IDC
Analyze the Future

WHITE PAPER

Meeting Mission-Critical Computing Requirements with Blade Servers

Sponsored by: HP

Jean S. Bozman Matthew Eastwood

February 2009

EXECUTIVE SUMMARY

Blades are the fastest-growing segment within the worldwide server market, and their compute density continues to grow. The increase in computing power on multicore processors, combined with virtualization software, is allowing IT managers to assign more mission-critical workloads to blade server systems. As enterprises work to break down information "silos" within the datacenter, and to improve IT flexibility, more mission-critical workloads are becoming concentrated on blades. Therefore, blades supporting such workloads would benefit from having as many reliability, availability, and serviceability (RAS) features as possible, because avoiding system downtime due to data errors and memory corruption is essential to ensuring uninterrupted access to key business applications.

Initially, blades were introduced within the enterprise datacenter because they provided a new way to deploy workloads onto IT infrastructure. The ability to add blades on an as-needed basis brought greater flexibility to IT infrastructure, and the ability to manage many blades within one chassis brought operational benefits. Now, customers deploying blades in their datacenter want to be able to mix blades that run different workloads and different operating systems. The ability to host multiple operating systems "under one roof" and the ability to manage the blades together are factors with the potential to reduce operational costs associated with IT staffing.

As datacenters work to bring a broader range of workloads onto a bladed server infrastructure, the ability to support mission-critical workloads with hardware specifically designed for business resilience will take a higher priority for blade server installations. The hardware itself, along with the software environment that runs on the hardware, will create a technology "stack" that must protect applications just as well as if they were supported on standalone tower or rack-optimized servers.

HP Integrity blades address these technical and operational issues with a platform that has been optimized to run multiple operating systems — HP-UX 11i (a Unix operating system), Microsoft Windows, Linux, HP OpenVMS, and HP NonStop — and to do so in a virtualized environment that offers management of virtualized workloads, as well as hardware and software support for high availability of the applications. This paper outlines the trends in the industry relating to blade server deployments, the challenges relating to systems management and power and cooling in that computing environment, and HP's approach to addressing those challenges.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

This paper provides an overview of the considerations for deploying and supporting mission-critical applications and a list of the key requirements for supporting mission-critical workloads. It also shows how the blade server form factor fits into customers' IT infrastructure — and how blade servers support workloads with flexible deployments. Finally, it examines the ways in which HP Integrity server blades are designed to support these types of mission-critical workloads, and it includes several "snapshots" of customer deployments of HP Integrity server blades as part of a wider blade server buildout within the enterprise.

SITUATION ANALYSIS

Mission-critical workloads reside within every organization, although the workloads that are identified as mission critical vary by type of organization, size of organization, and vertical market segment that best defines the work of the enterprise. Each enterprise needs to decide for itself which workloads absolutely must be available for access by its employee end users and end customers.

Once the nature of the mission-critical workloads is identified and established, the means by which those mission-critical workloads will be hosted and protected will vary as well. In summary, there is not only a spectrum of workloads but also a spectrum of expectations about how available the workloads must be for business to proceed. If those applications were to stop, or to slow, business results would be affected.

Until recently, to ensure high availability for applications, customers had to install or maintain scalable servers, and in many cases, they overprovisioned those servers and overpaid for their ongoing maintenance. Or, some customers who were consolidating datacenters found that there was literally not enough room in their new datacenter for all of the servers that had been running mission-critical workloads for many years. Today, customers have more choices about the form factor of the systems they acquire, the hardware architectures for those systems, and the operating systems that run on those systems.

Business Benefits of Blade Servers

IDC has been tracking the blade server market since the 1990s when early deployments tended to be used more heavily in high-performance computing (HPC) and Web-based processing, hosting stateless workloads, such as static Web pages, and refreshing them for access by client systems. In recent years, the growth rate of blade servers has accelerated, resulting in unit shipment growth at greater than 25% in 2006 and 2007 and accounting for more than 10% of all worldwide server unit shipments in 2008 — and more than 9% of worldwide factory revenue for servers. Within the datacenter, blade servers are taking on a more prominent role, as they represent increasingly higher percentages of all new servers shipped. In 2007, when 8 million servers were shipped worldwide, blade servers accounted for 10% of all unit shipments, in contrast to rack-optimized servers (58%) and non-rack-optimized, standalone servers (32%). By 2012, IDC expects that the mix will change; we forecast that 10.9 million servers will ship worldwide, with blade servers accounting for 29% of all server units shipped, in contrast to rack-optimized servers (48%) and non-rack-optimized, standalone servers (23%).

One of the key drivers for this change in the server form-factor mix is the increased flexibility blades bring to server deployments and capacity planning. Blades support IT flexibility because capacity, in the form of additional blades, is added as needed. This is, in effect, pay-as-you-go capacity expansion, driven by peaks of demand for computing resources, which can also be taken offline as needed. Multiple server blades are housed within one chassis, where they can be managed together, reducing operational costs compared with rack-optimized servers that have separate connections to storage and networking resources.

