

The AMD Opteron™ 6000 Series Platform: Revolutionizing Server Price/Performance

Key takeaway: AMD introduces a revolutionary 4P server pricing strategy

- >> Until now, the greater processing power and memory footprint that 4P servers provide came at a steep premium relative to the cost of 2P servers.
- >> In response to customer feedback about that, AMD is introducing the AMD Opteron™ 6000 Series platform, a single product that fits in both 2P and 4P systems—yet comes at a 2P price.
- >> For the first time, organizations with demanding workloads can buy a 4P server at a 2P price, because the exact same platform that installs in 2P systems also installs in 4P systems.
- >> This breakthrough pricing strategy can help hardware manufacturers to significantly lower server pricing in the 4P category.

Executive Summary

From a purely technical standpoint, 4P servers are usually a better choice than 2P servers for solutions requiring a balance of performance and scalability. Yet the traditionally high price of 4P platforms relative to the advantage they deliver in core count and memory footprint has made 2P servers the preferred option for many mid-market and enterprise organizations today. With workloads becoming more complex and servers being asked to support growing amounts of data and users, hardware makers and buyers alike need a way to leverage the processor, memory, and I/O capabilities of 4P servers without paying “the 4P tax.”

Recognizing this issue, AMD has responded with a new approach that helps make 4P servers more affordable than ever before. Historically, processor manufacturers have offered separate products in separate price bands for 2P and 4P servers, with 4P-capable processors coming at a healthy premium over their 2P-capable counterparts even if they were similar at the silicon level. When AMD first brought the AMD Opteron™ processor to market, it reduced this premium, delivering a higher-performing 4P processor with a lower price premium. However, truly revitalizing a 4P market that has shrunk over the years and truly giving customers a compelling reason to buy 4P instead of 2P requires a more dramatic step.

Accordingly, AMD is introducing the AMD Opteron 6000 Series platform, a single platform (processor and chipset) that offers support for both 2P and 4P server systems. In the past, the AMD Opteron 2000 Series processors supported 2P servers and the AMD Opteron 8000 Series processors supported 4P servers. With the launch of the AMD Opteron 6000 Series platform, however, one price band now supports both 2P and 4P products, enabling the exact same platform that installs into a 2P server to install in a 4P system, and vice versa. As a result, the 6000 Series platform is one of the first x86-based platforms that hardware vendors can use to deliver “value 4P servers” offering 4P capabilities at 2P economics.¹

¹ Based on standard power Six-Core AMD Opteron processor Model 2435 1ku pricing of \$989 as of 4/27/10 vs. standard power AMD Opteron processor Model 6172 (“Magny-Cours”) 1ku pricing of \$989 as of 4/27/10.



This white paper discusses the market dynamics driving demand for value 4P servers, explains how the AMD Opteron 6000 Series platform satisfies that demand, and explores the practical implications of the value 4P server's breakthrough combination of performance, scalability, and value for companies running high-intensity workloads such as virtualization, cloud computing, and high-performance computing.

Table of Contents

Executive Summary	1
Wanted: A Value 4P Server	2
Introducing the AMD Opteron™ 6000 Series Platform	4
Putting Value 4P Servers to Work: Core Usage Scenarios	6
Virtualization	6
Database	7
Cloud Computing and High-Performance Computing	8
The AMD Opteron™ 6000 Series Platform in Action: The Dell™ PowerEdge™ R815	9
Conclusion	11

Wanted: A Value 4P Server

Workloads that demand more processor cores, larger memory footprints, and greater I/O bandwidth are fast becoming commonplace in twenty-first century data centers. Consider, for example, the following:

