



SYBASE IANYWHERE
improves management
but lacks integration P_32

A TRULY MODULAR SAN?
Compellent's Storage Center
redefines flexibility P_36

NEW TCO SERIES:
Need to justify your WLAN?
Here's a cost breakdown P_72

Network Computing

For IT By IT

DATA CENTER DIET PLAN

Server racks bulging at the seams?
Today's blades are slim but pack
plenty of processing muscle. Steven
Hill evaluates three systems P_42

CISCO VS. MERU

With allegations flying about bugs
and standards violations, we asked
each vendor to explain its side P_67



STORAGE & SERVERS

REVIEW

Analysis: Next-Gen Blade Servers

Data Center Diet Plan

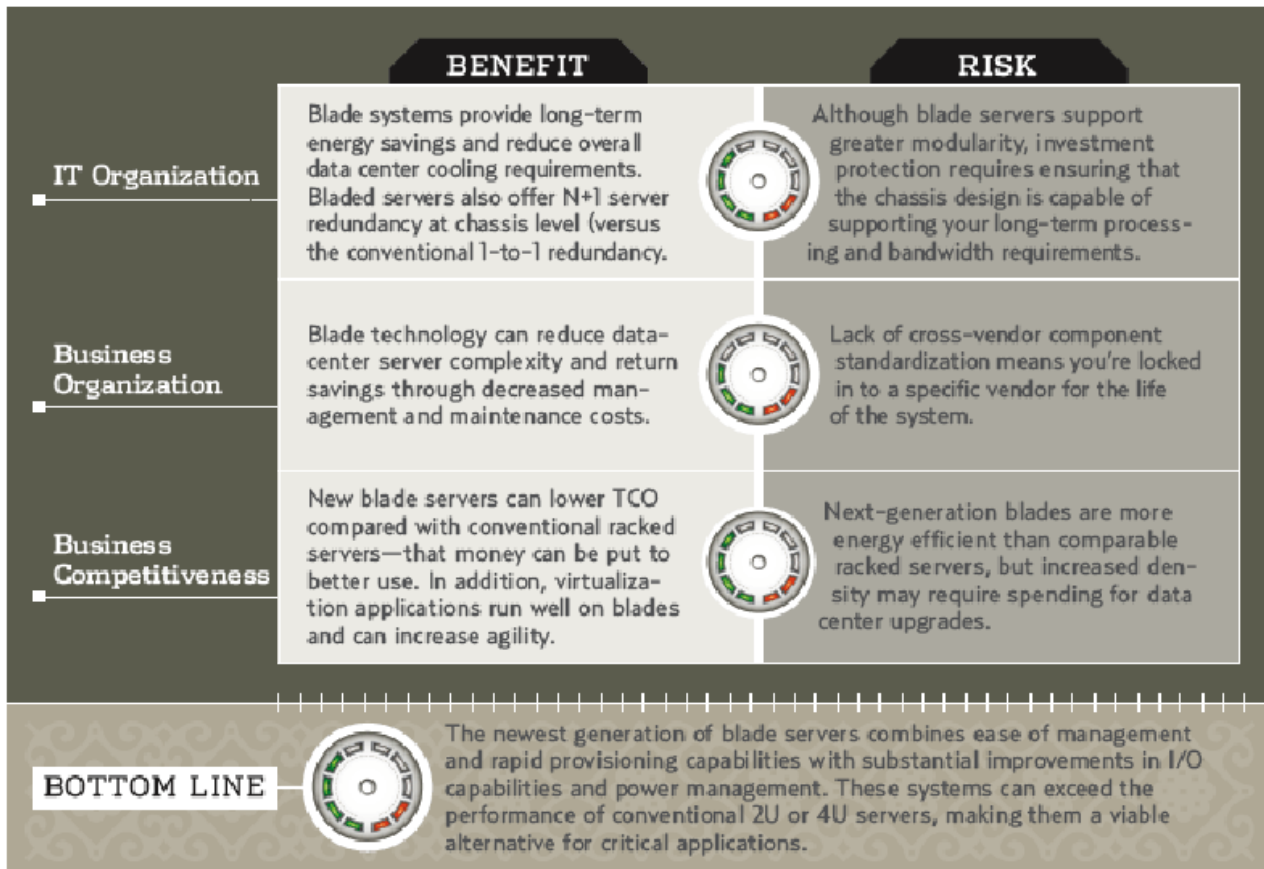
Should you slim down with the newest generation of blade server? Vendors claim the latest systems offer improved flexibility and promise to reduce data center bloat. They're right -- if your infrastructure can support the power and cooling requirements. We examine the latest technology.

Dec 7, 2006 - By [Steven Hill](#)

If the name of the [data center](#) game is getting more computing power for less, blades should be the hottest thing since South Beach. They're more manageable and deliver better [TCO](#) than their 1U counterparts--our latest testing shows as much as a fourfold increase in [processor](#) density combined with 20 percent to 30 percent power savings.

So why did Gartner Dataquest put this year's blade shipments at an anemic 850,000 units this year, just 10 percent of total server sales?

IMPACT ASSESSMENT: BLADE SERVERS



Today, blade server designs have improved, with enough midplane throughput and modularity at the chassis to provide investment protection for their three-to-five-year lifespan. Processor density has increased, and power consumption is lower than you might expect.

They also deliver incredible flexibility. Instead of limiting the blade system to particular types of I/O--interconnect, network or storage--vendors overprovision the I/O channel, providing sufficient bandwidth for targeted environments, or let IT allocate I/O as it sees fit. Seems vendors are following the lead of core switch vendors: Make the frame big enough to jam in just about anything you want for the foreseeable future.

Enterprises are finally catching on. Blade shipments will rise to 2.3 million units by 2011, to account for almost 22 percent of all server purchases, according to Gartner Dataquest. Although blades are still more expensive than conventional 1U servers, you should see operational savings in the 30 percent range, according to Imex Research. Those changes make the newest generation of blade servers excellent candidates for high-demand core and server-virtualization applications.

That's the blade server story vendors want you to know about. The less flattering side: Even while delivering power savings, the energy demands of these high-density systems will still tax the overall infrastructures of many older--and even some newer--data centers. You may be able to quadruple the processor density of a rack, but can you power and cool it?

Many data-center managers are saying no. By 2011, 96 percent of current data-center facilities are projected to be at their power and cooling capacity limits, according to a recent survey of the Data Center Users' Group conducted by Emerson Network Power. Forty percent of respondents cited heat density or power density as the biggest issue they're facing. We examine HVAC and electrical requirements and offer design suggestions in our "[This Old Data Center](#)" special issue at and our Data Center Power Issues analyst report at [nwcanalytics.com](#).

The Players

Most first-tier server vendors now offer blade lines. Gartner Dataquest, IDC and our readership agree that IBM, Hewlett-Packard and Dell, in that order, are the top three players in the market. IBM holds about a 10-point lead over HP, with Dell a distant third at less than half HP's share. No other single vendor is in the double digits, though judging by our testing, Dell should be watching over its shoulder for Sun Microsystems.

In an unlikely coincidence, IBM is flexing its muscle by proposing to standardize blade system modules and interconnects around its design. Although standardization would benefit the enterprise, we're not convinced IBM's proposal is the best solution; see "[IBM and the Quest for Standardization](#)" page 50).

We asked Dell, Egenera, Fujitsu, HP, IBM, Rackable Systems and Sun to send a base system, including a chassis and four x86-based server blades with Ethernet connectivity, to our new Green Bay, Wis., Real-World Labs®, where we had EqualLogic's new PS3800XV [iSCSI SAN](#) array and a Nortel 5510 48-port [Gigabit](#) Ethernet data-center [switch](#) online for testing.