Enterprises have reported a number of business benefits related to the adoption of blade computing, including the following:

- Improved resource utilization, due to support for consolidated workloads
- Improved energy efficiency, which reduces power and cooling costs
- Reduced use of datacenter floor space, due to increased server density
- Reduction in cabling to LANs and SANs, reducing capital expenditures and unplanned downtime, as storage and networking resources are shared among the blade servers
- Better use of IT staff time, due to more efficient management of server resources

MARKET TRENDS

Blade servers have been rapidly adopted in recent years, and according to IDC supply-side data, they now account for more than 10% of all server units shipped worldwide. Growth rates for blade servers, in terms of both units and factory revenue, are in the double-digit range — and blade servers are the fastest-growing type of server, by form factor, in any given calendar quarter over the past three years.

IDC Workloads studies, based on customer surveys at 1,000+ sites each year, have found that the range of workloads (application types) running on blades has changed over the years — and has broadened to include more enterprise applications over time. IDC research shows that business processing, decision support, and collaborative workloads are increasingly deployed on blade platforms; several years ago, the mix was more weighted to IT infrastructure and Web infrastructure workloads. The business processing category includes online transaction processing, ERP, CRM, and other line-of-business (LOB) workloads; decision support includes data-based analysis and business intelligence (BI); and collaborative includes email and groupware accessible by end users throughout an organization, where email and its delivery is also considered to be mission critical. Figure 1 shows the distribution of blade server revenue and units shipped by workload.

FIGURE 1

Blade Workload Segments

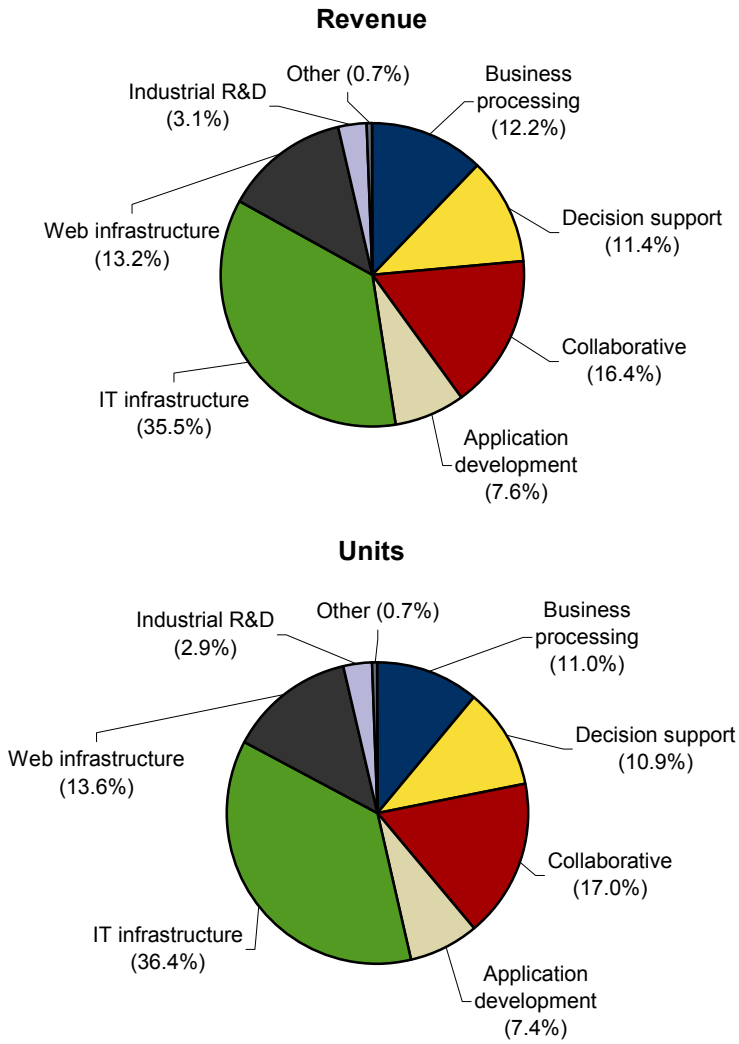

Source: IDC's Workloads Study, 2008

Business processing and decision support are the leading workloads deployed on server blades, but, as Figure 2 shows, their presence is even more pronounced on Itanium-based (EPIC) server blades. The side-by-side comparison shown in Figure 2 demonstrates that these two types of workloads represent a larger percentage of the total workloads for Itanium-based server blades than of the general population of server blades.

