Cloud-Ready High Performance Computing

Taking HPC to the Clouds

WHITEPAPER

suse.com

Adding the cloud to your HPC configuration can save you money and time, but it also pays to be careful and plan ahead. The ideal configuration depends on your budget, your workload and the location of your data.

High performance computing (HPC) was once the preserve of universities and government research institutions, where it offered the compute capacity necessary for modeling hurricanes and studying the Big Bang. In today's universe, however, HPC is equally at home in the enterprise. The vast computing power unleashed through HPC, and the accompanying tools and programming techniques of the HPC environment, are changing the rules for IT departments. HPC means any company can explore the possibilities of artificial intelligence, machine learning and advanced risk analysis.

Future-focused corporations are currently working to build HPC into their plans, for better:

- **Engineering:** HPC systems model new designs and simulate environmental forces.
- **Banking:** HPC offers state-of-the-art forecasting and fraud protection.
- **Medicine:** HPC systems provide on-demand diagnoses and push the envelope for better drug design.
- **Infrastructure:** HPC is at the heart of the drive for smart cities, traffic control and power grid optimization.
- **Retail:** Vendors depend on HPC for customer studies, sale modeling and supply chain management.

HPC is also the engine driving the advanced analytics (aka Big Data) revolution. Almost any company can improve efficiency and add value by employing Big Data techniques to derive valuable insights from sales records and internal process information.

Organizations around the world are well aware that HPC is in their future. According to Hyperion Research, worldwide HPC revenue is expected to reach over US$19.5 billion by 2022. One of the biggest drivers for this growth is a new generation of applications that are much more complex than traditional enterprise applications. This new generation includes data-intensive applications for AI, machine learning and advanced analytics that require both high performance and extreme scalability in order to produce timely results. This need, in turn, drives the need for an HPC infrastructure that helps to manage the workloads within a parallel cluster environment.

When building the HPC infrastructure, enterprises can choose to leverage hardware specifically designed to support the demands of HPC applications around memory, compute power, network communications and storage needs. In some cases, a combination of HPC hardware and cloud platforms can provide the best overall environment. According to the 2018 Gartner report, "Top 5 Ways to Successfully Deliver HPC Cloud Strategies," by 2023, the number of organizations adopting HPC cloud will increase to 30 percent from 5 percent today, resulting in higher infrastructure agility in HPC environments.

But what are some guidelines for determining how to balance your HPC infrastructure needs across on-premises and cloud? Solutions vary depending on needs and budget, and enterprises are increasingly looking to cloud computing as a cost-effective alternative to handle the new generation of HPC applications.

Cloud computing provides several benefits for the HPC environment. The cloud reduces overhead and offers unmatched flexibility for adapting to changing needs. Organizations are quickly becoming aware of the benefits of cloud computing in the HPC space. According to Hyperion Research, 70 percent of all HPC sites now run jobs in the public cloud (up from 13 percent in 2013). Overall, 10 percent of all HPC jobs now run in the cloud. Hyperion predicts that spending for HPC in the cloud will increase 83 percent by 2022.

This paper offers some insights on how to integrate cloud computing with your HPC systems.

HPC in the Cloud

Cloud computing reduces deployment costs and ongoing expenses, enabling the local team to focus on software development and analytics and minimizing the complexity of the hardware environment. The two most popular cloud scenarios for HPC are:

- **All-cloud:** The entire HPC configuration resides in the cloud environment.
- **Cloud bursting:** The base HPC configuration is located on-premises, with additional capacity in the cloud for times of peak activity.

The all-cloud configuration eliminates local hardware expense and maintenance and minimizes complexity, although it could potentially lead to a higher cost per compute cycle. The cloud-bursting configuration enables the organization to retain some local capacity, but allows for a more streamlined local presence without sacrificing the ability to run big jobs when the need arises.

All-Cloud: Minimize the Local Footprint

In its most basic form, the all-cloud option calls for all components of the HPC infrastructure—including the head node, compute node and storage nodes—to reside in the cloud environment. The user connects to the HPC cluster using some form of remote access technology, such as SSH (Figure 1). The user configures and launches the HPC workload. The cloud-based cluster springs into action, providing a result that is then displayed or downloaded back to the user.

