Transforming Your Enterprise

Synchronizing business and IT
to capitalize on change

March 2006

High-Performance Computing Special Edition

invent


change
hp
The spinoff that acted like a startup.
Avaya, a global leader in communication software, systems and services, spun off from Lucent with a legacy IT infrastructure that, while efficient, wasn't nimble enough to be a competitive advantage. HP partnered with Avaya to implement IT Service Management and HP OpenView, effectively re-deploying existing technology assets. Today, IT spending is down 30%. Millions have been saved by finding unused capacity. And Avaya answers whenever opportunity calls. www.hp.com/adapt
Solutions for the adaptive enterprise.
invent
©2004, 2005, 2006 Hewlett-Packard Development Company, L.P.

Transforming Your Enterprise

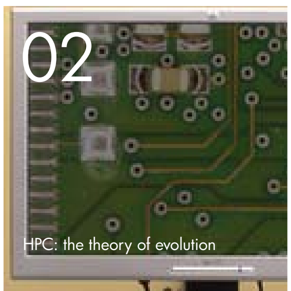

HPC: the theory of evolution

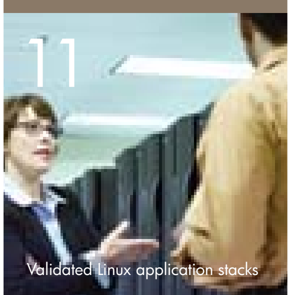

Validated Linux application stacks

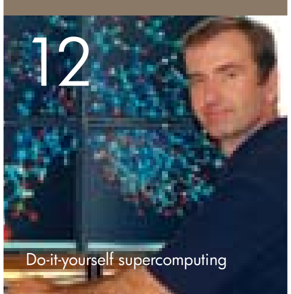

Do-it-yourself supercomputing

Table of Contents

Easing the pain of progress

Wherever there's change, there's challenge, and the high-performance computing (HPC) market is no exception. HPC has undergone a transformation over the past 10 years, with traditional proprietary high-end systems such as vector supercomputers (Cray) giving way to industry-standard servers running UNIX®, Linux® and even Microsoft® Windows®.

IDC estimates that scientists and engineers increased their purchases of clustered standards-based systems from about $285 million in 1999 to $2.6 billion in 2004*. In addition, non-technical environments are seeing increasingly demanding workloads and are turning to clustered HPC solutions. Thus, parallel computing, once only used by scientists and engineers, has become widely adopted in the commercial world.

But this two-pronged shift in the technical computing marketplace leaves many IT departments—in both technical and non-technical organizations—with unanswered questions. How do we manage change? Will our storage requirements be met? What options do we have when purchasing these types of systems? Do we need grid computing? Will our existing applications run in a clustered environment?

HP has created this publication to answer those questions—and many more. We recognize it's crucial for our HPC customers to have access to trends and data to stay one step ahead of future requirements and, just as importantly, the competition.

Your feedback is welcomed at www.hp.com/go/transformHPC.

For additional information follow the links inside, or contact your HP sales representative, your HP channel partner or the HP Welcome Center at 1.800.282.6672, press 5 and mention code BLDG.

Winston Prather, VP and GM, High-Performance Computing Division, Hewlett-Packard Company

Subscriptions/Feedback — To subscribe to receive a future complimentary printed or electronic copy of *Transforming Your Enterprise*, and to provide feedback on the content in the publication, visit: www.hp.com/go/transformHPC.
To unsubscribe to *Transforming Your Enterprise*, visit: www.hp.com/go/transformHPC.

*** Introduction, Page 3** — An IDC white paper sponsored by HP: "HP's Strategy for Delivering Cluster Technology to Technical Computing Environments," IDC, November 2005. "IDC Estimates 23% Growth in Technical Computing in 2005: Third Consecutive Year of Double-Digit Growth," IDC Doc #34611, December, 2005. *** Page 5** — "Smart Chip, System and Data Center Enabled by Advanced Flexible Cooling Resources," by Chandrakant D. Patel, Cullen E. Bash, Ratnesh Sharma, Abdlmonem Beitelmal, Christopher G. Malone, Hewlett-Packard Laboratories, December 2004. *** Page 8** — Information Week, December 5, 2005. *** Page 10** — "HP Makes Its Utility Computing Offerings More "Adaptive" to Customer Needs," by David Tapper, IDC, December 2005.

Printed in Canada.

