**High Performance Computing (HPC)**

# High Performance Computing (HPC)

White Paper

People, Process, and Technology | **One Firm**


**SBA CERTIFIED | DUNS: 83074292 | CAGE: 5H6B4 | TIN: 270175433**

**Fusion PPT**
8245 Boone Boulevard
Suite 200
Vienna, VA 22182
www.fusionppt.com

## 1.0 OVERVIEW

When heterogeneous enterprise environments are involved, there are complex internal and external factors to consider when seeking high performance computing (HPC) options. Decisions are further complicated as the industry continues to innovate while your plans are still being made. Navigating the market and analyzing possible solutions is a time-consuming process that can take your organization's focus off of the core mission. In addition, your organization may not have the latest experience and background that is required in order to be successful in examining and comparing the latest features and product capabilities.

**Fusion PPT Can Pinpoint Your Best HPC Approach.**

As an established leader in system integration services, Fusion PPT's team of HPC industry experts can leverage their experience with HPC environments to help your organization make the right choices. In the Fusion PPT Innovation Lab, we evaluate these environments and test performance against marketed features and objectives. Our HPC industry experts have extensive hands-on experience working with the latest vendor products. We understand how to evaluate a product, validate requirements and ensure functionality in your environment. We also have the opportunity to formally evaluate many products prior to their general availability within our Innovation Lab. Fusion PPT industry experts provide a level of insight into market trends, and gaps in product features and capabilities. Our Innovation Lab also provides us with the opportunity to solve specific enterprise problems using a best-in-class approach to tools and products.

**Focus on Your Mission. We'll Find Your Best HPC Solution.**

At Fusion PPT, our vendor independent expertise provide leveraging both new and established HPC technologies to solve unique business and operational challenges for our clients. With our assistance, your team will spend less time researching, analyzing, deploying, and managing complex technologies, and more time on serving the core needs of your organization.

HPC origins began in the academic arena for science and engineering applications as a means of migrating off of high-end supercomputer models of computation (cf. vendors such as Sequent and Cray). Over the years, it has also been applied to business and finance, military, and Internet applications that have massive requirements for memory, processing, and I/O.

HPC involves aggregating computing power in ways that provide superior performance in comparison to what can be achieved by the use of traditional desktop or server-based computing. HPC environments arrange computers into clusters or grids; Grid computing and HPC are terms that are often used interchangeably.

HPC clusters are comprised of multicore nodes (symmetric multiprocessors) and are networked via very high speed interconnects, using technologies such as InfiniBand. The Message Passing Interface (MPI) specification is typically used as the protocol to manage the parallel

communication between nodes. The resultant clusters provide an aggregate computational, memory, and networking bandwidth that greatly exceeds that of traditional standalone servers by orders of magnitude. The HPC cluster can solve problems and deliver solutions that no single server could solve due to resource limitations.

Applications vary widely in terms of the specific aspect of HPC that they primarily utilize (e.g. compute, storage, interconnect). For example, some applications are highly-parallel, but do not require high levels of network interconnectivity. Others will require high interconnect speeds for low latency and high throughput connectivity. Still others have high I/O requirements that see storage as the bottleneck. And at the top tier are applications that need to capitalize on all aspects - compute, storage, and interconnect.

## 2.0 LEADING VENDORS

In the traditional on premise server-based HPC market, the big three vendors with the largest HPC market share are Hewlett-Packard, IBM, and Dell, respectively.

As for cloud-based HPC providers, a number of large and some smaller vendors have specialized cloud computing designated just for HPC:

- IBM
- Amazon
- Penguin Computing
- R-HPC
- Sabalcore
- Gompute

In addition, Red Hat provides the Messaging Realtime Grid (MRG) that supports on-premise and cloud-base solutions

Red Hat MRG is a HPC clustering infrastructure based on a customized Linux distribution. Linux kernel has been optimized for parallel computing, high-speed messaging and cluster scheduling. The MRG infrastructure also supports the HTCondor software framework supporting high-throughput computing jobs that can run for months and years.

The MRG grid supports the management of physical servers or virtual machines, up to 10,000 nodes. Interconnect options range from Ethernet (1 GigE and 10GigE) up to Infiniband (20 Gigb). MRG is primarily aimed at the departmental level cluster and includes a cluster management toolkit to simply the management of the HPC solution.

## 3.0 PROS & CONS

There are a number of factors can influence the decision of using on premise or cloud-based HPC solutions.