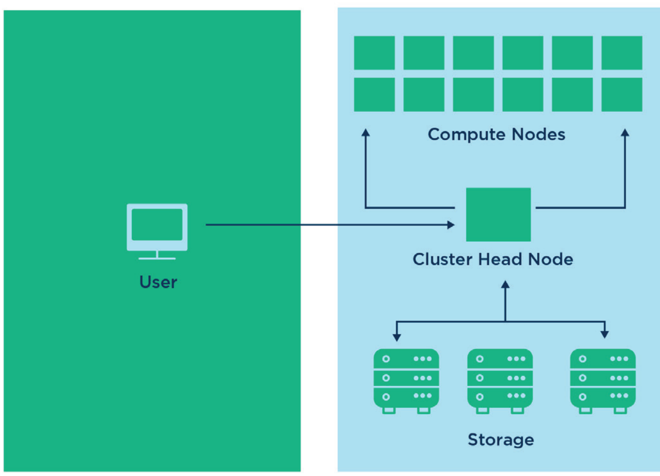

Figure 1: *The all-cloud option: The whole cluster is in the cloud and the user interacts with the system from a local network.*

The all-cloud option is ideal for scenarios where:

- You need to get an HPC system set up quickly, without waiting for hardware purchase, assembly and configuration.
- The utilization of the HPC system is not sufficient to warrant the expense of setting up an on-premises system.
- Network speeds and infrastructure are not suited for the bursting option.

Several vendors offer support for HPC in the cloud. Check the details carefully to ensure that you are truly getting what you expect. The network medium and system resources for the compute nodes within the data center could have a significant effect on latency and performance. Storage requirements could also determine whether all-cloud is a suitable option. Some organizations prefer to store their data locally. In other cases, the need for independent access to the data or the cost of cloud storage prohibits the use of cloud-based storage nodes.

Cloud Bursting: Expanding from Local

The bursting option is growing in popularity with the rise of virtual machine (VM) orchestration technology and fast internet connections. In a bursting scenario (Figure 2), a head node and cluster nodes form a core on-premises HPC environment. Additional compute nodes located in the cloud accommodate peak demand.

If you already have an HPC system and you are looking for a cost-effective way to upgrade without starting over, bursting is a powerful option. On the other hand, even if you are building a new HPC system, bursting enables you to save money while still maintaining a local HPC environment. You can design your local HPC presence for less than your expected maximum load and scale up into the cloud as needed.

Of course, the simple scenario depicted in Figure 2 is only part of the picture. Figure 3 shows a more detailed view. A bursting configuration requires a safe and suitable internet access method for integrating cloud nodes with the local infrastructure. Network access is typically achieved using a VPN gateway or other high-bandwidth network connection.

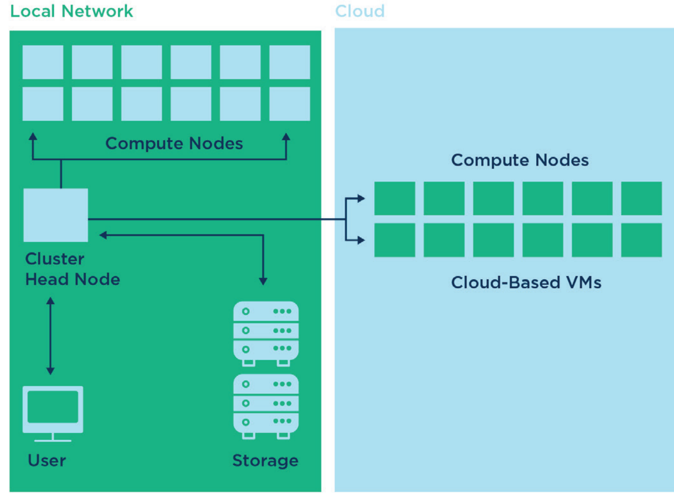

Figure 2: *Cloud bursting: A job runs on the local cluster and then bursts to the cloud when additional capacity is needed*

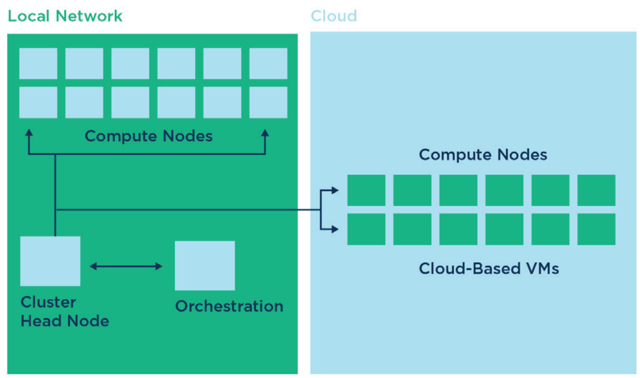

Figure 3: *A closer look at cloud bursting. An orchestration layer deploys VMs in the cloud. The cloud is linked to the local network through a high-bandwidth VPN or other low-latency connection.*

Bursting requires an orchestration layer for managing and provisioning compute nodes in the cloud. For instance, the OpenStack orchestration system is a popular alternative for initiating and managing the bursting process. OpenStack is a leading open source cloud management platform that offers cost-effective support for a public or private cloud environment. Other proprietary tools can also provide orchestration services in a bursting scenario.

