

DDR3 memory technology

Technology brief, 3rd edition

Introduction	2
DDR3 architecture	2
Types of DDR3 DIMMs	2
Unbuffered and Registered DIMMs	2
Load Reduced DIMMs	3
LRDIMMs and rank multiplication	3
DDR3 memory speeds	4
Memory power consumption	5
Low-voltage DDR3 memory for ProLiant G6 and G7 servers	6
HP SmartMemory for ProLiant Gen8 servers	6
Core DDR3 Technologies	7
Fly-by topology	7
On-die Termination	9
Address parity checking on RDIMMs	9
Integrated DIMM temperature sensor	9
DDR3 memory and NUMA systems architectures	10
Older server architectures	10
DDR3 and NUMA systems architecture	11
DDR3 memory and 2P ProLiant servers	12
DDR3 and 4P ProLiant architecture	12
Intel 4P architecture and DDR3	13
AMD 4P architecture and DDR3	13
Memory throughput with DDR3	14
System memory bandwidth	15
DDR3 latency	15
Achieving highest performance with DDR3 memory	16
Maximizing system throughput	16
Minimizing memory latency	16
Using balanced memory configurations	16
Conclusion	17
For more information	18



Introduction

DDR3, the third generation of Dual Data Rate (DDR) Synchronous DRAM memory, delivers significant performance and capacity improvements over older DDR2 memory. HP introduced DDR3 memory with the G6 and G7 ProLiant servers, coinciding with the transition to server architectures that use distributed memory and on-processor memory controllers. DDR3 continues to evolve in terms of speed and memory channel capacity, and the new HP ProLiant Gen8 servers fully support these improvements. This paper provides a detailed look at the core technologies that enable DDR3 memory and its benefits as well as the integration of DDR3 with the current server architectures.

Beginning with ProLiant Gen8 servers, we are introducing HP SmartMemory technology for DDR3 memory. HP SmartMemory DIMMs have passed our qualification and testing processes. Gen8 servers configured with HP SmartMemory DIMMs deliver extended performance and manageability that is not supported using third party DIMMs.

DDR3 architecture

DDR3 uses the same basic DRAM configuration and architecture as previous DDR implementations. Each DIMM consists of ranks of 9 or 18 DRAMs that deliver 72 bits—64 bits of data and 8 bits of ECC (error correction code)—in parallel to the memory bus to the CPU. DDR3 natively supports addressing up to eight banks, or ranks, of memory on a given memory channel. An individual DIMM module may support one, two, or four banks of these DRAMs, creating single-, dual-, or quad- ranked DIMM modules. The capacity of the DRAMs used and the number of ranks determine the overall capacity of a DIMM. DDR3 defines up to an 8 Gigabit DRAM, which will eventually lead to an individual DDR3 quad-rank DIMMs with a capacity as large as 64 GB.

Types of DDR3 DIMMs

DDR3 initially supported two types of DIMM memory—Unbuffered DIMMs (UDIMMs) and Registered DIMMs (RDIMMs). ProLiant Gen8 servers will support new third type memory we call Load Reduced DIMM, or LRDIMM.

Unbuffered and Registered DIMMs

With Unbuffered DIMMs (UDIMMs), all address and control signals, as well as the data lines, connect directly to the memory controller across the DIMM connector. Without buffering, each additional UDIMM that you install on a memory channel increases the electrical load. As a result, UDIMMs are limited to a maximum of two dual-rank UDIMMs per memory channel. In smaller memory configurations, UDIMMs offer the fastest memory speeds with the lowest latencies.

Registered DIMMs (RDIMM) lessen direct electrical loading by having a register on the DIMM to buffer the Address and Command signals between the DRAMs and the memory controller. The register on each DIMM bears the electrical load for the address bus to the DRAMs, reducing the overall load on the address portion of the memory channel. The data from an RDIMM still flows in parallel as 72 bits (64 data and 8 ECC) across the data portion of the memory bus. With RDIMMs, each memory channel can support up to three dual-rank DDR3 RDIMMs or two quad-rank RDIMMs. With RDIMMs, the partial buffering slightly increases power consumption and latencies.

Fully Buffered DIMMs (FBDIMMs), which buffer all memory signals (address, control, and data) through an Advanced Memory Buffer (AMB) chip on each DIMM, are not part of DDR3. FBDIMM architecture supported more DIMMs on each memory channel, but FBDIMMs were more costly, used more power, and had increased latency. The greater number of memory channels in server

architectures beginning with HP ProLiant G6 and G7 servers eliminated the need for an FBDIMM solution with DDR3.

Load Reduced DIMMs

Load Reduced DIMMs (LRDIMMs) are a new type of DDR3 memory to go alongside of the Unbuffered (UDIMM) and Registered (RDIMM) DDR3 memory. LRDIMMs buffer the memory address bus and the data bus to the memory controller by adding a full memory buffer chip to the DIMM module. Unlike the Fully Buffered DIMMs (FBDIMMs) from DDR2 days, Load Reduced DIMMs deliver data to the memory controller in parallel rather than using a high-speed serial connection.

Because LRDIMMs are completely buffered, they have both advantages and disadvantages when used in a server. They include:

- **Support for the maximum amount of memory per channel.** You can install as many as three LRDIMMs on a single memory channel. More important, with LRDIMMs you can now install three quad-ranked DIMMs on the same channel. You can only install two quad-rank RDIMMs.
- **Increased power consumption.** The addition of the buffering chip and the higher data rates increase DIMM power consumption.
- **Latency.** For equivalent memory speeds, buffering requires adding clock cycles to memory reads or writes, increasing the latency of LRDIMMs relative to single and dual-rank RDIMMs. This condition is offset by the fact that on fully populated memory channels LRDIMMs operate at a higher speed than quad-rank RDIMMs, resulting in a comparatively lower overall latency.

Table 1 compares the operating characteristics of a quad rank LRDIMM with those of an equivalent RDIMM.

