



Dell Expands HPC Access and Adoption with Intel Cluster Ready Program

Addison Snell

White paper sponsored by: Dell

OVERVIEW

High productivity computing (HPC) is being adopted at increasing rates across a variety of organizations with users ranging from experienced veterans to HPC novices. Growth is coming both in traditional HPC markets and in non-traditional “Edge HPC” applications as more organizations are realizing the productivity benefits of HPC usage and adoption. The tremendous performance and continuing standardization of HPC components have made HPC more accessible and more viable for a growing range of workloads.

However, many organizations that could benefit from the use of HPC clusters find that it is complicated to get the systems up and running because of limited IT resources or the complexities of the clusters themselves. The disaggregation of the HPC stack has led to challenges for system administrators in the deployment and maintenance of HPC clusters because many middleware and tool components are no longer standard, integrated parts of the server or operating system. End users might install software and middleware components from different sources at various stages of release schedules without these components being tested as part of an integrated solution. The resulting incompatibilities can be difficult to identify, isolate, and repair, especially for smaller or newer HPC users.

The Intel Cluster Ready program, for which Dell was an original partner, seeks to address this challenge by simplifying the purchase, deployment, and management of HPC clusters for entry-level and mid-range HPC users. The Intel Cluster Ready program is based upon the Intel Cluster Ready Specification, a common basis for building clusters and registering applications for HPC. The specification, which includes requirements for hardware, software, manageability, and functionality, helps to ensure that each cluster component conforms to industry standards or, if no standard exists, best-in-class practices. Intel provides the specification to hardware vendors, software vendors, and systems integrators, so they can build certified Intel Cluster Ready clusters and develop registered Intel Cluster Ready applications. The key to these certifications is the Intel Cluster Checker software management tool, which analyzes configurations to identify problems and provide detailed diagnostic information.

The Intel and Dell partnership has several benefits for the companies and for users:

- Users can buy certified HPC clusters from Dell that result in faster implementation, simpler management, better system health, and improved overall productivity.
- Dell addresses growing segments of the HPC user community with an integrated, total solution for HPC clusters.
- Intel, Dell, and other partners help the industry re-consolidate software stacks with long-term productivity benefits for HPC users.

Dell HPC solutions with Intel Cluster Ready are designed to enhance user productivity. The true value of Intel Cluster Ready for users is the overall reduction in time to results that can come from

more productive HPC installations. The result is increased productivity for end users because they can worry less about HPC deployment and management issues and can focus on getting work done.

MARKET DYNAMICS

The New HPC User

An important dynamic in the growth of the HPC industry is the adoption of high-performance clusters by new users. As HPC has transitioned to industry-standard technologies, the price/performance improvements have tipped the ROI equation in favor of adoption, and the systems are capable of conforming to IT standards.

New HPC usage comes in three main categories. In the most straightforward case, an organization spreads the usage of HPC to new groups within the same organization. In this case, knowledge transfer can ease the expansion, especially when there is any sharing of IT infrastructure. In the second type of new usage, established HPC applications are adopted by organizations that had not previously done digital simulations. For these new users, the applications may be available off-the-shelf, but there may be challenges in creating digital prototypes or building out systems to scale for the first time. Both of these types of new user expansion have gone on for decades, but growth rates have been higher in recent years.

The third – and most significant – type of new usage is the adoption of HPC technologies in application areas outside the traditional realm of science, engineering, and analytics. In its demand-side studies and frequent interactions with the user community, Tabor Research has identified several new and emerging non-traditional application areas that have begun to adopt HPC technologies in significant quantities. Tabor Research refers to these as the “Edge HPC” market (as opposed to “Traditional HPC”), whose size, growth, and set of customer dynamics will play significant roles in the adoption rates of new products, services, and/or initiatives.