Building and deploying the orchestration layer can require additional overhead and configuration time. However, a well-designed orchestration layer also provides efficient automation and helps you optimize cloud usage.
Depending on the cloud vendor, you can base your configuration on a number of bursting scenarios, including:

- **On-demand:** Pay for compute resources on a per-hour or per-second basis. This option is good for occasional spikes in workload.
- **Reserved:** Reserve a block of time to run a specific, pre-scheduled job. Some cloud vendors offer a lower price for pre-scheduled jobs, which allows them to plan ahead and reserve the resources in advance. This type of billing is best for big projects of known size that will run within a well-defined time window.
- **Spot:** The user bids for unused capacity on the cloud network. This option often offers minimum pricing for jobs that can run at off-peak hours..

As demand changes and your HPC configuration evolves, the cloud provides the flexibility you will need to adapt and optimize.

All-Cloud or Bursting?

Should you run all your HPC workloads in the cloud or build your HPC presence around a local cluster with additional cycles through cloud bursting? Finding the right solution requires you to consider several factors:

- **Storage:** Data movement can be a major bottleneck for HPC operations. Minimizing data movement often means that, if the data is stored in the cloud, the workload should run in the cloud. And if the data resides on the local network, your best option is to provide your primary HPC resources locally and use the cloud for bursting in times of peak load. Financial risk assessments that consume a lot of market data are I/O intensive and could require ultra-low latency. I/O-bound workloads like these might not be as amenable to cloud environments.
- **Network:** Network-bound workloads require a lot of bandwidth. Bursting is less efficient with slower internet connections and works better for compute-intensive workloads that minimize network traffic. If extensive node-to-node communication is required, as in deep learning applications, going “all-in” to the cloud might not be ideal.

SUSE Linux Enterprise High Performance Computing

A highly scalable, high performance open source operating system designed to utilize the power of parallel computing for modeling, simulation and advanced analytics applications. SUSE Linux Enterprise HPC provides a parallel computing platform for high performance data analytics workloads such as artificial intelligence and machine learning. Fueled by the need for more compute power and scale, businesses around the world today are recognizing that a high-performance computing infrastructure is vital to supporting the analytics applications of tomorrow. This product is built for easy adoption across a cloud environment and lower-cost Arm-based hardware all the way up to the largest supercomputers in the world – with a focus on flexibility and providing SUSE-supported capabilities for today’s HPC environments.

- **Compute:** If node-to-node communications and data movement are less of an issue, and the applications are more compute intensive, there could be a strong case for running those workloads in the cloud. Compute-intensive applications such as genome sequencing could benefit from an extreme-scale environment provided by the cloud.
- **Throughput:** Certain workloads can benefit from an extreme-scale environment that clouds can provide. If there is data parallelism, meaning that there is shared data across concurrently running jobs, without a need for communications between the jobs, a scalable environment such as the cloud could be ideal.
- **Cost:** Building and maintaining a local hardware infrastructure requires overhead, which becomes part of the total cost of using the system, but the contribution of these start-up costs to the average per-cycle expense diminishes as usage increases. On the other hand, the all-cloud option might have a higher per-cycle operational cost that is relatively constant with increased usage. A local cluster with bursting is, therefore, more efficient for systems that are heavily utilized. In addition, all-cloud is a better option for organizations that expect lighter usage.

As a general rule of thumb, it is more cost-effective to establish an HPC environment in the cloud if your HPC workloads meet the following criteria:

- **Not node-to-node communications-intensive**
- **Not I/O-intensive**
- **Not bound to on-premises data movement**
- **Uses shared data across many concurrent jobs**

In some cases, small HPC clusters can benefit from running in the cloud, regardless of the workload or data profile.

SUSE for HPC

SUSE, the world's largest independent open source company, specializes in agile, enterprise-grade open source solutions. Collaborating with partners, communities and customers, SUSE delivers and supports enterprise-grade Linux, software-defined infrastructure and application delivery solutions to create, deploy and manage workloads on-premises or in the cloud.