- **Virtualization**—A highly memory-intensive workload, server virtualization is currently one of the most sweeping and pervasive trends in IT. Indeed, almost 50 percent of server workloads will be running on virtual machines by the end of 2012, according to analyst firm Gartner Inc.² What's more, adoption is spiking especially sharply among midsize companies. By the end of 2010, Gartner analysts predict, there will be greater penetration of virtual machines among companies with 100 to 999 employees than among Global 500 enterprises.³
- **Database**—Spurred partly by the proliferation of digital video, social media, photos, and audio, the world is currently undergoing an unprecedented explosion in data volumes. Global stores of digital data totaled 487 billion gigabytes as of May 2009 and will double every 18 months going forward, according to analyst IDC.⁴ Processing all of that information and mining it for actionable insights consumes ever-increasing amounts of server capacity. Furthermore, database workloads have in recent years grown beyond order processing systems and data warehouses to include Internet-driven technologies such as social networking, e-commerce, and cloud computing. That is putting more stress on database servers as they field higher volumes of concurrent user requests.

²<http://www.networkworld.com/news/2009/102009-gartner-server-virtualization.html?src=netflash-rss>

³<http://www.eweek.com/ca/Virtualization/Server-Virtualization-Adoption-Growing-Rapidly-Gartner-821995/?kc=rss>

⁴<http://www.infoniac.com/hi-tech/digital-content-today-to-double-every-18-months.html>

- **Cloud computing**—Adoption of cloud-based solutions that deliver software, data, and services over the Internet is growing at a fiery pace, fueling an explosion of Web-based transactions. In fact, IDC expects worldwide spending on IT cloud services to grow from \$17.4 billion in 2009 to \$44.2 billion in 2013.⁵
- **HPC/Technical computing**—Organizations are increasingly using hundreds or thousands of inexpensive 2P x86 servers in large scale-out installations to solve large problems by breaking these down into small pieces that can be distributed across many servers. Whether decoding the human genome, rendering video animation, or searching for new deposits of oil and gas, high-performance compute (HPC) clusters are helping drive innovation. They're also hungry for more cores and more memory as they harness massive parallel resources to complete their tasks.

Aside from their rising popularity, however, workloads such as these have another quality in common: Running at acceptable performance levels often requires server hardware with heavy-duty processing capacity, large memory footprint, robust I/O resources, and headroom to grow as needs dictate. In today's market, 4P servers generally fit that description

better than 2P servers. But 4P servers have traditionally occupied the upper “niche” of the server market, acting as workhorses for database and mission-critical applications. With the introduction of the AMD Opteron™ processor in April 2003, the integrated memory controller and 64-bit memory addressing functionality formerly limited to UNIX-based RISC servers became available in mainstream 4P x86-based servers. Soon, these 4P x86 servers were successful in penetrating into the “high-end, high performance” servers, offering specific benefits in areas such as core count, memory footprint, and I/O capabilities over the volume market of 2P x86 servers.

However, while AMD technology-based 4P servers eliminated many of the performance and scalability limitations of the x86 server, 4P servers in general still imposed a significant price penalty. Like similar products from Intel Corporation, AMD Opteron processors have historically come in separate price bands for 1P, 2P, and 4P servers. The gap between the DP (2P) and MP (4P) price bands, in particular, has been disproportionate to the performance boost that 4P servers deliver. By offering high performance at a lower price premium in many cases than comparable products, the AMD Opteron processor has made strides in closing that gap. Yet even so, the “4P tax” continues to cause many

