

---

# Memory for Dell PowerEdge 12<sup>th</sup> Generation Servers

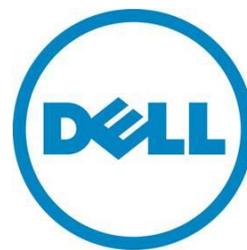
---

*This technical overview explains the memory options available for the Dell PowerEdge 12<sup>th</sup> generation servers and how these options will affect your environment*

ESG Memory Engineering

ESG Advanced Engineering

ESG Marketing



**This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.**

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

February 2012 | Rev 1.0

## Contents

Introduction .....	4
Memory subsystem: What does it do and how does it work?.....	4
Standardization .....	5
Reliability .....	5
Power .....	6
DIMM organization .....	8
Channels, slots, and ranks: How to build a memory subsystem .....	9
Channels, slots, ranks: How memory speed is affected .....	10
PowerEdge 12 <sup>th</sup> generation server memory features and benefits .....	11
What's new and improved?.....	11
Memory organization: Intel Xeon E5 processors .....	11
Memory channel usage and pairing .....	12
BIOS and system configuration .....	13
Reliability, availability, serviceability (RAS) features.....	13
DIMMs .....	14
Memory subsystem.....	14
New or advanced technologies .....	14
BIOS enabled.....	15
Memory roster .....	15
Summary .....	16

## Tables

Table 1. DIMM relative power consumption for various states .....	7
Table 2. Number of DIMMs per channel .....	9
Table 3. PowerEdge memory speeds by type and loading .....	10
Table 4. Feature comparison between the generations.....	11
Table 5. RAS features .....	13
Table 6. Sample memory configurations for R720 with 256GB.....	16
Table 7. Memory roster .....	16

## Figures

Figure 1. PowerEdge servers generational memory comparison.....	12
---	----

## Introduction

The Dell™ PowerEdge™ 12<sup>th</sup> generation servers significantly expand the memory subsystem capabilities of the previous generation servers. This technical overview discusses the memory subsystem capabilities of the newest generation of PowerEdge servers, highlighting the improvements from previous generations.

## Memory subsystem: What does it do and how does it work?

**Note:** This section is an introduction to DDR3 memory. If you already have a general understanding of DDR3, feel free to bypass this section.

In any computing environment, the processor executes a series of tasks operating on a collection of data, makes decisions or calculations, and produces results based on those decisions or calculations. While disk storage systems are used for long-term storage of these elements (task instructions and associated data), during processing, a high-speed, random-access memory, based on silicon technology, is used to provide fast and reliable short-term storage of these elements. The fundamental component used in this storage is the dynamic random-access memory (DRAM).

Each DRAM is capable of storing a finite number of data bits. In a PowerEdge server, a typical DRAM is capable of storing between 2 and 4 thousand million bits (Gigabits or Gb). While this may sound like a large number of bits, today's servers are capable of holding 512 Gigabytes (GB) of memory or more. At 8 bits per byte, 512GB of memory would require more than 2048 2-Gigabit DRAM devices.

In order to package large quantities of DRAMs into a single server, DRAM packages are mounted on a module capable of holding up to 72 DRAM devices. These modules are called dual in-line memory modules (DIMMs) due to their connector pin configuration—two rows of pins on a single-wide connector with DRAMs mounted on one or both sides of the module. With the ability to mount up to 72 DRAMs on a DIMM, we can see it would take only 32 DIMMs to provide 512GB of memory.

With so many DRAMs on a memory module, DIMMs have access to several DRAMs simultaneously, and in fact, will access 64 data bits within a single memory cycle. These 64 bits, or 8 bytes, may come from 8 or 16 DRAMs depending on how the DRAMs are organized.

As we mentioned earlier, today's server memory uses DRAMs that are either 2- or 4-Gigabit density. If we want to offer DIMMs that have capacities between 2 and 32 Gigabytes each, then we need to organize the DRAMs such that each DRAM of a given density has the proper width and depth to meet the desired DIMM capacity.