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

SUSE Linux Enterprise High Performance Computing delivers more than the operating system. It includes a bundle of popular HPC tools, such as Slurm for workload management and scheduling, Ganglia for workload performance monitoring and much more. In addition, all HPC tools are supported by SUSE. Full system images of SUSE Linux Enterprise High Performance Computing are available in the Microsoft Azure Marketplace and the AWS Marketplace, for those who wish to establish a strong HPC-optimized compute platform in the public cloud.

SUSE offers a full range of products and services to support your IT landscape—and to integrate your HPC resources with the rest of your infrastructure. For instance, SUSE Enterprise Storage is a Ceph-based virtual storage environment for cost-efficient data management. In addition, SUSE Manager offers secure and efficient network management with configuration and auditing.

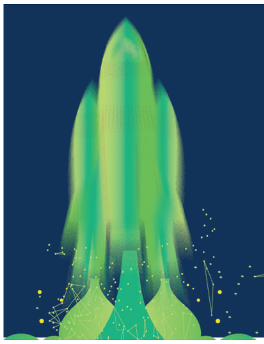

Dr. Herbert Huber, Department Head of High Performance Systems at the Leibniz Supercomputing Centre (Leibniz-Rechenzentrum, or LRZ), explained, "We have used SUSE® solutions at LRZ for over two decades now. For us, the overriding benefit of their solutions is the compatibility they offer. We first implemented SUSE Linux Enterprise Server on our very first general-purpose Linux cluster system in the late 1990s because we deemed it to be one of the best operating systems for performing standard HPC workflows: a quality which SUSE has reliably upheld over the years."

LRZ: Leibniz Supercomputing Centre

The Leibniz Supercomputing Centre (Leibniz-Rechenzentrum, or LRZ) is located on the campus area of the city of Garching, near Munich, Germany. LRZ is the IT service provider for all universities in the Munich area, as well as a growing number of research organizations throughout Bavaria. The LRZ is a world leader in its field, providing HPC and data center resources for scientific research across the region of Bavaria. On top of this, LRZ is one of the three key players in the Gauss Centre for Supercomputing (GCS)—Germany's foremost supercomputing institution—which is committed to creating a consolidated HPC infrastructure that can be applied to a broad range of scientific and industrial research projects.

LRZ is utilizing HPC in the cloud in a number of ways, from integrated container development, remote visualization of data and simulations.

Dr. Herbert Huber, Department Head of High Performance Systems at LRZ, explained "On the supercomputing level, we first deployed SUSE Linux Enterprise High Performance Computing in 2006, due to its support for very large shared memory nodes, as well as its seamless interoperation with the HPC software stack and many commercial application software packages. The seamless compatibility of SUSE Linux Enterprise Server with our HPC software stack was and is a key factor in our decision to employ the SUSE operating system on our HPC systems."

"Our future plans involve even greater collaboration with GCS and further development of cloud computing resources to aid with storage and processing of scientific data," said Dr. Huber. "We also have plans to expand our high performance hardware landscape once we secure the budget to do so. We are even in the concept stages of building innovative new hardware solutions, including quantum computing and exascale systems, which we hope to implement within the next decade.