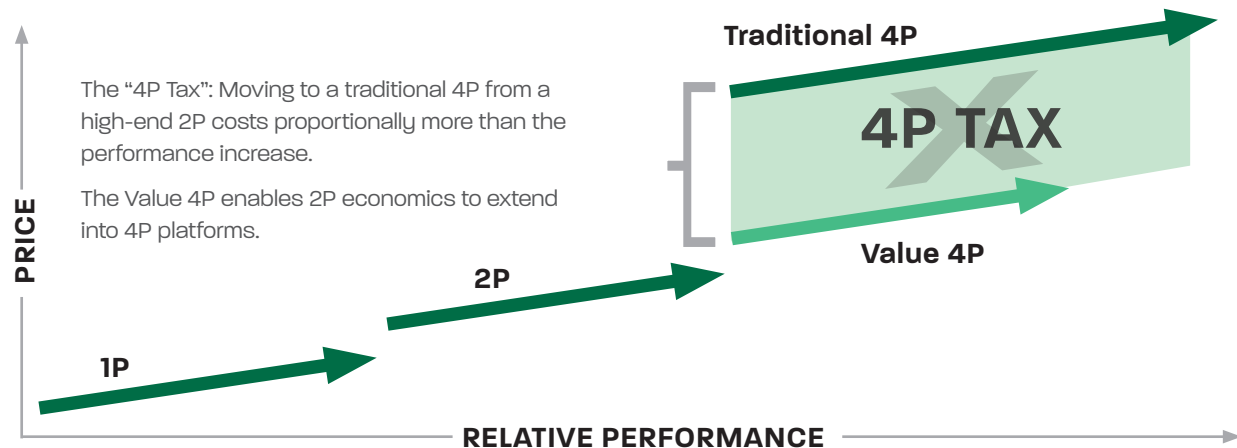


Figure 1: Thanks to the “4P tax,” 4P servers have traditionally often cost too much to justify the performance gains they deliver relative to 2P servers.

⁵ http://news.cnet.com/8301-13846_3-10372446-62.html?part=rss&subj=news&tag=2547-1_3-0-20

companies—especially in the mid-market space—to assign workloads better suited to 4P servers to lower-priced 2P value servers instead. For example, in 2009 less than five percent of x86 servers sold were 4P servers.

However, many industry watchers believe that would quickly change if businesses had access to a “value 4P server” that delivered 4P capabilities more economically.

“Virtualization, high-intensity database applications, and other heavy-duty workloads are driving greater demand for processing power and memory capacity in today’s data centers,” says Matthew Eastwood, group vice president of IDC’s enterprise platform group. “A value 4P server that delivered 4P performance at 2P economics would enable organizations to meet that demand far more efficiently and cost-effectively than is currently possible.”

Introducing the AMD Opteron™ 6000 Series Platform

With the launch of the new AMD Opteron 6000™ Series platform, such value 4P products are now available for the first time. Equipped with up to twice as many cores, up to two times the memory bandwidth or greater⁶, and almost twice the I/O bandwidth⁷ as previous-generation AMD Opteron processors, the AMD Opteron 6000 Series platform delivers the processing capacity organizations need to tackle challenging workloads efficiently. Yet it also accommodates both 2P and 4P systems in one price band; more succinctly, the 2P- and 4P-capable processors are now the same thing. A 4P AMD Opteron 6000 Series platform can help cut memory expenses, since the greater quantity of DIMM slots on 4P servers can allow you to run larger memory configurations using smaller, less expensive DIMMs.⁸ In addition, servers based on the AMD Opteron 6000 Series platform can help reduce management overhead, help reduce power costs, streamline software images, and conserve valuable data center floor space when used in conjunction with a consolidation strategy.

“The AMD Opteron 6000 Series platform came about as a result of something interesting that we learned from our conversations with customers,” says Patrick Patla, vice president and general manager for the server and embedded division at AMD. “Businesses that need high-performance server hardware were using several 2P products instead of 4P systems, not because the 2P servers were the best overall fit for the job but because 4P servers were comparatively too expensive. We realized that by creating a server platform offering 4P capabilities — in other words, by making the same processors work in both 2P and 4P systems — we could give organizations a cost-effective way to meet their performance and scalability needs while also helping reduce hardware footprints in the data center.”

The upshot is a breakthrough development for organizations with difficult workloads. Here are some of the AMD Opteron 6000 Series platform’s most important features:

⁶ Based on quad channel DDR3-1333 for AMD Opteron™ 6100 Series processor vs. dual channel DDR2-800 for Six-Core AMD Opteron™ processor.

⁷ I/O comparison based on 3x HyperTransport™ technology links @ up to 4.8 GT/s (up to 57.6 GB/s total bandwidth) for Six-Core AMD Opteron™ processor vs. 4x HyperTransport™ technology links @ up to 6.4 GT/s (up to 102.4 GB/s total bandwidth) for AMD Opteron™ 6100 Series processor.