For example, suppose we want a 2GB DIMM. Using 2-Gigabit DRAMs, we would need no more than 8 DRAMs. Since the DIMM word width is 64 bits (8 bytes), we would use one DRAM per byte and the result would require a DRAM that is 8 bits wide and 256 million words deep ( $8 \times 256M = 2Gb$ ). Using 8 of these devices would provide our 2GB DIMM. What if we used 4-Gigabit DRAMs? In this case, we would need a DRAM that is 16 bits wide and 256 million words deep. Four of these would work if we could use a DRAM that is 16 bits wide. As we will explain later in this paper, 8 bits wide is the largest width a DRAM can be on PowerEdge server DIMMs.

In addition to memory size and capacity, another important factor in today's servers is how fast we can access this memory. With the next generation of servers, Dell offers extremely fast memory that operates at speeds of up to 1600 megatransfers per second (MT/s or Mbps for megabits per second).

## Standardization

In order to access memory this quickly, well-defined access methods are required. In addition, the sheer quantity of DIMMs used in these systems means that these memory modules (and the DRAMs mounted on them) absolutely require industry design standards, and the primary industry standard for this generation of memory is based on JEDEC<sup>1</sup> DDR3<sup>2</sup> set of specifications. These specifications define the electrical interface, packaging, and operating protocols for the memory used in Dell PowerEdge 12<sup>th</sup> generation servers.

Standardization ensures interoperability and continuity of supply of memory modules from multiple suppliers. Standardization also helps keep memory costs in check by keeping the memory as a commodity volume based product. The PowerEdge 12<sup>th</sup> generation servers rely on both DDR3 and DDR3L standards. We will explore the differences between these two standards later in this paper. For the moment, however, we will simply say that these two standards define a 240-pin DIMM that is the basis for the 12<sup>th</sup> generation memory subsystems. This DIMM has a data bus width of 64 data bits with an additional 8 bits that provides an error-correcting code (ECC).

## Reliability

With the amount of memory used in the Dell PowerEdge 12<sup>th</sup> generation servers and the business critical applications in which these servers are used, memory reliability is an utmost concern. With DRAM semiconductor processing technology ranging down to 3x nanometers, and moving well into even lower 2x nanometer technology during the lifespan of these servers, the susceptibility of randomly occurring bit errors is increasing. These bit errors are due to many naturally occurring factors such as alpha particles or spurious electrical events.

Given the number of memory data bits in a PowerEdge 12<sup>th</sup> generation server, the probability of an error occurring is measurably high and a means to detect and, if possible, correct any of these errors is necessary. This is where the extra 8 bits of ECC data comes in.

With 8 bits of ECC, stored along with the 64 data bits, any possible single-bit error cannot only be detected, but can be corrected as well. Memory reliability and availability is significantly enhanced due to these added bits and the fact that error correction occurs in real-time.

The additional bits also enable the servers to detect multi-bit errors, which at least prevent the server from using bad data. These 8 extra bits brings the DIMM word width to 72 bits or 9 bytes. Consequently, the DRAM organization used on the DIMMs must be an integral sub-multiple of 72. The 16-bit word width of a 2GB DIMM requires 5 DRAMs wide to provide 72 bits of data, which means the 4-Gigabit DRAM in the earlier example would not work.

With  $5 \times 16 = 80$  bits, the result would be 8 wasted bits. Consequently, 12<sup>th</sup> generation server memory is organized with either 8- or 4-bit word widths on the DRAMs, which is defined by the DDR3 standard.

---

<sup>1</sup> JEDEC: Joint Electron Devices Engineering Council

<sup>2</sup> DDR3: Double Data Rate Type 3; synchronous dynamic random access memory

## Power

As with practically any electronic device, memory modules consume power. The amount of power consumed is generally a result of the memory semiconductor technology, memory operating voltage, memory speed, memory utilization, and any power management practices that are applied to the modules.

Before going into detail on power consumed by a module, a basic understanding of the nature of a dynamic random-access memory cell is necessary. The significant advantage of a dynamic memory cell is that each bit of data can be stored with a single transistor coupled with a capacitor, allowing very high density storage. The capacitor is the actual storage element, selectively storing electrons to be represented as one or zero stored in the memory cell. The transistor provides the means to read or write to the capacitor. To write to a cell, the transistor is turned on and a charge of electrons is either pumped into the capacitor (if writing a one) or bled from the capacitor (if writing a zero).

