INTEL HPC ORCHESTRATOR: SOFTWARE FOR A NEW ERA

Addison Snell

White paper
August 2016

EXECUTIVE SUMMARY

High performance computing (HPC) is inextricably linked to the relentless pursuit of advancements in scientific discovery, and until some far-off day when we declare we have solved all there is to know about science, the industry will continue to demand improvements in the performance and scalability of HPC. This is true not only in public sector research but also in business, as commercial markets continue to be the majority users of HPC resources and the driving engine behind the industry's growth in upcoming years.

This continual push to do more with HPC has led not only to growth in investment, but also to occasional wholesale changes in approach. Changes in processor architectures are changing the scalability model in HPC. Even within the standard microprocessor paradigm, the move to multiple cores on a chip has introduced a new tier of parallelism at the socket level. In addition, many-core architecture options have proliferated, including options such as Intel Xeon Phi. Intel Xeon x86 processors are the dominant microprocessor in HPC today, but users are facing challenges inherent in returning to a world of specialization. The state of the market in HPC middleware is markedly diverse, and new workloads exacerbate the challenge. Big data and analytics have already introduced new environments for data-intensive workloads, and research in machine learning is making its way toward mainstream commercial HPC users, such as those in finance and manufacturing.

Against this backdrop of complicating HPC dynamics, Intel is attempting to help ease HPC users into the future. In late 2015 the company participated in introduction of the OpenHPC community, which has become a "shared repository" for HPC middleware. With OpenHPC as a basis, Intel is now taking the next step in creating scalable software environments for HPC users. Intel HPC Orchestrator is a modular stack of Intel-supported software components based on OpenHPC, which can be tailored to particular HPC environment preferences. This stack becomes the software underpinning of Intel's Scalable System Framework.

Intel HPC Orchestrator extends beyond OpenHPC in three critical ways:

- *Advanced Testing*—Intel does additional testing and validation of the integrated OpenHPC system software components before bundling them into Intel HPC Orchestrator configurations, ensuring cross-compatibility and reliability.
- *Premium Software*—Intel incorporates additional licensed software components that are not necessarily freely available as part of OpenHPC, such as the Intel Parallel Studio XE Cluster Edition suite of compilers and tools.
- *Intel Support*—Intel provides Level 3 support for the integrated components included in Intel HPC Orchestrator, giving end users a single focal point in managing their middleware environments, as provided through system vendors.

By providing a common, supported software environment with Intel HPC Orchestrator, Intel aims to protect end users' HPC investments, while still allowing them to chase new levels of scalability and performance. ISVs can spend more of their time on new features and functionality, and administrators can more easily optimize their systems' performance for their specific workloads. In the march toward exascale computing, Intel HPC Orchestrator will play a critical role in connecting new levels of scalability in hardware to the true advancements that are the bellwethers of HPC.

MARKET DYNAMICS

A New Era in HPC

Throughout its history, the High Performance Computing (HPC) industry has been marked by occasional changes in the constant pursuit of one consistent goal: ever-increasing computational performance. This relentless drive comes from HPC's intrinsic link to advancements in science, engineering, and—to an increasing extent—business. The pursuit of scientific advancement and its role in our lives does not diminish: there will always be new cures to discover, new products to design, new energy to harness. And where HPC is applied, it can always be deployed with greater fidelity; models can be refined for more realism, more degrees of freedom, more predictive ability.

This is true not only in public sector scientific research, but in business as well. Commercial organizations make up 55% of HPC usage by revenue worldwide and represent higher forecasted growth than government or academia over the next five years.[1] In a study conducted by Intersect360 Research for the U.S. Council on Competitiveness, 86% of respondents—all representatives of U.S.-based, commercial HPC users—agreed with the statement, "HPC is critical to the future direction of our business." 68% of respondents said their most challenging applications could consume 10x more performance over the next five years; 37% said their most challenging applications could consume 1000x more.[2]

Projected Need for Additional HPC Performance over Next 5 Years
Survey of 100 U.S. Commercial HPC Users
Source: U.S. Council on Competitiveness, Intersect360 Research, 2014