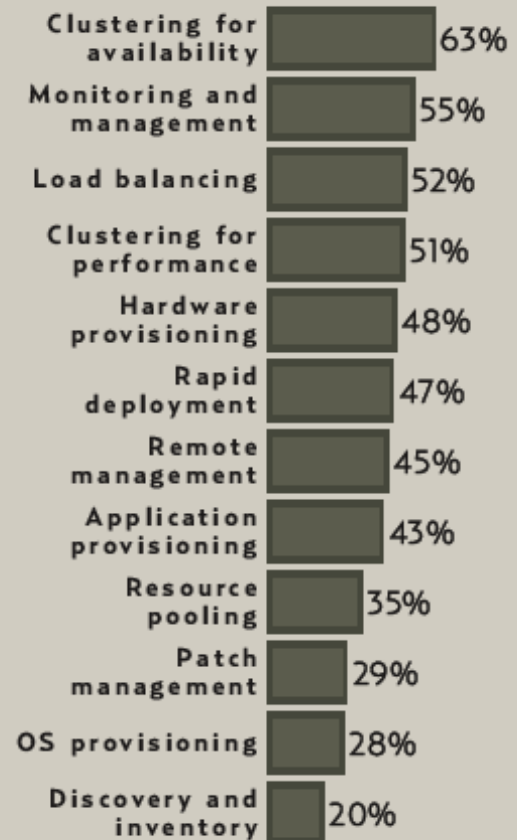
We were surprised when only HP, Rackable Systems and Sun agreed to participate. See our methodology, criteria, detailed form-factor and product-testing results at [nwcreports.com](#); our in-depth analysis of the blade server vendor landscape and poll data can be found at [nwcanalytics.com](#).

Not only are the cooling and power improved in our new lab digs, it's actually possible to see daylight while testing, if you set your chair just right. It may have been the bright light, but once we got going, the differences in our three systems came into sharp focus.

Three For The Money

READER POLL

Which of the following blade server features are important to your organization?



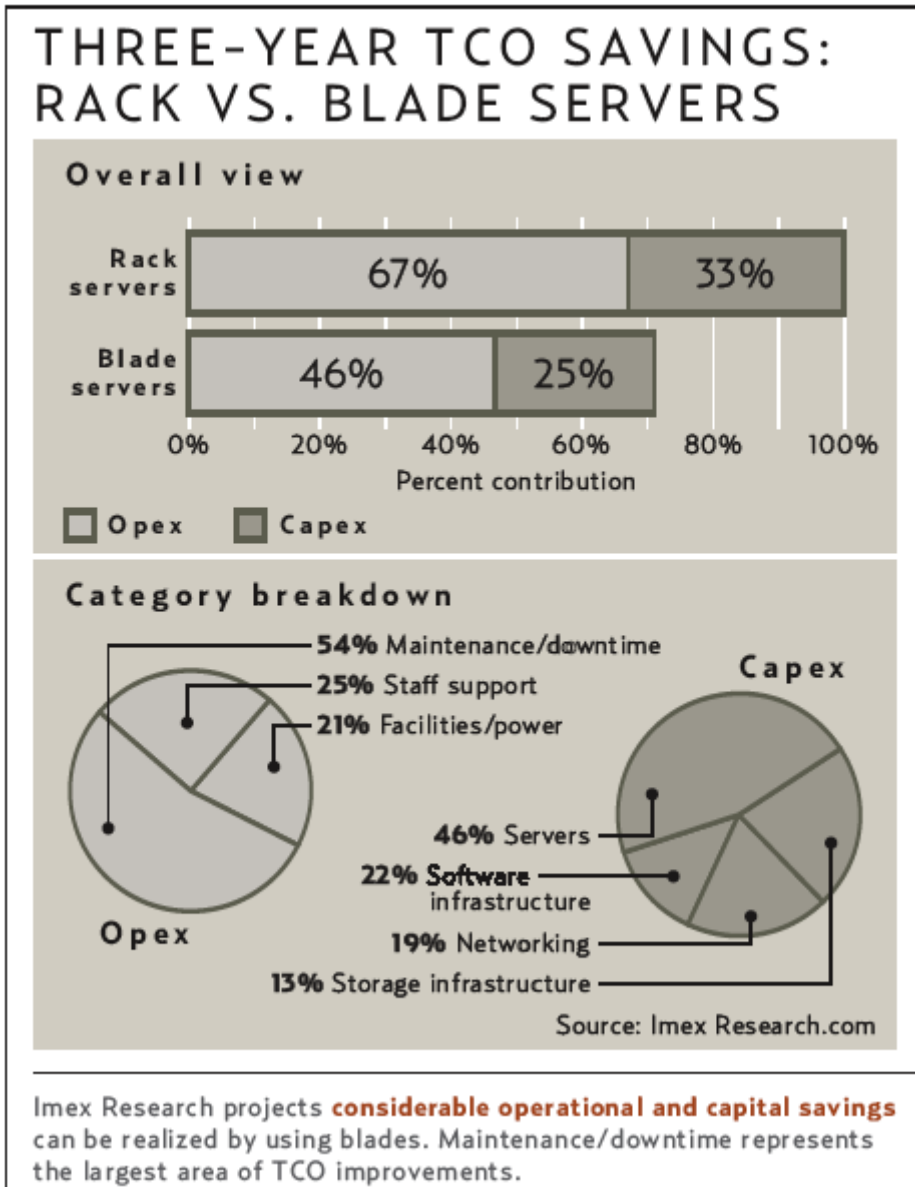
Source: NWC Poll, 527 respondents

HP submitted its new 10-U C-Series BL7000c enclosure along with two full-height and two half-height ProLiant BLc server blades. Sun sent its recently released 19-U Sun Blade 8000 Modular System with four full-height Sun Blade X8400 Server Modules. Rackable Systems provided five of its Scale Out server blade modules.

Rackable blew away the competition in sheer node count. Rather than basing its design on an eight- or 10-blade chassis, Rackable goes large with a proprietary, passive rack design that supports 88 blades and has an intriguing focus on DC power. If you're in need of processors, and lots of 'em, Scale Outs may be just the ticket.

What excited us most about the blades from both HP and Sun was their potential for future expansion to the next generation of processors and high-speed I/O interfaces. The midplane bandwidth offered by these systems will be perfectly suited for the rapidly approaching 10 Gigabit Ethernet and 8-Gb/10-Gb Fibre Channel adapters and switches (for more on blades for storage, see "Storage on the Edge of a Blade," page 52). These bad boys have clearly been designed to provide investment protection.

Review: Blade Servers: Data Center Diet Plan



IMEX Research Note:
[Click Here for additional Information.](http://imexresearch.com/reports/EXEC_SUMMARIES/IMEX - VZ Exec Summary.pdf)
http://imexresearch.com/reports/EXEC_SUMMARIES/IMEX - VZ Exec Summary.pdf

HP also is on the verge of introducing its new Virtual Connect technology for Ethernet and Fibre Channel. Exclusive to HP and the BladeSystem C-Series, Virtual Connect modules allow as many as four linked BladeSystem chassis to be joined as a Virtual Connect domain that can be assigned a pool of World Wide Names for Fibre Channel or MAC and IP addresses for Ethernet. These addresses are managed internally and dynamically assigned by the Virtual Connect system to individual blades within those chassis. By managing these variables at chassis level rather than at blade or adapter level, individual blades can remain stateless, making it easier for system administrators to swap out failed modules or assign hot spares for failover.

Simplify My Life

[Clustering](#) for availability as well as monitoring and management are the top two blade functions, according to our poll. And indeed, vendors tout their systems' ability to dispense with third-party management software and/or KVMs as a key selling point.

Physically configuring, installing and cabling conventional servers can be a massively time-consuming process that must be done during a scheduled downtime to ensure other servers in the rack are not disrupted accidentally. Even something as simple as replacing a failed NIC or power supply can be risky business.