Tabor Research divides Edge HPC into four segments:

- *Complex Event and Business Processing* – Applications that are driven by continuous data feeds generated by real world events. The applications enter the HPC realm as requirements to maintain low response latencies in the face of growing data volumes exceed the capabilities of standard solutions. Often times these applications have real-time response requirements and demand highly distributed infrastructure.
- *Process Optimization* – Applications with profiles that mirror traditional HPC workflows. These applications often make use of technology above and beyond standard enterprise solutions, either in architecture, software, or system management. Applications tend to be data and/or computationally intensive, but run in operational time (i.e., in synchronization with the operational timing of the organization).
- *Virtual Infrastructure and Environments* – Applications that implement computer network based business and social structures. These structures range from online gaming environments to multi-person/system training environments to virtual economies to virtual social environments.
- *Ultra-Scale Computing* – Applications that require systems to be specially designed or configured to effectively manage node counts that significantly exceed those provided by industry-standard products. Ultra-scale requirements appear across multiple application spaces, and thus this segment may be viewed, in part, as representing an “Other

Applications” category. Some broadscale, corporate IT infrastructures fit this profile, as do many back-end server farms for Internet search and Web 2.0 applications.

In many ways, the growth in Edge HPC can be seen as the co-evolution of technology requirements driven by HPC on one hand and large-scale datacenter or Web 2.0 applications on the other. This dynamic has created a growing pool of new “cross-breed” users who have the performance needs associated with traditional HPC but the risk-averse nature of mission-critical enterprise computing. In addition, these new users do not tend to associate themselves with “HPC,” which they may view as esoteric or experimental.

Dell and Intel Cluster Ready can empower HPC adoption for this type of user. Pre-tested, reliable systems from a trusted vendor are the types of scalable systems that an evolving enterprise user can acquire and manage.

The Cluster Management Challenge

Tabor Research studies show growth across all HPC segments. This growth has been fueled by transitions to standard components, led in part by companies like Dell and Intel. For the last decade, HPC systems have increasingly been built from standardized technologies that are ubiquitous in normal IT usage. Users select servers, microprocessors, and interconnects based on price/performance and other factors, and they add appropriate middleware and software to transform the pieces into HPC clusters. Because the modern hardware is both more familiar and less expensive than proprietary or custom architectures, it has enabled more organizations to become first-time HPC users and allowed existing users to expand their HPC infrastructures. The hardware barrier to HPC entry or expansion has been significantly lowered.

Software, on the other hand, has remained as a source of frustration for users, either as a barrier to entry or as an impediment to productivity. Even for seasoned IT veterans, the software stack for HPC clusters offers challenges in configuration, installation, optimization, and management.

The primary driving factor behind this frustration is disintegration of the software stack. The democratization of HPC extends beyond hardware to include software components such as operating systems, MPI layers, compilers, job schedulers, and other tools. Users commonly acquire different versions of each of these components; some are purchased, some are individually downloaded for free, and some are bundled together with systems, processors, or application software. In some cases, users may utilize software components that are written and maintained in-house. The result of this piecemeal approach can be suboptimal performance or application failure due to unforeseen incompatibilities between software components obtained from disparate sources.

A primary factor amplifying this trend is the rapid evolution and adoption of the Linux operating system. The two most prevalent Linux families – Red Hat and Novell SUSE – not only have multiple kernel and distribution incarnations, but also have spawned numerous overlays and optimizations to streamline them for particular server platforms or levels of scalability. As a result, systems acquired at different times or from different vendors are likely to show significant variability in the software versions installed and in performance. Even if an administrator sets up one cluster optimally, there is no guarantee that the second will behave the same way. The addition of multiple versions of various commercially-available applications compounds the permutations of potential incompatibilities.

A research study conducted by Tabor Research in 2007 found that the greatest source of software complexity for HPC users was the transition to processors with additional cores (i.e., single-core to dual-core, dual-core to quad-core, etc.). In addition to making sure release versions match, the software and applications need to be capable of taking advantage of the available performance. Additionally, when users found existing tools to be inadequate for their purposes, they were most likely to seek freely available tools to replace them. This combination of usage

and development trends has contributed to the rapid proliferation of software components that are not well-tested or qualified with each other.