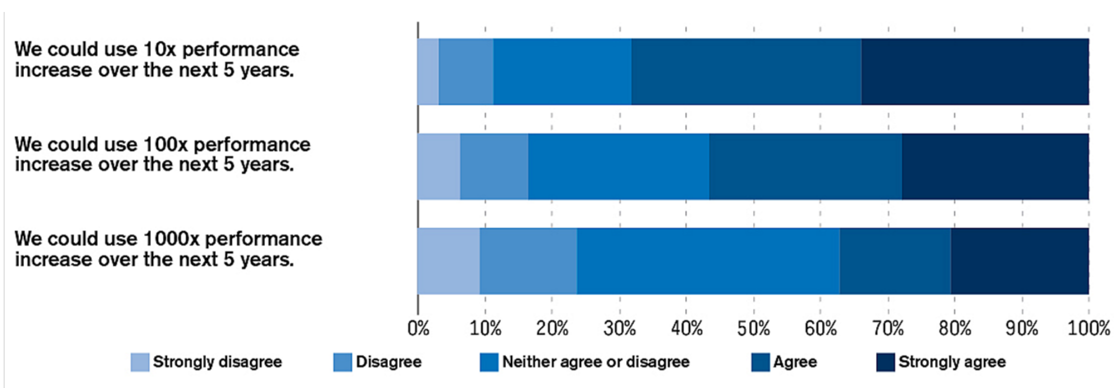

This continual push to do more with HPC has led not only to growth in investment, but also to occasional wholesale changes in approach. In the 1970s and '80s, most supercomputers were based on vector processors. The late 1980s brought a revolution, in which HPC systems could be built on scalar RISC processors in shared-memory, symmetric multiprocessor architectures running UNIX. This became the dominant paradigm until "Beowulf" clusters burst onto the scene in the second half of the 1990s. These systems ganged together the superior price/performance of x86 processors, mostly from Intel, into distributed memory systems running open-source Linux. A primary point of Beowulf clusters' appeal was the portability that came

1 Intersect360 Research, "HPC 2015 Total Market Model and 2016-2020 Forecast," June 2016.
2 U.S. Council on Competitiveness, "Solve. The Exascale Effect: The Benefits of Supercomputing Investment for U.S. Industry," October 2014, http://www.compete.org/reports/all/2695.

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

with standardization. Users who migrated their applications into a message passing interface (MPI) programming model could then run those applications on any cluster, regardless of where the components came from. Over time, these clusters would see upgrades, such as faster interconnects, blade form factors, and 64-bit processing architecture, but the Beowulf ideal of portability remained.

Today clusters are still the dominant HPC system architecture, but changes in processor architectures are starting to change the scalability model again. Even within the standard microprocessor paradigm, the move to multiple cores on a chip has introduced a new tier of parallelism at the socket level. In addition, many-core architecture options have proliferated, including options such as Intel Xeon Phi—available either as a standalone processor or as co-processor, using x86 programming tools—as well as GPUs, which can be attached to server processors over the PCIe interface. FPGAs (field programmable gate arrays) can be customized to particular applications. Some users have experimented with DSP (digital signal processing) technologies, or with low-power options such as ARM or Intel Atom. And some RISC processors, such as IBM Power, are still marketed for HPC workloads.

To be sure, Intel Xeon x86 processors are still far-and-away the dominant microprocessor in HPC today, but with this pending diversity in processing architectures, users are facing some challenges inherent in returning to a world of specialization. 88% of HPC users predict they will deploy multiple processor architectures, assigning applications to the types of processing elements that best suit them. Intel Xeon Phi plays a major part in this. 24% of HPC users report already having accelerated x86 elements in place for some applications, with an additional 48% saying they have plans to test or evaluate them.[3]

Why a Stable HPC Software Platform Matters

Against this diversity in hardware, end users must balance how much specialization they undertake. Putting too much effort into one arena could lead to a loss in competitive advantage if it proves to be the wrong path. HPC users are grappling with this dynamic as they explore architectural differentiation. In their evaluation of processing architectures, 87% of HPC users say that “availability of a compatible software stack (OS, support libraries, compilers, etc.)” is “very important” or “extremely important,” while 75% gave that same rating to “compatibility with existing applications, without modification.”[4]

