ExpressRoute connections don't go over the public internet. They offer greater reliability, faster speeds, consistent latencies and higher security than typical internet connections.

IMPROVED PERFORMANCE

SUSE and Microsoft collaborated to deliver the first enterprise Linux kernel optimized for Azure. The Azure-tuned kernel is included in SUSE Linux Enterprise High Performance Computing. It provides faster boot speeds and has a smaller memory footprint.

With a 25 percent faster network throughput and a 23 percent reduction in average latency, the Azure-tuned kernel can significantly improve the performance of HPC infrastructure.

For your most demanding HPC workloads, our partnership with Cray allows customers to get dedicated supercomputers on their Azure virtual network.

SEAMLESS ORCHESTRATION

In an all-cloud HPC system, Azure offers two tools to help you avoid the time and expense of re-architecting your applications for the cloud. You can simply extend the HPC resource environment that you already have to Azure by using:

⁴ Top500 Supercomputer Report, June 2019

⁵ Top500: Exabyte Measures Linpack Performance Across Major Cloud Vendors
<https://www.top500.org/news/exabyte-measures-linpack-performance-across-major-cloud-vendors/>

- **Azure Batch.** Batch enables you to manage your application workflows. Developers can define the Azure compute resources necessary to execute your applications in parallel or at scale, without having to manually configure or manage the infrastructure.
- **Azure CycleCloud.** While batch enables access to HPC resources by running jobs in Azure, CycleCloud runs compute clusters. IT can provide access to compute clusters for various business groups or departments. CycleCloud also offers tools for managing the flow of data and jobs between clusters.

IMPROVE TCO WITH RESERVATIONS

SUSE Linux Enterprise High Performance Computing is eligible for Azure Reservations. Azure Reservations enable you to control costs through up front purchases of Azure with SUSE HPC.

Microsoft partnered with SUSE before any other open source distribution to offer 1-year and 3-year reservations at significant discounts, as detailed here: suse.com/c/azure_reservations_for_suse_linux/.

Given these benefits, it's not surprising that over 70 percent of HPC sites run some of their jobs in public clouds, up from only 13 percent in 2011.³ But, how do you use the cloud for HPC? How can you ensure the

Every dollar invested in HPC brings an average of \$44 in profits or in cost savings.⁶

Solution Components

MICROSOFT AZURE

The Microsoft Azure cloud is designed for HPC workloads. It offers HPC-optimized infrastructure components, including high-memory virtual machines and the latest generation GPUs. It also includes high IOPS storage and mission-specific, field-programmable gate arrays (FPGAs) for everything from AI training

jobs, accelerated computing and visualization. You can also get a fully managed, dedicated supercomputing service on Cray.

Azure is also the only public cloud with support for InfiniBand.

SUSE LINUX ENTERPRISE HIGH PERFORMANCE COMPUTING

SUSE provides an enterprise-class, high-performance, highly scalable open source operating system built to use the power of parallel computing for modeling or simulation, artificial intelligence and advanced analytics applications.

The key feature of the solution is the HPC module, a supported set of popular HPC tools and utilities that make managing and monitoring parallel computing environments easier.

The HPC module is updated independently of the underlying operating system, ensuring that you can access the latest HPC innovations as quickly as possible. In addition, multiple service life options (including Extended Service Pack Overlap Support and Long-Term Service Pack Support) enable you to update the operating system at the best time for you, to avoid interrupting vital HPC projects.

Master the HPC Cloud Era

HPC often requires a massive number of cores, features a peak-and-valley nature to the jobs and is a field of constantly evolving technology. That's what makes HPC a natural fit for the cloud. With HPC expertise from SUSE and Microsoft Azure, you can take advantage of this new way of powering advanced applications and capture the flexibility, performance and cost savings of HPC in the cloud.

⁶ Hyperion Research Update: Research Highlights In HPC, HPDA-AI, Cloud Computing, Quantum Computing, and Innovation Award Winners, Hyperion Research, 2019



Thank You

SUSE
Maxfeldstrasse
90409 Nuremberg
www.suse.com

For more information, contact SUSE at:
+1 800 796 3700 (U.S./Canada)
+49 (0)911-740 53-0 (Worldwide)