For the IT professional or HPC user, identifying and resolving incompatibilities can be a taxing chore. It can be difficult to determine which components are combining to limit performance. As a result, the user may not know where to turn for assistance. If each component works in isolation, who do they work with when the panoply fails?

Some users are combating this complexity by investing in additional system administrators, but this solution is not ideal, particularly in many smaller operations. The talent pool for additional staff can be limited or costly, so hiring can be a challenge. Furthermore the traditional jurisdiction of system administrators is operating systems, while the management of variegated, unbundled middleware components requires a broader scope of expertise. Therefore, the addition of administrative staff is suboptimal and often unviable as a solution. The real need for many HPC environments is to simplify the overall HPC solution stack.

DELL AND INTEL CLUSTER READY

To combat the widespread problems associated with software incompatibilities or inefficiencies in HPC clusters, Intel partnered with Dell and Platform Computing to create the Intel Cluster Ready program. The Intel Cluster Ready program targets entry-level and mid-range HPC environments that need ease of installation and system maintenance in their HPC deployments.

The Intel Cluster Ready Program

To address frustration with HPC cluster management among the growing population of new HPC users, Intel introduced the Intel Cluster Ready program. Through this program, Intel works with hardware vendors to designate certified Intel Cluster Ready configurations: systems that are built according to the Intel Cluster Ready Specification and tested for component interoperability. Intel also works with software vendors to designate registered Intel Cluster Ready applications, which have been validated for use on certified Intel Cluster Ready systems. Because registered Intel Cluster Ready applications are verified to run on clusters with the same base specification, users are able to run multiple registered applications on the same certified Intel Cluster Ready system without rebuilding the software stack or reconfiguring the hardware.

To maintain the health of the cluster over time, end users can utilize the bundled Intel Cluster Checker. Intel Cluster Checker is the tool suite cluster providers use in order to validate Intel Cluster Ready certification. By running these tools themselves, users can reduce the time spent troubleshooting while enhancing system reliability and maintaining optimal performance levels. Intel Cluster Checker continually monitors the health of the cluster and alerts administrators to problems with a straightforward "traffic light" dashboard interface. Because Intel Cluster Checker simplifies troubleshooting, HPC users may also avoid the need to add specialized support personnel.

Although any HPC user might benefit from the Intel Cluster Ready program, it is most attractive to three primary HPC segments:

- *Entry-level or mid-range users.* Users with smaller HPC organizations are less likely to possess significant levels of HPC-specific system administration expertise. The acquisition of systems that are certified to work without intrusive software intervention is a potentially attractive concept for these users.
- *Production users.* Organizations or divisions who are working on production systems, as opposed to research or development systems, are intolerant of suboptimal performance, system failure, or downtime. They are also likely to strive to maintain consistency in their

software environments, with less tolerance of diverse, unqualified versions. The combination of certified clusters and registered applications is therefore a good fit for production users.

- *New HPC users.* Just as a new cook is more likely to work from a recipe than a well-seasoned chef, users who are implementing their first or second generations of HPC clusters can derive more benefit from Intel Cluster Ready than experienced HPC users, who may have already established particular combinations of hardware and software for their organizations. The growing Edge HPC market includes many new HPC users.

These three market segments overlap, and users at the intersection of these groups – production users of first- or second-generation, entry-level or mid-range HPC clusters – are likely to find the greatest benefits from Intel Cluster Ready solutions.

Of course, even users who benefit the most from certified solutions will invariably seek some level of customization, with additional hardware, an in-house program, or some potentially obscure piece of third-party software that is critical to their particular workflow. Dell HPC hardware is Intel Cluster Ready capable, so additional nodes, processors, or memory can be added to the certified cluster. The Intel Cluster Checker tool can be used to check the certification status of the cluster and to verify compatibility of new components, to diagnose and isolate problems, and to monitor system health on an ongoing basis.