FIGURE 2

Workload Segments for Itanium-Based (EPIC) Servers

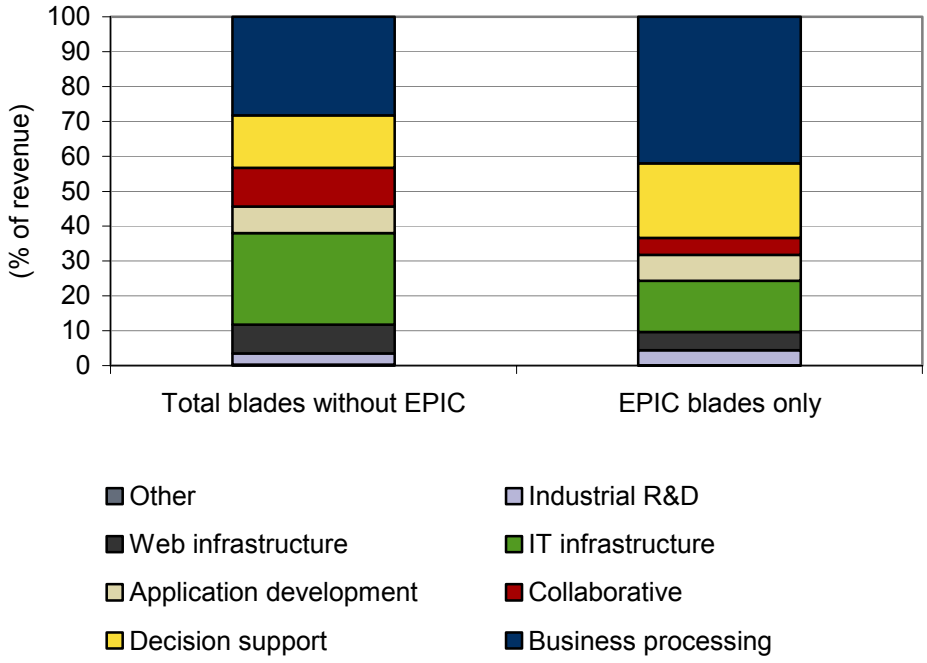

Source: IDC's Workloads Study, 2008

Mission-Critical Workloads

Some business processes would simply come to a full stop if the mission-critical workloads supporting them were interrupted for any reason. Although such outages are relatively rare, they must be avoided — and recovery should also be quick and efficient, allowing applications to restart and production data to be accessed. The ability to avoid interruptions and to continue running in spite of small errors in memory is critical to avoiding costly downtime. Historically, fault-tolerant systems and highly available systems were expensive, as were large, standalone systems in the datacenter. Now, customers have choices about how, and where, to install such capabilities — even within a scale-out infrastructure that includes rack-optimized servers or blade servers.

RAS features are all vital to support mission-critical workloads that cannot fail, or cannot be interrupted, for long periods of time. One way to look at this is that RAS features prevent small errors from becoming causes of a general failure of a system, while other aspects of the system support high availability of applications and data — even in the event of unplanned downtime due to electrical outage, network outage, or natural disaster.

Some examples, by vertical market segment, follow:

- **Telecommunications.** Telecommunications companies typically deploy rows and rows of rack-optimized servers, along with bladed servers, to support services to customers of wireless communications devices. These services include cell phone service, automated number lookup, and support for 911 dispatching. The scale-out infrastructure of a telecommunications company must be reliable because many government services (e.g., police, fire) and end customers rely upon it. Interruptions in these services are widely noted and cause disruptions to all the businesses that subscribe to them.

- **Logistics/distribution.** These companies are very aware that any interruption in computing will have a ripple effect through their distribution systems on a local, national, or global basis. Minimizing downtime is a high priority for companies within this vertical market — and supporting mission-critical workloads is at the heart of many of the companies' business processes.

- **Transportation.** For airlines and railways, computing outages translate into delays and into snarled schedules and unhappy customers. Some of the mission-critical systems support scheduled flights or train departures, while others support the mechanical systems that manage the transportation equipment.

- **Manufacturing.** Manufacturing companies look to end-to-end applications to manage supply chain interactions, as well as to ensure that process control within factories is running smoothly. Multiple business applications must be kept online, and outages must be repaired quickly, before major disruptions take place.

- **Healthcare.** Systems supporting healthcare, such as those used in hospitals and clinics, have historically been "siloed," according to the departments they support (e.g., customer records, laboratories, radiology, emergency room operations). The move toward integrated healthcare demands a more unified approach to information processing, and new IT infrastructure is being put into place that makes it easier to share data — and to store it — in an accessible way. Lack of access to important medical data hampers patient care and can delay scheduled medical procedures.

The Importance of RAS Features

RAS features provide an important foundation for mission critical workloads, such as LOB applications and databases. These features are important for environments in which downtime would stop business operations, causing revenue loss for the business and impacting profitability.

Hardware platforms that have RAS features have been optimized to support a range of mission-critical workloads, and the operating systems running on those platforms have been optimized to redirect data to alternate resources, and avoid shutdown, even if one or more hardware components presents an error or fails. Designs that support RAS features improve operational efficiency through continued support for important business processes, which in turn benefits business users.

