

# Thermal Management & Server Density: Critical Issues for Today's Data Center

*Creating a Cooler Future through Innovations in Rack-Mount  
Server and Processor Design*

---

## I. CLUSTER COMPUTING: NEW CHALLENGES FOR THE DATA CENTER

The data center, a fixture in enterprise IT environments since the introduction of the mainframe computer 40 years ago, today faces its biggest operational challenge ever. Propelled by the skyrocketing popularity of x86-based cluster computing—which is replacing the proprietary RISC/Unix servers that, a decade ago, largely displaced mainframes—many data center managers are struggling to provide their facilities' two most essential commodities: power and cooling.

The culprit is, quite simply, heat. The two top sources of heat in the data center are servers' power supplies and processors, with processors producing 75 to 80 percent of the heat that must be dissipated. The processor also drives the server's overall power draw requirement; the faster the processor, the greater the power draw, and the more heat generated.

Recent technology advances have enabled dramatic improvements in PC server density, with rack-mount systems packing more, ever-faster processors into a smaller footprint. This in turn allows more computing power to be housed in relatively fewer square feet, concurrently increasing the amount of power required to operate the server and the amount of heat that must be dissipated. Thus, while delivering the flexibility and scalability that enterprises and High Performance Computing (HPC) environments are clamoring for, clusters of dense servers raise power consumption levels and amplify thermal management challenges.

### ***Hitting the wall: Data center limitations***

It's not just antiquated data centers that are being taxed by computing clusters. Facilities built as little as two years ago are ill-equipped to hold today's high-density server racks, due to inherent limitations in their power utilization and cooling output. When data centers are constructed, they are subject to a physical limit as to how much power can be delivered to the site and stored. Similarly, data centers' air conditioning systems are limited in how many BTUs of cooling power can be delivered in a given space and, in public data centers, a predetermined level of wattage per cabinet that can be dissipated. The problem is so pervasive that through year-end 2008, heat and cooling requirements for servers will prevent 90 percent of enterprise data centers from achieving the maximum theoretical server density (0.7 probability).<sup>1</sup>

Server density and thermal management issues present more than just a facilities management challenge. They are omnipresent problems that adversely impact the three business dynamics that drive the data center: reliability, maintenance and serviceability, and cost. Lowering the amount of heat produced by servers can dramatically improve all of these factors:

**Reliability:** If a server can be made to run cooler, a large number of high-failure parts (such as power supply, hard drive, case fan and CPU fan) will have longer lifetimes, increasing the machine's reliability and reducing downtime. Other parts that tend to have lower mean time between failure (MTBF), such as CPUs and memory, can experience much higher failure rates if server temperatures become excessive.

**Maintenance and serviceability:** Higher server reliability also translates into lower overall lifetime maintenance and service costs, with longer MTBF.

**Cost:** Higher reliability yields a third related financial benefit, lowering the cost of replacement components associated with failed machines. It also impacts another large line item: power. Only about half of total IT costs are related to hardware and software; the rest are in administration and overhead costs. Power, and the cost of managing power elements within the data center, comprises a large part of operational costs.

### ***Rackable Systems and AMD: Definitive solutions***

Clearly, lowering the heat produced, and power consumed, by high density servers is a foremost goal of all designers and manufacturers of servers and processors. Rackable Systems and AMD today offer breakthrough technology that minimizes power consumption and heat generation, allowing IT managers to maximize the amount of cluster computing power they can extract from existing data center facilities.

In 1999 Rackable Systems designed the industry's first-ever half-depth, back-to-back rack-mount server designs. Rackable Systems further enhanced these dense servers' thermal management capabilities with the 2003 introduction of Distributed DC Power Supply technology. Developed by Rackable Systems, It significantly improves system availability, allows servers run to cooler and more efficiently, and delivers impressive cost savings.

---

On a per-system level, Rackable Systems servers fitted with DC Power supplies and the AMD Opteron™ processor HE consume up to 40 percent less power, and produce up to 40 percent less heat than servers equipped with competitive processors.

---

Distributed DC Power also enables Rackable Systems' servers to easily accommodate the latest high-speed processors while still running cool. These include the low-power AMD Opteron processor HE, which can ameliorate data center thermal management issues and allow more powerful clusters to operate in the same amount of space.

With greatly improved thermal management and lower power requirements, the combination of Rackable Systems' designs and the AMD Opteron processor HE presents a quantum improvement in allowing more computing power to be housed in existing data centers. For enterprises hungry for more computing power, but challenged to acquire it within their facilities' constraints, Rackable Systems-AMD Opteron processor HE solutions present some of the most exciting solutions available today.