The Intel Cluster Checker tools evaluate system wellness and behavior through a variety of tests, including some widely recognized HPC benchmarks. Some of the highlights of Intel Cluster Checker are:

- The STREAM benchmark, which tests system-wide memory performance.
- “Ping Pong” tests of networking functionality and performance.
- Intel’s own MPI benchmarks to test network performance inside the cluster.
- Version checking for compilers and other middleware.
- Consistency checking for MPI and network images.

The Intel Cluster Ready program isn’t for everybody. Experienced R&D experts at national supercomputer centers are notorious tinkerers who prefer to create, find, and eventually solve their own problems. But for buyers who value simpler, pre-tested, certified HPC solutions, Dell HPC clusters and the Intel Cluster Ready program provide a compelling alternative to unbundled, potentially incompatible hardware and software components.

Dell Promotion of Intel Cluster Ready

The success of the Intel Cluster Ready program relies on cooperative partnerships with major hardware and software providers. Users must find trustworthy solutions that are inclusive of components that are desirable for their installations. Potentially the most important role is that of the server platform vendor because this is the company and brand that the user identifies as having provided the cluster.

In this sense Dell could be Intel’s most significant partner on the Intel Cluster Ready program. Dell was one of Intel’s two co-developers of Intel Cluster Ready. (Platform Computing, the maker of cluster provisioning and management software such as LSF, was the other.) Dell has been involved since the project was conceived, a year before its launch, and has continued to work closely with Intel engineering to help further refine the Intel Cluster Checker tool and to develop a comprehensive distribution methodology.

TABOR RESEARCH ANALYSIS

Dell brings recognition and value to the Intel Cluster Ready program because of its particular profile in HPC and in general-purpose computing. Dell is a tier-one hardware vendor with a breadth of product offerings, so end users have a wide array of configuration options from which to choose. Dell has a reputation as a solid IT partner and is a trusted provider of IT solutions worldwide. As such, Dell is strong in commercial production settings. Furthermore, Dell's well-established sales channels make it easy for new users to acquire Dell systems.

Users will come to Dell for complete, supported solutions that they expect to work right out of the box. If and when problems do arise, users may feel confident in turning to Dell for support. The Dell-Intel partnership on Intel Cluster Ready is mutually beneficial, while providing supportable configurations to a significant class of HPC users.

Challenges and Opportunities

Tabor Research predicted that partnerships and bundled offerings would begin to combat system administration challenges for HPC customers, and Intel Cluster Ready, with close collaboration from Dell, is the industry's first large-scale manifestation of a collaborative solution. However, other competing partnerships may emerge with different standards for pre-configured solutions.

A challenge for the Intel Cluster Ready program is for users to perceive its value in a positive fashion. While the simplest way to explain Intel Cluster Ready is in terms of "ease of use," such positioning statements necessarily begin from the negative viewpoint of "what if it doesn't work." Furthermore, Tabor Research studies have shown that although users are frequently frustrated with cluster management challenges, they do not place significant value on concepts such as "comprehensive cluster management" in their system evaluation. This does not mean users cannot benefit from or perceive value in Intel Cluster Ready. Although comprehensive cluster management is not a primary feature that tends to drive acquisitions, users do perceive high value in productivity gains.

This is the true value of Dell HPC and the Intel Cluster Ready program to the end user. If buyers can install systems and get them up and running more quickly and consistently, with fewer headaches and less downtime, they will have more productive environments in which more work gets done. In this way, Dell and the Intel Cluster Ready program could help take users to new levels of engineering leadership, product design, and innovation.

For more information about Dell HPC solutions and the Intel Cluster Ready program, visit <http://www.dell.com/hpc>.