But with blade systems, modules are designed to be hot-swapped without the need to bring down the entire chassis. The time required for server maintenance can drop from hours to minutes, and the unified management interfaces offered by most blade systems can dramatically simplify the process of server administration from a software standpoint.

All three systems we tested offer onboard, blade-level network management interfaces as well as front-accessible KVM and USB ports for direct connections.

The tools Rackable provided offered a degree of control over any number of remote servers, but the absence of more powerful, integrated features, such as remote KVM and desktop redirection over Ethernet, is a noticeable omission in Rackable's Scale Out design when compared with chassis-based blade systems. And with its chassis-less design, Rackable doesn't offer a unified management interface.

Conversely, HP and Sun provide extremely detailed, Web-based management and monitoring capabilities at both server and chassis level. Both systems offer integrated lights-out management interfaces, for example, enabling us to do status monitoring and system configuration across all the blades in our chassis.

But HP's ProLiant was clearly the class of the management category. HP went the extra mile by adding to its Web-based controls a local multifunction LCD interface at the base of each chassis that supports all the management features found in the Web interface, without the need to connect a laptop or KVM to the chassis. At first we were unimpressed with yet another LCD display, but we were won over by the elegant simplicity and surprising flexibility offered by that little screen. The HP Insight [Display interface](#) is identical remotely or locally and offers role-based security, context-sensitive help, graphical displays that depict the location of problem components and a chat mode that lets technicians at remote locations interactively communicate with those working on the chassis. We could easily imagine how useful two-way communication capabilities would be in a distributed enterprise.

Go Green

The costs of power and cooling in the data center are at an all-time high, and it's not just the person signing the utility checks who's taking notice. In July, the House of Representatives passed House Bill 5646 on to the Senate. This measure directs the Environmental Protection Agency to analyze the rapid growth and energy consumption of [computer](#) data centers by both the federal government and private enterprise. Sponsors of the bill cited data-center electricity costs that are already in the range of \$3.3 billion per year, and estimated annual utility costs for a 100,000 square foot data center at nearly \$6 million.

Denser systems run faster and hotter, and the cost of running a server could well exceed its initial purchase price in as little as four years. Because the actual energy use of any blade system is dependent on a number of variables, such as processor, memory, chipset, [disk](#) type and so on, it's virtually impossible to establish a clear winner in this category. All three vendors told us their systems offer an estimated 20 percent to 25 percent reduction in energy costs over similarly equipped conventional servers, and all provide temperature-monitoring capabilities.

We were intrigued by Rackable's DC-power option, which could provide the greatest savings for enterprises that have implemented large-scale DC power distribution throughout the data center. Most of the power efficiency provided by the chassis-based blade systems from HP and Sun stems from the ability to consolidate power and cooling at rack level. But HP takes this concept a step further with its Thermal Logic technology, which continually monitors temperature levels and energy use at blade, enclosure and rack levels and dynamically optimizes airflow and power consumption to stay within a predetermined power budget.

Expensive Real Estate

Another key benefit of blade systems is the ability to pack a lot of [processing](#) power into the least amount of rack space. Of course, how well a vendor accomplishes this goes beyond number of processors--we looked at how blades were grouped and how well they're kept supplied with data.

Rackable's side-by-side and back-to-back rack provides by far the greatest processor density per square foot of floor space--as many as 88 dual-processor servers equate to 176 CPUs per rack, twice that if you count [dual-core](#) processors, putting Rackable well ahead of HP and Sun in this category.

Coming in second in the density challenge is HP. Four 10U BladeSystem c7000 enclosures fit in a single rack; each enclosure holds 16 dual-processor half-height blades, squeezing 128 CPUs into a conventional rack. The 19U Sun Blade 8000 chassis fit two to a rack, and each chassis can handle 10 four-processor server modules, for a total of 80 processors per rack.

But processor density is only part of the story.

When comparing blade systems, it's important to note the difference between two- and four-processor blades. The dual-processor options from Rackable and HP are the basic equivalent of conventional 1U and 2U servers, while each quad-processor Sun Blade X8400 matches a traditional 4U box. This configuration supports the assignment of much larger tasks on a blade-by-blade basis and makes the Sun system a better candidate for demanding applications.

Check Out Those Pipes

The modern IT environment is as much about data throughput as processing capacity. That means for blades to be competitive against conventional servers, they must keep up with high-speed fabrics, such as 4-Gb Fibre Channel, 4x InfiniBand and 10-Gb Ethernet, while still supporting multiple GbE links for management and basic network traffic.

When it comes to total backplane bandwidth, we found little difference between HP's and Sun's designs. PCIe, Fibre Channel, InfiniBand or Ethernet--it's all [serial](#) in nature. It really comes down to who's got the pipes, and in this case it's Sun.

What makes this an important IT issue is the fact that, for now, a decision to buy a given blade system locks you into a single vendor's hardware [platform](#) for a period of time. Ensuring that the chassis you purchase can accommodate future high-speed interfaces and processors provides investment protection.

Did we mention the Sun Blade 8000 has huge pipes? Its midplane can handle up to 9.6 terabits per second in combined throughput. According to Sun, this equates to 160 Gbps in usable bandwidth per blade when you add in [protocol](#) overhead and other factors.

HP's BladeSystem C-Series offers substantial 5-Tbps midplane bandwidth when measured at pure line speed, more than enough to support multiple high-speed fabrics from each blade. HP also offers additional high-speed cross-connects between adjacent blade sockets, designed to improve performance for multiblade clustered applications, as well as to support plans for future storage-specific blades.

In the case of Rackable Systems, the 2-Gbps offered by the dual GbE ports is perhaps the weakest [link](#) in the Scale Out design--this much lower bandwidth potential is a shortcoming that will probably limit use in many high-performance, high-bandwidth applications.

The other side of the I/O issue is port diversity and flexibility. Blade systems can potentially be greater than an equivalent sum of racked servers, thanks to their ability to share support of multiple fabrics and reduce cabling with integrated switch modules. The Sun Blade offered the most potential here by virtue of the remarkable bandwidth of its PCIe midplane architecture. But HP's BladeSystem C-Series currently provides by far the greatest [port](#) diversity when it comes to available backplane switches and pass-through modules.

What'll All This Cost Me?

Comparing pricing on a purely apples-to-apples basis turned out to be a fruitless quest because the variety of approaches taken by our three vendors made direct comparison difficult. To award our pricing grade, we created a formula that pitted two dual-processor, Opteron 285-based server blades from HP and Rackable against a single four-processor [Opteron](#) 885-based blade from Sun. We made sure to price each system on a blade level with similar memory, [SATA storage](#) and dual-port Gigabit Ethernet connectivity, without consideration for the differences in chassis/rack capabilities. Not perfect, but at least we're in the produce family.

Rackable's Scale Outs came in at \$9,780 for two dual-processor blades, making them the least expensive option on a processor-by-processor basis. The HP BladeSystem landed in the middle, at \$11,768 for a pair of ProLiant BL465c half-height blades. An equivalent Sun Blade would run \$24,885 apiece when you include the cost of the two PCIe Express GbE modules required to match the port counts of the other two systems.

This was no surprise: The Sun X8400 blade is a quad-processor system, and it was a foregone conclusion that it would be more expensive than its dual-processor counterparts. In all fairness to Sun, this was like comparing a pair of two-way, 1U servers to a single four-way, 4U system; even though the processor count is the same across configurations, it costs a lot more to make a four-CPU server.

See more pricing and warranty details in "[Bullish on Blades](#)" at [nwcreports.com](#).