In IDC's 2008 Workloads Study of end users, 63% of all Itanium-based servers (EPIC servers), as measured by revenue, are running business processing and decision support workloads, as shown in Figure 2. These two workload categories support a high percentage of mission-critical workloads, such as LOB applications and business analytics. In comparison, the share for the overall server market, in supporting these types of workloads on all platforms, is about 20% lower, at 43% of server revenue.

Further, IDC's 2007 Server Virtualization Survey found that more than half of all virtualized servers are deployed in support of business processing, database, and decision support workloads. In addition, a higher percentage of blades are virtualized versus other server form factors (rack-optimized and non-rack-optimized, standalone servers). In short, for many virtualization projects, blades are a preferred platform for hosting, and managing, virtualized applications.

Although Itanium-based blades account for a small percentage of all blade servers worldwide currently, IDC supply-side data shows that adoption of Itanium-based blades has been ramping quickly in terms of unit shipments and revenue on a worldwide basis. Itanium-based server blades grew from less than half a percent of total market share in 2005 to nearly 3% in 2008. This is a compound annual growth rate (CAGR) of 81%, when considering growth from 2005 to 2008.

Overall, Itanium-based servers, including all form factors, have been growing in the marketplace in recent years, hosting more than 13,000 applications on the Itanium-based (EPIC architecture) platform. Itanium-based servers grew 31% in 2007, in terms of factory revenue, while unit shipments grew 36% year over year. In comparison, the worldwide server market grew 3.3% in terms of factory revenue in 2007, while unit shipments grew 6.6%.

Blade Server Usage Models

Blade server deployments accomplish several important goals for IT managers, including the following IT deployment benefits:

- They improve flexibility, with the ability to add capacity on a pay-as-you-go, as-needed basis. The ability to mix and match resources, by adding different types of blades (e.g., different processor types, different operating systems) within a blade server chassis housing all of the blades, helps to break down the IT silos that were often seen in traditional datacenter deployments of standalone (pedestal-based) servers and arrays of rack-optimized servers.

- They support unified management of multiple server blades because multiple blades are housed within the same chassis — and a holistic management system can be applied to all blades in the chassis. Granularity of control over individual server blades and the workloads running on them (via virtualization) is another aspect of improved management through blade deployments.

- They share outside resources, including access to storage and to network devices. The ability to attach cabling to the chassis, rather than to each computer, reduces wired cabling, which would otherwise need to be connected to each individual server. This aspect of bladed server deployments leads to the buildout of a "fabric" of interconnected server devices and storage devices.

- Their deployment often coincides with virtualization/consolidation projects that gather up multiple workloads to run on fewer server "footprints" within the datacenter. This trend is seen in IDC customer-based research that shows blades are often a catalyst for IT transformation projects within the datacenter.

Business benefits associated with blade server deployments include the following:

- When deployment of blade servers is compared with deployment of rack-optimized servers, a reduction in operational expenditures (opex) associated with more efficient management of server blades can be realized. Opex has been climbing rapidly in recent years, largely due to rocketing "soft costs" involving maintenance and management and IT staff time, along with rising energy costs associated with power and cooling.

- Power and cooling costs can be further reduced through the consolidation process. This is aided by the lower power and cooling requirements of bladed servers, which are compute-dense systems housed within a blade chassis framework.

- Server blades require less floor space than other form factors, addressing budgetary limits on building new datacenters in the economic downturn and limited space to add capacity within existing datacenters.

Business Value of Blade Deployments

More than 10% of servers sold today are blade servers. IDC believes this will be a fast-growing market segment for the next five years and expects that by 2012, more than 25% of all server shipments will be blade servers. The increasing adoption of this dense server technology has been driven by datacenter managers looking to maximize real estate and improve physical space limitations. The improvements in blade technology over the past two to three years have also helped improve customer acceptance. While initial generations of blade servers were limited in terms of memory capabilities and were identified as creating hotspots in datacenters, more recent introductions of blade servers have larger memory footprints and supplemental cooling advancements.

Businesses look to bladed servers as a way to approach the server sprawl that has been driving operational costs higher in recent years. By provisioning workloads onto compute-dense form factors, such as server blades, businesses can preserve datacenter floor space — and contain power and cooling costs within the limits imposed by local power utilities.

Operational costs are rising in the datacenter; they are driven by increasing costs for power/cooling of server infrastructure, management of workloads across the infrastructure, downtime caused by system outages, and IT staff costs related to system administration and management. Technologies and products that address this important need to reduce operational costs will contribute to company revenue and profitability.

As shown in Figure 3, the top-of-mind challenges for datacenter managers — nearly all of them related to operational costs — are:

- **Power/cooling.** Reducing power/cooling costs through improved energy efficiency has emerged as the top priority for many datacenter managers. All datacenter servers, storage, and networking devices must be powered within a certain "power envelope" — and "power capping" has become common in the largest cities around the world. Importantly, the use of closely packed servers within the datacenter leads to "hotspots" that must be controlled, meaning that datacenter servers must also operate within a "thermal envelope" — or operational costs will rise even further.