Table 1: Operating characteristics of Quad-rank LRDIMM versus RDIMM

	32 GB Low Voltage Quad-rank RDIMM (ProLiant G7)	32 GB Low Voltage Quad-rank LRDIMM (ProLiant Gen8)	LRDIMM Advantage
Max. DIMMs per channel	2	3	50%
Max. Capacity per channel	64 GB	96 GB	50%
Max. Bandwidth per Channel	6.4 GB/s (@ 800 MT/s)	8.5 (@ 1066 MT/s)	33%
Idle Latency	80 ns	74 ns	7%
Max. Power	6 W	9W	-50%

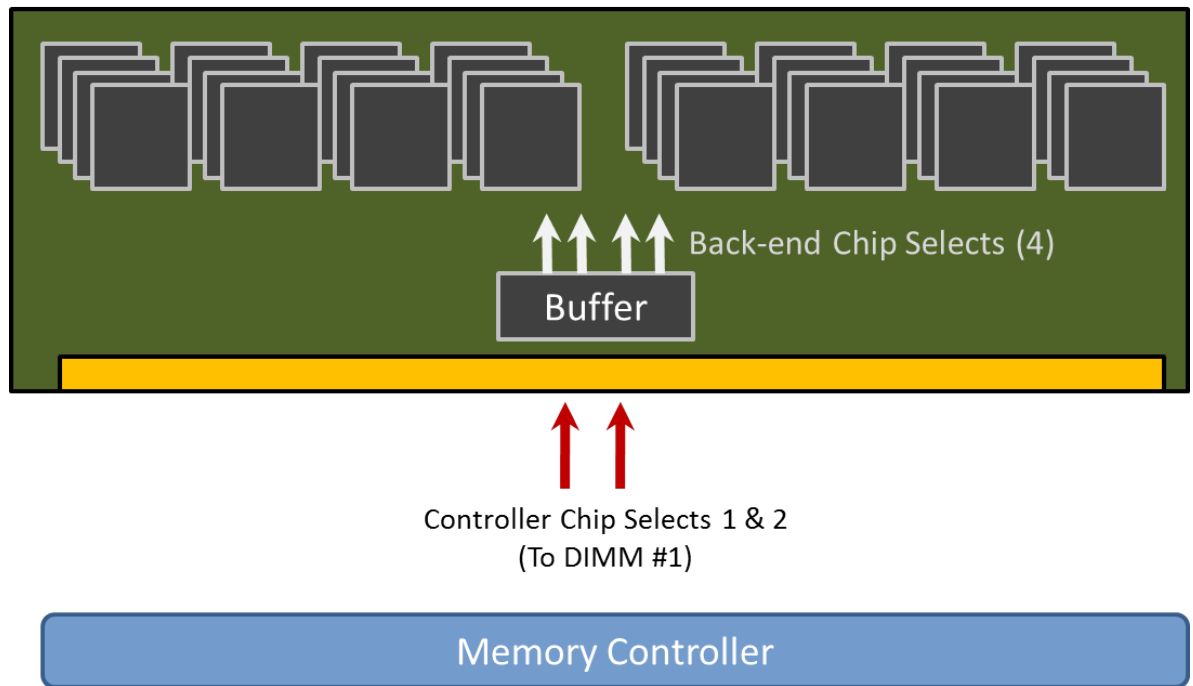
LRDIMMs allow you to configure large capacity systems with faster memory data rates than quad-rank RDIMMs. ProLiant Gen8 servers support LRDIMMs as well as UDIMMs and RDIMMs, although you cannot mix DIMM types within a single server. G6 and G7 servers do not support LRDIMMs.

LRDIMMs and rank multiplication

Load reduced DIMMs are the first DDR3 DIMMs to support operation of three quad-ranked DIMMs on a memory channel. DDR3 memory architecture can only natively support a maximum of eight ranks per channel, not the 12 ranks required for three quad-rank DIMMs. LRDIMM architecture solves this issue using by using a design known as rank multiplication.

Figure 1: The illustration shows LRDIMM rank multiplication using a quad-rank DIMM.

LRDIMM Rank Multiplication



With rank multiplication, the system memory controller sees each quad-rank DIMM as a dual-rank DIMM. The controller uses the normal external chip select signal to address one of the six ranks that it sees. The memory buffer chip on the selected Load Reduced DIMM then acts as an intermediary, intercepting the external chip select signal and mapping it to one of four back-end chip select lines located on the DIMM module. The buffer determines the correct back-end chip select line to activate based on both the external chip select and the row and column addresses that the memory controller asserts for the given memory operation.

DDR3 memory speeds

DDR3 memory uses faster clocking and data rates than DDR2 memory. The DDR3 specification originally defined data rates of up to 1600 mega transfers per second (MT/s), more than twice the rate of the fastest DDR2 memory speed (Table 2). G6 and G7 ProLiant servers support a maximum DDR3 DIMM speed of 1333 MT/s.

Table 2: Comparing DDR3 memory speeds

JEDEC Name	Common Name	Data Transfer Rate	Maximum DIMM Throughput
PC3 – 14900	DDR3-1866	1866 MT/s	14.9 GB/s
PC3 – 12800	DDR3-1600	1600 MT/s	12.8GB/s
PC3 – 10600	DDR3-1333	1333 MT/s	10.6 GB/s
PC3 – 8500	DDR3-1066	1066 MT/s	8.5 GB/s
PC3 – 6400	DDR3- 800	800 MT/s	6.4 GB/s