invent

HPC: the theory of evolution

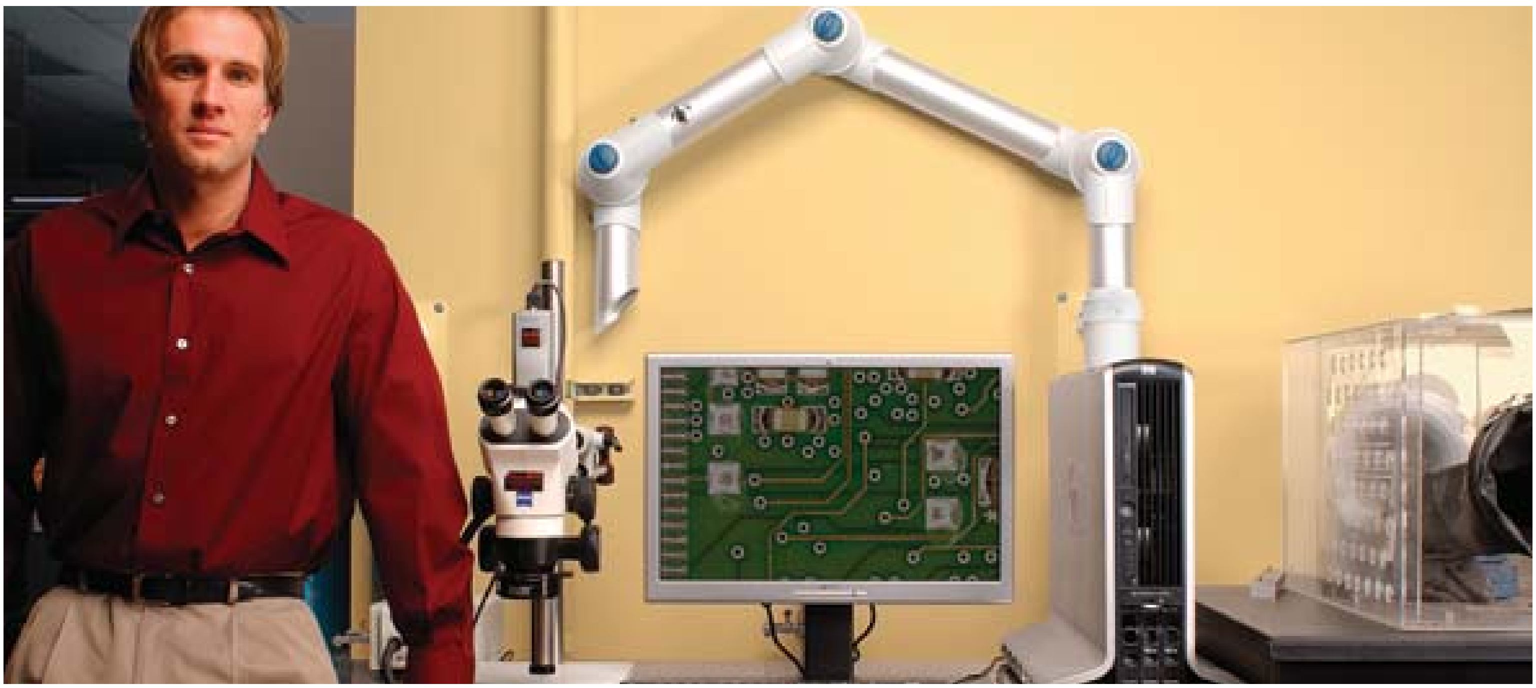

From Cray to Gates, technical computing is headed downstream in a big way.

You know there's been a big shift in the technical computing marketplace when Bill Gates headlines Supercomputing 2005, and talks about embracing the ever-accelerating trend of low-cost, high-performance compute clusters. “Technical computing is crucial to the many discoveries that impact our quality of life,” said Gates. He was preaching to the converted, of course: technical users and vendors have known this all along. But although market growth has been propelled by UNIX® and Linux®-based servers (often run in clusters) from HP and others, the idea of mini-supercomputers running Microsoft® Windows® on people's desks isn't that far from reality—thanks to rapid advances in hardware and increasingly sophisticated software (see “Do-it-yourself supercomputing,” page 12).

How quickly supercomputing zealots will embrace Windows over their beloved UNIX or Linux remains to be seen. The fact is, however, lines are blurring between business and scientific computing. “Industry-standard open source environments are making technical computing affordable for people who could never afford to buy these kinds of systems before,” observes Ed Turkel, manager of the product and technology marketing group for HP’s High-Performance Computing (HPC) Division. “Today, a two-person engineering shop can have its own cluster sitting in a corner. Or a research team can have their own system, rather than being dependent on some big shared system within an organization.” Indeed, the largest and fastest-growing areas of the HPC market are the low end—departmental and workgroup segments which include systems priced at less than $250,000*. These make up more than half of the overall market, and according to IDC's initial estimates, are expected to grow at roughly 35 percent from 2004 to 2005, while the classic supercomputer "capability systems" market has reached a mature stage and is expected to see flat growth over the same period*.

While this may come as news to many in the HPC segment, the shift hasn't happened overnight. For decades the industry has been moving away from monolithic multimillion-dollar proprietary supercomputers toward lower-cost configurations of VAX clusters in the mid-80s, then RISC/UNIX, and most recently, Linux and

Windows on industry-standard processors from Intel and AMD (see "The collaborative effort...," page 11).

Led by Digital Equipment Corp. in the mid-1980s, clustering is the biggest hardware trend in HPC today. The reason it's taken a while to gain acceptance is that not all supercomputer-based programs translate well onto the clustered model. “Some application vendors have to catch up,” says Turkel. For instance, in computer aided engineering (CAE), fluid dynamics applications run well on clusters but structural analysis applications tend not to. HP and others have been involved in government-funded research to create de facto standards such as Message Passing Interface which creates code to allow various end-user applications to run in parallel on multiple machines. The HP-MPI implementation is helping independent software vendors (ISVs) to provide applications support for clusters, a major factor in clusters moving into the mainstream.

Vendors are also working to improve cluster storage performance (see “Fast storage: the forgotten ingredient of high-performance computing,” page 16) as well as cooling technology (see “Staying cool in the data center,” page 4), because when you put too many nodes in a rack, well, things can get hot. These days, it's not uncommon for a customer to install 1,000 nodes in one shot. “A 1,000-node cluster a few years ago was considered huge,” says Turkel. “You couldn't do it without a major customer project. Now it's—I wouldn't say commonplace—but it's something we can do with standard configuration rules.” HP has built multiple 1,000-node and larger clusters, and is bidding on more all the time. “The size and scope are growing as the technology matures,” Turkel observes.

This up-scaling trend is important because traditional technical computing environments require ever-more processing power, in a more flexible package. A leading aircraft manufacturer that used to run large symmetrical multi-processor (SMP) systems for its structural analysis, fluid dynamics and impact analysis applications, has been shifting to several clusters of 96 HP Integrity server nodes running HP-UX. “In the past they would have considered this stuff that the lunatic fringe (our euphemism for high-end) does,” says Turkel. “But we're taking this technology once thought of as bleeding-edge and making it industry-standard.” To see how, turn the page.