This mechanism, while permitting high density, does impose one requirement on the memory design however. The capacitor, used to store the data, is not perfect and leaks electrons over time. Consequently, in order to maintain the data stored in each cell over a period longer than the time it takes to leak electrons, the memory design must perform refresh cycles that are used to restore charge into the capacitors. These refresh cycles take place over a relatively brief interval by human standards, typically 64 milliseconds, but a relatively long interval by memory access standards (at 1333 MT/s at 80% bandwidth utilization, approximately 550MB of data could be read within this interval).

Memory for the PowerEdge 12<sup>th</sup> generation servers has several levels of operating states, all of which consume some amount of power. Some power states, such as burst writes and burst reads consume the most power, while others, like memory idle, consume less. The total memory subsystem power is the summation of all the power consumed over time as the subsystem transitions between various states. As we shall see, the dynamic nature of the memory and how it is managed at the system level has a very strong influence on power saving.

As mentioned in the first paragraph of this section, the contributing factors influencing power consumption of memory include:

- Memory semiconductor technology, primarily the size of the transistor geometries
  - > In Dell PowerEdge 12<sup>th</sup> generation servers, the latest available memory technologies are used. As a rule, this means that Dell is offering the lowest power devices available from all of Dell's suppliers.
- Memory operating voltage
  - > The majority of Dell's memory offerings are based on DDR3L standard, which allows the memory to operate at either 1.35 volts or 1.5 volts, and Dell's PowerEdge 12<sup>th</sup> generation servers operate the memory at the lowest possible voltage to meet your performance needs.
- Memory operating speed
  - > The majority of memory offered on Dell's PowerEdge servers operates at 1333 MT/s, even when run at the lowest possible voltage. In cases where maximum performance is required, Dell offers a comprehensive set of 1600 speed DIMMs.

- Memory utilization or how much of the available channel bandwidth is used for transfers
  - > This is largely a function of applications running on the server; however, with the comprehensive range of DIMMs offered by Dell, good balance of power versus available capacity is achieved.
- Memory subsystem power management practices
  - > Dell PowerEdge 12<sup>th</sup> generation servers utilize some of the most aggressive power management algorithms available in the industry, which is a primary driver for overall power savings.

The first three factors contributing to system power consumption and related savings are largely controlled by diligence and design rigor—a well-designed memory subsystem to operate at maximum speed while remaining at the lowest voltage, and the best quality parts in your supply base to reliably deliver this capability. While memory utilization is a function of customer applications, the comprehensive range of devices offered by Dell provides the means by which you can optimize this factor.

Finally, addressing the power management practices, if we examine any memory subsystem, especially a large system where power management is most critical, we can easily see that memory utilization per DIMM is likely to be below 50%—this is simply because with two or more DIMMs per channel only one DIMM can be accessed at a time. In these cases, the remaining DIMMs are either idle or refreshed. Hence, the power savings can be significant if we are able to keep a DIMM in its lowest possible power state during this idle or refresh time.

Fortunately, the DDR3 standard defines operating states that, when properly implemented, can result in significant power reduction during this idle time. Without exhaustively analyzing each possible power state, from a high level, a typical DIMM (8GB, 1333MT/s) would have the relative power consumption values, (normalized to operating at full speed and 80% bandwidth utilization) shown in Table 1.

Table 1. DIMM relative power consumption for various states

Condition	Relative power consumption
Full access @ 80% utilization	100%
Active idle	63%
Self-Refresh/S3	7%

As shown in Table 1, power consumption when simply idle, while significantly lower, is almost nine times higher than when the memory is operating in a state called Self-Refresh/S3. What is this state?

The requirement to refresh a memory cell every 64 milliseconds establishes a need for something to perform this refresh cycle. In older generation servers, the system’s memory controller performed a refresh through an Active Idle state, a condition necessary for the system to send commands to the memory. With DDR3, memory placed in a Self-Refresh state performs its own internal refresh cycles whereby significant power savings can accrue.

Unfortunately, managing transitions in and out of Self-Refresh involve critical timing coordination and, in the recent past, the industry has struggled to provide a robust implementation. This is where the design rigor and diligence that Dell has invested in its PowerEdge 12<sup>th</sup> generation servers really benefits you—Dell has developed the most comprehensive test and verification methodologies in the industry to ensure safe and reliable memory power management.