⁸ Based on comparison of 4P AMD Opteron™ 6000 Series processor-based system @128GB (12 DIMMs per CPU; 32 x 4GB) vs. 2P Intel Xeon (“Westmere”) processor-based system @128GB (16 x 8GB). Memory price for AMD processor-based system: \$5,888. Memory price for Intel processor-based system: \$12,000. Prices as reported on crucial.com as of 12/17/09.

- **Heavy-duty processing capacity:** Offers up to 48 physical processing cores in a 4P server, delivering the raw power you need to take on computationally demanding or heavily parallel workloads.
- **Massive memory capacity bandwidth:** The platform supports up to 512GB of DDR3 memory today and is planned to support higher memory capacities with future processor releases.
- **Dramatically enhanced I/O performance:** The AMD Opteron 6000 Series platform's integrated chipset comes with PCI Express® Gen 2 technology that nearly doubles I/O performance.⁹
- **Improved processor-to-processor speed:** With HyperTransport™ 3 (HT3) technology, four links run up to 6.4 gigatransfers per second, bringing greater processor-to-processor speed and efficiency to your infrastructure than with previous generations. The greater number of links gives OEMs the ability to create a fully-connected 4P server in which every processor and memory location is one hop away.
- **Expanded processor cache:** 12MB of L3 cache, twice as much as previous generation AMD Opteron processors, enables faster data access and optimal performance.
- **AMD-P power efficiency features:** The AMD-P suite of power management features dynamically enables reduced power usage of the overall processor, individual cores, and the logic within each core, depending on application needs and server workloads. Even with as many as twice as many cores, the new 12-core processors can fall within the same power and thermal envelopes as the previous 6-core model.
- **AMD Virtualization™ (AMD-V™) technology:** The AMD-V™ suite of virtualization features are hardware-based technologies that help reduce virtualiza-

tion software overhead. Additionally, the increased number of cores and memory channels in the AMD Opteron 6000 Series platform means that virtual machines can have much more robust configurations, giving virtualized environments near physical performance.

- **Low-power memory support:** Low-voltage DIMM support enables low overall power draw for high energy efficiency.
- **A wide range of models without feature compromises:** Each AMD Opteron 6000 Series processor has the same full set of features regardless of the power band, so there's no compromise on cache, memory speed, and interconnect speeds.
- **RAS features to help keep applications online:** Reliability-availability-security (RAS) features include the ability to turn off cores and bad memory cells, and to enable automatic switching to online spares in the event of problems.

In addition, the AMD Opteron 6000 Series platform offers the efficiencies and infrastructure stability that can come with platform consistency. From entry-level systems to HPC clusters, all AMD Opteron processors feature a common architecture and can be paired with the same chipset. This can help you meet all of your data center needs with fewer software images, fewer drivers, and reduced pre-production testing, all of which helps reduce one of the largest pieces of total cost of ownership—management costs. Plus, processors in the AMD Opteron 6000 Series platform are socket compatible with planned upcoming AMD Opteron 6200 Series processors (scheduled for release in 2011), so you can seamlessly take advantage of the next generation of AMD Opteron processor technology without changing platforms.¹⁰

⁹ I/O comparison based on 3x HyperTransport™ technology links @ up to 4.8 GT/s (up to 57.6 GB/s total bandwidth) for Six-Core AMD Opteron™ processor vs. 4x HyperTransport technology links @ up to 6.4 GT/s (up to 102.4 GB/s total bandwidth) for AMD Opteron™ 6100 Series processor.

¹⁰ BIOS update may be required.

Putting Value 4P Servers to Work: Core Usage Scenarios

The strength of a value 4P server is easier to appreciate in the context of real-world applications. The remainder of this white paper explores specific usage scenarios in which the compute power, scalability, and value offered by value 4P servers like those based on the AMD Opteron™ 6000 Series platform are particularly beneficial.