### *3.1 Cloud Perks*

Cloud-based HPC solutions enable customers to take advantage of the many benefits that cloud computing has to offer such as rapid scalability, virtually unlimited storage, resiliency and fault tolerance, and an on-demand metered pricing model.

In fact, cloud solutions now rival some of the established pre-existing HPC solutions. For example, the Amazon HPC solution supports instance clustering and cluster networking. The C3 instances (recently superseded by the C4 instance) falls in the TOP500 ranking of most powerful computer systems in the world.

### *3.2 Cloud Bursting*

Cloud computing also offers another inherent capability which makes it a viable option for HPC workloads – bursting. Cloud bursting can enable use cases where internal resources are overused or in cases where testing and small experiments are needed.

### *3.3 Management*

Some of the challenges in the cloud-based virtualized HPC environment involve management. Due to the sheer volume of computing resources involved in HPC solutions, the management of VMs could present challenges – for example VM "sprawl." Implications span capacity planning and network congestion as well. Proper tools and management will become essential. Licensing models must also be taken into account, again in dealing with the large numbers of VMs potentially involved.

### *3.4 Data Transfer*

Another consideration that has cost impact for cloud HPC, but not as relevant for on-premise HPC, involves data transfer. Cloud costs can increase when data transfers are considered, particularly transfers outbound from the cloud. When dealing with HPC loads, large amounts of data can be involved so users should be wary and plan accordingly based on data transfer estimates.

### *3.5 Consideration for Applications*

The large computational capacity in HPC environments is of minimal benefit without the applications that can fully harness the capabilities designed into the HPC environments. With the maturity of on-premise HPC over the years, many of the applications have been tuned to natively take advantage of capabilities within the on-premise hardware. This might even include applications that bypass the OS kernel.

Many applications in the HPC domain are highly numerical, for example the use of numerical simulations and modeling, data analytics, and computational fluid dynamics. As the HPC environment has largely used bare metal to date, the applications have been developed and coded to capitalize on the computational, memory, and network characteristics of physical servers.

In addition, parallel software is not necessarily easy to port between different platforms. The coding that capitalizes on one inherent capability of Platform A may lose that capability on Platform B.

### *3.6 HPC Storage*

The open source community has been a big champion in providing solutions to support the demanding storage needs of HPC. HPC storage options are now available to support on-premise and Cloud-based solutions. For example, a number of open source fault tolerant, high performance, scalable distributed file systems are available to enable HPC storage options:

- Moose File System (MooseFS)
- Hadoop Distributed File System (HDFS)
- Lustre
- CephFS
- GlusterFS

In particular, the Lustre file system (a blend of the words Linux and cluster) provides support for distributed file systems, scale out storage, and cluster computing, ranging from the workgroup cluster to the multi-site cluster. Due to its capabilities and the GNU General Public License open licensing, Lustre is used by a majority of the top 100 computers in the TOP500 list of supercomputers in the world (including the #2 Cray Titan and the #3 IBM Sequoia). In fact, Lustre supports tens of thousands of nodes, each potentially containing tens of petabytes (PB) of storage and very high throughput exceeding 1 terabyte per second (TB/s) of aggregate throughput.

Ceph is another distributed and highly-available platform scalable to the exabyte level, providing object storage, thin-provisioned block storage, and a file system (CephFS) that sits atop the object or block store. The Ceph RESTful API interface also enables the block and object stores to be supported in cloud-based environments.

In addition, the Gluster file system (GlusterFS) supports HPC by providing a petabyte-scalable parallel network-attached file system. The GlusterFS consists of storage servers aggregated via Infiniband or Ethernet.

While Gluster, Lustre, and Ceph remain publicly available via General Public Licenses, a number of acquisitions have moved some of the lead software development to Red Hat and Intel. Gluster was acquired by Red Hat in 2011 and now incorporates the technology into RHEL. Red Hat now offers the Red Hat Storage Server (based on Gluster) that supports private, public, and hybrid cloud options. Intel purchased Whamcloud (the purveyor of Lustre) in 2012, continues to develop the software, and provides a cloud-based software edition. Most of the Ceph development is now done internally at Red Hat after the acquisition of Inktank Storage (developer of Ceph) in 2014.

## 4.0 SUMMARY

While sales of high-end HPC equipment have shown decline in recent years, market-research firm IDC expects an expansion over the coming years. It's yet to surface if the cloud-based HPC market will affect these projections. A battle may lie ahead in which we see the entrenched on-premise

HPC environment with its customized applications vie against a cloud-based model promoting many benefits over bare metal such as flexibility, availability, scalability, and cost.