To get started in a clustered environment customers can opt for HP's Cluster Platform Express, aimed at the small to mid-sized cluster user from four to 32 nodes. Any user can go to the website and get an online quote for a cluster by entering key requirements (memory, processor, software, etc.). Price range: $40,000 to $500,000. For information visit www.hp.com/go/transformHPC2

Staying cool in the data center

Scaling out your server environment in order to quench your high-performance computing thirst can be like trying to squeeze 150 people into one seat to watch a symphony. It can create heating issues so dire, whole systems could crash. Fortunately, there are ways to alleviate this acute thermo-mechanical problem.

There was a time when you could install as many servers as you wanted into a data center. Since most centers were originally built to house monolithic mainframes of an era gone by, they would contain special heating, ventilation and air conditioning as well as chilled water lines designed to combat any heat dissipation.

That's no longer the case, particularly with the advent of server blade technology and clusters, which are not only compact and powerful, but generate enough heat to push the thermal envelope to the limit. Dan Cox, a manager of High Performance Computing Programs at HP, says the problem is not only about power in a room, but also the ability to extract the heat. Many data centers, he adds, have reached the law of diminishing return.

Conditions are also much different than they were in the 1980s and 1990s when air-cooled Unix®-based proprietary RISC machines were the dominant hardware choice in the data center. Then, the infrastructure, design and management practices could keep pace with the power and cooling demands, but that's no longer possible as the physical size of servers continues to decline, and their numbers increase. As an example, two years ago it was possible to fit an average of four to six servers into a single rack at any given time; the typical size of each was 5-7U, and the number of electronic watts generated was anywhere from 1,500-3,000. Today, the average number of servers per rack is 12-16, with an average server size of 1U (1U = 1.75 inches), and the power generated per rack has skyrocketed to between 5,000 and 6,000 watts. Looking ahead, power densities of 1,000 watts per U are likely. With density that high, data center personnel have very little time to react in the event of a cooling system failure.

Ken Baker, an Infrastructure technologist with HP, suggests that it's critical that organizations employ best-in-class practices on the data center floor so they do not waste cooling capacity. "You can continue to add cooling capacity, but because you have a finite fixed geometry within that enclosure, how and where you arrange those resources, and how they perform, is going to diminish over time," he says. "Those considering the deployment of additional clusters into their existing set-up should first evaluate their capabilities to environmentally support them," Baker cautions.

In order to mitigate the potential risk of an A/C meltdown, HP Services currently provides a static cooling service that is designed to reduce energy use and energy costs. Created at HP Labs, the Smart Cooling service uses computational fluid dynamics to create a 3D model of temperature distribution throughout a data center. Based on this model, it then recommends where computing resources and air conditioning should be placed. HP researchers believe the analysis — which includes written reports based on a thorough on-site inspection — could reduce cooling expenses by as much as 25 percent for a data center with 1,000 racks requiring 10 megawatts to power the infrastructure, and an additional 5 megawatts to dissipate and remove the heat*.

According to Chandrakant Patel, who oversees the thermal technology research program at HP Labs, and is co-inventor of the Smart Cooling service, the data center is now akin to a symphony hall in which every seat holds 150 members of the audience. "For such power density, energy balance alone in sizing air conditioning capacity in a data center does not suffice," he says. "One has to model the airflow and temperature distribution to assure inlet air temperature to systems."

When an animation firm, for example, has a cluster of machines executing rendering, that cluster may generate more heat in one area of the data center. This asymmetrical loading can cause problems, but with static smart cooling a center is analyzed and optimized so that for any particular configuration the air conditioners are well balanced.

The next evolutionary step in this process will be the introduction of HP's Dynamic Smart Cooling. Now in development, this service uses numerical modeling techniques to statistically provision the cooling resources based on a given layout of heat loads to avail energy efficient heat removal. Patel contends that when it comes to proper cooling techniques, the average home can be more sophisticated than a $100 million data center. "In many centers, air conditioners run at a fixed rate and somebody goes in and dials in the temperature. There is no flow control and certainly not much temperature control based on need," he says.

HP plans to include a pervasive sensing infrastructure in it's Dynamic Smart Cooling service, which is being designed so that air going to the servers is measured and used as a proxy for control. "If you have a blade environment and a fully-loaded rack, we plan for sensors to be at the front and at the back of the rack. The system is being designed so that it collects all that information about current heating and cooling conditions and an HP algorithm analyzes it and sends new flow and temperature points to the air conditioner," says Patel

If you want highly available, highly reliable ultra dense installations — especially in the realm of large clusters and scientific computing — this is a field that cannot be ignored, says Baker: "It is critical that you engage the proper professionals to analyze your environment prior to deploying the hardware so that you do not make a very expensive mistake."

For an HP brief on optimizing data centers for high-density computing, visit: www.hp.com/go/transformHPC

To view a *Scientific Computing* webcast "Power Up/Cool Down" visit: www.scimag.com/coolpower

Grid myths debunked

Understanding the evolution from distributed computing to the ubiquitous grid, and where to begin.

A cursory glance of recent computer tabloid headlines will tell you grids are the next big thing to hit IT. From promises of utility-style computing to talk of mega-grids, just about every vendor has a definition or technology to sell.

But before you get caught up in the hype, there's something to keep in mind: a grid is not a product, it's an evolution, says Sara Murphy, HP marketing manager for Grids.

"Grids are built, not bought," stresses Murphy. "When we talk about grid computing, we're really talking about laying the groundwork for the evolution of your current enterprise-wide distributed computing environment, not a magic black box that you order and deploy."