## DIMM organization

Earlier we discussed the organization of DRAMs and how DIMMs can mount up to 72 DRAMs in a single module. We will explore this in more detail to fully understand how memory, at a system level, is organized and why some memory combinations are constrained.

Going back to our initial example of a 2GB DIMM, we recall that this could be implemented with eight 256M x 8 DRAMs plus one more for ECC data bits; nine devices in all. If we want to use the same DRAM but increase the DIMM capacity to 4GB, we can simply add another row of DRAMs to the DIMM, yielding 18 DRAMs on the DIMM. If we have two rows, however, how do we access one row versus the other?

The JEDEC DDR3 DIMM specification provides a set of signals that determines the row of a DIMM selected for access. These signals are called  $S_n[0:3]$  and, as you might infer, there are four such signals. Does this mean we can have up to four rows of DRAMs on a DIMM? Yes, however, there are limitations to the type of DIMM that can have more than two rows.

In JEDEC terminology, ranks are each row of DRAM. In our first example, the 2GB DIMM 1Rx8 describes one rank by an 8-bit wide organization. A 4GB module would be 2Rx8 organization. Carrying this one step further, it is possible to have a 4Rx8 8GB module using the same 2GB DRAM. As stated earlier, there are some limitations to this implementation and other factors that would make other possible implementations preferred.

When the number of DRAMs increases on a module, the number of signal loads also increases proportionally. With a single row of eight DRAMs, the loading is tolerable. Each DRAM ties to their own data I/O lines and all share common clock, address, and command lines. With two ranks, the clock, address, and command line loading would be sixteen loads and the data I/O lines between the two DRAMs is shared, one per row.

Carrying this further would again double the loads on the individual signals reaching an unwieldy electrical configuration with very poor signal integrity. The DDR3 standard addresses these loading conditions by defining three different types of DIMMs:

- **Unbuffered DIMM:** UDIMM, does not buffer either add, control, clock or data I/O lines
- **Registered DIMM:** RDIMM, buffers add, control, clock lines but does not buffer data I/O lines
- **Load reduced DIMMs:** LRDIMMs, buffers both add, control, clock and data I/O lines

UDIMMs are like the simple examples we used earlier, where we have eight or sixteen DRAMS on a module, and the DRAM signal pins are connected directly to the module connectors. This is a simple implementation but introduces the heaviest signal loading on the memory bus.

RDIMMs, by buffering and registering the address, control, and clock lines, removes these loads from the system memory bus and isolates these to the well-defined signal paths on the DIMM itself. While this register introduces some latency to buffer these signals, the pipelined nature of the memory

access allows full speed access once the transfer begins. Rather than having eight or sixteen loads on these signals, the system memory bus only sees one load, that of the register. However, the data I/O signals on an RDIMM remain connected directly between the connector and the individual DRAMs (up to four loads on a 4-rank DIMM).

LRDIMMs are a new technology, introduced on the PowerEdge 12<sup>th</sup> generation servers. LRDIMMs buffer the data I/O signals between the DRAMs and the system memory bus by using a buffer device.

Consequently, even on a 4-rank DIMM, the system memory bus only sees one load on all of the DIMM connector pins. The expected numbers increase as memory scales upward in the future.

## Channels, slots, and ranks: How to build a memory subsystem

Now that we know something about DIMMs and DRAMS, how does a PowerEdge server put these together to build a memory subsystem?

Looking at the processor, in order to access memory, there must be at least one memory channel provided to connect between the processor’s memory controller and the memory module(s). A memory channel has one or more slots containing DDR3 DIMM connectors. The more slots we add to a channel, the more memory we can add to the system.

There is a limit to how many slots we can add to a channel. In PowerEdge servers, the maximum number of slots we can add to a channel is 3. Fortunately, in Dell’s high-end PowerEdge servers, the processors provide 4 memory channels, permitting up to 12 slots per processor. With Dell’s four-processor PowerEdge servers, there are 48 memory slots.

Going back to our different DIMM types with three slots, we can now think about how much memory we can add to a channel. Since each DIMM type represents different loading and selection characteristics, every channel has a corresponding limit on the number of DIMMs it can support, given the respective DIMM type. Table 2 shows the maximum number of DIMMs of each type supported by PowerEdge servers.