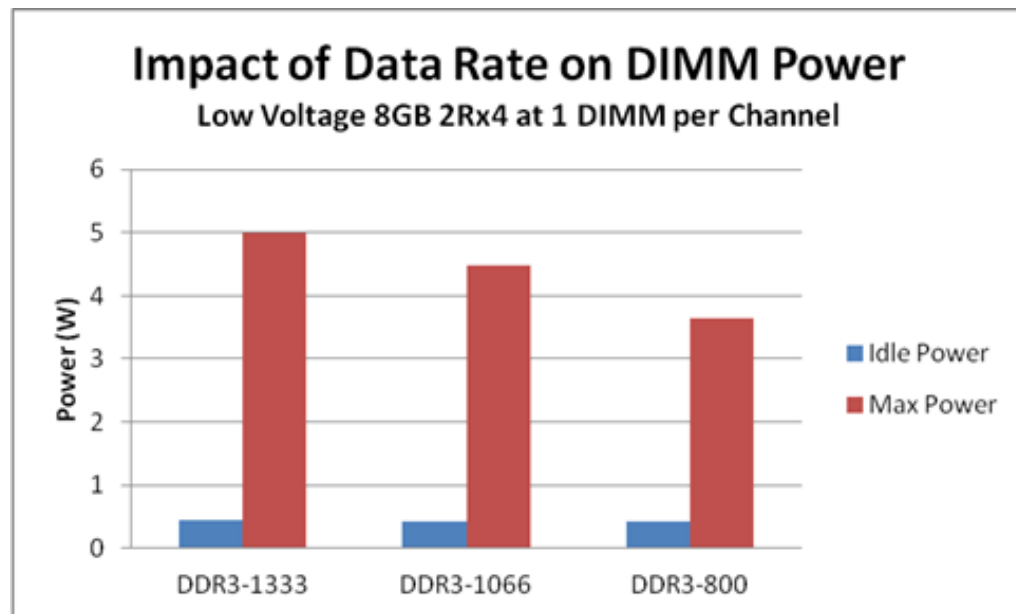
JDEC has extended the DDR3 specification to define additional memory speeds of 1866 MT/s and 2166 MT/s. ProLiant Gen8 servers will initially support memory speeds up to 1600 MT/s. We have designed the ProLiant Gen8 platform architecture to run at memory speeds up to 1866 MT/s once processor chipsets supporting it are available.

The faster data rates for DDR3 result in maximum possible throughput rates that are significantly greater than those of DDR2 are. The maximum bandwidth represents the amount of data that can move between a memory controller and the DIMMs if data transfers occur on every transfer cycle. With DDR3-1600 memory, the maximum bandwidth per memory channel is 12.8 GB/s.

Memory power consumption

DDR3 is more power efficient than DDR2. DDR3 DIMMs operate at 1.5V compared to 1.8V for DDR2 memory. DDR3 also has for low-voltage DIMMs that operate at 1.35V, further lowering power consumption. Figure 2 shows the idle and loaded power consumption for an 8 GB dual-rank DDR3 RDIMM in a ProLiant Gen8 server.

Figure 2: Power consumption by memory speed for a low-voltage 8 GB dual-rank RDIMM.



DDR3 memory running at 1333 MT/s uses about 25% more power than the same memory running at 800 MT/s. As a general guideline, DDR3 memory uses about 30% less power than DDR2 memory

when running at the same speed (800 MT/s). At 1333 MT/s, DDR3 uses about the same power as DDR2 memory while providing an improvement in maximum bandwidth.

The new 1600 MT/s DDR3 RDIMMs for the ProLiant Gen8 servers use 1.5V and consume the most power. Running at 1600 MT/s, DDR3-1600 memory uses about 15% more power than the equivalent DDR3-1333 memory in G6 and G7 servers.

DDR3 also has two power savings modes that the memory controller can employ to reduce memory power consumption further—CKE power down and self-refresh mode. With CKE power down, the memory controller performs a look-ahead at the queue of pending memory operations. If none are scheduled for a particular DIMM, then the controller places the DIMM in a lower power state and only re-awakens it for refreshes. For additional power savings, the memory controller will place the DRAMs in self-refresh mode whenever the system CPU enters a power down state (C6 for Intel processors). When this occurs, the DIMM goes into a still lower power state and performs its own refresh cycles.

Low-voltage DDR3 memory for ProLiant G6 and G7 servers

DDR3 memory is also available in low-voltage DIMMs, which operate at 1.35V instead of the 1.5V for standard DDR3 DIMMs. Low-voltage DDR3 DIMMs consume 10 to 15% less power than standard DDR3 DIMMs. A dual-rank 8 GB low-voltage DDR3 DIMM running at 1333 MT/s consumes about 8.5 watts compared to the 10 watts that a standard voltage DIMM consumes. At lower voltage, electrical loading and signal integrity are greater challenges. For these reasons, low-voltage DDR3 DIMMs are available only in single- and dual-rank configurations.

The low-voltage DDR3 DIMMs for G6 and G7 servers cannot operate at full speed when the number of DIMMs installed per channel increases. Tables 3 and 4 summarize low-voltage DDR3 DIMM operation in G6 and G7 ProLiant servers.

Table 3: Operation of low-voltage DDR3 DIMMs in AMD-based ProLiant G7 servers

	1 DIMM per channel	2 DIMMs per channel	3 DIMMs per channel
Single Rank	1333 MT/s @ 1.35V	1333 MT/s @ 1.35V	1066 MT/s @ 1.35V
Dual Rank	1333 MT/s @ 1.35V	1333 MT/s @ 1.35V	800 MT/s @ 1.35V

Table 4: Operation of low-voltage DDR3 DIMMs in Intel-based 2P ProLiant G6 servers

	1 DIMM per channel	2 DIMMs per channel	3 DIMMs per channel
Single Rank	1333 MT/s @ 1.35V	1333 MT/s @ 1.35V 1066 MT/s @ 1.35V (UDIMMs)	800 MT/s @ 1.35V
Dual Rank	1333 MT/s @ 1.35V	1066 MT/s @ 1.35V	800 MT/s @ 1.35V

As these tables illustrate, low-voltage DDR3 DIMMs for G6 and G7 servers reduce power consumption, but they operate at lower data rates in larger configurations. In some cases, these speeds are lower than the data rates at which standard voltage DIMMs operate.