Blade Servers By the Numbers:

0.7%	Worldwide market share currently held by RLX Technologies, which helped pioneer the blade server concept. Source: Gartner Dataquest
8 kilowatts	Power level above which it generally becomes difficult to cool a standard rack, because of airflow limitations in most data centers. Source: Forrester
25 kilowatts	Potential per-rack power draw if populated with dense servers or blade chassis. Source: Forrester
\$0.07	Average U.S. per-kilowatt-hour cost for electricity. Source: U.S. Department of Energy
40 to 1	Reduction in cables possible by using blades rather than 1U rack servers. Source: Burton Group
7%	SMBs that use blades, compared with 25% running traditional racked servers. Source: Forrester

IBM And The Quest For Standardization

Buying into a blade system has long meant purchasing servers and interconnect modules from only one vendor. How much this worries you depends on your point of view--many companies prefer to deal with a single vendor for the support, service and pricing benefits offered by volume server purchases. Still, the perception remains that blade systems are more constrictive than their conventionally racked counterparts.

This year, IBM launched Blade.org, which it describes as a "movement" to provide a standardized design for blade system modules and interconnects that would theoretically lead to the commoditization of blade parts. The only drawback for competing vendors: These standards are based on IBM's BladeCenter concept.

Standardized hardware improves competition and makes life better for IT, and there's plenty of precedent--PCI, ATX, SATA, SCSI. But while we applaud the idea of design specifications for blade server components, we're not convinced IBM's BladeCenter system represents the best-of-breed benchmark for long-term standardization.

For our money, the new high-performance concepts behind the Sun Blade 8000 Modular System deserve serious consideration. Taking I/O modules off the blade, moving them to hot-swappable locations on the backplane and using vendor-independent PCIe Express Module technology for [communications](#) and interconnects is a more future-proof methodology than perpetuating the existing blade concept. And, it would engender more industry cooperation than the Blade.org's IBM-centric solution.

Regardless, 60 [software](#) and component vendors, including Citrix, Brocade and Symantec, have signed on to the program. It's not surprising--they have nothing to lose--and everything to gain.

Storage On The Edge Of A Blade

Given the smaller form factor of blade server modules, early-generation blades relied on internally mounted 2.5-inch laptop drives for onboard storage. But these desktop-class drives offered neither the performance nor the reliability of their 3.5-inch counterparts.

Today's blade customers favor boot-from-SAN configurations, but the growing popularity of 2.5-inch enterprise-class [SAS](#) and higher-capacity SATA drives have brought the convenience of externally accessible and hot-swappable drives to blade systems.

Because the Rackable Systems Scale Out blades we tested support full-size 3.5-inch drives, they could be fitted with as much as 1.5 TB of SATA disk per blade using dual, internally mounted 750-GB drives. Sun's blades offer two hot-swappable 2.5-inch disks per module, and Hewlett-Packard's c-Class system supports two hot-swappable 2.5-inch drives on its half-height blades and four drives on full-height blades.

In 2007, HP plans to introduce c-Class storage blades that will hold six SAS drives, supporting up to 876 [GB](#) per half-height blade and linked to an adjacent blade slot with a dedicated x4 Serial link. We can't wait.

To qualify for this review, we asked vendors to submit a base blade chassis, all software required to manage chassis and blades, and four matching server blades with extended 64-bit x86 processors and Gigabit Ethernet (GbE) connectivity for conventional network throughput and iSCSI storage.

PARTICIPATING VENDORS

Hewlett-Packard, Rackable Systems and Sun Microsystems

TESTING SCENARIO

To test blade servers, we built a dedicated network using our Nortel 5510 48-GbE data center switch. We installed Windows Server 2003 R2 on all systems under test. For our networked storage, we created an iSCSI SAN based on a combination of EqualLogic's SAS-based PS3800XV and SATA-based PS400E storage arrays, and assigned iSCSI logical unit numbers to each individual server module.

SCORING CRITERIA

Energy Management: 20% Rates voltage options, power-supply redundancy, published efficiency statistics, thermal monitoring features and user-configurable energy-management tools.

System Management: 20% Evaluates integrated management interfaces at both blade and chassis level, and rates them on ease of operation, level of detail, reporting features, security and systemwide unified-management capabilities.

Cost: 15% Our formula compares a pair of dual-processor, Opteron 285 server blades from HP and Rackable Systems against a single, four-processor Opteron 885 blade from Sun. From this baseline we established a sample price for each system with similar raw processing capabilities, system memory, SATA storage and dual-port GbE connectivity.

Features: 15% See our features chart at nwcreports.com.

I/O Bandwidth: 15% The ranking for I/O bandwidth was a combined comparison that took into account blade-level I/O capabilities, total chassis-level bandwidth, port diversity, and availability of third-party switch and IO modules.

Processor Density: 15% Scoring is based on the maximum number of processor sockets that could occupy a full 42/44-U server rack.

RESULTS

Our Editor's Choice is HP's c-Class BladeSystem. This well-rounded offering provides a winning combination of system performance, module flexibility and price, backed up by a strong service and support structure. We found the HP management system the most user-friendly and energy-conscious, and it's obvious HP has attempted to solve many of the problems of earlier generations of blade servers. It earned consistently high scores across all grading areas. The beefy Sun Blade 8000 Modular System earned a perfect score for I/O bandwidth and could be a solid alternative to conventional servers for core-level applications, such as large Oracle or SAP environments, as well as high-performance I/O and compute-intensive apps. Its combination of four-way blades and massive bandwidth at both blade and midplane level sets a new standard for x86 blade performance and throughput. Still, the Sun Blade 8000 is somewhat pricey and currently offers only one I/O module option for its four Network Express slots.

In contrast, the cost-effective Rackable Scale Out blades would be well-suited to edge applications that benefit from extremely high server density. This system packs capacity for 88 blades into a passive rack design that can support DC power. However, the Scale Out lagged in I/O performance and flexibility.

Steven Hill is an NWC technology editor. Write to him at shill@nwc.com.

NWC Reports: Bullish on Blades

For our latest round of testing, we asked seven leading blade server vendors to send a chassis and four x86-based server blades with Ethernet connectivity to our new Green Bay, Wis., Real-World Labs®. Hewlett-Packard, Rackable Systems and Sun Microsystems agreed. Fujitsu and Dell declined; IBM said it had no available hardware; and Egenera begged off, saying it "doesn't connect to iSCSI storage."

HP sent its ProLiant c-Class BladeSystem, which for our testing comprised a 10U ProLiant c7000 chassis, two 3.2 GHz Xeon ProLiant BL460c blades, two dual-3.2 GHz Xeon ProLiant BL480c blades and four GbE (Gigabit Ethernet) 2c Ethernet blade switches for the c-Class BladeSystem.

Rackable Systems provided five of its Scale Out Blade Servers--four dual 2.2 GHz Low-Power Opteron 275 HE Scale Out Blade Servers and a dual 2.6 GHz Opteron 285-based Scale Out Blade Server.

Sun delivered its recently released Sun Blade 8000 Modular System, which included a 19U Sun Blade 8000 Chassis, four full-height quad-2.6 GHz Opteron 885 Sun Blade X8400 server modules, two Sun Blade 8000 GbE 20-port Network Express modules, four PCIe Dual-GbE ExpressModules, two PCIe Dual Infiniband ExpressModules and two PCIe QLogic dual 4-Gb Fibre Channel ExpressModules.

When we first reviewed blade servers, in June 2003 (see "Pitching Blades"), we were looking to answer the question: Is managing blades easier than dealing with a stack of 1U or 2U servers? The answer, then as now, is an unqualified yes. But other aspects of blade servers, including design, energy efficiency, throughput and fabrics supported, have improved significantly.

FORM UP

We found significant architectural variations among the three blade systems tested.

Rackable's Scale Out system is targeted at high-density environments looking to provide hundreds, even thousands of individual system nodes. Its hot-plug design lets IT load one rack with as many as 88 server modules in a side-by-side and back-to-back configuration. With an optional 19-inch mount kit, conventional racks can hold five vertical chassis in a 7U space.

