Blade Servers: Driving the standard

By Lin Nease and Craig Yamasaki

HP Blade Servers are a series of CompactPCI-based dedicated servers and related peripherals that brings blade technology to enterprise and service provider customers, and at the same time enhances the offerings to telcos. These servers use an extension of the CompactPCI specification called OpenBlade, which adds a dedicated management capability, including a port for a dedicated management LAN, as well as Fibre Channel support for access to storage. The result provides great density, lower power consumption, easier management, and better performance than previous CompactPCI-based systems. The result for customers will be a lower TCO, greater flexibility better security, and lower risk.

Hewlett-Packard's launch of the blade servers provides an important new area of growth for the technology and for the CompactPCI standard. Our servers extend the existing CompactPCI standard, provide a new standards-based line of products, and allow existing users of CompactPCI technologies to leverage their hardware investment and expertise while growing their capabilities. Soon, we expect to see additional vendors joining to create a wealth of new products meeting the CompactPCI standard.

The new HP Blade Servers extend open standards-based CompactPCI blade server technology to new markets, including service providers, corporate and enterprise data centers, in addition to the more traditional carrier-grade markets. Because of the widely varying needs of these new classes of users, and because of the demand for increased performance among traditional users, HP's Blade Servers are available to meet a wide variety of requirements, and to meet several levels of performance demands.

Especially exciting are the new types of CompactPCI blade servers that are now available. These include:

- Complete Web hosting appliance blades
- General-purpose server blades
- Storage blades
- Ethernet switch blades and Fibre Channel arbitrated loop and switch blades
- Expansion I/O blades
- Management blades

In effect, using HP Blade Servers will allow service providers or enterprise customers to bring new levels of density, flexibility, serviceability, manageability, and efficiency to their infrastructure, while gaining the security and technology innovation of a standards-based computing environment.

Edge and enterprise customers typically need high performance computing and storage services that can be accessed via the Internet or through an internal enterprise network or intranet. Servers are expected to provide Web pages and related data, as well as database access and application services on demand. The traffic demands on these servers can be quite heavy, and as Web and intranet usage becomes more mission-critical, the demand for data delivery and storage grows accordingly.

Currently, edge, enterprise, and commercial users are faced with choosing between 1U dedicated servers or proprietary blade servers. The 1U dedicated servers are relatively inexpensive and provide good redundancy, can be added to the data center as needed and use familiar technology. However, they are complex to manage, consume significant amounts of power and demand some physical access for troubleshooting, upgrading and maintenance.

Some early blade server offerings could arguably make better use of floor space in the data center than 1U servers, but they can still be complex to manage, and have the downside of being based on proprietary technology. This means that customers must make the decision to buy into a specific company for a significant amount of their computing needs. Proprietary technology also limits upgrade possibilities to only those offered by their vendor. It eliminates the ability to add new capabilities from other vendors without adding a different proprietary solution and increases the chance of incompatibilities. Many competitors are designing proprietary, ultra dense servers, but because of this density their servers are thermally limited to under-powered processors and limited memory.

The HP Blade Servers offer a third alternative. By using the existing CompactPCI standard, plus some extensions, users can gain the density benefit of blade servers, while maintaining an industry-standard environment. Users of HP Blade Servers can realize a gain of twice the computing resources for the same space requirements when compared against a similar volume of 1U servers. In addition, these users will find that their demand for power and cooling is significantly less than a rack full of 1U servers.

The OpenBlade specification

CompactPCI has historically addressed the needs of the communications and control equipment world. In this value chain, equipment providers are the solution integrators. They are essentially designing hardware out of CompactPCI suppliers' components.

OpenBlade is a specification based on CompactPCI (both present and future), but augments the specifications of CompactPCI in order to address the critical needs of enterprise end-user customers. Enterprise customers need to use blades in a manner that is very different from equipment-provider OEM customers.

The OpenBlade specification was designed to enable rapid solution development for customers. It is HP's contention that blades offer a unique new economy in system integration – an economy that can only be unleashed with an interoperable environment of complimentary products. Only after whole solutions can be built out of interoperable blades will the promise of blades be fulfilled.