- **Availability of applications and data.** Having highly available servers is critical in a business that must operate on a 24 x 7 x 365 basis. Data and applications must be accessed on a 24-hour basis, and that access must be ensured from any time zone around the world. If unplanned downtime occurs, workloads must be transferred to other servers for continued processing, using high-availability software to restart the processing.

- **Space utilization within the datacenter.** Crowding in the datacenter contributes to power/cooling problems — and customers are looking for new ways to deploy servers to use the existing datacenter space more efficiently while supporting more energy-efficient operations. The use of multicore processors and virtualization software has greatly contributed to the ability to consolidate workloads in the datacenter. IDC expects this trend to continue during the economic downturn, resulting in more cost-efficient use of available datacenter space.

FIGURE 3

Top Challenges in the Datacenter

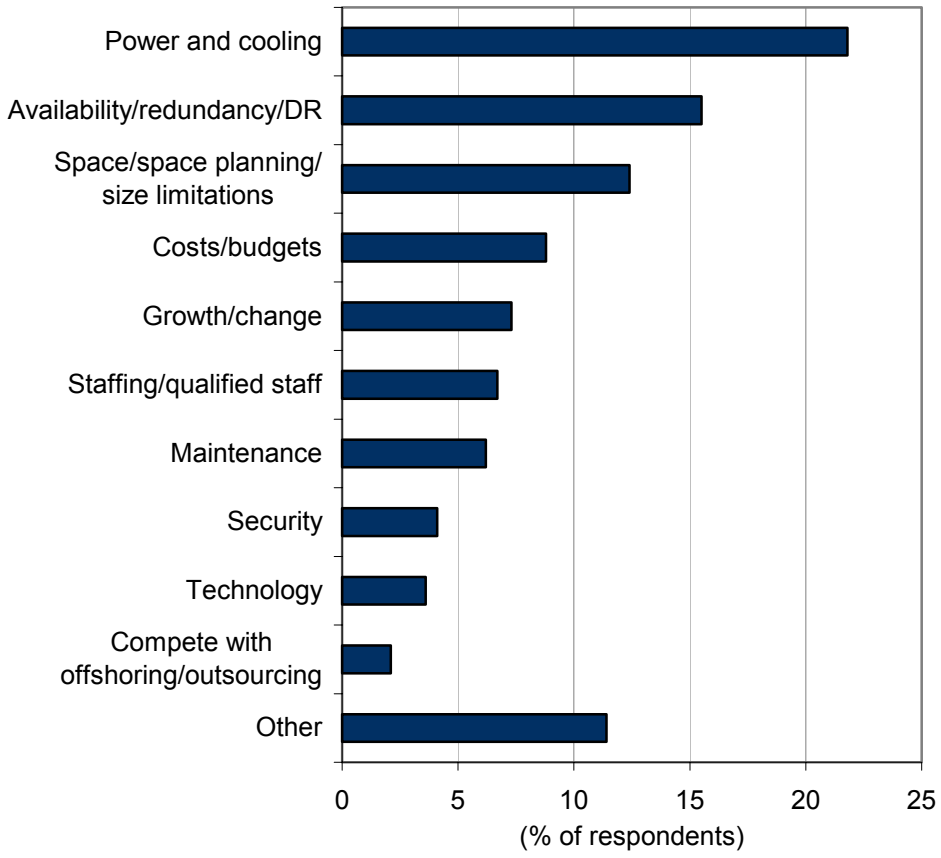

Source: IDC, 2008

Importantly, blade servers support business value in a time of economic constraints. They do so in several ways: by reducing power and cooling costs, reducing management costs, and reducing IT staff time — all of which contribute to improved operational costs. The chief contributors to improved operational costs include IT staff costs, power and cooling costs, and the costs associated with building out and operating a datacenter. Blade deployments clearly address the power and cooling costs and space-related costs within the datacenter, due to their compact designs, including features that minimize power and cooling requirements for each individual server blade. Reducing IT staff costs is just as important as reducing other types of operational costs. Staff time is expensive, and training of IT staff is also expensive. The ability to leverage fewer IT staff to control many more servers — that is, improving the ratio of system administrators to total number of servers managed — is a primary way to reduce opex within a constrained IT budget.

Combining blades with virtualization technology makes the blade deployments a more powerful platform for workload consolidation. These virtualization tools allow multiple virtual servers to coexist on a single physical server and have been widely adopted as a means to reduce the sprawl of physical servers within a datacenter. IDC research shows that customers report average ratios of virtual machines to physical machines of 4:1, and increasingly, many customers are reporting 10:1 ratios.

HP'S INTEGRITY SERVER BLADE SOLUTIONS

This section provides technical specifications that are specific to the HP Integrity blades product line, including HP Virtual Connect for server I/O and HP Thermal Logic for power and cooling management, and it describes the software, including the HP Virtual Server Environment (VSE), for virtualizing HP's Integrity server blades.