Key to understanding the promise of grid computing, she adds, is the realization that a grid — as it relates to IT — is a software environment based on open standards and protocols that make it possible to share disparate, loosely coupled IT resources across organizations and geographies. In a grid, resources such as computer cycles, storage capacity, databases, applications, files, sensors or scientific instruments, can be dynamically provisioned to the users or applications that need them.

One reason for confusion in the marketplace, is that different organizations are at varying stages of the grid evolution, depending on their own set of unique circumstances, adds Niraj Srivastava, a senior technologist for the Advanced Solutions and Development Practice of HP Services. For example, they might begin with clusters and load balancing, or they could focus on administrative portals, data services and security issues.

No matter where you begin, pulling together components and protocols to build a grid is much like designing a house, says Srivastava. "You don't just order a house; you work with an architect to design the house that best suits your lifestyle and budget."

For some organizations, the driving factor

Instead of each research lab building its own cluster of computational resources and then watching it sit idle for weeks or months at a time, they can deploy a grid strategy to share resources across boundaries on an as-needed basis and with enhanced security.

is the need to deploy a new IT infrastructure in support of a new business line. Others are looking at grids as a way to gain computational power they wouldn't be able to otherwise acquire. And still others are building grids to enable geographically dispersed users to collaborate. Determining those requirements up front should be the first step in any grid exercise.

“The components and protocols that make up a grid aren't always new technologies,” explains Srivastava, referring to resource managers, portals, schedulers and clusters. “Building a grid infrastructure is about pulling them together in a modular framework based on standards.”

From an IT perspective, one of the key benefits to deploying a grid infrastructure is the ability to pinpoint under-utilized resources. In large pharmaceutical companies, for example, grid technology is successfully being used to allow researchers located in different countries to share computational resources. As Srivastava explains, instead of each research lab building its own cluster of computational resources and then watching it sit idle for weeks or months at a time, they can deploy a grid strategy to share resources across boundaries on an as-needed basis and with enhanced security. IT is able to get much more efficiency out of its infrastructure and researchers are able to bring drugs to market in a more timely fashion.

Other examples of grid computing are found in the scientific and technical research arena, including projects like CERN's Large Hadron Collider Computing Grid, used by a worldwide community of scientists; the Southern Partnership for Advanced Computational Infrastructures (SPACI), the first grid-enabled research and development IT environment in Europe; and, Canada's Shared Hierarchical Academic Research Network (SHARCNET), a computing grid serving 11 academic institutions.

As a grid pioneer, and major contributor to these and other grid projects, HP is committed to bringing the benefits of grid computing to the enterprise. HP has ported, tested and tuned the Globus Toolkit — the open source grid infrastructure implementation — to HP server platforms, and works closely with independent software vendors such as United Devices, DataSynapse, Platform Computing, Altair Engineering, NICE and TurboWorx to ensure third-party grid software works well on HP platforms. It also provides the management software and storage architectures that enable flexible, scalable grid solutions.

In addition, HP Consulting and Integration Services provides a comprehensive range of services to help plan, design and implement a grid computing infrastructure, including Grid Orientation, Grid Assessment Workshop, Grid Knowledge Transfer Workshop, Grid Design and Deployment Services, and Grid Optimization Service.

“Grids are no longer a vision for the future,” says HP's Murphy. “They're here today, but it's still early on in the evolution.”

For more information on HP grid technologies and services, and to download an HP grid white paper and solutions brief, visit: www.hp.com/go/transformHPC

To view an HP webcast discussing the problems solved by grid, including example use cases, visit: www.hp.com/go/transformHPC4

Feature

IT on tap!

Pay-per-use computing with enhanced security is now a reality for performance-hungry environments.

Like the idea of buying compute power on an as-needed basis, but don't like the potential loss of control? Take another look, because the pay-per-use model has come a long way since its inception in the late 1990s. In fact, Gartner estimates the worldwide market for IT utility-computing services will grow from $15 billion in 2004 to about $80 billion in 2008, and approximately 50 percent of all data-center assets will be purchased as a service by 2010*.

Under this model, companies in high-performance computing sectors like oil and gas, manufacturing and financial services can plug into offsite computing power for one-on-one use at a per-hour rate. “Utility computing is as important to high-performance computing as Linux® was when it first came out,” says Chris Smith, business development manager for HP Utility Computing Services. “It allows companies access to the compute power when they need it and not be left

holding that cost when compute power isn't called for."

Often compared to the approach taken by public energy utilities, the IT-as-a-service concept also echoes the timeshare principle of the mid-20th century. Back then, companies couldn't justify the expense of buying a mainframe, nor could they fully utilize or maximize the assets. So they shared the time, and cost, with other companies. These days the cost of hardware is far less prohibitive, but high-performance environments are nevertheless constantly at risk of being constrained by technology limitations. Companies — particularly those that process terabytes of data in intensive cycles — end up playing leapfrog and incurring enormously high IT overhead as a result. "It's almost impossible for any company today to fully meet the needs of their engineering group," observes Smith, "even though that same group forms the core of a company's success or failure by their ability to create the next great product faster than their competitors."

One answer might lie in HP's Flexible Computing approach, launched last November. From its offsite data centers, HP hosts customers' applications and/or sells compute cycles on an as-needed basis. It's an especially suitable model for companies that need a large amount

IDC's David Tapper: "By taking responsibility for the capital burden, HP can reduce total costs and lower the investment risk for customers*."

of extra capacity at certain times of the year to run compute-intensive programs like oil reservoir simulation, product prototyping and crash analysis. As well, the company offers access to the latest versions of pre-integrated vertical applications in conjunction with leading independent software vendors (ISVs). Small to mid-sized customers — for example, the ISVs themselves — are prime candidates, because their design, development and prototyping cycles peak at certain times.