However, in the enterprise computing world, solution integration is performed by IT personnel or systems integrators

Reprinted from CompactPCI Systems / April 2002

who buy network computing products off-the-shelf. The demand to mix and match vendors' offerings has created the need for an extremely high level of interoperability. Ethernet equipment is an example of what a low-touch integrator or end-user customer needs in terms of interoperability, because Ethernet products of all types interoperate reliably. Likewise, if blades don't possess a similar level of interoperability and ubiquity, their advantages in manageability, cost, and density will be meaningless.

However, CompactPCI addresses a range of issues that are not applicable to blades. For example, the very specification CompactPCI leverages, PCI, is a standard for I/O adapters. The multitude of issues associated with I/O bus layouts and signaling are not central issues for blades. Rather, the central issue for blades is access to heterogeneous solutions. These solutions include server blades (known in CompactPCI as CPU blades), storage blades, network switch blades, management blades, appliance blades, and many more. Thus, OpenBlade is really about Ethernet and Fibre Channel – not PCI – to enable interoperability and ubiquity.

OpenBlade augments CompactPCI by enabling:

- Interoperability akin to the Ethernet world
- Solution integration from commercial building blocks, rather than OEM components
- Enterprise data center capabilities such as Fibre channel and management
- Strict focus on network-connected blades, rather than I/O adapters

The changes to the CompactPCI specification required to become OpenBlade compliant are minimal. Starting from CompactPCI 2.16, OpenBlade's management and storage pinouts are derived from the H.110 TDM bus pin-out, but the pins assigned to -48VDC signaling have been changed to ground. The remaining pins, which are used for the two Fibre Channel loops and the management LAN use previously unpopulated pins.

The diagram in Figure 1 shows that the changes to the existing CompactPCI pinout specification take place exclusively in the P4 / J4 combination. The detailed pinout chart in Figure 2 shows the changes in the existing specification, and the pin locations for the new communications requirements.