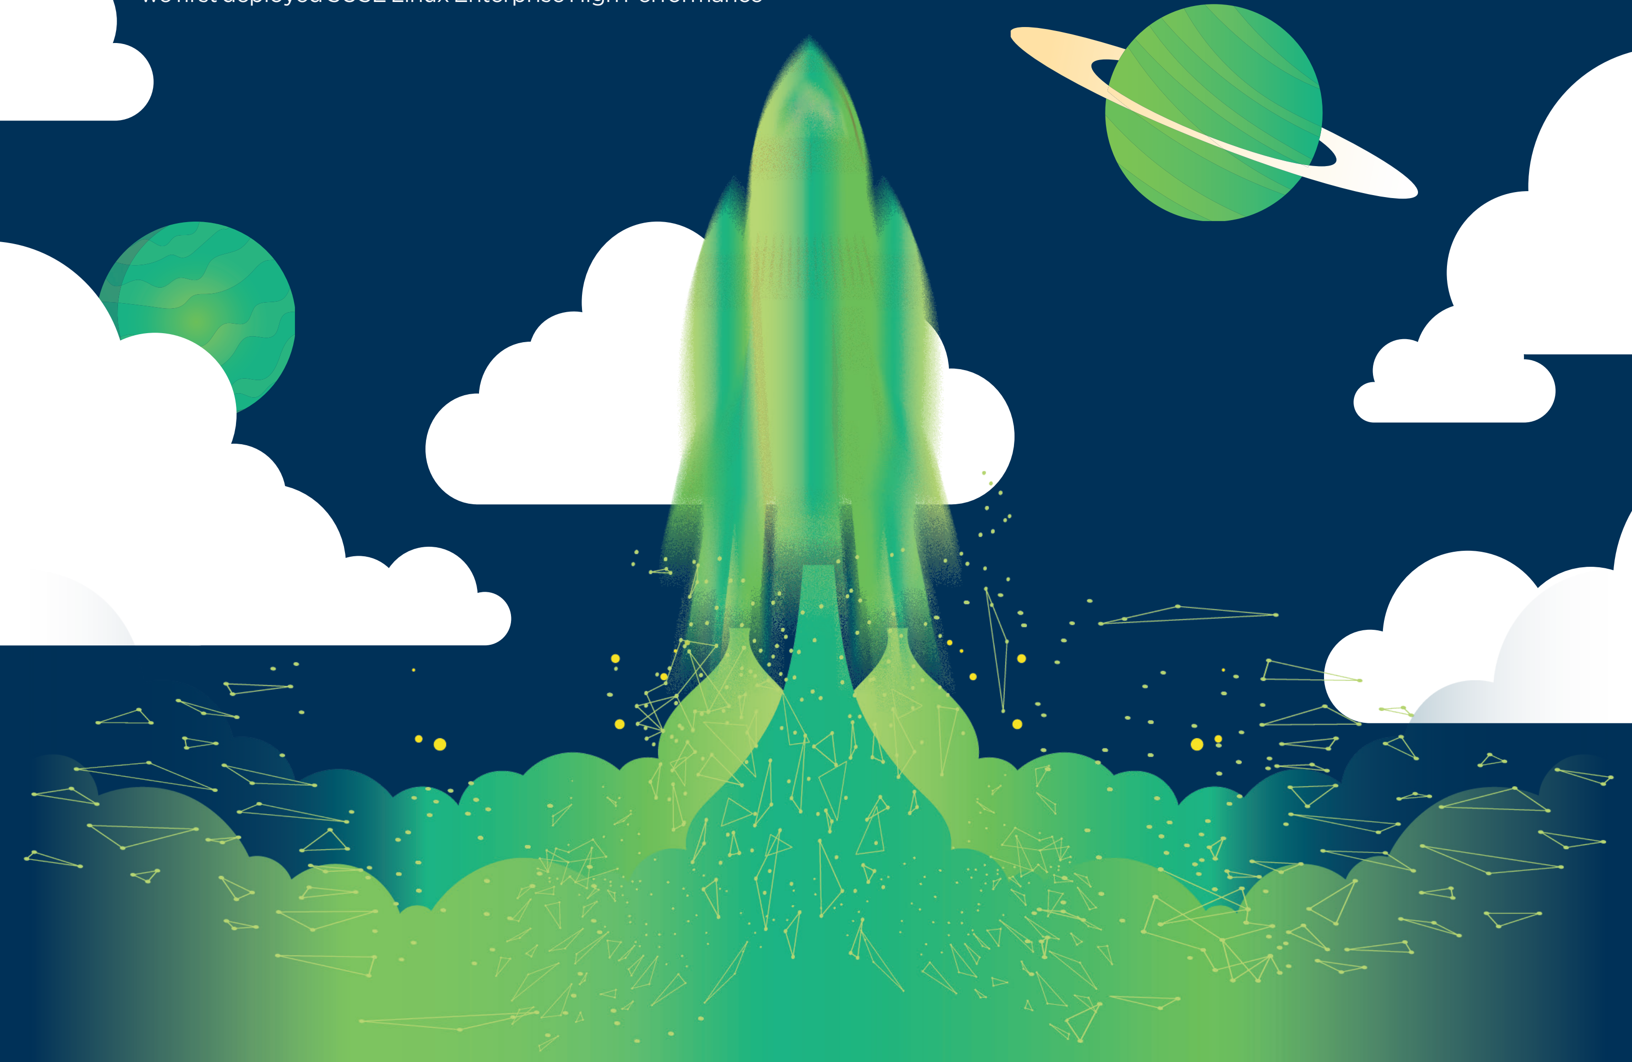

For more information, contact your local SUSE Solutions Provider, visit us online or call SUSE at:

1-800-796-3700 (U.S. and Canada)
1-801-861-4500 (Worldwide)

SUSE
Maxfeldstrasse 5
90409 Nuremberg
Germany

suse.com