Advancing scalability and performance
with four-socket Dell PowerEdge servers

By Armando Acosta and Robert Bradfield

Four-socket Dell™ PowerEdge™ R810 and PowerEdge R910 rack servers and PowerEdge M910 blade servers offer highly scalable processing power and memory capacity to help cost-effectively boost performance and expand virtualized environments, while other features help increase reliability and simplify administration.

Enterprises running large-scale databases or virtualized environments need servers that let them scale processing power and memory capacity flexibly and cost-effectively. Because servers in these types of environments run everything from enterprise productivity tools to mission-critical business applications, they must provide outstanding reliability. And to help optimize worker productivity and enhance business agility, servers must also be quick to deploy, simple to manage, and easy to maintain. The Dell PowerEdge R810 and PowerEdge R910 rack servers and PowerEdge M910 blade servers are designed to meet these requirements, combining scalable processing performance and tremendous memory capacity with a dense server design, reliability features, and streamlined management.

Achieving scalable compute performance

Dell PowerEdge R810 and PowerEdge R910 rack servers and PowerEdge M910 blade servers are designed to provide exceptional scalability within each server, helping organizations to quickly and cost-effectively expand database environments or add virtual machines to each physical host. In the past, scaling processor performance frequently meant adding servers. With the PowerEdge R810 rack server and PowerEdge M910 blade server, however, organizations can begin with a two-socket configuration and then scale to four sockets when they need to increase application performance or extend a virtualized environment (see Figure 1). Organizations can therefore double the processing performance within a server without needing to provide

Dell PowerEdge R810 rack server

Dell PowerEdge M910 blade server

Dell PowerEdge R910 rack server

	Dell PowerEdge R810 rack server	Dell PowerEdge M910 blade server	Dell PowerEdge R910 rack server
Sockets			
	Two or four	Two or four	Four
Form factor			
	2U rack server	Full-height blade server	4U rack server
Processors			
	Intel Xeon processor 7500 or 6500 series	Intel Xeon processor 7500 or 6500 series	Intel Xeon processor 7500 series
Memory slots			
	32	32	64
Maximum memory capacity			
	512 GB	512 GB	1 TB
PCI Express (PCIe) slots			
	6 PCIe slots	Up to 4 mezzanine card slots	Up to 10 PCIe slots

Figure 1. Four-socket Dell PowerEdge rack and blade server configurations

Maximum performance

Forrest Norrod, vice president of the Dell Server Group, answers questions on the latest Intel Xeon processor–based Dell servers and how they can help IT organizations increase performance, reliability, and efficiency for mission-critical applications.

youtube.com/watch?v= Y7OCYT9WSms

the extra power, cooling, or data center space associated with additional servers.

Each of the four sockets in these PowerEdge servers can be filled with powerful multi-core Intel® Xeon® processors that offer significant performance enhancements compared with previous-generation PowerEdge servers. The PowerEdge R810 rack server and PowerEdge M910 blade server can be equipped with the Intel Xeon processor 7500 or Intel Xeon processor 6500 series; the PowerEdge R910 rack server can be equipped with the Intel Xeon processor 7500 series. Using this generation of Intel Xeon processors, the PowerEdge R910 rack server has achieved up to 219 percent higher performance compared with the previous-generation PowerEdge R900 rack server, and the PowerEdge M910 blade server has achieved up to 76 percent higher performance compared with the PowerEdge M710 blade server in a four-socket configuration.[1]

By using the latest multi-core Intel Xeon processors, these servers can dramatically increase the number of processor cores per rack compared with previous-generation PowerEdge servers. In its four-socket configuration, the PowerEdge R810 server can be equipped with 32 processor cores in a single 2U enclosure—up from a maximum of 12 cores in previous-generation 2U PowerEdge servers. Both the PowerEdge R910 4U rack server and the

[1] Based on SPECjbb2005 performance testing by Dell Labs in March 2010; actual performance will vary based on configuration, usage, and manufacturing variability. For the latest SPECjbb2005 benchmark results, visit spec.org.

219%

The Dell PowerEdge R910 rack server has achieved up to 219 percent higher performance than the PowerEdge R900.

76%

The Dell PowerEdge M910 blade server has achieved up to 76 percent higher performance than the PowerEdge M710.

1/3

The Dell PowerEdge M910 blade server provides double the memory capacity of previous-generation 4U Dell servers in one-third the space.

PowerEdge M910 full-height blade server also support up to 32 cores per server, compared with a maximum of 24 cores in the PowerEdge R900 rack server and 12 cores in the PowerEdge M710 blade server. In addition, the thread-level parallelism provided by Intel Hyper-Threading Technology enables highly efficient use of these processing cores—providing a substantial performance boost for multi-threaded software compared with the same number of cores without this technology.

Packing additional cores into each server enables organizations to support large databases and large-scale virtualized infrastructures while helping to keep power, cooling, and space requirements under control. For example, the PowerEdge M910 server enables the deployment of up to eight four-socket blade servers in just 10U of rack space. Compared with previous-generation 4U Dell servers, the PowerEdge M910 server is designed to

provide double the memory capacity in one-third the space.