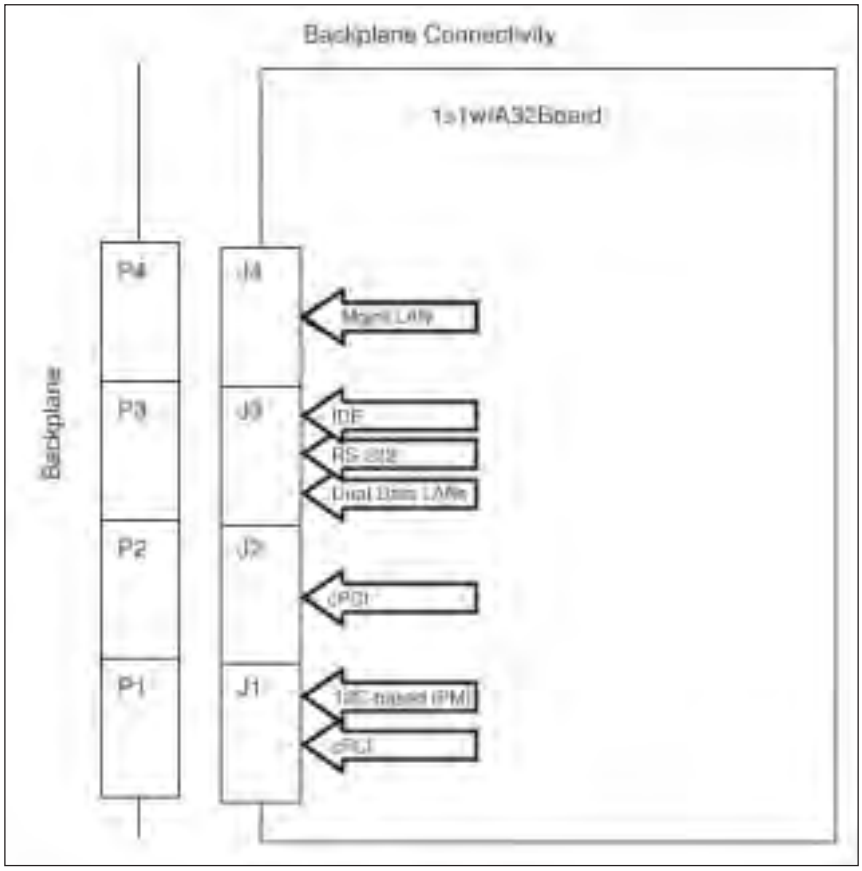

Figure 1

The changes to CompactPCI for the proposed OpenBlade specification allow interoperability between blades meeting that specification and today's CompactPCI specification. Because this specification is well-known and thoroughly debugged, it opens the door for new manufacturers to create new CompactPCI/OpenBlade products. The result is more flexibility and greater opportunity for both manufacturers and their customers. Integrators will be more likely to use OpenBlade and CompactPCI products because they can be assured that their solutions will work and be widely supported. Customers can be assured that they won't be stranded by single-vendor specifications that go away when the vendor changes its marketing plans or goes out of business.

The new capabilities of the OpenBlade specification can result in products that are more useful to a broader set of customers. For example, the management support designed into HP's Blade Servers means that customers can manage their networks centrally, rather than having large numbers of on-site technicians constantly looking at each dedicated server.

HP Blade Servers and peripherals

HP Blade Server suite includes server blades, storage blades, a management blade, an Ethernet switch blade, and optional Fibre Channel blades. Each of these blades includes all CompactPCI standard services, including power, signaling, and network connections. In addition, the extensions on HP Blade Servers include Gigabit Ethernet (defined by PICMG 2.16), dual fibre channel and an additional Ethernet management network. These servers will work in existing blade enclosures with existing CompactPCI connectors, but the extended services will not be available in enclosures that do not support the OpenBlade extensions. Likewise, existing blades will work in the HP Blade Server chassis. The details of the extensions will be discussed below in the section on the OpenBlade specification.

The HP Blade Server bc1100 (Figures 3 and 4) is the initial server product and is a complete server with a processor, memory, disk storage, and network connections included. This blade is designed to mount in the HP Blade Server bh7800 chassis, which will hold up to 16 blade servers. HP's blade server solution currently supports the Intel Pentium III processors running the Linux operating system. The three distributions of Linux currently supported are: Debian 2.2r3 kernel 2.2.19; Red Hat 7.0 kernel 2.2.16; or SuSE 7.1 kernel 2.4.0. HP Blade Servers will, in the future, be

available with the PA-RISC processor and Intel's Itanium processor family. HP-UX and Microsoft Windows will also be supported in the future.

The HP Blade Server bc1100 currently uses a 700 MHz Intel Pentium III processor. The blades each come with 512 Mbytes of ECC PC100 memory and a 30 Gbyte IDE hard disk. Both PICMG 2.16 10/100 Ethernet ports are supported. The blade is CompactPCI hot-swap level 3 compliant and it supports SNMP. For legacy users, each of these server blades supports two USB ports, an RS-232 port and an SVGA port. As a part of the extended CompactPCI specification, each blade also includes an additional Ethernet port intended for out of band management as preferred in xSP, corporate, and enterprise data centers.

Any of these operating systems can use the standard Ethernet connections through the mid-plane, thus significantly reducing the cabling complexity associated with three LANs per server. In addition, the HP Blade Server solution will support IDE, SCSI, and FC. At initial release the server blades will have access to local storage through the embedded IDE hard disk on the server blade itself and by pairing the server blade with an optional storage blade with two IDE hard disks for additional storage. The dual Fibre Channel interfaces also routed through the mid-plane are enabled by an FC-AL SAN blade (which creates a standard FC arbitrated loop storage area network), or a Fibre Channel switch blade (which would create switched fabric to each server slot in the chassis).