Table 2. Number of DIMMs per channel

DIMM type	Ranks (per DIMM)	Maximum number of DIMMs in channel	Comments
UDIMM	1 or 2	2	
RDIMM (1 or 2R)	1 or 2	3	
RDIMM (4R)	4	2	Limited to 8 ranks in a channel
LRDIMM	up to 4	3	

Before we can determine the maximum capacity we can add to a system, we need to dig a bit more into ranks, DRAM organization, and DIMM design.

We mentioned at the beginning of this paper that DIMMs can hold up to 72 DRAM devices and established that DRAMs are organized either 4- or 8-bits wide. If we use 4-bit wide devices, we have 18 devices per rank with a 4R DIMM using the maximum of 72 DRAMs. At 4 Gigabit per DRAM, this

establishes the maximum amount of memory per DIMM as 32GB. The extra ECC bits are not counted when considering capacity. Note that 72 8-bit wide DRAMs would require 8 ranks, exceeding the maximum amount of ranks allowed in our table above.

Since 32GB of memory would have 4 ranks, from the table above, we see that if we use 32GB RDIMMs, we can only put 2 DIMMs per channel, whereas, if we use 32GB LRDIMMs, we can put 3 DIMMs per channel. With 4 channels per processor, this would yield 384GB per processor, which is 1.5TB in select Dell PowerEdge four-socket servers.

Note that in order to mount 72 DRAMs on the small DIMM circuit board (30 x 133.35 mm = 1.18 x 5.25 inches) special packaging is required. When more than 36 DRAMs are on a single board, a dual die package (DDP) made of two DRAM die are placed in a single surface mount device and two chip selections on the device pins select one of the two die.

Note also that mixed DIMM types within the same system is not supported. The entire system must be populated with the exact same type of DIMM.

### Channels, slots, ranks: How memory speed is affected

The Dell PowerEdge 12<sup>th</sup> generation servers offer completely new performance in their memory subsystems with offerings that include 1600 MT/s speed. Critical design and testing elements ensure that Dell server memory achieve the desired speeds.

As you might expect, the loading within a channel significantly affects memory speed and, looking at Table 3, you can see the maximum achievable speed based on memory loads.

Table 3. PowerEdge memory speeds by type and loading

DIMM type	DIMM ranking	DIMM rated voltage, speed	Generation memory POR (socket R, 3DPC platforms)						
			1 DPC		2 DPC		3 DPC		
			1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	
RDIMM	SR/DR	DDR3 (1.5V), 1600MT/s		1600		1600			1066
RDIMM	SR/DR	DDR3L (1.35V/1.5V), 1333MT/s	1333	1333	1333	1333			1066
UDIMM	SR/DR	DDR3L (1.35V/1.5V), 1333MT/s	1066	1333	1066	1333			
RDIMM	QR	DDR3L (1.35V/1.5V), 1066/1333MT/s	800	1066	800	800			
LRDIMM	QR	DDR3L (1.35V/1.5V), 1333MT/s	1333	1333	1333	1333	1066		1066

- Green boxes are defaults for performance/performance per watt
- White boxes (with lower voltage) can be chosen to save power but run slower
- White boxes (with higher voltage) can be chosen through custom configuration in the bios

Table 3 confirms that installation of a third DIMM in the channel results in a dramatic drop in speed. This is unavoidable due to loading, but for maximum performance, the Dell PowerEdge 12<sup>th</sup> generation servers offer great performance per watt at 2 DIMMs per channel.

## PowerEdge 12<sup>th</sup> generation server memory features and benefits

We will now examine the features and benefits of the Dell PowerEdge 12<sup>th</sup> generation server memory.

### What’s new and improved?

There are currently more options than in previous generations—more capacity choices and more frequency choices. Also available are increasingly intricate tradeoffs based on the DIMM choices and how the DIMMs are populated including reliability, availability, and serviceability (RAS), performance, cost, and power. There is interoperability between some of the memory used in previous generation PowerEdge servers, but this is limited to UDIMMs and RDIMMs. Table 4 shows a comparison of the PowerEdge 11<sup>th</sup> and 12<sup>th</sup> generation servers.