HP offers two Integrity server blades, as follows:

- **BL870c:** The HP Integrity BL870c server blade is a four-socket, full-height server blade that supports three different models within the Intel Itanium 9100 series processor line, and it is supported by up to 192GB memory (via 24 DIMM slots). The BL870c is designed to address server and application consolidation, virtualization, and demanding mission-critical workloads.
- **BL860c:** The HP Integrity BL860c server blade is a two-socket, full-height blade server that supports different models within the Intel Itanium 9100 series of processors, and it is supported by up to 48GB memory (via 12 DIMM slots). The BL860c is a versatile platform designed to be used with mission-critical application tier workloads and transactional workloads.

Both the HP Integrity BL870c and BL860c server blades coexist with HP ProLiant server blades and StorageWorks storage blades within the HP c-Class enclosure.

Supporting Mission-Critical Workloads

As platforms for mission-critical workloads, Integrity blades can be viewed as a catalyst for reducing server footprints within the datacenter. While they support mission-critical workloads, they are often deployed in a new and different way, compared with rack-optimized servers and standalone scalable servers. They are often deployed to support workload consolidation projects, leveraging virtualization technology to reduce server footprints within the datacenter. The business objectives for these kinds of projects include reduction in power and cooling costs and reduction in IT maintenance costs associated with management of large groups of rack-optimized servers — each of which houses a separate link to storage and network resources.

The HP Integrity BL870c server blade, which is based on Intel Itanium 9100 series processors, supports demanding, mission-critical solutions that would otherwise run on multiple rack-optimized servers. Based on dual-core Itanium processors, these systems are shipped with up to 192GB of memory and the high-performance HP zx2 chipset, including RAS features of the HP Integrity server blades (on a system level), Itanium features (on the processor level), error correcting code (ECC), and memory protection — including cache memory protection and support of I/O and data transfer.

HP Integrity Blade Technologies

- **Operating systems for Integrity server blades.** HP Integrity blades support a wide range of operating systems that support mission-critical workloads in the datacenter, including HP-UX 11i Unix, Windows, and Linux, along with two longstanding HP operating environments for the datacenter: OpenVMS and HP NonStop Kernel (NSK). This variety of operating systems allows customers to bring forward mission-critical applications that already run in these environments — and to consolidate them on HP Integrity blades so that they can be managed together within HP's BladeSystem c7000 and c3000 blade chassis enclosures.

- **Virtualization technologies.** IT managers and system administrators have a number of options when it comes to supporting virtualization on Integrity systems, including Integrity blades. Integrity blades can exist within HP VSE, which is optimized for virtualizing mission-critical applications. HP VSE includes HP Integrity Virtual Machines that can support virtual machines running HP-UX 11i, as well as Microsoft Windows or Linux distributions (Red Hat or Novell SUSE). It is important to note that HP VSE supports multiple types of virtual machines, including Windows Virtual Machines and Linux Virtual Machines. Finally, HP Integrity systems running the NonStop fault-tolerant operating environment were announced, and shipped, in 2008 (see HP Integrity NonStop BladeSystem sidebar on page 13).

- **Enhanced virtualization software for HP VSE.** In January 2009 HP announced VSE 4.1, which introduced expanded capabilities to bring the flexibility of virtualization to physical Integrity blade servers. It supports the creation and management of "logical server" images, or templates, that can be easily moved across physical server blades. The profile includes information about the number of cores supported by the logical server, the amount of memory, the MAC addresses for the network connections, and storage information for server HBAs and the SAN array controller. Logical servers can be deactivated and saved offline, or created and stored as templates, and then can be quickly activated when needed. They can also be used to support "planned downtime" in which logical servers are moved away from a physical server that is scheduled for repair, software upgrade, or replacement — and then moved back to the new or repaired server, once the maintenance is completed.

- **Managing the virtualized environment.** Within a VSE, the virtualized environment running on an HP Integrity Blade can be monitored and managed with the same tools that are used for all Integrity servers. The VSE management software provides a unified view of all physical (P) servers and virtual (V) servers running on a virtualized blade. This means that system administrators are able to allocate resources and balance workloads as needed. Importantly, this VSE management software provides a unified point of control that gives administrators a 360-degree view of all workloads and their status. If some workloads need more computing resources, additional resources can be allocated to the workload, or the workload can be assigned to one or more blades, within the blade server chassis. A graphical user interface makes it easier to identify which workloads should be shifted to other hardware resources — and speeds redeployment of the logical server, as a "template" that can be moved as a drag-and-drop object, to a new position within the bladed server infrastructure.