The powerful processing capabilities and compact designs of these servers contributes to exceptional price/performance ratios. At the same time, supporting additional processors in a reduced amount of space helps organizations minimize power consumption. In four-socket configurations, these PowerEdge servers also deliver outstanding performance per watt compared with previous-generation PowerEdge servers.

Breaking the server memory capacity barrier

Many organizations running large enterprise applications and virtualized environments find that server memory capacity is a primary limitation for scaling. These four-socket Dell PowerEdge servers accommodate significantly greater memory capacity than previous-generation PowerEdge servers, enabling organizations to break through those limitations. The PowerEdge R810 and

PowerEdge M910 servers both offer 32 memory slots for up to 512 GB of memory in each server; for the PowerEdge M910 server, that is more than twice the maximum memory capacity of previous PowerEdge blade servers. The PowerEdge R910 server provides 64 memory slots for up to 1 TB of memory.

These large memory capacities enable organizations to increase the number of virtual machines per server or enhance the performance of memory-bound database applications. They can also capitalize on the large number of memory slots to help reduce costs: populating these slots with smaller, less-expensive memory modules enables organizations to provision sufficient memory without incurring the cost of larger-capacity modules.

In addition, Dell FlexMem Bridge technology allows organizations to take advantage of large memory capacities without having to scale the number of processors. This innovative, patent-pending

technology in PowerEdge R810 and PowerEdge M910 servers enables organizations to access the capacity of all 32 memory slots while using only two of the four available processor sockets. While other servers typically require populating all four sockets to gain access to this type of large-scale memory capacity, Dell FlexMem Bridge technology enables organizations to achieve a suitable mix of memory and processing performance for their particular environments without incurring the expense of purchasing and running large servers. And because virtualization software licensing costs are often calculated per socket, this technology can help organizations avoid the fees associated with those extra processors—enabling them to increase virtual machine density while controlling licensing costs.

Boosting reliability and simplifying management

These four-socket Dell PowerEdge servers also incorporate a range of features designed to enhance reliability for mission-critical workloads. For example, they include dual internal Secure Digital (SD) modules for embedded hypervisor redundancy. The embedded hypervisors can accelerate virtualization deployment, and including two modules enables failover to help minimize downtime even if there is a hardware problem with one SD module.

Using advanced reliability, availability, and serviceability

(RAS) technology features, processors from the Intel Xeon processor 7500 and Intel Xeon processor 6500 series are designed to automatically monitor, report, and recover from hardware errors to maintain data integrity and keep mission-critical applications online. In addition, the physical design and production of these servers help ensure reliability and increase uptime. PowerEdge M910 blade servers, for example, capitalize on the multiple power, cooling, and networking infrastructure redundancies provided by the PowerEdge M1000e modular blade enclosure to help avoid unplanned outages.

Simplifying server management is also critical to enhancing IT efficiency. These PowerEdge servers incorporate key management capabilities introduced in 11th-generation PowerEdge servers to help reduce the time and effort required for a range of administrative tasks. The Dell Lifecycle Controller in each of these servers is an integrated chip that helps simplify tasks by enabling administrators to access a comprehensive set of management functions—including driver installation, firmware updates, hardware configuration, and diagnostics—in a pre-OS environment. Using a single intuitive interface called the Unified Server Configurator, administrators can get to work as soon as the server is powered up, without needing additional media or even requiring an OS to be installed—helping

speed up server deployment, save time on maintenance, and minimize downtime.

Enhancing resiliency and performance

Dell PowerEdge R810 and PowerEdge R910 rack servers and PowerEdge M910 blade servers are designed to increase the performance of large database-intensive applications and scale virtualized server environments rapidly and cost-effectively. From the integration of Dell FlexMem Bridge technology to the inclusion of dual SD modules for embedded hypervisor redundancy, many of the advances included in these PowerEdge servers resulted from discussions with IT professionals who sought ways to enhance scalability, increase reliability, and simplify management while controlling costs. These PowerEdge servers are designed to boost the performance of core business applications, expand database environments, and ultimately increase the overall efficiency of the enterprise. PS

Armando Acosta is a senior product line consultant at Dell and has more than 12 years of experience in the IT industry.

Robert Bradfield is a senior product line consultant at Dell and has more than 15 years of experience in the IT industry.

Learn more

Dell PowerEdge servers:
dell.com/poweredge

PowerEdge R810

Armando Acosta gives a peek inside Dell FlexMem Bridge and other key features of the PowerEdge R810 rack server.

youtube.com/watch?v=
KoMIxtDvK68

PowerEdge M910

Robert Bradfield provides a rundown of the design features of the PowerEdge M910 blade server, including the network fabrics and dual internal SD modules.

youtube.com/watch?v=
iUnc1o0xLXE

PowerEdge R910

A Dell team discusses the architecture of the PowerEdge R910 rack server and the advantages it offers—including manageability, ease of access, and efficient power and cooling.

youtube.com/watch?v=
oulHU7hGRDM