The lower ejector handle on each blade initiates a power-down cycle to allow safe removal and replacement in a running chassis. In this manner, which is compliant with the PICMG 2.1 standard, the blade

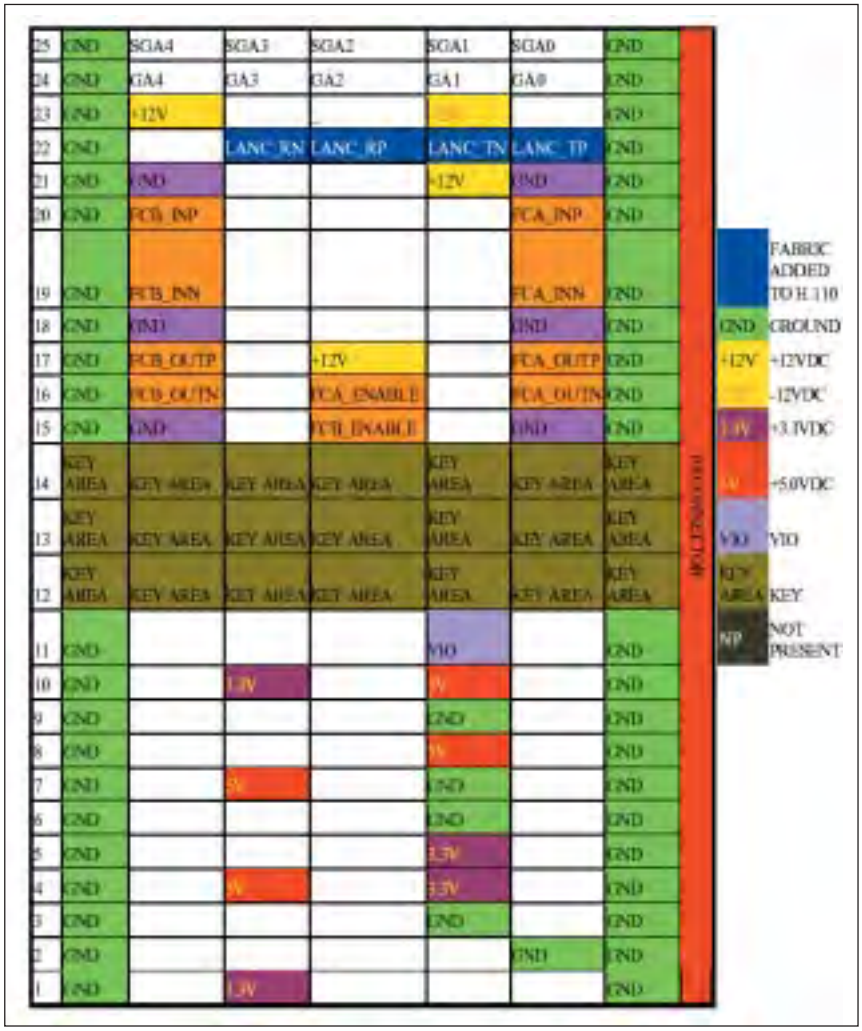

Figure 2

Figure 3. HP Blade Server bc1100

Figure 4. Connector pins for the HP Blade Server bc1100

can shut down safely, while also notifying the Network Operations Center (NOC) that the blade is being removed. This function can also be initiated at the NOC. Each blade includes environmental sensors that

communicate with the management blade which, in turn, controls fan speeds, slot power, etc., and which can alert the NOC of abnormal conditions or events. This combination of features means that blades are less likely to suffer environmental faults, and are easier to service.

Because of its ability to function as an independent, dedicated server, as well as to exist in a high-performance Gigabit Ethernet or SAN environment, the HP Blade Server is an ideal solution for companies that need a high-density solution. The bc1100 exceeds the demands of companies with space constraints and power needs by meeting those needs and also delivering high levels of manageability and flexibility. A future two slot-wide, two way IA-32 blade will utilize the 100 watts available to obtain the highest performance and flexibility to meet the demands of the enterprise customers while improving their manageability of servers. The key to this manageability and flexibility are the PICMG standards, the OpenBlade extensions to the CompactPCI standard as well as the FC and management blades that enhance the HP Blade Server offerings so that customers can get a more complete server solution.