- **Virtual Connect.** The HP Virtual Connect architecture provides enhancements in server I/O that improve overall aggregate bandwidth into and out of a bladed server environment supporting virtualization. HP Virtual Connect Flex-10 supports a high-speed "backplane" for a wide range of I/O types, including 10 Gigabit Ethernet (GbE), Fibre Channel, and InfiniBand I/O links. This Virtual Connect fabric results in the deployment of fewer extra NICs, switches, and cables — reducing IT-related operational costs as well as capital costs. HP Virtual Connect Enterprise Manager (VCEM) is used to simplify the management of bladed server systems, managing server-to-network connections across the datacenter. Due to the modular design of server blades, Virtual Connect can be used to quickly reprovision server blades within the blade chassis, through the use of workload-balancing software and high-availability software. P servers and V servers running on the blades can also be recovered by leveraging Virtual Connect connections, supported by VSE software (see earlier reference) that allows administrators to rapidly deploy "logical servers" from one physical server blade to another. When combined, these technologies support highly virtualized IT infrastructure and, as such, are foundational technologies for next-generation datacenters leveraging bladed server designs.

HP Integrity NonStop BladeSystem

HP Integrity NonStop blades, supporting the NonStop Kernel (NSK) fault-tolerant operating system, were introduced in 2008, delivering high levels of reliability and supporting Availability Level 4 (AL 4), which is the highest level on IDC's four-point Availability Level Spectrum. (AL4 denotes support for continuous [fault-tolerant] processing.) When NonStop software is deployed in the server blade form factor, it provides near-linear scalability, allowing workloads to scale up by adding blades on an as-needed basis. By shipping NSK on blades, HP reduced the point of entry to this style of computing, which is generally found on midrange or high-end servers based on Integrity server technology.

The idea driving the introduction is that HP has been fostering a "blade everything" approach to next-generation datacenter infrastructure. That is, it is bringing a wide array of workloads onto the same bladed server infrastructure, wherever possible, for the purposes of more flexible deployment and more unified management than was possible when different types of workloads ran within separate "information silos" throughout the Web, application, and database "tiers" of the datacenter.

Integrity NonStop blades can be used to consolidate workloads that have been running elsewhere within the datacenter — or they can be "mixed in" with other types of blades, including HP blades based on x86 architecture (Intel or AMD x86 microprocessors), so that all workloads can be managed from a single control point. Examples would include deploying a database on the NonStop blade, which is then accessed by requests running on x86 server blades — or running business intelligence workloads on the NonStop blades to analyze transactional data generated by other servers in the datacenter, including other blades in the blade server chassis.

Telecommunications companies, including companies that provide wireless cell phone service, can leverage Integrity NonStop blades to reduce the "footprint" that they currently have in-house, compared with older HP NonStop systems they have installed over the years. HP ships carrier-grade NonStop Integrity blades that conform to the NEBS standard for telecommunications carrier-grade bladed infrastructure. Financial services companies currently running their ATM, POS, or other transactional applications on older NonStop models will also be able to move their applications onto the new Integrity NonStop blades, allowing them to leverage the higher performance and smaller "footprint" of the Integrity NonStop blades. Other types of companies will use Integrity NonStop blades to move data-gathering systems supporting decision support. Often, they are deployed in retail companies or financial services companies that collect transactional data and then want to see patterns in the data through the use of business intelligence software.

- **HP Solution Blocks.** HP BladeSystem Solution Blocks are a set of example configurations, based on HP best practices or testing, for a range of enterprise workloads, including SAP ERP, Oracle PeopleSoft, BEA's SOA solution, Siemens PLM Software, and IBM WebSphere for Application Integration. Several configurations are available for each solution (e.g., small, medium-sized, and large [or full feature]). The intent is to provide customers with a solid starting point upon which specific customizations can be made depending on specific customer requirements. The value to the end user is in the cost savings associated with enclosure optimization and the speed of deployment, including related storage components.

- **Management software.** HP Integrity blades are managed by a combination of HP Systems Insight Manager (SIM) and, on an enterprise datacenter level, the HP Business Technology Optimization (BTO) suite of IT management products (the next generation of HP's OpenView enterprise management suite). This combined approach simplifies overall management challenges related to managing the hardware itself, the blade systems within the chassis, and multiple computing systems across the datacenter. At the same time, both physical and virtual servers are monitored and managed, providing a unified view of all computing components.

- **Addressing energy efficiency and flexible deployments.** Integrity blades now have access to several HP technologies that were originally developed for deployments of x86 server blades within the HP BladeSystem chassis. These technologies, including HP Thermal Logic and HP Virtual Connect, are being leveraged to provide a consistent set of products and technologies across bladed server infrastructure that includes blades based on x86 processors and those based on Itanium processors. In addition, both the BL870c and BL860c Integrity blades support the Integrated Lights-Out 2 (iLO2) Advanced Pack, which is designed to provide automated remote server management as a standard feature of the server solution.

- **Services for blade server deployments.** HP offers a variety of services to companies that are evaluating blade servers for future deployments. This is important because they include HP startup and deployment services, assessment and design services, deployment management services, and education services. HP also offers support for mission-critical workloads on a 24 x 7 x 365 basis through its Proactive 24 service and its Support Plus 24 service, which is customized for specific workloads.