HP's Flexible Computing Model consists of three tiers: Infrastructure Provisioning Service (secure processing services from HP data centers on Linux, HP-UX, and Microsoft® Windows® 32-and 64-bit systems); Infrastructure Provisioning Service Plus (this adds grid and workload/scheduling management and/or compiler software); and, Application Provisioning Service (sits on top of IPS or IPS+). Pricing ranges from about 55 cents an hour per 32-bit Intel® Xeon®-based server to $1.50 per processor per hour for Intel® Itanium®, while pricing for 64-bit Xeon and AMD Opteron™ falls somewhere in between.

It is a true turn-it-on, turn-it-off solution…and the customer has complete and exclusive control.

HP's strategy should help change the way customers purchase technology — as a service rather than a fixed asset. As IDC's David Tapper says: "By taking responsibility for the capital burden, HP can reduce total costs and lower the investment risk for customers*."

One such customer, Schlumberger Information Solutions, is working with HP to help provide its own customers with the extra compute cycles and improved performance needed to run large numbers of reservoir simulations. This should help Schlumberger deliver faster turnaround times to customers, eventually allowing them to speed up the cycle from exploration to production.

What's the downside? Clearly, risk is a sensitive subject for many high-performance computing customers, who historically shied away from the utility model because they refused to send company-sensitive data over the Internet and outside their own firewalls. "That's still a challenge for a lot of companies today," concedes Smith. But he points out that network security has improved to the point where, if the right solutions are in place, it's practically impenetrable. HP, for example, not only enhances the security of its data center assets with advanced intrusion detection, VLAN extranets and secure cell architecture environments; it also treats the assets as though they were on the customer's own premises. Unlike traditional outsourcing, HP gives customers full control over the IT assets, including a root password to the system. Once a VPN connection or secure shell has been opened, even the staff on the data center floor can't access or even view the data within those systems. Smith says HP's IT personnel need biometric hand scanners combined with password and badge control just to get through the door. "I can't even get into this facility."

Okay, so security is a surmountable problem. But is pay-per-use computing the right solution for your company? IT directors and company management should ask themselves the following questions: if high-performance computing is considered a mission-critical asset within your organization, could the needs of your engineering group be better served? Would it be cheaper to use a utility service, versus buying or leasing the assets? Finally, think through your usage model. If you're running 50 compute nodes 24 hours a day 365 days a year, at some point it would be cheaper to buy or lease than to implement a utility model. But let's say your company wants to deploy 256 nodes for 12 to 16 weeks per year during its product design cycle. At roughly $36,000 per week, that's far more cost-effective than purchasing 256 nodes you can't turn off, at $50,000 to $100,000 apiece, over the lifetime total cost of ownership for each node.

"It is a true turn-it-on, turn-it-off solution," says Smith. "It enables customers to make it a virtual extension of their own IT assets, and during the time they use the utility, they have complete and exclusive control."

Still worried about those control issues? Thought not.

Companies wanting to test drive the Flexible Computing model can sign up for HP's Flexible Computing Club, a low-risk, low-cost way to experiment with these types of services. Club membership is $5,000 and provides customers with startup and consulting services, training, a dedicated home node server and a 48-hour pilot project. If you place an order of more than $20,000 within 90 days, then the fee is credited back. To have a representative contact you, visit: www.hp.com/go/transformHPC5

The collaborative effort to push Linux® clusters to top performance

Standardized and validated solutions are pushing Linux® clustering performance ahead and reducing management complexity.

There are many business and technical benefits to deploy a Linux® cluster, including an enticing price point, enhanced compute speed and robust processing. But the most alluring aspect is the emergence of a standardized, qualified Linux clustering solution that will minimize management and support complexity.

A Linux-based cluster is not a plug-in, problem-free platform. Application vendors build solutions on specific Linux flavors. Interoperability issues hit when a cluster platform doesn't match a vendor's chosen flavor, and complexity rises as more applications are laid down. Very often, the increased maintenance and support work needed to manage the cluster ends up siphoning off the cost efficiency gained from open source technology. "For some of our customers, Linux has proven to have hidden costs related to deployment and management," says Barbara Hutchings, director of strategic partnerships for Fluent, a commercial computational fluid dynamics (CFD) vendor. "It has also been a challenge to work with as a vendor, due to all the permutations and related trouble shooting needed."

Yet, despite the management headaches, Fluent saw a significant up tick in 2005 of its installed base moving to a Linux-based clustered environment. The primary reason is that Fluent was built specifically to work in a distributed memory environment and scales extremely well on Linux clusters. The increasing customer interest in standardized high-performance computing (HPC) is why Fluent is working with HP to ensure its application meets standardized software stack requirements.

Last year HP and Novell announced a joint agreement to certify and support the SUSE LINUX Enterprise Server on HP ProLiant servers and the HP BladeSystem. A team — comprised of HP, Novell and a select group of software providers — built a validated solution suite for high-performance and grid computing called the Validation Suite. It features industry standard EM64T, AMD Opteron™ servers, Novell SuSE SLES 9, and Scali Manage and MPI Connect for management needs, as well as vertical compute libraries such as the MPI Library. Components within each validated stack work together to provide ease of operations.

Novell, HP, and their team of software providers have integrated and tested the components to deliver a wide range of specifically configured and verticalized stacks on Linux. Validated configuration data — including performance test results and best practices for deployment — is available from Novell, HP and the software providers with hardware and software components for each stack available from HP. The Validation Suite is complementary to the HP-developed XC software stack which is based on RedHat.

"We're pleased to work closely with HP's technical team, and a critical component is the HP-MPI library to help ensure that the software we send out is compatible with industy-standard software stacks," says Hutchings. Fluent offers very robust support for SUSE SLES 8 and 9 and is adopting HP-MPI in the upcoming release of its software this year.