Unfortunately, the state of the market in HPC middleware is markedly diverse. In the Beowulf cluster era, standardization in hardware meant that end users would seek differentiation in software. Of the over 300 middleware packages identified by HPC users in an annual survey, 62% are mentioned only once, and the top 20 middleware packages combined account for less than half of all total mentions.[5]

New markets and workloads exacerbate this dynamic. Big data and analytics have already introduced new environments for data-intensive workloads. And since beginning to track “ultrascale internet” as a distinct HPC segment in 2007, Intersect360 Research has monitored the growth and evolution of what is now known as the hyperscale market, now tracked as a separate market from HPC, in which the top tier of buyers all spend in excess of $1 billion per year in IT and have dramatic influence in application development and software adoption across the enterprise IT landscape.[6] Research in machine learning is fueled by hyperscale organizations like Google, Facebook, Amazon, and Baidu, and today mainstream commercial HPC users, such as those in finance and manufacturing (the largest commercial HPC vertical markets), are seeking to adopt machine learning applications.

---

3 Intersect360 Research special study, “Processor Architectures in HPC,” 2016.
4 “Processor Architectures in HPC.”
5 Intersect360 Research, “HPC User Site Census: Middleware,” 2015.
6 Intersect360 Research, “The Hyperscale Market: Definition, Scope, and Market Dynamics,” April 2016.

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

If these weren't enough challenges for HPC environments, increasing numbers of HPC users are moving parts of their technical workloads off-premise. Although the total proportion of spending on cloud computing in the HPC market remains low—2.4% of total product and services spending in 2015[7]—more than one-fourth of HPC users utilize cloud for at least some portion of their workflow[8], and cloud is expected to have the highest spending growth rate of any product segment over the next five years.[9]

As users seeks to optimize their increasingly diverse workloads across a hardware environment that is specializing to embrace new architectures and off-premise resources, it will be beneficial for many to create stability and centralization in the software platform. Creating a central repository for software that is certified to run across a panoply of resources in an integrated fashion will assist many users in continuing to pursue an optimized environment without a loss in productivity or a prohibitive increase in system administration costs.

INTEL HPC ORCHESTRATOR

Against this backdrop of complicating HPC dynamics, Intel is attempting to help ease HPC users into the future by shepherding what it calls a "shared repository" for HPC middleware and by incorporating supported elements into a family of products for HPC users. The first part of this endeavor is accomplished through OpenHPC, a community of partner OEMs, independent software vendors (ISVs), and public sector HPC users tasked with the stewardship of open-source and commercial software in the new era of HPC. The second part integrates selected, Intel-supported components from OpenHPC with advanced features and testing to create Intel HPC Orchestrator, a trusted, high-performance middleware layer that forms the software foundation of Intel's Scalable System Framework strategy.

OpenHPC

OpenHPC was introduced in late 2015 as an endeavor of the Linux Foundation, the titular arbiters of open-source advancement, with dozens of prominent inaugural members from across the HPC community. The goal of OpenHPC is to create a centralized crèche of open-source software development, incubating amongst qualified "reliable, relevant versions" of Linux distributions, developer tools, file systems, and resource managers.

OpenHPC succeeded immediately in attracting the dominant OEM vendors of HPC systems (including HPE, Dell, Lenovo, Cray, Atos, SGI, Fujistu, Sugon, Penguin, Inspur, and NEC) as well as notable commercial software vendors (Allinea, Altair, Ansys, Dassault Systèmes – Simulia, MSC Software, and Nimbix). Just as importantly, OpenHPC has participation from many leading supercomputing sites worldwide. This matters not only in giving OpenHPC credibility and critical, highly scalable testbeds, but also because these academic and government research labs are often the creators and maintainers of critical open-source components.

The net result of OpenHPC is to normalize an environment for HPC software. The effect is seen in three areas: in programming, where common compilers, developer tools, and distributions make it easier to create scalable applications; in deployment, where pre-qualifications of supported components speed up installations, with fewer conflicts; and in optimization, where users can now find greater synergies between components in seeking greater performance at scale.