HP SmartMemory for ProLiant Gen8 servers

HP SmartMemory DIMMs enable extended performance and manageability features on HP ProLiant Gen8 servers. For the 2P ProLiant Gen8 servers using the Intel E5-2600 series of processors, HP

SmartMemory DIMMs support extended performance when compared to third party memory. This is true for several DIMM types and configurations. Table 5 summarizes these performance extensions.

Table 5: Extended performance for HP SmartMemory DDR3 DIMMs in 2P ProLiant Gen8 servers

	1 or 2 DIMMs per channel	3 DIMMs per channel
1600 MT/s RDIMMs	1600 @ 1.5V	1333 @ 1.5V (HP SmartMemory and RBSU option) 1066 @ 1.5V (3 rd Party)
1333 MT/s RDIMMs	1333 @ 1.35V (HP SmartMemory) 1333 @ 1.5V (3 rd Party)	1066 @ 1.35V (HP SmartMemory) 1066 @ 1.5V (3 rd Party)
1333 MT/s LRDIMMS	1333 @ 1.35V (HP SmartMemory) 1333 @ 1.5V (3 rd Party)	1066 @ 1.35V (HP SmartMemory) 1066 @ 1.5V (3 rd Party)
1333 MT/s UDIMMs	1333 MT/s (HP SmartMemory) 1066 MT/s (3 rd Party)	Not Supported

The HP DDR3 Memory Configuration Tool provides additional assistance with configuring DDR3 memory (including HP SmartMemory) in HP ProLiant servers. It is available at http://h18004.www1.hp.com/products/servers/options/tool/hp_memtool.html

Core DDR3 Technologies

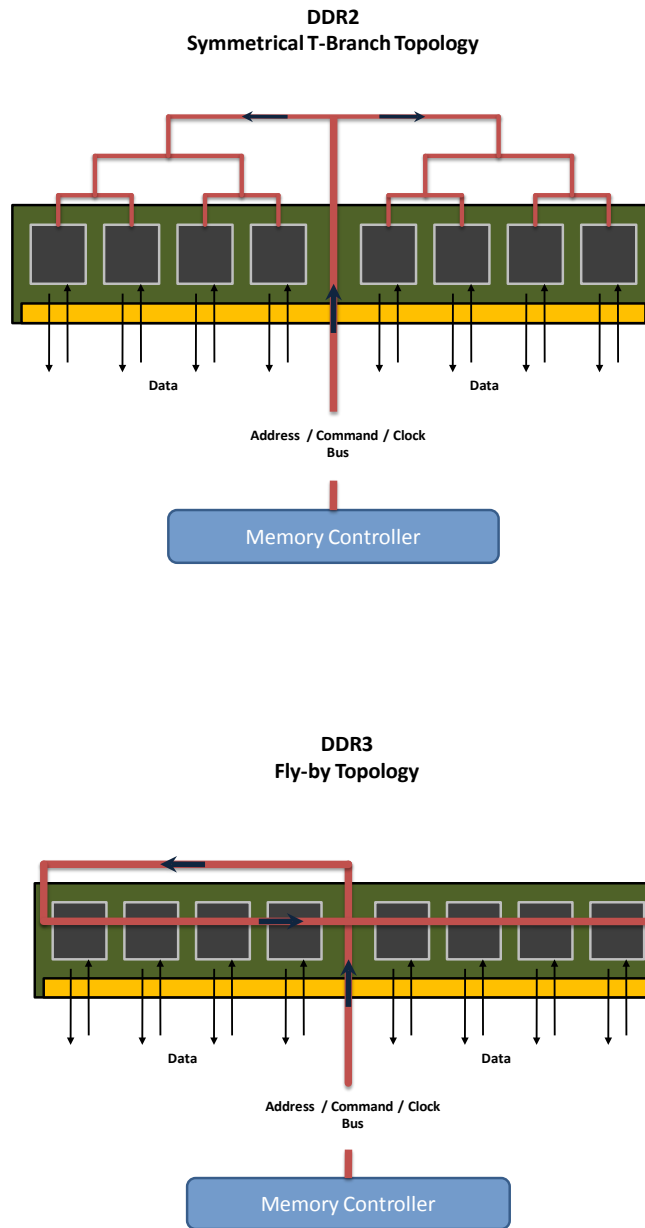
DDR3 memory specifies data transfer rates that are greater than twice that of DDR2 memory. Achieving these rates required significant engineering work to improve electrical signal integrity as well as new technologies to address the increasingly small timing tolerances. DDR3 memory has also added features to improve overall reliability and manageability.

Fly-by topology

Fly-by topology is one of the key technological innovations that allow DDR3 DIMMs to achieve twice the speed of DDR2 memory. In general, it refers to how address and command lines run to the DRAMs on the DIMM module and the timing adjustments required on the memory controller to compensate for it. To understand the workings of Fly-by topology, we will compare it to the T-Branch topology used in DDR2.

Figure 3 shows how one of the address and command signals run to the DRAMs in an Unbuffered DDR2 DIMM compared to a DDR3 DIMM. The symmetrical T-branch topology of DDR2 ensures that command and address signals arrive at all the DRAMs as close to simultaneously as possible. With this design, all of the DRAMs present their data (during a read) to the memory bus at the same time. The memory controller then reads the set of parallel bits. The window of time when the data bits can be read is known as the data eye, and the size of the eye shrinks as memory clock speeds increase.

Figure 3: The charts compare DDR2 and DDR3 address and command signal topologies.



Fly-by topology solves the problem of the shrinking data-eye by eliminating the need to deliver the data signals simultaneously to each DRAM. With Fly-by topology, each command and address signal flows along a single path that goes from DRAM 0 to DRAM 8. This simpler topology helps increase signal integrity. But it guarantees that the command and address signals won't arrive at each DRAM at the same time. If the signals arrive at DRAM 0 at time N , then they should arrive at DRAM 1 at time $N+1$, DRAM 2 at time $N+2$, and so forth. The result is that, on a read, each DRAM presents its data to the memory controller at a slightly different time.