Hutchings notes that the cluster development trend of the past saw customers acquiring different compute pieces from multiple vendors. "They'd get the compute nodes from one source, an OS from another, the interconnect from another, and then struggle to assemble a working system." Today, she says, customers are increasingly appreciative of a validated and standardized software stack with everything bundled and tested. It's likely going to be less expensive in the long run, she says.

HP, Novell and their team of software providers continue to improve and extend the Novell Validation Suite. HP offers a free service to software providers and customers that are interested in working on the Validation Cluster suite. For details, e-mail: askforhelp@hp.com

For more information on Fluent engineering solutions visit: http://www.fluent.us/

Technologies

Do-it-yourself supercomputing

Compute clusters and parallel programming suites deliver supercomputing capabilities once reserved for the privileged few.

Financial modeling, computational fluid dynamics, molecular modeling, computational chemistry, structural stress analysis and other tasks requiring supercomputing assistance have in the past necessitated two types of individuals: the researchers, engineers and scientists who understand the science and comprehend the problems and dynamics of these complex application areas; and, the technology experts and programmers who must implement the requirement digitally, and analyze and test them in the computer realm.

According to Dan Cox, manager of High Performance Computing Programs for the Industry Standard Servers group at HP, this dynamic is changing. "Several developments are enabling the research world at-large in some very exciting ways," he notes. "Thanks to affordable computer clusters and new parallel programming suites, small companies, universities and other emergent research organizations now have access to supercomputing capabilities that once were reserved for the most prominent enterprises, governments and education institutions."

The possibility of what could be termed "do-it-yourself supercomputing" was born from the creation of dual core CPU chips, which enabled the aggregation, or clustering, of several computing systems based on multiprocessor architectures. These clusters deliver supercomputing capabilities at a much lower cost and complexity than traditional, standalone systems.

"Industry-standard compute clusters now can be purchased affordably and easily placed in a cubicle, delivering the same performance that previously necessitated a multimillion dollar supercomputer that filled a large room," exclaims Cox.

Affordable hardware clusters meet only half the challenge, however. The multiprocessor architectures on which compute clusters operate require parallel programming approaches, such as multithreading, openMP and message passing interface (MPI). According to Cox, parallel programming using these current industry standard approaches is difficult at best for the engineer, researcher and scientist who have not had to deal with this technology before. "There is a shortage of trained personnel who can effectively develop, maintain and enhance parallel applications. Even the experts are not very productive in these tasks," says Cox. Furthermore, software developers seek to create parallel multiprocessor applications that can run efficiently on all of the multiprocessor platforms.

"Creating an effective parallel version for even one platform is extremely difficult," says Martin Schultz, founder of Scientific Computing Associates Inc. (SCAI), an organization that pioneers the commercial use of parallel and distributed computing. "Creating parallel versions for multiple platforms is a daunting task that few developers are prepared to face." Which leads us back to the reliance of researchers, engineers and scientists on the select few technology experts and specialist programmers. In what has become a significant bottleneck, the productivity, efficiency and sheer availability of parallel application developers has threatened to stall the do-it-yourself supercomputing movement.

Auspiciously, a new parallel programming approach resolves this dilemma and boosts the viability of commodity multiprocessor clusters. It's based on the generalization of the current fashion of processor virtualization, Schultz indicates, but goes one step further. Instead of the virtualization of a single CPU, this new approach considers the virtualization of entire multiprocessor systems. "As an alternative to using different parallel programming approaches to different architectures and being forced to maintain and enhance four 'parallelized' versions of each application, the developer only has to create a parallel code for one architecture," explains Schultz. "In other words, we create a single parallel source code, and when we wish to retarget the application to different parallel architectures, we recompile the code with a compiler for that architecture." Representing this new approach to parallel programming, SCAI's Linda Virtual Machine software suite (LindaVM) is widely used by scientists, researchers and non-programming professionals for parallel application development and processing. HP and SCAI are collaborating to deliver LindaVM tools on HP cluster platform solutions. The result is cost-effective supercomputing with greatly simplified parallel programming, making it possible for most developers, scientists and researchers to succeed, regardless of budget or programming expertise.

"Industry standard clusters with SCAI software tools provide an excellent alternative to difficult and time-consuming message-passing programming for parallel and distributed computing projects," Schultz continues. "They save both programming time and hardware costs when working with applications that previously were designated only for supercomputer class machines. Developers and scientists are able to focus on their code and computational solutions — not expensive and complex systems integration efforts."

"Smaller entities, enterprises and research organizations now have access to supercomputing capabilities for the first time," says Cox. "Bringing supercomputing to the masses could very well lead to an explosion of research, learning and knowledge in much the same way that personal computing empowered the everyday productivity worker and revolutionized the ways in which we utilize and benefit from IT."

To download a white paper on the HP Unified Cluster Portfolio, visit: www.hp.com/go/transformHPC

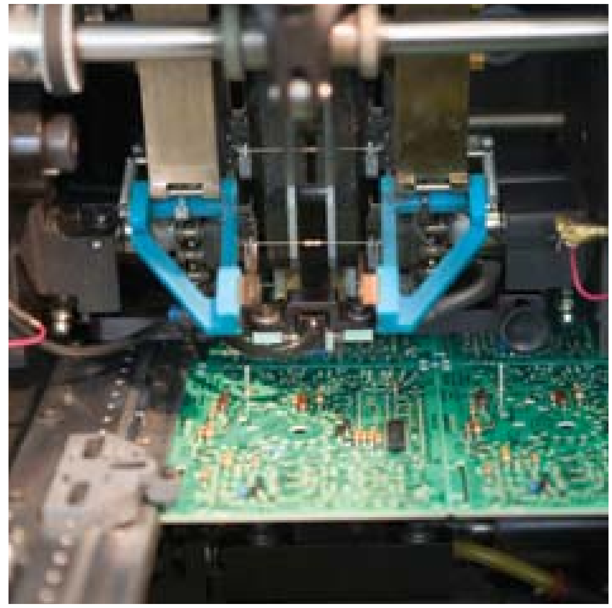

Boosting EDA with Linux®

High-performance computing clusters can enhance EDA efficiency, reduce costs and improve time-to-market.