Intel HPC Orchestrator

With OpenHPC as a basis, Intel is now taking the next step in creating scalable software environments for HPC users. Intel HPC Orchestrator is an Intel-supported modular stack of software components based on OpenHPC,

7 "HPC 2015 Total Market Model and 2016-2020 Forecast."
8 Intersect360 Research, "HPC User Budget Map: Industry Averages," 2015.
9 "HPC 2015 Total Market Model and 2016-2020 Forecast."

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

which can be tailored to particular HPC environment preferences. This stack becomes the software underpinning of Intel's Scalable System Framework.

Intel HPC Orchestrator extends beyond OpenHPC in three critical ways:

- *Advanced Testing*—Intel does additional testing and validation of the integrated OpenHPC system software components before bundling them into Intel HPC Orchestrator configurations, ensuring cross-compatibility and reliability.
- *Premium Software*—Intel incorporates additional licensed software components that are not necessarily freely available as part of OpenHPC, such as the Intel Parallel Studio XE Cluster Edition suite of compilers and tools.
- *Intel Support*—Intel provides Level 3 support for the integrated components included in Intel HPC Orchestrator, giving end users a single focal point in managing their middleware environments, as provided through system vendors.

Fully laid out, Intel HPC Orchestrator represents the complete operating, programming, and runtime environment, inclusive of the OS and all middleware components, such as provisioning, workload management, compilers, libraries, and utilities. ISV applications that are not sold by Intel can be qualified to fit into the environment.

Intel HPC Orchestrator Component Stack
Source: Intel

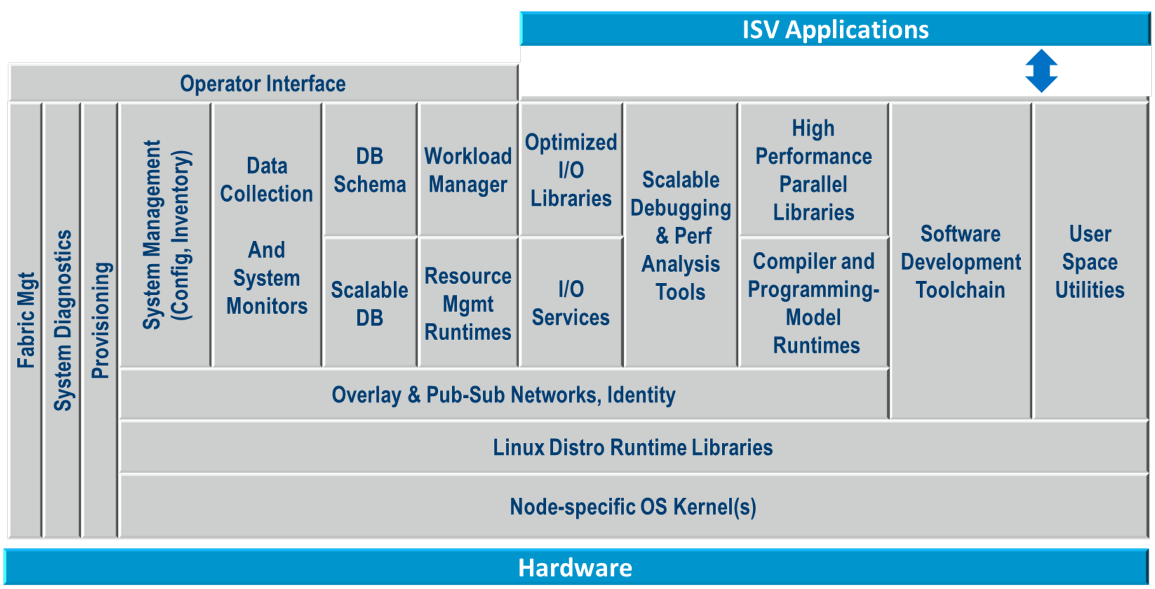

Over time, Intel HPC Orchestrator will extend into a family of products focusing on different constituents within the HPC market. This initial deployment, Intel HPC Orchestrator – Advanced, targets the midrange through high end of the HPC spectrum—users whose capabilities might place them on the TOP500 list of supercomputing sites, but not in the top 50. In the future Intel plans to roll out different iterations that will more closely map to the needs of the ISV community and the elite-scale supercomputing sites.