VIRTUALIZATION

Though virtualization can help companies save money and increase efficiency, it can also impose some serious CPU and memory demands on host hardware. This is particularly the case when leveraging virtualization to consolidate servers, strengthen their disaster recovery capabilities, and streamline desktop management and security. Many businesses have limited their virtualization efforts to 2P servers, due to the cost constraints associated with traditional 4P servers. However, the value 4P server now changes the game. With up to 12 cores, four memory channels per processor, and up to half a terabyte of memory support, the AMD Opteron 6000 Series platform offers the level of resources needed to support numerous virtual machines as robust as the standalone servers they are replacing.

Server Consolidation

Of the many use cases for virtualization, server consolidation remains the most popular and compelling. Moving applications from underutilized servers onto virtual machines has allowed organizations to reduce their total number of physical servers, resulting in simplified admin-

istration, reduced power and cooling costs, and freed-up data center floor space. However, many organizations have limited their initial consolidation efforts to 2P server platforms and are now taking the next step by moving more demanding, production-level applications onto their virtualized clusters.

In theory, 4P servers make great host devices for virtual machines, as they offer the generous CPU and memory resources needed to lower physical server counts while retaining and sometimes even improving overall computational and memory capabilities, supporting large numbers of users, and accommodating virtual machines running more robust applications that require above-average CPU and memory capacity. In addition, using value 4P servers to host virtual machines can help you:

- Reduce the number of host virtualization software images you must maintain, patch, and upgrade
- Improve power efficiency, as fewer servers can mean fewer power supplies
- Cut memory expenses, since the greater quantity of DIMM slots on 4P servers can allow you to run larger memory configurations using smaller, less expensive DIMMs¹¹

Value 4P servers based on the AMD Opteron 6000 Series platform enable organizations to take the next steps to virtualize and consolidate their infrastructure without making exorbitant upfront investments in new hardware.

¹¹ Based on comparison of 4P AMD Opteron™ 6000 Series processor-based system @128GB (12 DIMMs per CPU; 32x4GB) vs. 2P Intel Xeon ("Westmere") processor-based system @128GB (16 x 8GB). Memory price for AMD processor-based system: \$5,888. Memory price for Intel processor-based system: \$12,000. Prices as reported on crucial.com as of 12/17/09.

"As organizations increasingly turn to virtualization for help in consolidating their most critical and demanding workloads, we're seeing a real need for server hardware with greater processing power, memory bandwidth and power efficiency," says Microsoft's Eric Jewett, director server marketing. "With the new AMD Opteron 6000 Series platform and Windows Server® 2008 R2 with Hyper-V, customers can take advantage of the advanced throughput, scalability, and cost-effectiveness to support greater server density in the data center and help reduce their power footprint."

Disaster Recovery

If money were no object, every production server in your data center would have a backup equivalent on round-the-clock standby. With IT budgets tighter than ever today, though, few businesses can afford to maintain a comprehensive duplicate hardware structure. As a result, many mid-market organizations in particular are inadequately prepared for major server outages.

Value 4P servers provide an answer to this dilemma, as they offer the memory and CPU capacity needed to support large numbers of "failover" virtual machines at cost-effective prices. That can empower even midsize businesses to recreate their data center using fewer, larger backup servers.

Desktop Virtualization

Consolidation makes as much sense for desktops as it does for servers. Indeed, virtualizing and centralizing desktops can help organizations strengthen security and improve support.

However, virtual desktop environments require server hardware with enough compute power and memory to serve multiple client sessions simultaneously without

increased user response time. While 4P servers can meet this requirement, they have historically cost more than many organizations have been willing to pay relative to the performance advantage they deliver. Value 4P servers, though, can enable even mid-market companies to afford powerful host servers with the cores, memory, I/O bandwidth, and throughput required to keep virtual clients speedy and responsive.