Table 4. Feature comparison between the generations

Feature	11 <sup>th</sup> generation	12 <sup>th</sup> generation
<b>DIMM type</b>	U DIMM ECC R DIMM	U DIMM ECC R DIMM LR DIMM
<b>Transfer speed</b>	800 MT/s 1066 MT/s 1333 MT/s	800 MT/s 1066 MT/s 1333 MT/s 1600 MT/s
<b>Voltage</b>	1.5 V (DDR3) 1.35V (DDR3L)	1.5 V (DDR3) 1.35V (DDR3L)

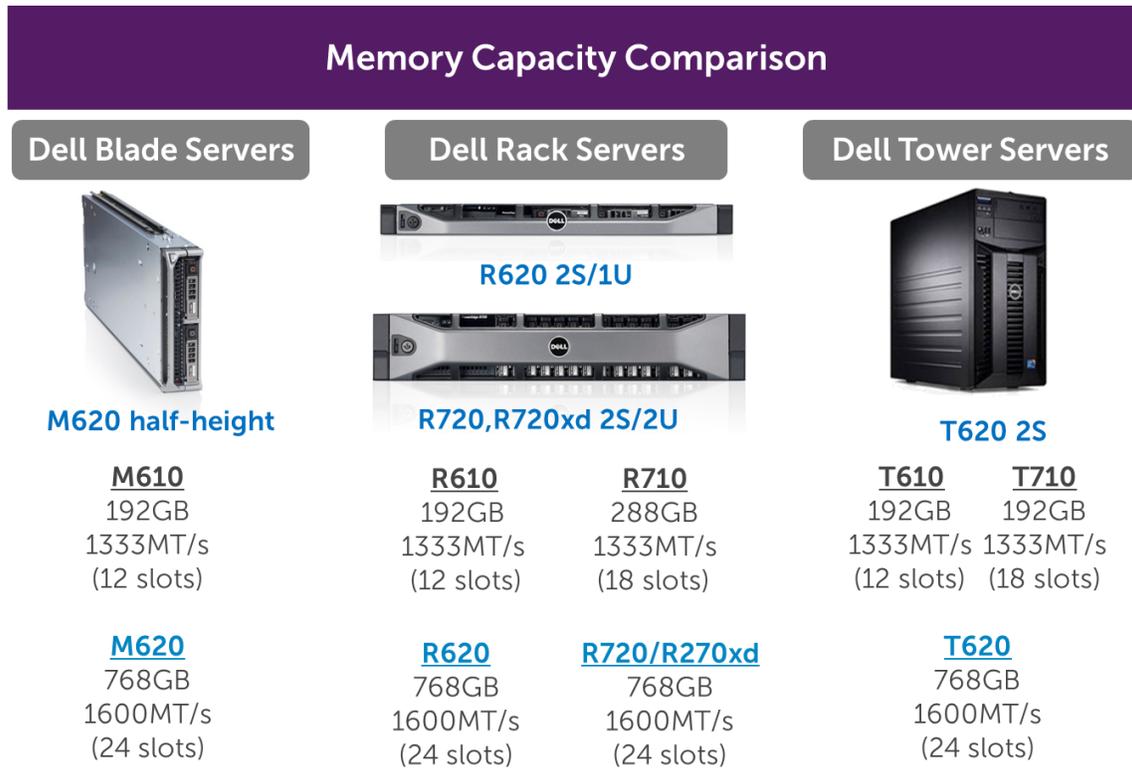
### Memory organization: Intel Xeon E5 processors

Dell PowerEdge 12<sup>th</sup> generation servers, based on the Intel<sup>®</sup> Xeon<sup>®</sup> processors E5-2600 and E5-4600 product families are available in a wide range of rack, tower, and blade 2- and 4-processor configurations. As such, a wide range of memory capacities is available for these systems.

The fundamental building block for any of these systems is the Intel Xeon E5 processor with each processor offering 4 memory channels and supporting up to 3 DIMM slots per channel.

Figure 1 shows a comparison between PowerEdge 12<sup>th</sup> generation servers and 11<sup>th</sup> generation servers in terms of memory.