Rackable's affordable Scale Out blades come in both single- and dual-processor versions and can be populated with Intel Xeon or AMD Opteron processors. On-blade storage is provided by one or a pair of internal, 3.5-inch SATA or SCSI drives, and server I/O options are limited to dual GbE ports. But what makes Rackable unique in this market are its multiple power alternatives--letting Scale Out blades be used in AC or DC environments--and the fact that its blades are based on an open-architecture server module designed to accommodate a wide range of standard motherboards and hard disks. If you need the flexibility to tailor blade configurations based on specific hardware requirements, Rackable has you covered.

Meanwhile, the Sun Blade 8000 Modular System is built for speed, with the performance to handle demanding applications, such as Oracle or SAP. Its 19U chassis can hold 10 four-processor blade modules, each equivalent to a conventional 4U racked server. Storage is provided by one or a pair of 2.5-inch SAS drives per blade. Sun has chosen to align with AMD, so it offers only Opteron processors. This may not sit well with some IT managers, but we can only hope that the days of name-brand prejudice will come to an end eventually.

What sets Sun's servers apart are its exclusive use of four-processor blades and its gutsy decision to move network components off the blades and on to the chassis backplane. This breaks the traditional blade-design mold, where each blade must be configured with the desired combination of NICs or HBAs at the time of purchase. The Sun Blade 8000 server modules have no NICs or HBAs, and all communication from the blade is kept strictly PCIe. Each blade has four x8 PCIe and two x4 PCIe links that run through the passive midplane of the chassis. These provide substantial bandwidth to two individual x8 PCIe Express modules per blade, as well as four Network Express modules that provide shared connectivity across all ten blades. This unique methodology allows hot-pluggable I/O customization of individual blades using optional Ethernet, FC or InfiniBand PCIe Express modules, while simultaneously supporting as many as eight additional GbE connections per blade. Finally, the new HP ProLiant c-Class BladeSystem offers the most balanced approach of the three systems we reviewed. Targeted at enterprise data centers and server consolidation, the 10U c7000 enclosure supports as many as eight full-height dual-processor blades, 16 half-height dual-processor blades, or any combination thereof. HP offers blades with either Intel Xeon or AMD Opteron processors, and its half-height blades can support two hot-swappable 2.5-inch SAS drives, while the full-height blades can handle as many as four drives each.

Although not quite as beefy as the Sun Blade when it comes to chassis-level I/O bandwidth, the c-Class surpasses its predecessors by offering the flexibility of multiple, integrated GbE NICs on each blade, as well as the option to add two additional NIC or HBA modules--known as mezzanine cards--on the half-height blades and three more on the full height. The internal GbE NICs and mezzanine cards are mapped to eight bays on the chassis backplane that can be populated with a wide choice of switches and pass-through modules, using up to four different fabrics.

MANAGE ME

We weighted system management at 20 percent of our grade. Rackable's optional Roamer management system provides Ethernet or serial connectivity and an LCD status panel. The text-based management interface gave us remote access to basic power cycling, BIOS configuration, system reset and temperature monitoring, but little else in terms of system status information, event logging or automated problem reporting.

The Sun Blade 8000 System offers on each Server Module an Integrated Lights Out Manager (iLOM) on each Server Module that combines with the Chassis Monitoring Module to provide secure, browser-based and on-site access to a broad spectrum of status monitoring and system configuration tools. The hardcore Sun worshipper will be pleased to find a DMTF-compliant CLI as well.

Sun also offers its optional Sun N1 System Manager software, which supports an even broader set of features, including bare-metal discovery, OS provisioning, and the ability to manage hundreds of systems across multiple chassis and racks from a single console. We believe N1 System Manager would be a worthwhile addition for large environments, but the software is available only for Solaris.

HP's Integrated Lights-Out 2 (iLO 2) management interface, now in its fourth generation, teams with HP's Onboard Administrator module to provide single-session management of all blades and I/O devices in the chassis. We found many similarities in the ILO management capabilities of the HP and Sun systems, but HP goes the extra mile, providing additional features with its new HP Onboard Administrator. This handy app integrates with existing iLO 2 features on all components, plus provides both local and remote control over the enhanced management tools added to the BladeSystem c-series; these include such unique features as real-time control over power and cooling, visual indicators of fault conditions and configuration errors, and simplified diagnostics and graphical instructions for troubleshooting.

We were also impressed with HP's elegant and flexible local interface. A multifunction LCD display, located at the base of each chassis, gave us access to all Web interface management features, including role-based security, context-sensitive help and graphical displays of problem components. A chat mode allows technicians at a remote locations to communicate. Handy stuff.

Building on these integrated tools are a wealth of high-level software products from HP that offer expanded control of the server environment; these include the Insight Control Data Center Edition management suite for rapid provisioning, software installation and license management, and the ILO Select Pack for additional security and reporting features. And at the top of the food chain, there's the HP Open View line of enterprise-level management tools.

FEEL THE HEAT

Rackable's Scale Out blades support three power options. The least efficient is an all-AC configuration, using blades with conventional AC power supplies. The second setup combines DC-powered blades with AC/DC rectifiers for power conversion at rack level. These rectifiers can be configured for [redundancy](#) and mounted either on-rack or outside the data center for better heat management. In addition, Rackable's Scale Out blades are mounted in a back-to-back configuration at rack-level, so exhaust heat is ducted to the top of the rack rather than to the back of the system and into the aisle or rack behind it.

Perhaps the most energy-efficient option would be for data centers wired for redundant -48 DC. These shops can use DC-powered blades in combination with Rackable's DC power-distribution system, which allows the rack to connect directly to a central DC power supply.

Rackable's DC-powered blades incorporate a DC [power supply](#) with no moving parts that generates much less heat--while being several times more reliable--than conventional power supplies. One liability to the Scale Out blade design, however, is its incorporation of two high-speed fans in each blade module. This setup works well to manage heat generated by the blades themselves, but it offers no redundancy or hot-swap capabilities; not to mention a full rack of 88 Scale Out blades would have 176 fans running at all times.

Most of the power efficiency provided by the chassis-based blade systems from HP and Sun stems from their ability to consolidate power and cooling at rack level. While supplying redundant power to conventional servers requires that dual, active power supplies be mounted in each server, the blade chassis designs from HP and Sun use a maximum of six power supplies to handle the entire rack. This provides both AC redundancy and N+N power-supply redundancy. Cooling fans are handled in the same manner, with both chassis offering redundant, high-efficiency, multispeed, hot-swappable fan modules that automatically respond to variations in the cooling needs of all blades in the chassis.

A feature unique to HP is Dynamic Power Saver, which lets an administrator choose between AC redundancy, power-supply redundancy and a power-saving mode that allows some power supplies to shut down during off-peak periods. One of the main problems with AC power supplies is that they're inefficient when running under reduced demand, so the Dynamic Power Saver automatically puts unnecessary power supplies on standby, letting the system run fewer power supplies at higher and more efficient utilization levels.

Administrators can actually set a hard limit on the amount of AC power a chassis or rack can consume. This affects only new loads and not cooling, so there's no risk to systems already on line. This is not a hard limit on the power, but rather a threshold setting. The number of available power supplies and their redundancy level also determine the power available to the chassis; the system will automatically limit the number of server blades that can come online based on those guidelines and the power draw of the system. It's a nifty high-level feature.

I/O BANDWIDTH

We evaluated all three vendors' designs to see whether they'd be able to keep up with high-speed fabrics, such as 4 Gb FC, 4x InfiniBand and 10 GbE, while still providing GbE links for management and network use.