## II. COOL BY DESIGN: RACKABLE SYSTEMS HIGH DENSITY SERVERS

Rackable Systems' patented server design is at the heart of its ability to build high performance, high-density rack-mount servers. The company solves traditional servers' density and thermal management shortcomings by mounting its servers back-to-back. In doing so, the company has patented an innovative way to address the problems caused by traditional rack-mount servers, which are typically stacked in a single row. These problems include:

- lower density per rack
- the creation of "hot spots" in the data center as heat is pushed out the back of the servers and congregates in one area
- complications in servicing the equipment since cables are snaked out the back, making it difficult to quickly identify, remove and service a failed server.

Rackable Systems' approach entails minimizing the number of components inside the server, and putting input/output (I/O) mechanisms such as cables and ports on the front of the server. Not only is density doubled in the rack, but heat from the servers is vented out the top of the cabinet, providing immediate cooling benefits for the data center. Serviceability is also greatly improved, since cabling and the LCD display are visible and accessible in front.

Rackable Systems' thoughtful approach extends to the design of each server. Rather than building "general-purpose" servers, the company designs systems with a focus on minimizing power requirements. When specifications are drawn up, Rackable Systems' engineers strategically choose power supplies to lower the necessary power draw as much as possible. For example, a 250-watt power supply is used for a server with a 200-watt requirement; other manufacturers may use a power supply of up to 500 watts to satisfy the same 200-watt need. While this overage is not significant in small installations, in large clusters it drives up power consumption considerably, as well as the amount of heat that must be dissipated.



Rackable Systems' patented "back-to-back" rack mounting technology enables:

- twice the CPU density
- improved serviceability
- optimized thermal efficiency

## ***DC Power: Improving business drivers in the data center***

Rackable Systems is the first server provider to manufacture standard rack-mounted servers with Distributed DC Power technology and make them generally available. The DC Power Supply solves the heat generation problems found in data centers fitted with many servers, each with a dedicated alternating current (AC) power supply. By distributing redundant direct current (DC) power to each server—and replacing the standard AC power supply with a more reliable DC power supply—approximately 20 to 40 percent of the thermal load is successfully shifted outside the server to AC-to-DC rectifiers, which can easily be installed either at the top of the cabinet or outside near an HVAC (cooling) unit. The DC power supplies are 93 percent efficient and do not generate a significant amount of heat.

### **RACKABLE SYSTEMS SERVER HIGHLIGHTS**



- Maximized density: Up to 176 processors per cabinet
- Easy to service with front-facing I/O
- Highly flexible configurations with the latest processor technology
- Integrated, out-of-band, lights-out remote management
- Ideal for multi-server deployments
- Choice of AC or DC input power

Announced in November 2003, this advance is extremely timely, as demand for cluster computing solutions and ongoing innovation in processor technology create new challenges and opportunities. Specifically, as CPU power consumption and thermal output reaches and surpasses 100 watts per processor, the primary high-density server design challenge has shifted from merely achieving desired density levels to also keeping the servers cool. Today's servers with the highest-speed processors frequently require higher-wattage AC power supplies—typically, the higher the wattage power supply, the less its efficiency. This results in more heat inside the chassis, leading to higher system failure rates and greater power consumption.

Compatible with both AC- and DC-based data center environments, Distributed DC Power cluster solutions squarely address these challenges and extend the advantages of Rackable Systems' server designs. As a result, Rackable Systems' Distributed DC Power technology improves the three business drivers that govern data center operations:

**Reliability:** The DC power supplies in the individual servers have no moving parts (unlike AC supplies) and have an MTBF of over 7 million hours, versus the typical 100,000 hours of a normal AC power supply. This 70-fold increase in reliability at the power supply level increases overall server reliability by approximately 27 percent. Because of these factors, DC-based servers have much lower failure rates in comparison to similarly configured AC-based systems.

An even larger factor in the increased availability and reliability is that Rackable Systems' Distributed DC Power solution provides similar improvements as N+M redundant power supplies, but without the extra cost or extra heat. This is possible because of the redundancy at the rectifier level. Spare rectifiers allow single rectifiers and even entire AC power circuits to fail without affecting system power.

**Maintenance and serviceability:** Distributed DC Power solutions alleviate formidable challenges in the areas of reliability and CPU density, while significantly reducing data center maintenance costs. The reduction in heat increases overall server reliability, thereby cutting maintenance costs. The use of Distributed DC Power also improves configuration flexibility and reduces related maintenance requirements.

**Cost:** Because DC power is inherently more efficient than AC, power savings of up to 40 percent can be achieved in Rackable Systems servers equipped with Distributed DC Power supplies.

Combined with Rackable Systems' superior Remote Management solution, which delivers complete lights-out management capability through a secure serial interface, Rackable Systems Distributed DC Power solutions enable all data centers to achieve the flexibility, serviceability, management and reliability that is critical today.