Figure 1. PowerEdge servers generational memory comparison



With one exception, all of these platforms offer the same set of memory capabilities. The one exception, the M620, due to its extreme compact size, has some memory population constraints when the highest processor speed bins are used. It is important to note that on the M620 blade server, the heatsink for 115W or 135W thermal design point processors obstruct two of the DIMM slots for each processor and the number of usable DIMM slots drops below the 24 slot maximum. See Table 6 for exact slot count based on configuration.

### Memory channel usage and pairing

The Intel Xeon E5 processor, with its four memory channels, has four memory controllers that can run independently of each other or, in certain modes, run as two pairs of memory controllers. The channels are numbered 0-3 with channels 0 & 1 and 2 & 3 paired when run in paired modes. The physical location of the channels correspond to this pairing, where channels 0 & 1 are arranged on one side of the processor package and channels 2 & 3 are arranged on the opposite side of the processor.

As with some things in life, optimizing for one set of conditions creates trade-offs. For example, optimizing for performance could theoretically affect memory channel settings for more conservative reliability modes. There are more details on this in the RAS\_REF section, but specifically, if the channels are running independently and all channels populated in an identical fashion, then the system will achieve the highest level of performance. However, if error checking and recovery capability are of significant importance, choose the appropriate RAS modes and pair the memory channels to achieve the desired results.

The Dell PowerEdge 12<sup>th</sup> generation servers provide full flexibility in these choices and the 12<sup>th</sup> generation memory offerings nicely complement these options.

## BIOS and system configuration

With the Dell PowerEdge 12<sup>th</sup> generation servers, you can easily configure the server operating conditions through the system profiles. By selecting a particular profile in the BIOS setup, you can configure a system to operate at highest performance, highest performance per watt, high-density configuration, or custom configuration.

While the first two profiles are self-evident, the high-density configuration runs the system in an intentionally designed conservative mode when you are more interested in maximizing capability and reliability over performance. This mode has the most default influence on the memory subsystem by selecting more conservative clocking, operating at a higher voltage with more margins, and performing safety type actions such as running refresh cycles at twice the normal rate. All of these result in a platform that runs as reliably as possible even when populated with the maximum potential memory.

Custom configuration, of course, gives you full control of all of the system settings and can be oriented toward any behavior you prefer.

## Reliability, availability, serviceability (RAS) features

Of all the improvements made since the previous generation, the addition of several new RAS features make the new platforms the most reliable and available ever offered. The new RAS features could be an entire white paper onto itself, but to give you a taste of what is available, Table 5 summarizes memory-specific RAS features.

Table 5. RAS features

Category	RAS feature
DIMMs	ECC
	Register/PLL
Memory subsystems	SDDC
	Mirroring
	Rank sparing
	Demand and patrol scrubbing
New or advanced with 12 <sup>th</sup> generation	Memory buffer (LRDIMM)
	DIMM SPD error logging
	Correctable error threshold
	Memory page retire
	High-density profile
BIOS enabled	Corrupt data containment
	MCA recovery
	Device tagging

## DIMMs

While ECC memory advantages are numerous, one additional advantage to the register/buffer in both RDIMMs and LRDIMMs is the ability to perform parity checking on both addresses and commands, which is a feature lacking in UDIMMs. ECC detects improper addresses or commands, notifying the system that memory data corruption has occurred. Depending on other RAS operating modes, trapping this condition may permit error recovery in the form of retries or alternate DIMM/rank operations.

## Memory subsystem

### Single device data correction

Single device data correction (SDDC) organizes the data in the channel such that a single bundle of data, accessed in burst mode, is sourced from a series of single x4 or x8 DRAM words. Should corruption happen to any or all of the bits in a bundle, then corrections occur. The x4 condition is corrected on all cases. The x8 condition is corrected only if the system is in advanced ECC (lockstep) mode.

Dell PowerEdge 12<sup>th</sup> generation servers go beyond the memory enhancements of the previous generation of servers through more comprehensive memory offerings and an architecture preventing memory channel loss when advanced ECC mode is active.

### Mirroring

Mirroring memory from one channel into a paired channel provides 100% redundancy. 12<sup>th</sup> generation memory mirroring does not result in the loss of a complete memory channel as it did with the previous of generation servers.

### Rank sparing

Rank sparing reserves a rank of memory as a spare as long as there are at least 2 DIMMs or 4 ranks in a channel. The memory controller moves a rank exhibiting a high number of correctable errors to the spare rank; the spare rank is invisible to the operating system and applications. This feature was also available on the previous generation of servers.