To compensate for this skew, the memory controller must adjust its timing to lock in the bits from each DRAM at an appropriately delayed interval. This process is known as read leveling. These delays also vary slightly from one DRAM and DIMM to another. The memory controller must determine them empirically and then program them each time that the system reboots. This process is known as memory training. For memory writes, this scenario reverses, and the memory controller must delay the presentation of different sets of data bits to the bus to match the time when each DRAM is ready to receive them.

On-die Termination

Electrical circuits that carry signals need to be terminated with resistors to damp electrical reflections and to improve overall signal integrity. Earlier memory standards had memory termination on the system board. On-die termination puts the resistors on the DRAMs themselves, increasing their effectiveness by placing them at the end of the memory bus circuits. With DDR3 memory, the number of possible termination values is significantly greater than for DDR2. Just as important, the memory controller now empirically sets the termination values during POST based on the configuration of the DIMM module itself (number of ranks) and its position on the memory channel. Both of these refinements contribute to the signal integrity improvements necessary to support the faster DDR3 speeds.

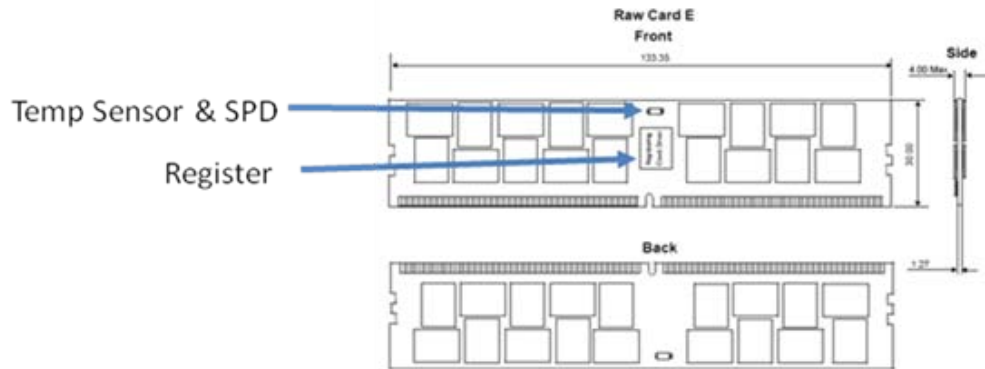
Address parity checking on RDIMMs

In DDR2, address parity detection was an optional feature. With DDR3, it's now standard. On DDR3 RDIMMs, the register chip performs a parity check on the DRAM address lines and compares it to the parity bit from the memory controller. This process detects potential addressing errors. Although address parity checking cannot correct addressing errors, it does stop the controller from writing data to an incorrect DRAM address, preventing silent data corruption. Unbuffered DIMMs do not support address parity checking because they do not have a register.

Integrated DIMM temperature sensor

DDR3 memory modules have built-in temperature sensors accurate to $\frac{1}{2}$ °C located in the center of each DIMM (Figure 4). HP engineers have taken advantage of the information these DIMM temperature sensors deliver by integrating it into Sea of Sensors fan control technology controlled by the iLO management processor found in all current ProLiant servers.

Figure 4: DDR3 DIMM with temperature sensor.



As a starting point, HP engineers have performed extensive modeling and testing to determine the operating temperature of each DRAM on a DIMM based on the readings from the DDR3 DIMM sensor. These values are determined by evaluating each of the following.

- The measured temperature from the DIMM sensor
- The relative location of each DRAM on the DIMM
- The direction of the airflow across the DIMM in a given server system

The iLO management processor in each ProLiant server collects this information from the DDR3 DIMMs and uses it, along with temperature data from other sensors in the server, to control fan cooling inside the server. This Sea of Sensors fan control technology ensures optimal cooling and helps prevent possible system failure while reducing power consumption by eliminating overcooling.


DDR3 memory and NUMA systems architectures

DDR3 is a stand-alone memory specification. But its use in servers goes hand-in-hand with the transition to new server architectures that use Non-Uniform Memory Access (NUMA). AMD Opteron™-based servers have used NUMA architecture since their inception, with DDR1 and later DDR2 memory. The AMD-based ProLiant G7 servers use an updated NUMA architecture that supports DDR3 memory. Starting in G6 and G7, Intel-based HP ProLiant servers began incorporating NUMA architecture along with other new features. The NUMA server architectures and DDR3 address memory throughput and latency issues that were limiting system performance under older architectures as system memory footprints continued to increase. All ProLiant Gen8 servers utilize NUMA architecture.

Older server architectures

Figure 5 shows the typical architecture for a two-processor (2P) server that used the traditional memory architecture. With this general design, known as uniform memory access, memory controllers and memory channels were located on a centralized system chipset. Each processor used the same pathway to access all of the system memory, communicating with the memory controllers across the front side bus. The controllers then accessed the DIMMs on the memory channels, returning the requested data to the processors. The architecture supported memory controller functions, each of which managed two memory channels for four memory channels per system. The system supported larger memory footprints by supporting up to four DDR2 FBDIMMs per channel.