The dual GbE ports in Rackable's Scale Out system provide 2-Gbps throughput; this lower bandwidth will limit use in many high-performance applications--however, Rackable's target audience seems to run more toward those needing high server density rather than fat pipes.

In contrast, the Sun Blade 8000 wowed us with its capacity. Its midplane can handle up to 9.6 Tbps in combined throughput, and Sun specs the systems at 160 Gbps of usable bandwidth per blade. HP's BladeSystem c-series is no slouch either, with 5 Tbps midplane bandwidth. We liked that HP provides additional high-speed cross connects between adjacent blade sockets--a boon for multiblade clustered and storage-specific applications.

System bandwidth aside; the major difference between the Sun and HP designs lies in the locations of their I/O devices. Sun has chosen to extend multiple PCIe interfaces from the blade, through a passive midplane, and to the chassis backplane, where it locates all I/O devices for FC, InfiniBand and Ethernet connectivity. On the other hand, HP's approach uses a passive serial midplane to connect specific fabrics that originate on the blade, with the backplane reserved for switch or pass-through modules.

ANY PORT IN A STORM?

We also evaluated port diversity and flexibility. Ideally, integrated switch modules enable blade systems to support multiple fabrics across multiple servers. If this is important to you, HP's BladeSystem c-series offers the greatest port diversity in terms of available backplane switches and pass-through modules.

On the HP c-series, GbE switch modules are available from both HP and Cisco, and FC adapters from either Emulex or QLogic. You also can use pass-through adapters that can provide direct connections between blade-level adapters and third-party switches. Aside from the dual, integrated GbE ports on the half-height and the four ports on the full-height blades, HP offers multifunction GbE mezzanine cards that support [TCP/IP](#) off-loading as well as iSCSI acceleration and RDMA (remote direct [memory](#) access) capabilities. When it comes to features, port diversity and overall system flexibility, we have to give

this round to HP. But it was a close call thanks to the remarkable bandwidth offered by the Sun Blade's PCIe midplane architecture.

Sun offers connectivity for most fabrics through the use of individual Ethernet, InfiniBand and FC PCIe Express modules connected by two individual x8 backplane slots per blade. But at this point, the only option available for Sun's four aggregated Network Express slots on the chassis backplane is a 20-port GbE pass-through module. Even though this offers the potential for eight individual GbE ports per blade--with room to spare for two x8 PCIe Express Modules--the absence of shared switching modules limits the Sun Blade 8000's overall flexibility. This makes Sun pretty much all dressed up with nowhere to go at this point, but we'll be keeping an eye on developments in this area.

On the Scale Out blades from Rackable Systems, dual onboard GbE ports are wired along with serial/Ethernet management communications through a blind-mating connector at the back of the blade; they terminate within the rack as standard [RJ-45](#) network cables. Because there is no dedicated chassis, there are no integrated switch modules; instead, space is left available in a full rack to allow for installation of third-party switching equipment. An open [PCI](#) slot located on the front panel of the blade is available for conventional FC or InfiniBand adapters, but this would necessitate routing individual cables to the front of each blade and finding room in the rack to accommodate additional switching.

PRICE DILEMMA

The disparate architectures of these blade systems made comparing price difficult. We created a formula that pitted two dual-processor, Opteron 285-based server blades from HP and Rackable Systems against a single four-processor Opteron 885-based blade from Sun. We priced each system on a blade level with similar memory, SATA storage and dual-port GbE connectivity, without consideration for the differences in chassis/ rack capabilities.

The Scale Outs from Rackable Systems came in at \$9,780 for two dual-processor blades, making them the least-expensive option on a processor-by-processor basis. The HP BladeSystem landed in the middle of the pack, at \$11,768 for a pair of ProLiant BL465c half-height blades.

An equivalent Sun Blade would run \$24,885 when you include the cost of the two PCIe Express GbE modules required to match the port counts of the other two systems. We expected this--the fact that its design requires 800-series rather than 200-series Opterons would make Sun about \$6,000 more expensive than rivals from the get-go. We think the premium is worthwhile, though, for companies that can make use of the Sun Blade's power. Our as-tested pricing for each system: HP, \$36,546; Sun, \$149,680 (both including chassis and modules). Rackable, \$9,780 (server blades only, no chassis or power supplies).

From a warranty aspect, both HP and Sun offer a three-year warranty on the chassis and blades with next-business-day response times for onsite service, while the Rackable Systems Scale Out servers have a one-year warranty combined with an RMA repair program.

All three vendors have both Web-based and 24/7/365 telephone accessibility and offer a number of optional, multi-year support plans that can provide up to 24/7/365 onsite service. Overall, HP offers the most complete matrix of service options for both hardware and software support, but Sun has upped the ante by launching a new pricing scheme that reduces maintenance costs by covering all devices within the entire chassis.

AND THE WINNER IS...

Our Editor's Choice is HP's ProLiant c-Class BladeSystem. This well-rounded offering provides a winning combination of system performance, module flexibility and price, backed by a strong service and support structure. We found the HP management system the most user-friendly and energy-conscious, and it's obvious that HP has put a great deal of thought into solving many of the problems associated with earlier generations of blade servers. It earned consistently high scores across all grading areas.

The beefy Sun Blade 8000 Modular System earned a perfect score for I/O bandwidth and could be a solid alternative to conventional servers for core-level applications, such as large Oracle or SAP environments, as well as high-performance I/O- and computationally intensive apps. Its combination of four-way blades and massive bandwidth at both blade and midplane level sets a new standard for x86-based blade performance and throughput. Still, as much as we liked its performance capabilities, the Sun Blade 8000 is somewhat pricey and currently offers only one I/O module option for its four Network Express slots.

In contrast, the cost-effective Scale Out blades from Rackable Systems would be well suited to edge applications that value extremely high server density over I/O performance.

Because earlier-generation blade servers were like fad diets--long on hype, short on delivery. Despite vendor promises, they didn't represent much of a savings over conventional devices. Most of the systems we evaluated when we reviewed blade servers in June 2003 were struggling with first-generation blues--an 8- or 10-blade chassis used the same amount of rack space as equivalent IU devices and suffered [I/O bandwidth](#) limitations between blades and backplanes, making them better-suited for [Web server](#) consolidation than running critical databases.

But even then, one fact came through loud and clear: Managing blades is substantially easier than dealing with individual racked boxes.

HEWLETT-PACKARD C-CLASS BLADESYSTEM

HP designed the new c-Class BladeSystem, released in June 2006 to replace the p-Class, with an eye toward providing enough system performance and flexibility to carry the platform well into the next decade. One thing that came through loud and clear in our testing is HP's focus on simplifying management. This, combined with outstanding energy efficiency and a substantial list of connectivity options, earned HP top marks for enterprise-class blade servers.



In an interesting departure from its earlier systems, the new c-Class offers a choice of full- or half-height server blade modules. This flexible design lets the eight-slot chassis support as many as eight full-height or 16 half-height blades, or any combination thereof. At this point, both full and half blades support a maximum of two processor sockets per blade, but the full-blade design appears to have plenty of room for four-socket full-height blades at some point in the future.

The system we tested came equipped with two full-height blades sporting a pair of 3.2 GHz dual-Core Xeon processors, and two half-height blades, each with a single 3.2 GHz dual-Core Xeon.

There are a wide variety of half-height blade options available for the c-Class BladeSystem; the Intel-based BL460c modules sport two CPU sockets and a choice of nine different dual-core Xeon processors, ranging from 1.6 GHz to 3.2 GHz, and also support as many as eight sockets of PC2-5300 DDR2 FB-DIMM memory.

The AMD-based BL465c blades offer a choice of seven different Opteron 2000-series processors and provide eight sockets for PC2-5300 DDR2-registered DIMM memory. Each half blade has room for dual 2.5-inch hot-swappable SAS or SATA drives and comes with an integrated HP SmartArray E220i RAID controller that supports 0, 1 and 0+1 arrays.