DATABASE

A database is a multi-threaded workload perfectly suited to servers with large amounts of cores, memory, and I/O capacity. In fact, enterprises have traditionally used 4P or larger servers to run their larger database engines and related business process applications for precisely that reason. But many of the departmental database servers were relegated to 2P platforms because of tight budgets. Today, however, with databases growing ever larger and more complex, value 4P servers like those based on the AMD Opteron™ 6000 Series platform offer the low price points and robust capabilities a broader range of organizations need to scale expanding databases up, scale them out, consolidate them, or mine them for business insights efficiently and cost-effectively.

Scaling Up and Out

As their data stores grow, organizations have one of two basic options: Scale their infrastructure up by replacing existing servers with more powerful models or scale their infrastructure out by adding servers and splitting databases across them. Value 4P servers are a wise choice for either strategy:

- In "scale up" scenarios, value 4P servers offer a potent blend of strength, scalability, and value that cannot be

matched by 2P servers, delivering the power to tackle compute-intensive database workloads affordably.

- In “scale out” scenarios, moving from a 2P server to a value 4P server can help organizations to save money by dividing large databases across fewer, more powerful devices. Not only can this help lower hardware procurement costs, but it can also help simplify management, reduce the number of software images, and conserve data center floor space. Plus, fewer servers can consume less power and require less cooling, helping reduce operational expenses.

Consolidating Databases

Another increasingly popular approach to expanding data volumes is to consolidate by placing multiple databases on the same host device. This can be done with or without virtualization technology. However, today's 2P servers do not always offer the core and memory capacity required to serve multiple databases within required service level agreements. To implement a consolidation strategy without sacrificing responsiveness, organizations need server hardware with the kind of robust processing and memory capabilities enabled by 4P platforms. Value 4P servers allow businesses to acquire such platforms cost-effectively and to consolidate databases effectively.

Supporting Business Intelligence Workloads

Business intelligence is a growing trend with companies of all sizes. The ability to gather, analyze, and disseminate data provides businesses with a powerful tool for making better decisions. One of the most efficient ways to architect a business intelligence solution is to house both your data and your analytical software on the same server. Say you are currently running the

database and analysis engines on separate 2P servers. These servers have to communicate to each other over the network and you have to maintain separate software images for these servers. Value 4P servers based on the AMD Opteron™ 6000 Series platform have a large memory footprint and may offer enough compute power to host both your data and your analysis software on the same hardware, enabling a more efficient environment with reduced network traffic and software management.

CLOUD COMPUTING AND HIGH-PERFORMANCE COMPUTING

Organizations that manage large cloud computing and HPC server clusters typically prioritize power efficiency above other criteria when evaluating server hardware. But even in these high-density environments, performance and memory are important requirements for certain tasks, such as running large back-end databases or performing ASIC-chip verification. In those cases, 4P servers usually make more technical sense than 2P systems due to their superior capacity and compute density. Yet until now the comparatively high price of 4P platforms has discouraged most operators of cloud and HPC environments from using them. The low price of 2P servers has led them to become the platform of choice for HPC. Yet in environments where the interconnect is not saturated, consolidating two 2P platforms down to a single 4P can really make sense from both an architectural and a financial perspective.

Value 4P servers like those based on the AMD Opteron™ 6000 Series platform can enable organizations to deploy scalable, high-density servers cost-effectively. Equipped with up to 48 cores, four memory channels, and up to half a terabyte of high-speed DDR3 memory

capacity, 4P servers based on the 6000 Series platform offer very high FLOP rates and STREAM memory bandwidth up to 110GB per second.¹² As a result, cloud and HPC data centers can now economically tackle computationally-demanding workloads with fewer servers, helping save money on procurement and administration even as they reduce the number of network connections needed and shrinking the overall hardware footprint. By reducing the number of nodes (i.e. replacing a pair of 2P nodes with a single value 4P server), organizations often will be able to get the same levels of performance at lower costs, lower power consumption, and, most importantly, with greater manageability because

there are fewer systems in the cluster to manage.