Figure 5: Diagram of a traditional 2P uniform memory architecture



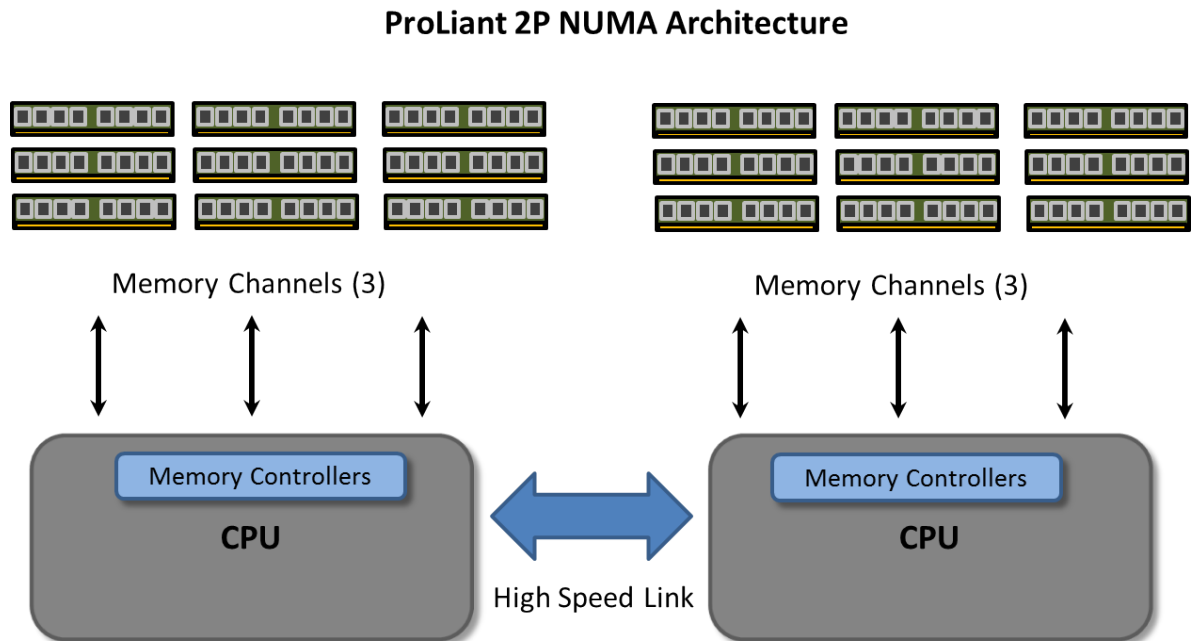
This architecture gave each memory channel a maximum raw bandwidth of 9.6 GB/s for systems supporting PC2-6400 fully buffered DIMMs. The memory channels of systems that use registered DIMMs can support a maximum bandwidth of 6.4 GB/s. With four memory channels per system, the theoretical maximum memory bandwidth for these systems is 38.4 GB/s and 25.6 GB/s, respectively. There are, however, factors that limit the achievable throughput:

- The maximum bandwidth of the front side bus became a performance choke point.
- Larger memory footprints required fully buffered DIMMs, which increased memory latency and decreased memory throughput and performance.

DDR3 and NUMA systems architecture

Although they vary slightly in their implementation details, NUMA architectures share a common design concept. With NUMA, each processor in the system has separate memory controllers and memory channels. In addition to increasing the total number of memory controllers and memory channels in the system, each processor accesses its attached memory directly. This eliminates the bottleneck of the front side bus and reduces latency. A processor accesses the system memory attached to a different processor through high-speed serial links that connect the primary system components. In Intel-based systems, it is the QuickPath Interconnect (QPI). For AMD-based systems, it is the HyperTransport technology. Beginning with the HP ProLiant G6, all HP ProLiant servers use DDR3 memory to help increase memory throughput. Figure 6 shows how the NUMA architecture looks for a typical ProLiant 2P server.

Figure 6: Typical G6, G7 ProLiant 2P NUMA Architecture



NUMA architecture solves the two related problems that emerged as system complexity grew:

- It eliminates bottlenecks in the memory subsystem that constrain system memory throughput
- It supports larger memory footprints without significantly lowering memory performance

DDR3 integrates with NUMA architecture to deliver significantly improved memory throughput. With a maximum transfer rate of 1600 MT/s, DDR3 memory can potentially deliver 12.8 GB/s of bandwidth per channel. ProLiant Gen8 NUMA architecture also supports more memory controllers and channels. For example, using DDR3 memory with six (eight for Gen 8) channels, the architecture for the Intel-based 2P ProLiant G6 servers has a maximum theoretical memory bandwidth of 64 GB/s, 65% greater than that of the older architecture using DDR2 memory.

DDR3 memory and 2P ProLiant servers

For 2P ProLiant servers, the memory architecture is relatively the same for both Intel-based and AMD-based designs (Figure 6). For G6 and G7 servers, each of the two processors has three memory channels supporting up to three DIMMs each.

For ProLiant Gen8 servers, the number of memory channels in 2P systems increases to four per processor. Memory speeds increase to support up to 1600 MT/s initially and 1866 MT/s later. Using quad-rank LRDIMM, the maximum memory capacity for the Gen8 2P systems will increase to 768 GB. This applies to both Intel-based and AMD-based server designs.

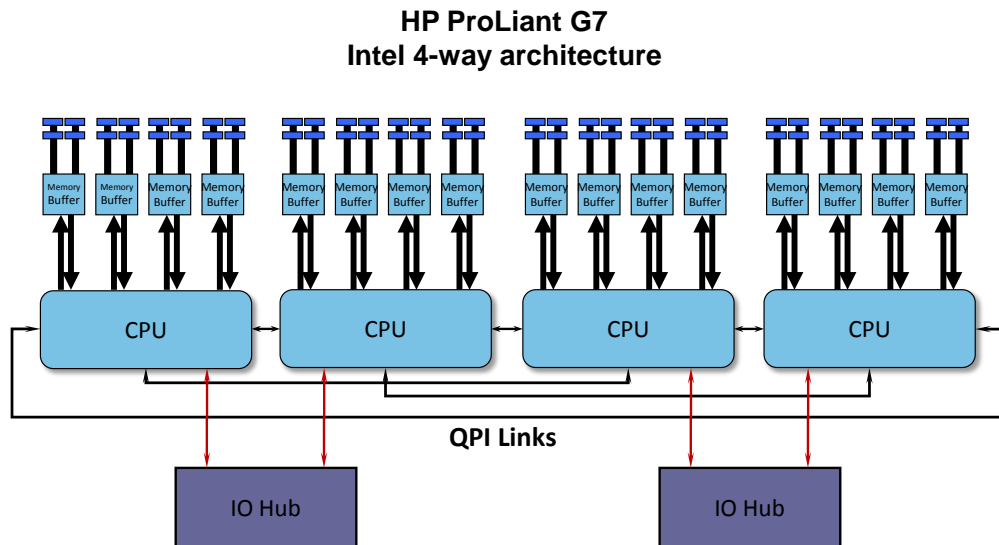
DDR3 and 4P ProLiant architecture

For ProLiant 4P server architectures, Intel-based designs take a slightly different approach to the memory subsystem than AMD-based systems. Each design approach has its own set of benefits, and both rely on DDR3 memory.