### **III. THE AMD OPTERON PROCESSOR HE: HIGH PERFORMANCE WITH ULTRA-LOW POWER CONSUMPTION**

The AMD Opteron processor is a top choice for cluster computing, providing a dramatic leap forward in compatibility, performance, investment protection and reduced total cost of ownership (TCO). The AMD Opteron processor HE is a low-wattage (55 watt) version of this versatile processor, delivering industry-leading performance per watt that is ideal for rack-dense servers, such as Rackable Systems', which are used in data center environments. Each low-power AMD Opteron processor offers equivalent performance to the full-power AMD Opteron processor with the same model number. For example, the AMD Opteron processor Model 146 HE provides equivalent performance at almost 40 percent lower wattage compared to the full-power AMD Opteron processor Model 146.

The consistent 940-pin design of the AMD Opteron processor HE variants allows for compatibility with the existing infrastructure used for corresponding full-power AMD Opteron processors. Rackable Systems servers can be interchangeably equipped with either the standard or low-wattage versions of the AMD Opteron processor, with no hardware modifications necessary.

#### ***Ideal for cluster applications***

An evolutionary processor, the AMD Opteron delivers two key benefits to users with cluster computing applications:

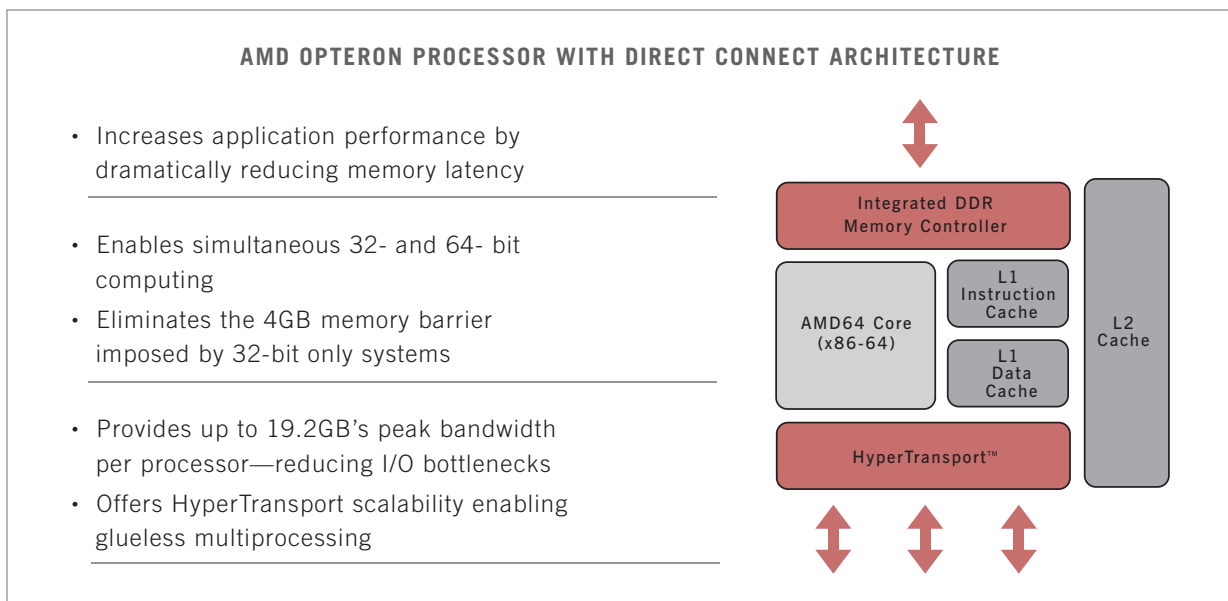
**Compatibility:** The AMD Opteron is the first x86-based processor that allows users to simultaneously run both high-performance 32-bit applications and 64-bit applications. This enables organizations to migrate their applications to 64-bit environments at their own pace, and does not require an either-or choice to be made between 32- and 64-bit.

The aforementioned interchangeability between the standard AMD Opteron processor and the low-wattage AMD Opteron processor HE provides an additional strong compatibility benefit.

**Performance:** An x86 processor that is 64-bit capable, the AMD Opteron processor readily accommodates high-performance 64-bit applications while giving a boost to 32-bit applications. The latter is due to its elimination of the four GB memory address ceiling in 32-bit systems, and by no longer requiring that the operating system run in the same partition as 32-bit applications. Both of these improvements can substantially improve the performance of 32-bit applications, while offering users a streamlined migration path to 64-bit computing.