Many cloud computing providers are using Memcached environments where they need large memory footprints distributed across multiple systems. In these environments, using fewer value 4P systems rather than larger numbers of low-priced 2P systems can allow customers to get to very large memory footprints without incurring the latency from a larger pool of systems. If CPU compute is not even a requirement, and memory addressability is the only requirement, then value 4P systems can deliver enhanced memory scalability at a very aggressive price point.

The AMD Opteron™ 6000 Series Platform in Action: The Dell™ PowerEdge™ R815

A powerful new rack server from Dell—the Dell™ PowerEdge™ R815—offers a compelling example of the groundbreaking value 4P products made possible by the AMD Opteron™ 6000 Series platform. Designed to support highly-threaded workloads such as virtualization and high-performance computing, the PowerEdge R815 packs the processing capacity, memory footprint, and I/O bandwidth of a 4P device into a 2U form factor with a 2P price point. This provides an innovative combination of high performance, compact size, and outstanding price for performance.

Sally Stevens, Dell's vice president of server platform marketing, credits the AMD Opteron™ 6000 Series platform with assisting Dell in achieving that milestone. "Customers

continually ask us for four things: More CPU power, more memory, more I/O, and more density," she says. "Thanks to the AMD Opteron™ 6000 Series platform, we're delivering on all of those requests in the new PowerEdge R815. The 6000 Series platform's high core count and substantial memory capacity were instrumental in enabling Dell to create the PowerEdge R815. And its breakthrough economics allow us to make all of the R815's power and efficiency available at a value normally associated with 2P systems."

Among the ways that Dell PowerEdge R815 servers and the AMD Opteron™ 6000 Series benefit companies with demanding workloads are these:

Higher Rack Density: Traditionally, x86-based 4P rack

¹² Based on measurements in AMD Labs as of March 17, 2010. Configuration: 4 x AMD Opteron™ processors Model 6174 in "Drachma" reference design kit, 64GB (16 x 4GB DDR3-1333) memory, SuSE Linux® Enterprise Server 11 64-bit.

servers have been available only as 4U devices. As a result, organizations in need of a 4P system's processing power and memory resources could fit no more than 10 such servers in a typical 42U rack. Thanks in part to the AMD Opteron™ 6000 Series platform's ample core count and memory footprint, however, the Dell PowerEdge R815 delivers 4P performance in a 2U case. Consequently, companies can install 21 R815s per rack and enjoy all of the capacity of a 4P system while more than doubling their rack density.

The resulting savings can be considerable. According to research by Dell, data center floor space currently costs as much as \$1000 per square foot on average. A 42U server rack enclosure usually occupies 17 square feet, plus another 10 square feet for associated power and cooling equipment. That means each 42U rack enclosure in an average data center potentially consumes up to \$27,000 worth of floor space. By extension, then, every time an organization replaces two racks containing 10 4P/4U servers each with a single rack containing 21 4P/2U Dell PowerEdge R815 servers it potentially frees up to \$27,000 worth of valuable data center real estate.

Greater Compute Density: 4P servers based on the AMD Opteron™ 6000 Series platform include up to 48 processing cores. A 42U rack enclosure bearing 21 Dell PowerEdge R815 servers, then, contains 1008 cores. A rack holding 10 previous-generation 4P/4U servers with six cores per processor, by contrast, contains just 240 cores. The PowerEdge R815, then, enables companies to increase compute density per rack over legacy 4P systems by up to 320 percent.

Furthermore, at 1008 cores per rack, the PowerEdge R815 delivers processing density rivaling that of blade

servers, but with greater flexibility. Blade chassis aren't interchangeable; buying one commits you to a specific server model as your capacity requirements increase. The PowerEdge R815, by contrast, deploys in standard server racks, so it can share space with any rack server you wish to use.

Enhanced Energy Efficiency: In addition to rack and floor space, the Dell PowerEdge R815 is designed to conserve power as well. For starters, the AMD-P suite of power management technologies built into all AMD Opteron processors enables the AMD Opteron™ 6000 Series platform to deliver 12-core performance within the same general power and thermal envelopes as earlier 6-core products. Secondly, the Dell PowerEdge R815 is a 2U server, so unlike 4P/4U systems it contains just two power supplies instead of four, resulting in lower power draw.