Common to both half-blade designs are a pair of integrated, multifunction GbE NICs that support TCP/IP off-load under Windows, as well as accelerated iSCSI and RDMA capabilities. For other connectivity options, there are two onboard mezzanine sockets that can be used for additional standard or multifunction GbE modules and 4x DDR InfiniBand, as well as 4 GB FC modules from Emulex or QLogic. All blades for the c-Class provide connectivity for HP's Integrated Lights Out 2 (ILO 2) system management interface without sacrificing a dedicated GbE port, and a high-density connector on the front of each blade uses a custom fan-out cable for direct KVM/USB connections.

The full-height blade we tested for the c-Class was the BL480c; like the BL460c, the 480 is available with 11 dual-core Xeon processor options. Although sharing the same processing capacity as its half-height brethren, the 480's full-height design leaves room for 12 total slots of PC2-5300 DDR2 FB DIMMs, two additional embedded GbE ports, three total mezzanine cards, and as many as four hot-swap SAS or SATA drives. Also included is an integrated HP Smart Array P400i SAS controller with 256 MB of cache and support for hardware RAID 0, 1 and 5. HP has also recently released the new BL685c blade, which is based on four Opteron 8000-series processors and offers as many as 16 DIMM sockets supporting DDR2 memory.

The c7000 chassis offers a substantial number of features that improve the system and energy management capabilities of the c-Class compared with HP's previous p-Class design.

At the heart of the 10U chassis is an Onboard Administrator Management Module that provides secure single-sign-on access to device configuration, power usage, system monitoring and temperature control over all components in single or multiple enclosures. It also serves as a single point of access to the ILO 2 systems on each blade, and supports virtual KVM services over Ethernet for ILO 2-enabled systems across multiple enclosures assigned to a common management domain. HP put a lot of effort into simplifying its management interface, and it shows. Along with detailed online management capabilities, a front-mounted LCD Insight Display at the base of each chassis provides for text-based troubleshooting, graphical error indication and a live-chat mode that allows a technician in a data center to communicate with an administrator working from a remote management terminal.

The c7000 chassis is available in single- and three-phase versions for U.S. or international power sources. Slots for as many as six front-accessible power supplies provide multiple levels of power redundancy.

HP offers one of the most detailed power management systems in blades today, and unique to HP's c7000 is new Dynamic Power Saver technology that allows the system to reduce the number of active power supplies during off-peak usage. Also unique is the ability of an administrator to set a hard power budget for single or multiple chassis, enabling the system to dynamically choose the amount of resources that will be allowed to come online.

System cooling is another integral part of HP's energy strategy, and its new Thermal Logic technology combines detailed inflow and outflow temperature analysis with a highly manageable, load-balancing cooling system. Based on the PARSEC (Parallel Redundant Scalable Airflow) architecture, the c7000 can use as many as 10 of the new high-flow, low-noise HP Active Cooling fans to provide independent zone-based cooling at blade level, with full redundancy across the chassis. This new fan module design is the quietest we've ever heard, which is especially, well, cool given its capacity--HP told us that a single Active Cooling module has the power to cool as many as five DL360 G4 1U servers.

The c7000 is rated for 5 Tbps in raw SerDes (serialization-deserialization) bandwidth. Because HP continues to keep storage and I/O modules on the blade, the chassis must be able to ensure that fabrics are connected properly through the midplane. To do this, the Administrator Management Module offers port-mapping capabilities that let us ensure that adapters mounted in mezzanine ports on the blades were properly connected to the interconnect bays in the backplane. A dedicated link between adjacent blade slots is available to support server clustering and connect to storage-specific blades in the future.

The last step in the I/O chain is connectivity at the backplane, and the c-Class offers the largest number of backplane options of the systems we tested. At the rear of the c7000 chassis are two bays for dual, redundant Administrator Management Modules and eight additional interconnect bays, each with 16 connections to the blade modules. Two of the interconnect bays are dedicated to redundant GbE switches; the remaining six can be populated with connectivity modules for Ethernet, FC, InfiniBand or SAS fabrics.

This design can offer fully redundant backplane support for as many as four fabrics simultaneously on a system filled with full-height blades. There are also pass-through adapters that can provide direct connections between blade-level adapters and third-party Ethernet or FC switches for companies that prefer to handle switching outside the chassis.

SUN MICROSYSTEMS SUN BLADE 8000 MODULAR SYSTEM

When Sun discontinued its Sparc/Athlon/Xeon-based Sun Fire B1600 Blade Platform in mid-2005, the company vowed to be back in a year with a whole new blade system. It wasn't kidding. The new Sun Blade 8000 Modular System, which bears little resemblance to its 3U predecessor, incorporates one of the most interesting technology decisions we've seen to date: By taking I/O modules off the blades and extending the PCI Express bus through the midplane, Sun can locate fabric connectivity modules at the backplane and provide for hot-swapping of PCIe Express modules without the need to take a blade offline.



From the start, it was obvious to us that the major focus of the Sun Blade 8000 is performance. This is perhaps best illustrated by the fact that only four-socketed blades are being built for the system. For the new Sun Blade X8400 Server Modules, Sun is partnering with AMD. At present, its blades are based on quad-940 sockets, which leaves three available Opteron processor options: the 870, 875 and 885. Of course, this is only a temporary limitation because blades based on the new 1207-pin Rev.F socket and 8000-series processors will support both current dual-core and future quad-core Opterons without a major blade redesign.

GbE, X4 DDR InfiniBand and 8- or 10-Gbps FC will all occur during the anticipated lifecycle of these blade systems, a factor that should make total system bandwidth a point of consideration to ensure long-term investment protection. The bandwidth and flexibility offered by the PCI Express-based design has provided Sun with some interesting and unique I/O options.

Of the four 8x PCIe channels available per server module, two slots per blade are dedicated to individual PCIe Express modules, based on the PCI SIG form factor and mounted at the top rear of the chassis. To fill the 20 available slots per chassis, Sun offers three possible PCIe Express modules: an Intel-based dual-port GbE NIC, a QLogic-based dual-port 4-Gbps FC HBA, and a Mellanox-based dual-port 4x InfiniBand host channel adapter. These hot-pluggable modules can be mixed and matched on a server-by-server basis, allowing for a great deal of granularity.

The remaining two 8x and two 4x PCIe channels are connected to four PCI Express Network Express Module (NEM) slots located directly below the individual PCIe modules. These four slots are designed to aggregate the links from all 10 server modules, and could be used for four different, or two redundant, aggregated I/O fabrics. At this point, the only NEM module available is a 20-port GbE device offering dual GbE NICs per server module; Sun says it will address this limitation in the near future.

In spite of the current lack of NEM options, we respect the leap of faith Sun took to adopt this I/O strategy. Given its adoption of an industry-standard PCIe-compliant architecture, it's hard to imagine that I/O device vendors will have difficulty fitting Sun's design specifications for the smaller, individual PCIe Express modules. The bigger challenge may be to find switch vendors with designs that will convert multiserver PCIe input to an aggregated, switched fabric output. A technology that seems ideally suited here is the anticipated development of I/O devices with hardware-enabled virtualization capabilities. An I/OV-enabled device would be capable of presenting multiple virtual interfaces to the combined PCIe connections provided by the NEM slots; for example, four I/OV-capable 10 GbE NICs could potentially be shared among all 10 server modules.

At 19.5 inches by 18.5 inches, the X8400 Server Modules are the largest that we tested. But Sun takes full advantage of that real estate and loads them up with four dual-core processors and sixteen sockets of system memory. Each blade has two hot-swappable 2.5-inch SAS or SATA drives and front-mounted sockets for conventional KVM/USB hookups. There's also an on-board iLOM service processor that offers remote KVM and storage capabilities and supports Sun's management interface as well as third-party IMPI 2.0-compliant management solutions.