The AMD Opteron processor with Direct Connect Architecture increases memory latency performance, provides more balanced throughput and I/O, and allows for more linear symmetrical multiprocessing. Specific features that allow for enhanced overall performance include:

- o *An integrated memory controller* that reduces latencies during memory access in a symmetric multiprocessor (SMP) server system, to yield fast computational processing for increased performance and productivity
- o *HyperTransport™* technology that provides high-speed I/O links that provides up to 19.2 GB peak bandwidth per processor, reducing I/O bottlenecks. Up to three coherent HyperTransport links can be supported, to deliver substantial I/O for current and future applications needs.
- o *Scalability* from one-way to eight-way across entire data or compute centers utilizing the same hardware and software infrastructure, allowing for maximum flexibility in IT infrastructure.
- o *256 TB of memory address space*, creating a significant performance benefit for applications in which large (or many) datasets are held in memory.



### ***Improved thermal management and lower power consumption***

As noted at the beginning of this paper, processors are a leading source of server heat generation within cluster environments. The AMD Opteron processor HE resolves specific heat and power issues, providing high performance at a lower wattage—enabling data centers to better manage costs associated with server space, thermal control and power consumption. Rackable Systems servers equipped with the AMD Opteron processor HE allow more powerful clusters to be housed in the same amount of data center space, while improving productivity and budget efficiency.

## **IV. SUMMARY**

The data center, the site of enterprise computing innovation for four decades, is suddenly challenged to support ever-increasing amounts of computing power. With the rapid expansion of growing numbers of high-density computing clusters, many data center managers are struggling to provide their facilities' two most essential commodities: power and cooling. Facilities built as little as two years ago are ill-equipped to hold today's high-density server racks, due to inherent limitations in their power utilization and cooling output.

Server density and thermal management issues are more than just a facilities management challenge. They adversely impact the data center's three business dynamics: reliability, maintenance and serviceability, and cost.

With these business drivers in mind, lowering heat production and power consumption is a foremost goal of all designers and manufacturers of servers and processors. Rackable Systems and AMD today offer breakthrough technology that minimizes power consumption and heat generation, allowing IT managers to maximize the amount of cluster computing power they can extract from existing data center facilities.

Rackable Systems' half-depth, back-to-back rack-mount server designs enable increased density, largely eliminate "hot spots" in the data center and simplify serviceability. The company has further enhanced its dense servers' thermal management capabilities with Distributed DC Power Supply technology. Far more reliable and 20 to 40 percent cooler than AC power technology, Rackable Systems' Distributed DC Power appreciably improves system availability, allows servers run to cooler and more efficiently, and delivers noteworthy cost savings.

It also enables Rackable Systems servers to easily accommodate the latest high-speed processors while still running cool. These include the low-power AMD Opteron processor HE, which offers equivalent performance to its full power AMD Opteron processor counterpart, at almost 40 percent lower wattage. Low-power processors like the AMD Opteron processor HE can reduce data center thermal management issues, allowing more powerful clusters to operate in the same amount of space.

With greatly improved thermal management and lower power requirements, the combination of Rackable Systems' designs and the AMD Opteron processor HE presents a quantum improvement in allowing more computing power to be housed in existing data centers. For enterprises hungry for more computing power, but challenged to acquire it within their facilities' constraints, Rackable Systems-AMD Opteron processor HE solutions are, in fact, the most promising high-density cluster servers available today.

### **ABOUT RACKABLE SYSTEMS**

Founded in 1999 and headquartered in San Jose, Calif., Rackable Systems leads the industry in large-scale data center deployment solutions. Through innovative design, patented technology and commitment to customer service, the company offers a broad line of server products that can be rapidly deployed in data centers and other demanding IT environments. Serving leading-edge technology companies, universities and research & development organizations, Rackable Systems' customers include Cold Springs Harbor Laboratories, Deutsche Bank, Electronic Arts, Google, Lawrence Livermore National Laboratory, nVidia, Pfizer, Stanford Linear Accelerator Center, Sony America, Tellme Networks, Toshiba America Electronic Components, Inc., U.C.S.C. Human Genome Project, U.S. Department of Justice, Webex and Yahoo!.

To learn more about Rackable Systems, go to [www.rackable.com](http://www.rackable.com), call 408-240-8300, or write to [info@rackable.com](mailto:info@rackable.com).

### **ABOUT AMD**

AMD (NYSE:AMD) designs and produces innovative microprocessors, Flash memory devices and low-power processor solutions for the computer, communications and consumer electronics industries. AMD is dedicated to delivering standards-based, customer-focused solutions for technology users, ranging from enterprises and governments to individual consumers. For more information visit [www.amd.com](http://www.amd.com).

<sup>1</sup> "Gartner Predicts 2004 Future of Server Density," Gartner, J. Enck, G. Weiss, November 2003.