The HP Network Switch Blade is functionally identical to the HP Procurve 2524. This blade will support up to 16 10/100 Ethernet ports through the mid-plane. Two Network Switch Blades are required to support two Ethernet ports per server (one switch blade is required; a second is available for redundancy). In addition to the 16 mid-plane ports, this blade includes eight external 10/100 Ethernet ports for communications with external network devices, and a Gigabit Ethernet uplink port. This is a layer 2 managed switch with VLAN support.

The HP Fibre Channel SAN Blade is inserted into a patented 3U carrier/slot on the mid-plane. When two 3U FC SAN blades are mounted on the 6U carrier, two Fibre Channel ports are enabled to each server blade. The FC-AL blade creates an arbitrated loop within the mid-plane, and two FC-AL blades create two arbitrated loops. Likewise, the FC Switch blade would enable switched FC fabric to each server blade, and two of these blades would allow for two switched FC interfaces. Both FC blades have interfaces that allow for connections to external FC switches and/or storage devices.

The HP Management Blade supervises the operations of the bh7800 chassis and provides access to the management consoles of each installed server blades. This blade provides consolidated management of the chassis, all installed blades, as well as event reporting, configuration management, and hot-swap control. The Management Blade is designed to use the third, out of band management port in the OpenBlade specification. However, the management LAN connection from the Management Blade can also be connected via the Network Switch Blade to the two production LANs if "payload" traffic doesn't need to be separated from the management traffic.

The HP Blade Server bh7800 chassis

The HP Blade Server bh7800 is a unique solution for the service provider or the enterprise. The unprecedented density, the efficient power use, and the vast simplification of the cabling plant can reduce personnel costs, increase staff efficiency, and allow customers to build larger, more capable data centers without having to build more space, or even install more power (see Figure 5).

The bh7800 chassis is designed to be a NEBS-compliant 13U enclosure that has 38 slots, of which up to 16 can be Server Blades. It includes N+1 hot-swap power supplies, fans and power line inputs, and it has local front and rear display panels and keypads.

The bh7800's unique two-sided enclosure allows customers to access the enclosure and insert blades from either the front or back. This enclosure eliminated the traditional rear transition modules, instead employing a mid-plane design that allows full-height 6U blades to be inserted from either end. The mid-plane is not a shared PCI bus topology where the PCI signals are delivered to each slot. Using through hole technology for the pins in the mid-plane, the PCI bus for a given server is dedicated to the I/O expansion slot that is directly opposite the server from the mid-plane. Furthermore, server slots alternate from front to rear on adjacent slots, allowing for a maximum number of two slot wide high performance server blades. (Each server blade is its own independent server on a card, making it possible to mix and match blades in the chassis, i.e., IA-32, PA-RISC, and IPF.)

This enclosure also contains a management blade, which has many important cost-saving features, including inventory management down to serial numbers, and configuration management capability that can track and restore a blade's management IP address when it's been hot-swapped. The management blade monitors all chassis functions including individual blade logs, power utilization and compliance, and thermal compliance. The management blade can also increase the fan speeds when over-temperature conditions arise and return them to normal when appropriate. All of the management data can be reported to management utilities using SNMP, or to network managers such as HP Toptools or HP OpenView. Although central to the infrastructure, the individual blades and remaining infrastructure will continue to operate in the event of a management blade failure, which can also be hot-swapped.

The HP Blade Server bh3700 Chassis is a NEBS-compliant chassis that shares many characteristics with the bh7800. It also supports the OpenBlade specification, and supports enhanced capabilities, including Gigabit Ethernet and Fibre Channel. However, the bh3700 is a 10 slot, 6U single sided enclosure designed more for the traditional carrier installation than for the enterprise or service provider. But because the bh3700 is capable of sharing many of the blades as the bh7800, HP will leverage the volumes and best in class of both businesses thus providing customers with a lower total cost of ownership. The most important difference between the two chassis is that the bh3700 has an I/O backplane, while the bh7800 does not. The bh7800 can only accommodate Ethernet or Fibre Channel connectivity between blades.