Furthermore, the PowerEdge R815 includes Dell Energy Smart technologies such as these:

- **Dell Energy Smart System Design:** Careful positioning of system components and internal shrouding, along with high-efficiency voltage regulators and Dell's efficient high-capacity fan technology, help keep PowerEdge R815 servers cool, so they use less energy. Plus, the PowerEdge R815 contains sensors that automatically help decrease power consumption by automatically tracking and regulating internal thermal activity.
- **Dell Energy Smart Management:** Capabilities such as power capping, power scheduling, and device disablement further help Dell PowerEdge R815 servers limit energy usage.

Intelligent Management Capabilities: Like all 11th generation Dell PowerEdge servers, the PowerEdge R815 comes with a variety of sophisticated management features that help simplify administration and lower operational costs, including:

- **Lifecycle Controller:** Most server makers require IT managers to consult a CD or DVD before performing common provisioning functions such as deployment, configuration, or updating. Dell PowerEdge servers, by contrast, make all of those functions available via the Lifecycle Controller, an embedded chip installed directly on the motherboard. By sparing administrators the hassle of relying on multiple pieces of easily-lost CD/DVD media, Lifecycle Controller streamlines the server configuration process and helps reduce downtime.
- **Interactive LCD screen:** Positioned on the front of every PowerEdge server, this miniature monitor enables technicians to view error messages, execute selected boot-up options, see power consumption information, and more.

- **Dell Management Console (DMC):** A standard feature on all Dell servers, the DMC offers IT managers a consolidated view of their IT infrastructure, including deployment, inventory, monitoring, and updates.

Proven Reliability: All Dell PowerEdge servers are designed from the inside out to provide long-term value. Dell uses only the most durable materials, emphasizes redundancy in all components, performs rigorous testing and validation, and emphasizes ease of servicing, to help minimize downtime.

Maximum Hypervisor Flexibility: Server manufacturers often limit their customers to one choice of hypervisor. Dell PowerEdge servers and the AMD Opteron™ 6000 Series platform, however, both support a wide range of hypervisors, including products from VMware, Microsoft, and Citrix. As a result, organizations using the PowerEdge R815 as a virtualization host server are free to select the right hypervisor for their specific needs.

Additional information about the Dell PowerEdge R815—including white papers, product specs, and other resources—is available at www.dell.com/poweredge/amd.

Conclusion

With their greater processing power and larger memory and I/O bandwidth, 4P servers are generally a better fit for challenging workloads such as virtualization and multi-threaded database applications than 2P servers. Historically, however, processor manufacturers have supported 2P and 4P servers via separate products in separate price bands—and the gap between those bands has often been too large to justify the performance advantage that 4P platforms deliver relative to 2P systems. Though the AMD Opteron™ processor has helped narrow that gap, many enterprise and mid-market companies continue to use large numbers of 2P servers to handle compute-intensive tasks better suited to smaller, more manageable numbers of 4P devices.

In an effort to make such compromises a thing of the

past, AMD has introduced the new AMD Opteron 6000 Series platform. By supporting both 2P and 4P server systems with the same product, the AMD Opteron 6000 Series platform enables 4P performance at 2P economics, giving organizations with high-end workloads a cost-effective way to deploy the right hardware for their needs while simultaneously simplifying administration and conserving data center floor space.

Of course, performance and scalability aren't top priorities for every workload, and 2P servers remain an excellent choice for many mainstream processing tasks. But organizations in need of server hardware with heavy-duty compute, memory, and I/O capacity should investigate the AMD Opteron 6000 Series platform before settling for a 2P product instead.

© 2010 Advanced Micro Devices Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Opteron, AMD Virtualization, AMD-V, and combinations thereof are trademarks of Advanced Micro Devices Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Microsoft, Windows, and Windows Server are registered trademarks of Microsoft Corporation in the United States and/or other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners.

PID # 48537-A