Intel 4P architecture and DDR3

Figure 7 shows the processor and memory architecture for the 4P HP ProLiant G7 servers that use Intel Xeon 5600 series processors. While the basic NUMA architecture is evident, there are distinct differences between this and the design of 2P systems. The most significant difference for the 4P server is the existence of separate memory buffers between the CPU and the memory channels. These buffers use a proprietary, high-speed serial link to transport memory data between themselves and the CPU while providing a standard memory bus interface to the DDR3 DIMMs. Using this approach, each memory controller supports two memory channels of two DIMMs each. In addition, the 4P architecture uses four memory controllers per CPU rather than three. Taken together, these design choices allow the Intel-based 4P systems to support up to 64 DIMMs, or 2 Terabytes of memory using 32 GB DIMMs.

Figure 7: 4P memory architecture for Intel-based HP ProLiant G7 servers



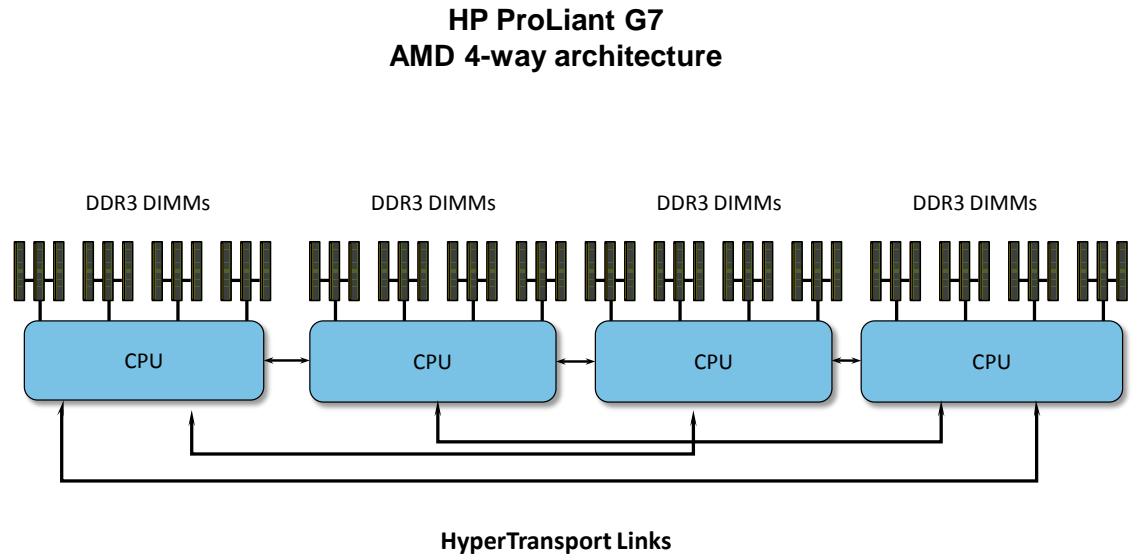
With the memory buffering used in this architecture, the system memory operates at 1066 MT/s for all memory configurations, including fully populated systems.

For ProLiant Gen8 servers, this architecture will remain relatively unchanged, although the memory speed will probably increase to 1333 MT/s.

AMD 4P architecture and DDR3

AMD-based HP ProLiant servers have used NUMA architecture since their inception. The ProLiant G7 servers are the first generation to use DDR3 memory. Figure 8 shows processor and memory architecture for an AMD-based 4P ProLiant G7 server with three DIMM sockets per memory channel.

Figure 8: 4P memory architecture for AMD-based HP ProLiant G7 servers



The 4P AMD-based ProLiant G7 servers use a slightly different memory architecture than the 4P Intel-based systems. Each processor has four memory controllers, each with a channel supporting either two or three DDR3 DIMMs, depending on the system. The three DIMM memory channels are a “T” electrically, with one DIMM installed at the center of the “T” and the other two DIMMs on the ends. This design provides improved signal integrity by keeping the lengths of the electrical paths to the DIMMs as close to the same length as reasonably possible. In order to help maintain this symmetry, you install DIMMs on both ends of the “T” before installing the third DIMM in the center.

This architecture allows the memory subsystem to support DDR3 operating at 1333 MT/s when memory channels are not fully populated. The absence of external memory buffering also results in lower overall memory latency. Without buffering, however, the architecture only supports a maximum of 48 DIMMs, or 1TB of system memory.

For ProLiant Gen8 servers, the AMD architecture does not change significantly. Without memory buffering, these systems will support DDR3 memory operating at 1600 MT/s in smaller memory configurations. With the availability of LRDIMMs capable of supporting three quad-rank DIMMs per channel, the maximum memory footprint will increase to 1.5 Terabytes @667 MT/s.

Memory throughput with DDR3

Next, we turn our attention to memory throughput for systems using DDR3 memory.

System memory bandwidth

By removing the front side bus and moving the memory controllers onto the processors, the newer system architectures eliminate some of the previous memory bottlenecks. The maximum theoretical memory bandwidth is unattainable in practice, because it represents an idealized scenario in which all memory channels operate at full throughput all the time. Using NUMA architectures, 2P ProLiant servers can achieve improved measured memory throughput relative to their theoretical maximums. See Table 6 for details.