Conspicuously absent on these modules are the blade-mounted Ethernet and storage adapters that we've come to know and love. All the I/O capabilities in the Sun Blade design are routed directly through the passive midplane using four x8 and two x4 [PCI Express](#) links per blade. This unique methodology supports 160 Gbps of total usable bandwidth from each individual blade and eliminates the need to pull a blade to replace a failed I/O module.

Like the X8400 blades, the Sun Blade 8000 chassis was the largest we tested. At a height of 19U, it weighed in at over 500 pounds and supports 10 server modules when fully configured. But, considering that a fully-loaded Sun Blade 8000 Modular

System offers the processing equivalent of 10 conventionally racked, four-socket 4U servers, the 8000 is an impressive, brushed aluminum tower of power. Like the HP c-Class, the Sun Blade 8000 chassis supports dual redundant CMMs (Chassis Management Modules), which provide a unified management interface for all components and servers.

The CMM offered us detailed monitoring capabilities for module status and temperature parameters and is designed to integrate with management tools for Sun's high-end X64-based systems. The 8000 chassis has six front-mounted power supplies that offer N+N redundancy for input and output power; it uses three to six single-phase 20A 220V circuits, depending on the level of redundancy desired. System cooling is provided with multispeed, hot-swappable fans. Three fans are dedicated to cooling the top PCI Express modules, six are for the power supplies, and nine are mounted at the back of the chassis to cool the server modules.

Perhaps the most impressive aspect of the Sun Blade 8000 Modular System stems from Sun's decision to use a passive midplane concept that passes the PCI Express system bus through the system, rather than using blade-based GbE, FC or InfiniBand I/O fabrics. This design offers 40 serial links per server--or a total of 400 serial links--and provides a massive 9.7 Tbps of overall SerDes chassis bandwidth that can be designated to any combination of fabrics and I/O modules that can be mounted at the backplane. Though this may seem like overkill at present, the adoption of high-speed interconnects like 10

RACKABLE SYSTEMS SCALE OUT SERVER SERIES

Rackable Systems introduced its Scale Out servers in 2004 as an alternative to conventional blade servers.

Rather than adopting the typical chassis/module concept, Rackable's design is based on providing the highest possible server density combined with a unique, full-rack-oriented design. By mounting Scale Out Blades in a side-by-side and back-to-back configuration, Rackable squeezes 88 servers in a single 88 in.-by-28 in.-by-44 in. cabinet, offering as many as 176 CPUs--or 352 cores. Now that's density.



Rackable has also been a long-term advocate of the adoption of DC power to improve overall energy efficiency and reduce the cooling burden in the modern data center.

Rackable prides itself on delivering what it calls "open blades"--which in this case means blade designs based on industry-standard components, such as standard-form-factor motherboards, 3.5-inch disk drives and a wide variety of processor selections from either Intel or AMD. Rackable could base a Scale Out Blade on almost any server-class ATX motherboard that can fit in its case, offering a great deal of flexibility for customers with specific needs.

have been able to provide practically any combination of processor we could have requested. When we opened up the Scale Out blade we found a conventional ATX dual-socket 940 motherboard with 8 DIMM sockets mounted with the I/O risers facing the front of the blade chassis. This allowed access to the KVM and USB ports through a panel on the front. The dual onboard GbE NIC ports were jumpered using Category 6 cabling to a blind-mating connector at the back of the chassis. By mounting the motherboard backward--relatively speaking--Rackable leaves room above the PCI slots to internally mount dual, 160-GB, 3.5-inch SATA drives; we could have incorporated SAS storage or as many as four 2.5-inch drives in either flavor. The power supply is mounted right behind the drives at the back of the chassis, and dual high-speed fans provide cooling for the entire module. Like conventional blade designs, the Scale Out modules are designed to be hot-pluggable without the need for tools.

For this review we requested an AMD system based on a pair of dual-core Opteron processors, but Rackable Systems would

Chassis? Who needs a chassis? Perhaps the biggest departure from the conventional blade model is Rackable's chassis-less approach. The closest thing we could find to a chassis is the custom rack-mounting system that the Scale Out blades require. Cooling air is drawn from both the front and back of the rack, and exhaust air is directed straight up through the top of the rack. Each blade has dual independent fans, which equates to 176 total fans in a full rack.

Each server module uses a relatively low-density blind-mating connector that provides power to the blade as well as I/O connectivity for the Ethernet and optional serial connections. In the middle of the rack these connectors tie into an integrated cabling system that terminates in RJ-45 connectors near the switch mounting bays. For companies requiring support for additional connectivity, there's access to an open PCI slot in the front of the blade, but this would necessitate routing cables from the face of any affected servers to the required switch.

Absent the chassis, there's no need for chassis-level management. But this also means that Rackable Systems offers no integrated, unified server management interface. The optional Roamer management interface available on the Scale Out blades provides individual control of basic power-cycling and BIOS settings while monitoring ambient case temperatures, but absent a centralized management interface, we saw limited out-of-the-box options for large-scale system management without resorting to third-party software. In this regard, the Scale Out Blades are identical to conventionally racked servers. Aside from extremely high processor density, what makes the Scale Outs stand out are the three powering options offered by Rackable Systems. For those choosing an AC-only power option at blade-level, the Scale Out Blades offer pretty much the same energy efficiency as conventional servers, and depend on a 1:1 ratio between AC power supply and blade. The real power efficiency lies in systems based on DC-powered blades that use systemwide DC or DC power that's converted at rack level. First, the internal DC power supplies for the Scale Out servers are more efficient, generate less heat and are far less prone to failure than conventional AC power supplies. And by supporting DC at blade-level, you can use multiple DC rectifiers

at rack level to provide the same N+N power redundancy found in chassis-based designs. For the ultimate in power efficiency, Rackable Systems advocates the adoption of data-center-level DC power, which could be converted outside the data center to decrease thermal load inside.

The biggest downsides of Rackable Systems' design are limited I/O bandwidth and options for the Scale Out blades we tested. These systems can obviously provide substantial processing capabilities, but would be bottlenecked for a number of applications by the limitations imposed by their dual-GbE NICs. Even though there is room for one additional, front-facing PCI card, this clearly wouldn't be an optimal solution in many data centers. We hope this is an issue Rackable addresses in the next generations of Scale Out systems. Given the company's "open blade" strategy, it's likely that rapid advancements in server-class [ATX motherboard](#) design will offer new alternatives in the future, but the relatively low density of its existing blind-mating connector will still lead to bandwidth limitations.

In all fairness, there are a number of applications that [function](#) perfectly well in a GbE-only environment. We've also found that many shops using Ethernet-based iSCSI for storage rarely stress their dual GbE connections, and most 1U and 2U racked servers come standard with only two GbE ports. With its affordable Scale Out blades, Rackable Systems is clearly targeting server-intensive environments, such as telcos and service providers, where massive scalability equals efficiency and affordability can be a key factor.

REAL-WORLD LABS REPORT CARD		BLADE SERVERS	
	HP BladeSystem	Sun Microsystems Sun Blade 8000 Modular System	Rackable Systems Scale Out Server Series
Energy management (20%)	4.5	4	4
System management (20%)	4.5	4	3
Cost (15%)	4	3	4.5
Features (15%)	4.5	4	2
I/O bandwidth (15%)	4	5	2
Processor density (15%)	4	3.5	5
TOTAL SCORE (100%)	4.30	3.90	3.40
<small>A>4.3, B>3.5, C>2.5, D>1.5, F<1.5 A-C GRADES INCLUDE + OR - IN THEIR RANGES. TOTAL SCORES AND WEIGHTED SCORES ARE BASED ON A SCALE OF 0-5</small>	A-	B	C+