Of course, the single most important feature of the HP Blade Server bh7800 is the CompactPCI standard. More important, it's the only blade standard that has momentum and an existing community of vendors, designers, and suppliers.

Equally important, using the CompactPCI standard for HP Blade Servers means that other vendors can take advantage of the features of the HP blades and provide additional capabilities, services and features. In today's marketplace, no single company can do everything, and with its CompactPCI-based blade servers, HP opens a wealth of potential applications to its customers.

The result is that customers will have the opportunity to select the products they need to meet their requirements, regardless of whether those products come from HP. CompactPCI helps assure customers that they will have a consistent upgrade path, either to more performance, or as

Figure 5. Three fully configured HP Blade Server bh7800 Chassis in a rack

solutions to problems and needs that don't currently exist. As a result, basing the HP Blade Servers on CompactPCI means that there will be an assured path to innovation, which is absent in a proprietary environment.

These advantages, of course, mean that the risk of adopting a blade server solution is reduced, as is the total cost of ownership. Likewise, new products will appear sooner, innovations will happen more quickly, and owners will have greater flexibility, and the comfort of knowing that they have an industry-standard solution.

HP Blade Server management

The HP Management Blade is the central component for managing the HP Blade Server bh7800 Chassis. This blade is designed to communicate with the central management station using a separate management LAN, to which all of the other HP Blade Servers are also connected. While this management LAN can be interconnected to the primary "payload" LAN through a switch or router, the design is intended to allow management traffic to remain separate.

The management LAN provides access to the management circuitry built into each HP Blade Server. In addition, it allows management of the operating system, which is especially important in applications where the "payload" LANs are open to the public, and where the internal operations of the OS can never be exposed. This results in a much more secure network than can ever exist where management must take place over the production network.

The functions of the HP Management Blade are as follows:

- Supervises the operation of other components in the chassis, including power and cooling, and the status of installed blades
- Reports failures to the central management system
- Provides switching for access to the console management for the management blade and the processor blades
- Provides initial boot configuration for the chassis
- Provides configuration for the network blades
- Keeps an inventory of the network blades
- Brings up, shuts down, and restarts blades
- Consolidates events from installed blade servers and transmits them to the management application
- Controls indicator lights for each blade, making faulty blades easy to locate in the enclosure

Each blade server includes a management controller that communicates with the management blade. This controller, which exists on a daughter board installed on each bc1100, provides console access and remote management functions. This remote management circuit (RMC) provides specific management information directly to the management LAN. The RMC provides:

- Text mode VGA access over the LAN
- Telnet access to the console
- Firmware update access
- Remote on and off through CompactPCI hot-swap support
- Remote reset and reboot

The information gathered by the HP Remote Management Blade is made available to management applications, such as HP's Toptools, and to management frameworks, such as HP OpenView. Either of these packages provides a centralized management solution that will monitor every blade server, the chassis, and in the case of HP OpenView, the operating system running on each blade server.

Because of the level of control, and the security provided by the separate management LAN, customers can manage and control nearly every aspect of the operation of these devices. This means that physical access to the blade servers is only necessary for physical events, such as the hot-swap of a blade. In addition, because the management blade is tracking errors as well as operations, intervention can take place as soon as an error is detected, meaning that server uptime is enhanced.

HP Blade Server Alliance Program

Because the OpenBlade specification is designed to allow other vendors to create blades that will interoperate with HP Blade Servers, HP created the HP Blade Server Alliance Program to encourage and assist companies that want to enter this exciting market.

The program includes a set of services that facilitates the integration of products on the HP Blade Server platform. This program will accelerate market acceptance of blade servers by speeding up blade development, increasing reach to customers through HP's partners and promote customer confidence in the interoperability of their blade solutions.

Services offered to partners through this program include hardware and software development kits, a hardware and software certification program, events and partner communication, joint marketing activities and promotional services.

The future of HP's Blade Servers

The HP Blade Server offering available today is just the beginning. The roadmap includes a Fibre Channel switch, blades with PA-RISC and IPF processors, support for Microsoft Windows and HP-UX as well as dual-and quad-processor blades.