CUSTOMER USAGE SCENARIOS

The following three examples offer "snapshots" of customers who have deployed HP Integrity server blades in their enterprises to support mission-critical workloads. In each case, the reason for the deployment, the types of workloads supported, and the specific computing environments (e.g., operating systems and applications) that run on the bladed server systems have been identified:

- **Manufacturing company in Italy.** Dierre SpA, an Italian manufacturer of doors located in northern Italy, used HP Integrity blades to consolidate HP OpenVMS workloads, reducing floor space use within the datacenter. Power and cooling costs and software licensing costs both were reduced by more than 30% annually. Importantly, OpenVMS applications, which this company has used for many years, were brought forward to new Integrity server blades, protecting previous investments in software and staff training.

- **Software development company in India.** iSOFT, a developer of healthcare software in India, installed server blades and storage blades at its sites in the cities of Chennai and Bangalore. Consolidation of workloads onto HP BladeSystems supporting VMware ESX Server virtualization software reduced

datacenter space requirements by 70% and reduced energy-related costs by 50%, compared with the previous rack-optimized server solutions. The company used HP Enterprise Virtual Arrays storage that was linked to HP c7000 blade chassis enclosures. By using iLO2 and HP Systems Insight Manager (SIM) software tools, in conjunction with HP Integrity blades, the company reduced IT staff costs associated with system administration by about 50%, addressing ongoing operational costs.

- **Healthcare provider in the United States.** Priority Health, located in western Michigan, sought to simplify operations of its HP-UX 11i Unix computing environment, and to make it more flexible, so that the infrastructure could be adapted in response to changing business requirements. Bladed servers allow applications to be shifted to available computing resources, and the overall bladed server environment accesses an extended storage resource based on HP XP24000 storage arrays. This change to blades preserved the overall computing environment, in terms of the software stack and storage components, but reduced the operational costs associated with the Unix server platforms.

CHALLENGES/OPPORTUNITIES

Bladed servers, like all major product types, are the subject of intense competition in the marketplace. That competition is likely to become even more focused on price, and price/performance, in the face of IT spending constraints associated with the current economic downturn. HP's bladed server offerings must compete with those of other vendors, including IBM, Dell, Sun, and Fujitsu. Each vendor must work to differentiate its offerings while continuing to enhance overall performance and usability of blade servers as a category of servers over time.

Further, all bladed servers must coexist with other types of form factors in the datacenter, including standalone servers and rack-optimized servers (which account for more than 40% of all unit shipments worldwide — about four times the percentage of servers known to be based on blade servers (now edging past 10%+ worldwide). However, IDC expects that bladed servers will become more prevalent over time, given their ability to host different types of workloads on separate blades — and yet be able to manage all server blades within each chassis in a unified way.

Several factors will improve adoption of bladed servers over time, leading to more widespread adoption:

- Support for a broader range of workloads, including high-performance computing, Web-enabled workloads; enterprise line-of-business applications; and databases

- Management tools that increase the number of servers that can be managed by a system administrator, providing ease-of-use software that offers a single-pane-of-glass view of all bladed servers under management

- Increased use of policy-based automation software

- Continuing improvements in power and cooling efficiency

These factors are included in IDC's assumptions about the forecast increase in bladed server unit shipments, as a percentage of total, by 2012, the end of the forecast period.

HP is well-positioned to provide the policy-based management tools and ease-of-use software described earlier — and to deliver server blades with energy-efficient power and cooling features and advanced RAS capabilities, as described in this paper. Due to HP's focus on bladed server environments, including HP's Virtual Connect and support for virtualization on server blades, the company is expected to remain one of the top providers for server blades worldwide. However, IDC notes that the blade server space will be a competitive market, as current blade vendors and new entrants are also expected to develop and ship bladed servers in the future.

CONCLUSION

Datacenter infrastructure is changing, becoming more flexible to support constantly evolving business requirements. Bladed servers address top-of-mind IT concerns while reducing operational costs associated with power and cooling, space utilization within the datacenter, and IT staffing costs. The process of migrating more workloads to bladed infrastructure is bringing more mission-critical workloads onto blades, increasing the IT requirements for high availability and reliability in the platform to ensure end-user and end-customer access to those important applications as well as business continuity.

HP Integrity blades are designed not only to support mission-critical workloads through the use of a range of operating systems that have traditionally supported enterprise applications and databases but also to support new workloads related to Web-enabled computing and collaborative software. Hardware that combines reliability features formerly found on standalone datacenter servers with support for the operating systems and thousands of ISV applications is now available in a blade server format. This includes HP NonStop Integrity blades supporting fault-tolerant workloads for applications that cannot be interrupted, along with HP Integrity blades supporting general-purpose business applications and databases. The ability of these blade server systems to manage a variety of workloads from a single point of control, and to move workloads from one blade to another, as needed, can be expected to support corporate initiatives to reduce both capital expenditures for blade server hardware and operational expenditures for the business organization's datacenter.