It's an unrelenting trend. The complexity of modern electronics continues to grow, while their size continues to shrink. This trend is advantageous for the consumers and businesses that utilize the electronics, of course, but proves increasingly problematic for the engineering teams that design and produce them. Not only do these teams need sophisticated Electronic Design Automation (EDA) software with sufficiently high computing power, but they also face the pressures of making better products, improving time-to-market and reducing operations costs.

"As design complexity intensifies and device geometries become smaller, verification becomes a big hurdle for getting quality products to market," notes Brian Lowe, EDA alliance manager for HP. "Development costs have surged because simulation cycles have increased tenfold. And functional errors at the system level are the leading cause of design revisions, which can cost millions in lost profits. Something has to give," he says. Lowe suggests design teams must improve existing methodologies with tools that scale across design complexity and multiple levels of abstraction. Needed are processor capabilities to more efficiently run the increasingly complex design codes, without overextending development budgets.

Alva Barney, an EDA project manager at HP, agrees. "At HP, we've found that we can improve all areas of product development, including EDA efficiency, operations costs, time-to-market and even product quality, by enhancing the underlying hardware environment," Barney claims.

According to Lowe, the old standard of proprietary, expensive mainframe hardware systems is now being replaced by open source, high-performance computing clusters. These Linux-based clusters, typically utilizing AMD Opteron™ or Intel® Xeon® and Itanium® processors, deliver greater compute power at a much lower cost than proprietary, UNIX®-based systems.

As Barney points out, the advantages of high-performance computing clusters extend beyond hardware affordability. "Linux-based clusters often deliver a 2X performance increase, which means we can conduct twice as many tests and simulations in the same amount of time," he explains. "This allows us to improve our design efficiency, boost time-to-market and ultimately create better products. In addition, conducting more tests in a given time period with the same number of EDA licenses reduces software licensing costs and the overall expenditure on product development and verification."

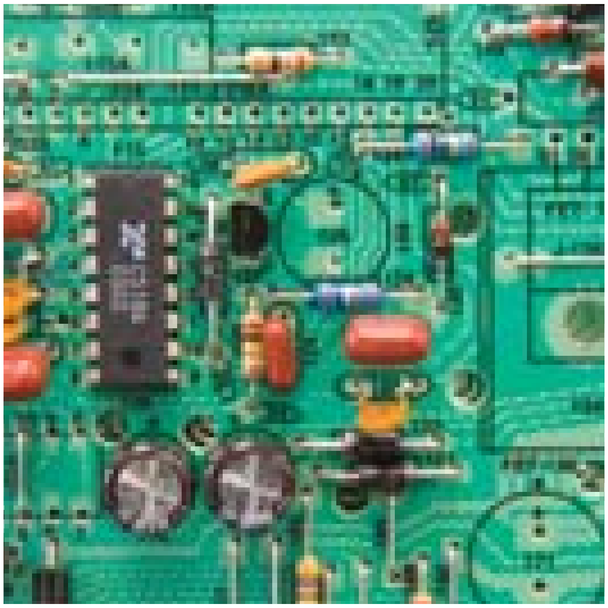

“By utilizing current HP workstations and servers with either Intel® or AMD processors, design teams can run their 32-bit and 64-bit applications with ease, reducing run times from weeks to hours.”

Linux provides a stable, open operating system environment, Lowe adds, which meets the requirements of global engineering teams that must share design configurations over far reaching geographies. Linux also offers significant deployment flexibility, with the ability to mix and match different generations of technology, including 32-bit and 64-bit nodes running on x86-based systems as well as 64-bit EPIC-based systems. “Linux accommodates both old and emerging technologies. Design teams working on multi-year projects don't have to worry about scrapping their existing systems. And EDA applications running on Linux can support future advances in CPU technology for servers and workstations,” indicates Lowe. “It's a no-brainer. HP's Linux delivers mainframe-like performance for a fraction of the cost of proprietary systems, and will likely meet advanced EDA requirements well into the future.”

While the decision to move to a Linux-based server and workstation environment may be a 'no-brainer,' there are notable considerations for engineering teams working with EDA applications, warns Norm Reini, an EDA segment manager at HP. “We have reached the physical limitations of clock speed on chips, so the industry is very clearly moving toward multi-core systems to improve performance,” Reini explains. “In order for design teams to get the most from their EDA applications, their software must work seamlessly with their hardware. If they employ single-threaded applications on multi-core computing clusters, they leave performance and money on the table.”

Reini advises organizations and design teams to consider systems with multi-core processors, and seek software vendors that are developing multi-threaded EDA applications. “By utilizing current HP workstations and servers with either Intel or AMD processors, design teams can run their 32-bit and 64-bit applications with ease, reducing run times from weeks to hours,” he suggests. “In addition, using the low profile, four-core processors not only saves valuable data center real estate, but also lowers licensing costs by reducing the number of CPUs employed.”

HP offers a range of computing platforms for deploying high-performance computing clusters. They include HP Integrity and HP ProLiant servers, blade systems, and operating environments such as Linux, HP-UX and Microsoft® Windows®. The company also provides a full suite of integration and consulting services to improve rack management, wiring, power distribution and cooling conditions for a variety of environments.

“Experience counts,” says Barney. “We utilize these systems every day for our own product development. I cannot recommend a better way to improve EDA productivity and keep costs down than by deploying Linux-based clusters with the latest processors. They certainly work for us, and we can help them work for others.”