Summary

HP's Blade Servers are an important new development for both the blade server market and the CompactPCI industry. The arrival of standards for blade servers is a critical step for their broad acceptance in the service provider and enterprise markets, and it's an important step in bringing new capabilities to carriers in a form that's compatible with their existing infrastructure.

The OpenBlade specification will help ensure the growth of CompactPCI, and of blade servers, because it will provide an

interoperability path for vendors and help ease the way for integrators and IT departments. In the meantime, the HP Blade Servers bring lower costs, better performance, greater security, and lower risk to customers.

The ability to leverage the advantages of blade servers while also providing easy access to the rest of the enterprise through industry standard storage solutions and management simply enhances the fact that both the vendors and the customers will win.

Lin Nease *is chief technologist for HP's UNIX server business. He is responsible for leading the commercial server business product portfolio process as well as the technology strategy. Lin has been with HP for 15 years, primarily in R&D, but also in marketing and business strategy.*

For more information, Contact Lin at:

Lin Nease
Hewlett-Packard
8000 Foothills Blvd., mailstop 5602
Roseville, CA 95747

Craig Yamasaki *is Product Planning Manager for HP's blades. He is chartered to leverage HP's vast experience with Windows, Linux, and UNIX servers to lead the definition of HP's Blade Server strategy. Craig has been with HP for 16 years, primarily in hardware R&D, but also in support, software R&D, and marketing.*

For more information, contact Craig at:

Craig Yamasaki
Hewlett-Packard
8000 Foothill Blvd. mailstop 5596
Roseville, CA 95747
E-mail: craig_yamasaki@hp.com
Web site: www.hp.com

GNP demonstrates Intel-based blade architecture is ready for high-availability applications

By Bruce Rostowfske

System manufacturers are touting a new type of server architecture called blade. Blade is a low-cost, low-power design that lets engineers fit hundreds of densely-packed computers into a standard rack by optimizing the available processing power of DSP and I/O boards. In addition to eliminating space needed for traditional separate host CPU and I/O boards, this super-slim server architecture enables multi-node clusters. For those applications that require scalability of processing power and I/O, blade is the next-generation architecture alternative.

GNP has demonstrated the effectiveness of blade's low-power/high-density design in a carrier-class telecommunication application powered by four Intel StrongARM-based computers. By loading its high-availability middleware, Continuant Cluster Suite, directly onto Intel's integrated 32-bit StrongARM SA-1110 I/O processing boards, GNP eliminated the need for system slot processors and provided a clustered platform for natural load-balancing. In a clustered blade configuration, these low-power processors support administration and application operations so that the system integrator does not need to design in a high-availability dual-host server to support other I/O computers.

Continuant Cluster Suite is ideal for implementing blade architectures for core network applications due to its combination of five-nines availability and scalability. Blades are commonly used to power applications running at the edge of a network, and are therefore not the standard default server architecture for high-availability applications. However, the blade configuration is ideal for core network applications with high-reliability requirements.

GNP's patent-pending Natural Clustering Technology, embedded in the Continuant middleware, keeps all nodes active and responsible for work distribution in an N+k configuration, where N is the number of blades required to maintain a desired performance level and k is the built-in reliability level. When in a blade environment, the middleware communicates between the application and the hardware framework, including OS and drivers, thus ensuring 99.999 percent uptime. Blade is a comprehensive operations framework that automates routine HA functions and provides out-of-band remote control and maintenance.

Blades are modular, scalable, low-power, low-cost computers, and GNP's demonstration proves that this architecture can function as part of a high-density server farm. When combined with effective clustering middleware and powerful embedded communications technologies like CompactPCI Packet Switched Backplane (cPSB), blade boards will emerge as a standard network application design. Whether as the primary application in a server farm, or one of many, blades support scalable high-availability applications for optimized computing power in limited space.

For more information, contact:

GNP Computers
555 E. Huntington Dr., Monrovia, CA 91016
Tel: 626-305-8484 • Fax: 626-305-9177
Web site: www.gnp.com