OPTIMIZING DDR3 MEMORY SETTINGS IN NEW 11TH-GENERATION DELL POWEREDGE SERVERS

By Paul Benson

Dell™ PowerEdge™ R610 and PowerEdge R710 rack servers, PowerEdge T610 tower servers, and PowerEdge M610 and PowerEdge M710 blade servers feature the Double Data Rate 3 (DDR3) memory architecture, which enhances flexibility by enabling up to three memory channels per processor. By understanding best-practices DDR3 configurations, administrators can optimize settings for targeted enterprise applications and virtualized environments.

Related Categories:

Blade servers

Dell PowerEdge blade servers

Dell PowerEdge rack servers

Dell PowerEdge servers

Dell PowerEdge tower servers

Memory

Servers

Visit DELL.COM/PowerSolutions for the complete category index.

The new 11th-generation Dell PowerEdge server family takes advantage of the Intel® Xeon® processor 5500 series, which utilizes Double Data Rate 3 (DDR3) memory. By providing the flexibility to configure up to three memory channels per processor, the DDR3 architecture allows administrators to optimize memory configurations to suit specific requirements for a wide range of enterprise usage scenarios.

DETERMINING MEMORY REQUIREMENTS

With its flexible design, memory can be configured to meet the demands of a variety of usage scenarios. *High-performance mode* may be appropriate for environments in which memory speed is a top priority—such as servers running enterprise applications and supporting a virtualization infrastructure. For example, if a Web server is primarily delivering user-specific content, then new pages must be created frequently. This scenario typically calls for a high-performance memory configuration because the pages are likely not cached. Application commands must be processed as efficiently as possible, but high-capacity memory is not required because there is little reason to hold on to the result in memory past the time it is displayed to users or stored on disk.

Balanced-performance mode is applicable in a wide variety of business scenarios. Designed for situations that require the best of both worlds—good performance and midrange capacity—this configuration allows for up to twice the memory capacity available in high-performance mode.

High-capacity mode is designed to provide up to three times the memory capacity available in high-performance mode, although the trade-off is low memory speed. High-capacity configurations may be suitable for Web servers handling pages that do not vary from one user to the next. Large-capacity systems are also generally appropriate for virtualized environments and for organizations that have large databases that they want to pull into memory to help avoid increased disk access times.

Mirror mode facilitates balanced performance while providing for the correction of memory errors. It provides a memory backup, so only half of the installed memory is reported to the OS. A data integrity configuration may be appropriate, for example, for 24/7 transaction processing scenarios in which an enterprise cannot afford to take its servers offline for anything other than scheduled maintenance.

DDR3 memory can also be configured for substantial power savings when heating and cooling usage in the data center is a critical concern. Unbuffered dual in-line memory module (UDIMM) technology is typically best suited for power-conscious organizations, as it is designed to consume 1 W less power per memory module than registered DIMM (RDIMM) technology.

SELECTING DDR3 OPTIONS FOR TARGET USAGE SCENARIOS

Each processor in the new 11th-generation Dell PowerEdge server family includes three separate memory controller hubs (MCHs) within the processor package. This design avoids the requirement to transfer memory transactions between the processor and an external device. However, the inclusion of three MCHs per processor introduces considerations about the appropriate way to populate and configure a system based on specific requirements for reliability, availability, and serviceability (RAS) features and speed (see the "DDR3 configuration in 11th-generation Dell PowerEdge servers" sidebar in this article).

The integrated memory controllers in the Intel Xeon processor 5500 series used in Dell PowerEdge R610 and PowerEdge R710 rack servers, PowerEdge T610 and PowerEdge T710[1] tower servers, and PowerEdge M610 and PowerEdge M710 blade servers support DDR3 technology, which is designed to provide a high-performance memory interface capable of low-latency response and high throughput. DDR3 memory is also designed to enable higher bandwidth with lower power usage than DDR2 memory.