For more information on HP’s EDA solutions and moving beyond 32-bit EDA Linux computing, visit: www.hp.com/go/transformHPC

Fast Storage: the forgotten ingredient of high-performance computing

Just as computing clusters replicated the performance, scalability and manageability of supercomputers at a much lower cost, so too can storage clusters.

What do digital content creators, governments, design automation enterprises, bioscience institutions, oil exploration companies, research entities and engineering organizations have in common? All require high-performance computing (HPC) to be successful.

Thankfully, the cost of high-performance computing has come down significantly in recent years, with computer clusters providing the performance and scalability once relegated to powerful—and tremendously expensive—supercomputers.

The picture is not all rosy, however.

Improvements in computer processing have resulted in the need to boost storage performance and capacity. The traditional barriers relating to processing speed have been replaced by barriers in storage speed. “Organizations often forget about the need for fast storage, but immense calculations per second don't mean a thing if storage and input/output (I/O) can't keep up,” says Kent Koeninger, product and technology marketing manager for HP's High-Performance Computing Division. “This has resulted in some very real problems for those seeking efficient, scalable and cost-effective high-performance computing.”

When processing speeds outpace the access and transmission of stored data, clusters can be hard to use. The result is an imbalanced system with poor throughput, explains Koeninger. This creates common file corruption problems, inefficient system performance, frequent application rewrites and under utilization of system resources.

To confront these challenges, many organizations add servers that utilize Network File System (NFS) protocols. According to Koeninger, however, NFS-based servers can present speed, reliability and complexity issues. “NFS is a partial solution to a bigger problem.” he explains. “It is an antiquated protocol that lacks the necessary caching validation to prevent file corruption when multiple nodes attempt to read and write the same file in real time.”

In addition, organizations are forced to add NFS servers over time to scale the system, and their loose connectivity and disjointed management introduce an inordinate amount of complexity. “The more NFS servers you add, the larger the headache over time,” quips Koeninger.

Fortunately, new high-performance computing storage architectures have been introduced that address the need for fast system storage and work around or bypass NFS' traditional shortcomings. Just as computing clusters replicated the performance, scalability and manageability of supercomputers at a much lower cost, so too can storage clusters, suggests Koeninger.

By clustering inexpensive storage devices that act as a single server with a single, scalable file management system, organizations can achieve exceptional I/O performance and reliability while easing system administration and reducing costs.

Furthermore, these high-performance computing storage architectures scale from 100 megabytes per second to more than 30 gigabytes per second, and from a few terabytes to multiple petabytes in size, all with a single file system that can be accessed by any size cluster reaching thousands of cluster nodes.

“By utilizing lots of small, affordable storage devices in a cluster with a common management platform, companies can realize outstanding I/O

“If a company is operating more than one NFS server or if its current NFS server is proving insufficient, there is a strong likelihood it could benefit from a high-performance storage cluster.”

performance and unlimited scalability without the file corruption, slow throughput and other limitations of deploying multiple small NFS servers,” notes Koeninger. “Better yet, there are a variety of these solutions that address almost every need and budget.” Sometimes scaling NFS is the best solution and sometimes newer technologies deliver features not possible with NFS.

Indeed, HP offers a spectrum of solutions to scale storage for high-performance clusters. One can consolidate and dramatically expand existing NFS service using HP StorageWorks EFS Clustered Gateway (based on PolyServe Matrix scalable-NFS technology) or one can bypass NFS with newer, more robust, NAS protocols such as Lustre, IBRIX, or Terrascale. The newer solutions can deliver dramatically higher performance and scalability at a much lower cost than NFS-based systems.

For example, the HP StorageWorks Scalable File Share (HP SFS) is a powerful file server that is easy to use and administer. HP SFS provides excellent price/performance across many terabytes of highly reliable, high-bandwidth storage. It uses a proven open-source protocol called “Lustre” that eliminates the file corruption and performance issues inherent in the NFS protocols and scales to more than ten times the bandwidth of even the newer PolyServe-based scalable-NFS solutions.

Koeninger notes that HP equipment is compatible with a wide range of storage solutions from leading software vendors, including IBRIX and Terrascale, which offer a different set of advanced proprietary features than those found in the more open NFS and Lustre protocols. These include faster transaction rates (IOPs for small IOs) and scalable metadata rates (GETATTR, READDIR, etc.). They include standard NFS and CIFS support in addition to their proprietary protocols.

HP provides consulting services to help identify and resolve processing/storage balance issues within an HPC system. With storage experts in each industry segment, full evaluation and installation services and seamless integration with common file sharing and management software platforms such as HP Clustered Gateway (PolyServe scalable NFS), HP SFS (Lustre) IBRIX and Terrascale, HP works with customers to improve the storage capabilities of their HPC systems.

“If a company is operating more than one NFS server or if its current NFS server is proving insufficient, there is a strong likelihood it could benefit from a high-performance storage cluster,” claims Koeninger.

For a white paper and solutions brief on HP StorageWorks Scalable File Share, visit: www.hp.com/go/transformHPC


Keeping up with new customers.
Kimberly-Clark Corporation is home to some of the world's most trusted brands—like HUGGIES® and KLEENEX.® With operations in 43 countries, the company needed to simplify its operations. HP helped them implement a standardized IT environment that streamlined the move into new markets and sped the integration of disparate technology into their operations. Today, Kimberly-Clark enjoys a 41% lower TCO of IT, and is able to stay ahead of its customers. www.hp.com/adapt
Solutions for the adaptive enterprise.
hp
invent
change
hp
©2004, 2005, 2006 Hewlett-Packard Development Company, L.P. Other company, product and/or service names and symbols mentioned herein are trademarks and are the property of their respective owners.