Table 6: Theoretical maximum versus measured memory throughput for 2P ProLiant servers

	Theoretical maximum memory bandwidth	Measured maximum memory throughput
Intel-based 2P ProLiant G5	25.6 GB/s (RDIMMs) 38.4 GB/s (FBDIMMs)	12 GB/s
Intel-based 2P ProLiant G6/G7	64 GB/s	40 GB/s
Intel-based 2P ProLiant Gen8	102.4 GB/s	88.6 GB/s

NUMA architecture also allows the 4P ProLiant G7 servers to deliver significantly increased memory bandwidth (Table 7).

Table 7: Theoretical maximum memory throughput for 4P ProLiant servers

	Theoretical maximum memory bandwidth
Intel-based 4P ProLiant G5	38.4 GB/s (FBDIMMs)
Intel-based 4P ProLiant G7	136.4 GB/s
AMD-based 4P ProLiant G6	51.2 GB/s
AMD-based 4P ProLiant G7	169.6 GB/s

DDR3 latency

Memory latency is a measure of the time required for the CPU to receive data from the memory controller once it has requested it. It is an important measurement of memory subsystem responsiveness. Retrieving data from the memory subsystem consists of several steps, each of which consumes time. Taken together, these times comprise the overall latency:

- Time memory request spends in the processor I/O queue and being sent to the memory controller
- Time in the memory controller queue
- Issuing of the Row Address Select (RAS) and Column Address Select (CAS) commands on the memory address bus
- Retrieving data from the memory data bus
- Time through the memory controller and I/O bus back to the requesting processor Arithmetic Logic Unit (ALU)

The setting of RAS and CAS lines determine which memory address will be accessed. The electrical properties of DRAM are such that setting them requires about 13.5 nanoseconds each. This time is

roughly the same for both DDR2 and DDR3 memory. This means that 27 to 28 nanoseconds of memory latency are relatively fixed. Any improvements in memory latency must come elsewhere. DDR3 achieves improvements in latency through its faster data rate and the elimination of Fully Buffered DIMMs.

There are two different measurements of memory latency for a system— unloaded latency and loaded latency.

- Unloaded latency is the fastest possible time for the processor to retrieve data from the memory subsystem. It is measured when the system is idle. The timing and electrical properties of the memory subsystem determine unloaded latency.
- Loaded latency is a measurement of the latency when the memory subsystem is saturated with memory requests. With loaded latency, many additional factors come into play, including the number of memory controllers in the memory subsystem; controller efficiency in handling queued requests; and memory interleaving. Loaded latency gives a truer representation of the memory subsystem’s capabilities in real-world environments.

Table 8 compares the unloaded and loaded latencies of DDR2 and DDR3 memory in Intel-based 2P ProLiant servers.

Table 8: Memory latency of DDR2 and DDR3 in Intel-based ProLiant servers

	2P ProLiant G5 DDR2 at 667 MT/s	2P ProLiant G6 DDR3 at 800 MT/s	2P ProLiant G6 DDR3 at 1333 MT/s	2P ProLiant Gen8 DDR3 at 1600 MT/s
Unloaded latency	126 ns	80 ns	70 ns	65 ns
Loaded latency	147 ns	140 ns	100 ns	124 ns

Achieving highest performance with DDR3 memory

DDR3 memory delivers improved performance over DDR2 memory. With the NUMA architectures, the way you install DDR3 DIMMs in the system is important.

Maximizing system throughput

The key to maximizing system throughput is to populate as many of the system memory channels as possible. This helps to ensure that the memory bandwidth of all channels is available to the system. With 2P ProLiant G6 servers based on the Intel Xeon 5500 series processors, this means installing a minimum of six DIMM modules, one in each memory channel.

Minimizing memory latency

You can optimize memory latency, particularly loaded memory latency, by running at the highest data rate. For systems that are capable of supporting the higher data rates, achieving this memory speed depends on the number of and the rank of the DIMMs installed in each channel.

Using balanced memory configurations

For nearly all application environments, the optimal configuration for DDR3 memory is to balance installed memory across memory channels and across processors. Balancing installed memory across memory channels on a processor optimizes channel and rank interleaving, ensuring maximum memory throughput.

Balancing installed memory across processors ensures consistent performance for all threads running on the server. If you have installed more memory on one processor, threads running on that processor will achieve significantly higher performance than threads on the other processor. A performance imbalance can degrade overall system performance, particularly in virtualization environments.

The white paper *DDR3 Configuration Recommendations for HP ProLiant Gen8 servers*, downloadable from <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c03293145/c03293145.pdf>, provides a detailed discussion of DDR3 memory configuration considerations for HP ProLiant Gen8 servers.

Conclusion

DDR3 memory is the next step forward in memory technology for industry standard servers. It provides improved raw memory transfer speeds compared to DDR2. Combined with the NUMA architectures, DDR3 helps to deliver major improvements in server memory throughput and latency in HP ProLiant servers. With the introduction of ProLiant Gen8 servers, DDR3 continues to evolve—supporting faster memory speeds and bigger server memory footprints than ever before. HP SmartMemory for ProLiant Gen8 servers delivers additional improvement in memory performance that is not available with 3rd party DDR3 memory.

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
Configuring and using DDR3 memory with HP ProLiant Gen8 servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c03293145/c03293145.pdf
Memory technology evolution: an overview of system memory technologies	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf
HP ProLiant Server Memory web page	http://h18004.www1.hp.com/products/servers/options/memory-description.html
HP ProLiant DDR3 Memory Configuration Tool	http://h18004.www1.hp.com/products/servers/options/tool/hp_memtool.html
Innovative Technologies in HP ProLiant Gen8 servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c03227849/c03227849.pdf

Send comments about this paper to TechCom@HP.com



Follow us on Twitter: <http://twitter.com/ISSGeekatHP>

© Copyright 2012 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel® and Xeon® are registered trademarks of Intel Corporation. AMD is a trademark of Advanced Micro Devices, Inc.

TC1201833, April 2012