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

Some of the software elements included in this first instantiation of Intel HPC Orchestrator are listed below. Intel tests for dependencies between components, including building any third-party libraries for each combination of compiler and MPI library, resulting in hundreds of modules in the repository.

• Base operating system:	RHEL/CentOS 7.2
	SLES 12
• Administrative tools:	Conman
	Ganglia
	Intel Cluster Checker
	Lmod
	LosF
	ORCM
	Nagios
	pdsh
	prun
• Provisioning	Warewulf
• Resource management	SLURM
	Munge
	PBSPro (planned addition)
• I/O services	Lustre client
• Numerical / scientific libraries	Boost
	GSL
	FFTW
	Metis
	PETSc
	Trilinos
	Hypre
	SuperLU
	Mumps
	Intel MKL
• I/O libraries	HDF5 (pHDF5)
	NetCDF (including C++ and Fortran interfaces)
	Adios
• Compiler families	GNU (gcc, g++, gfortran)
	Intel Parallel Studio (icc, icpc, ifort)
• MPI families	OpenMPI
	MVAPICH2
	Intel MPI
• Development tools	Autotools (autoconf, automake, libtool)
	Valgrind
	R
	SciPy/NumPy
• Performance tools	PAPI
	Intel IMB
	mpiP
	pdtoolkit TAU
	Intel Advisor
	Intel Trace Analyzer and Collector
	Intel Vtune Amplifier

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

HPC Community Benefits

Intel foresees multiple benefits of Intel HPC Orchestrator across the HPC community. For their OEM partners, Intel envisions it as a way to reduce the time it takes to release a new server product, alleviating the burden of each OEM qualifying its own middleware stack and moving differentiation to more value-added areas. For ISVs, Intel HPC Orchestrator reduces the myriad combinations of base configurations for them to qualify against, moving resources from testing and bug fixing to adding new features or greater scalability. And for administrators, it removes the headache of rooting out dependency issues between low-level components, freeing them to spend more time on performance optimization.

But for any initiative of this breadth, the true value must include the benefits to the end user, and here Intel HPC Orchestrator has value as well. Processing architectures are diversifying, and software is inextricably linked to this dynamic. By providing a common, qualified software layer in its Scalable System Framework, Intel gives users of its systems a stable foundation upon which to layer the applications that connect them to new levels of discovery and innovation.

INTERSECT360 RESEARCH ANALYSIS

As HPC continues to evolve, the primary question is not just theoretical performance, but of paramount importance is what can actually be done with the computer. For HPC to continue to be a tool for enabling advancement and scientific discovery, the industry must focus on how to bring together disparate elements into a truly scalable application framework.

As the industry moves forward toward exascale—and a higher volume of end users have access to petascale capabilities, either on premise or in the cloud—the biggest challenge is not in hardware but in software. At every stage of industry evolution, new innovations must be undertaken in software environments in order to harness new forms of parallelism, over the network and within the node.

This is where Intel's investment in OpenHPC and Intel HPC Orchestrator shines. Without the establishment of standards in software, the developer community spends a high proportion of its efforts maintaining support across an increasingly broad proliferation of operating environments. It is becoming too burdensome for each individual system provider, programmer, and administrator to test for dependencies, cross-certify components, and optimize clustered environments.

By providing a common, supported software environment with Intel HPC Orchestrator, Intel aims to protect end users' HPC investments, while still allowing them to chase new levels of scalability and performance. ISVs can spend more of their time on new features and functionality, and administrators can more easily optimize their systems' performance for their specific workloads.

Ultimately any new advancement in HPC is meant to support advancements in science, engineering, and business. When a new era introduces changes in hardware and parallelism, adjustments need to be made in programming and administration, as was the case with the introduction of clusters and the industry's migration to MPI and new middleware tools focused on resource and job management. Now the industry is changing again, and Intel is in the middle of the revolution. Intel HPC Orchestrator will play a critical role in connecting new levels of scalability in hardware to the true advancements that are the bellwethers of HPC.

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com