New 11th-generation PowerEdge servers support both error-correcting code (ECC) DDR3 RDIMMs and ECC DDR3 UDIMMs at bus speeds of 800 MHz, 1,066 MHz, and 1,333 MHz. PowerEdge R610, PowerEdge T610, and PowerEdge M610 servers have 12 memory sockets split into two sets of 6 sockets, with one set for each processor; PowerEdge R710, PowerEdge T710, and PowerEdge M710 servers have 18 memory sockets split into two sets of 9 sockets, also with one set for each processor. The memory interface supports memory demand and patrol scrubbing, single-bit error correction, and multi-bit error detection.

Offering capacities up to 8 GB, RDIMMs are suitable for a large amount of memory, extensive RAS features, and maximum expandability. RDIMMs enhance expandability because they allow for three DIMMs per channel, instead of two in UDIMMs. In addition, RDIMMs offer address parity, a RAS feature designed to halt the system if an incorrect address is detected. However, unless business requirements call for very high memory capacity, UDIMM ECC is typically a cost-effective alternative to a comparable RDIMM configuration.

The memory subsystem can run in one of three modes:

- **Memory optimized:** This mode uses all three memory channels in a processor, and populates all three channels identically. Although it enables the most DIMM population flexibility and system memory capacity of the three memory subsystem modes, this mode does not support memory mirroring, nor does it support single-device data-correction (SDDC) for x8-based DIMMs. However, it does support SDDC for x4-based DIMMs.
- **Advanced ECC:** This mode joins two controllers into a lockstep mode, thus creating a 128-bit data path. The advantage of this configuration is that it allows SDDC to work for both x4- and x8-based memory devices and supports SDDC features in both types. Memory modules must be identical in size, speed, and technology in corresponding slots.
- **Mirror:** This mode populates two channels per processor identically; the third channel is left unused. The two channels operate as mirrors of each other, with

DDR3 CONFIGURATION IN 11TH-GENERATION DELL POWEREDGE SERVERS

Each processor in 11th-generation Dell PowerEdge servers has three memory controllers, and up to three dual in-line memory modules (DIMMs) can be attached to each controller (see Figure A for an example of the processor and memory layout in a Dell PowerEdge R710 rack server). In models with 12 DIMM slots, each controller has two DIMM slots assigned to it. Both Dell and Intel strongly advise configuring systems in a balanced configuration, in which each channel is similarly populated. However, administrators should keep in mind that if the system has only a single processor, any memory installed in the slots assigned for the second processor cannot be accessed.

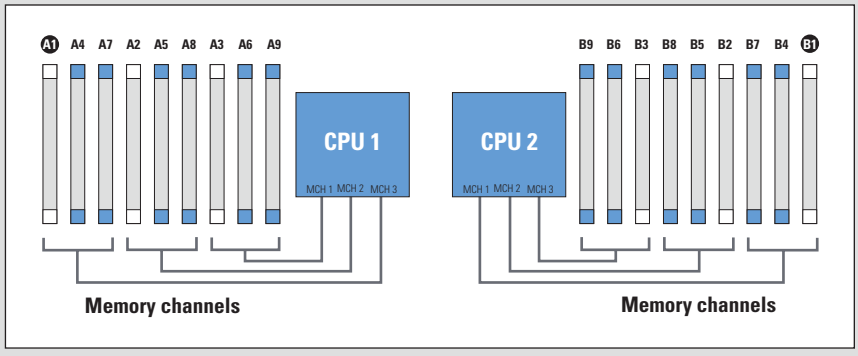

Figure A. *Dell PowerEdge R710 processor and memory layout supporting up to three DIMMs per channel*

[1] The PowerEdge T710 server was not yet released at press time; availability is planned for June 30, 2009.

writes going to both channels and reads alternating between the two channels. This mode is appropriate for typical enterprise applications. If excessive correctable errors occur or an uncorrectable multi-bit error occurs, the problematic channel is disabled and forthcoming memory accesses are routed through the functional channel. The user is informed of this event to service the system. Because the memory is mirrored, half of the installed memory is actually seen and reported by the OS.

This range of memory modes offers several potential advantages. Organizations can specify a single base server model to match specific applications or purposes, which helps simplify the server provisioning process. In addition, IT administrators have the flexibility to modify the server configuration for different applications or purposes at a later time—enhancing its long-term cost-effectiveness.

CONFIGURING DDR3 MEMORY FOR EACH USAGE SCENARIO

While specific requirements vary for individual organizations, understanding best-practices configurations for various DDR3 memory options in typical usage scenarios helps IT administrators achieve optimal performance and energy efficiency in targeted enterprise applications and virtualized environments.

As part of the memory configuration process, selecting a processor that supports required memory speed can be a key consideration. In high-performance configurations at the time of release, DDR3 memory at 1,333 MHz across all three channels per processor allows maximum bandwidth. Frequency limitations mean that this speed can support only one DIMM per channel, and thus a maximum of three DIMMs per processor. Dual-rank 4 GB RDIMMs are designed to provide peak capacity at 1,333 MHz and enable a system capacity of 12 GB per processor.[2]

System memory speed (the speed at which the memory is actually running) is set by the BIOS depending on processor capability, DIMM types used, and the number of DIMMs populated per channel. To run at maximum frequency, administrators should configure no more than one DIMM per memory channel. To help avoid errors in the memory subsystem, the system lowers the memory clock speed as DIMMs are added to each channel. Therefore, this approach supports the least amount of memory but at the highest possible speed among the configurations described in this article.

In a balanced-performance configuration, DIMMs should be installed across all three channels—but DIMMs at 1,066 MHz may be used and can be populated with two DIMMs per channel. Administrators may enable a capacity of 48 GB per processor by installing dual-rank 8 GB RDIMMs at 1,066 MHz.

As of the March 30, 2009, release date, 8 GB RDIMMs offer the highest capacity available for 11th-generation Dell PowerEdge servers. By using dual-rank 8 GB RDIMMs at 1,066 MHz, administrators may enable a capacity of 72 GB per processor for high-capacity configurations. In this case, the system is designed to automatically down-clock itself to 800 MHz because each memory channel will have three DIMMs installed.

For enterprises that require SDDC support for x8-based memory or mirroring in the data integrity configuration, only two channels from each processor should be populated. The memory channel farthest from the processor should not have any DIMMs installed—and it is vital that the DIMMs across the channels match each other.

Power-conscious IT departments can trade low capacity for reduced power consumption by using UDIMMS, which are designed to use less power than RDIMMs. However, UDIMMs are currently limited to a maximum 2 GB capacity and cannot support three DIMMs per channel. They also do not support address parity. Memory frequency is also a component in power consumption—for example, a 1,066 MHz configuration is specified to use less power than one set up as 1,333 MHz.

FINE-TUNING DDR3 FOR SPECIFIC MEMORY REQUIREMENTS

New 11th-generation Dell PowerEdge servers are designed to take advantage of the DDR3 memory architecture utilized by the Intel Xeon processor 5500 series. DDR3 is designed to support a range of performance requirements for enterprise applications and virtualized environments by enabling administrators to configure up to three memory channels per processor. Following best-practices guidelines for various usage scenarios helps administrators configure DDR3 memory options for optimal performance, capacity, data integrity, and energy efficiency.

Paul Benson was part of the 11th-generation Dell PowerEdge development team at Dell, with memory being one of his focus areas. He has a B.E. in Electrical Engineering, Computer Science, and Mathematics and an M.S. in Electrical Engineering from Vanderbilt University, as well as an M.B.A. from the University of North Carolina.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Efficient Enterprise Portfolio:
DELL.COM/Switch

Dell PowerEdge servers:
DELL.COM/PowerEdge

[2] Dell plans to continue looking into enhancing systems to support higher frequencies with additional memory.