### Demand or patrol scrubbing

The memory subsystem accepts automatically written corrections to errors once discovered during a read.

## New or advanced technologies

### Memory buffer (LRDIMM)

As with the parity checking provided by the register in RDIMMs, the LRDIMM provides parity checking on the address and command lines. The LRDIMM increases robustness of the system by buffering the data I/O lines and isolating the loaded data path to the DIMM module.

### DIMM SPD error logging

Each DIMM has a serial presence detect memory on it to tell the system about the DIMMs characteristics. This non-volatile flash memory saves a small OEM data in a small designated area. PowerEdge 12<sup>th</sup> generation servers use this space to record critical information gathered should a memory error occur, which will significantly reduce the time for failure analysis in the event of a failure. PowerEdge 11<sup>th</sup> generation servers have a limited form of this capability.

### Correctable error threshold

The BIOS in the 12<sup>th</sup> generation servers provides a more robust correctable error tracking capability. Each DIMM has a record in memory to keep track of frequency and location of correctable errors. A dramatically increased rate of correctable errors may indicate a failure is imminent. The memory subsystem accepts and escalates these error logs to the system management subsystem.

### Memory page retire

This is a completely new and Dell exclusive feature on the PowerEdge 12<sup>th</sup> generation servers. With the coordination of a running hypervisor, memory faults are monitored, and should certain regions, such as memory page, produce recoverable errors beyond a certain threshold, then the faulty page will be retired and effectively removed from the system so that it is no longer accessed. The event is logged so that memory can be replaced during scheduled service time.

### High-density profile

As discussed earlier, this feature allows you to configure your system to operate with more conservative operating parameters, thereby improving the overall robustness and availability of the system.

### Corrupt data containment

This feature is contained within the Intel chipset but coupled to the Dell BIOS. An uncorrectable error detected during a memory read is deferred and reported only when the processor actually consumes the corrupt data value. Cached faulty data causes no system impact, thereby improving the overall system performance.

Note that the mechanism to defer the error reporting uses the “Error Containment Bit” which is also known colloquially as “Poisoned” data.

## BIOS enabled

### Machine check architecture recovery

Machine check architecture (MCA) recovery is a general class of error recovery mechanisms in which BIOS intelligently processes various machine check conditions and attempts to recover, or, at a minimum, log the faulty events. The system automatically shuts down and restarts in a worst-case scenario.

### Device tagging

If one of the DRAMs starts having errors, the system can move the contents of the failing data DRAM to a DRAM normally used to maintain the ECC bits, thus degrading the error correction and checking protection but eliminating the failures. It is most useful for the elimination of hard errors and stuck bits until a replacement DIMM is available.

## Memory roster

One of the most significant changes in memory has little to do with technology. The changes are largely operational and change the way Dell offers memory to its customers. In previous generations, Dell established fixed configurations of memory of a specific type and capacity at the system level. With the 12<sup>th</sup> generation servers, Dell offers a comprehensive selection of DIMM choices and allows you to configure your system according to your specific needs. Table 6 lists a sampling of configurations for 256GB of memory in a PowerEdge R720 server.

Table 6. Sample memory configurations for R720 with 256GB

Qty	Size	Type	Feature
16	16GB	4Rx4 1066 RDIMM	Lowest cost point, but slowest speed (800)
16	16GB	2Rx4 1333 RDIMM	Medium cost point, good speed (1333)
16	16GB	2Rx4 1600 RDIMM	Medium+ cost point, great speed (1600)
8	32GB	4Rx4 1333 RDIMM	High cost point, decent speed (1066), good scalability
8	32GB	4Rx4 1333 LRDIMM	Highest cost point, good speed (1333), great scalability

Table 7 lists the memory supported by for Dell PowerEdge 12<sup>th</sup> generation servers.

Table 7. Memory roster

DIMM speed	DIMM type	DIMM capacity (GB)	Ranks per channel	Data width	SDDC support	DIMM voltage
1600	RDIMM	2	1	x8	Advanced ECC	1.5
1333	RDIMM	2	1	x8	Advanced ECC	1.35
1333	UDIMM	2	1	x8	Advanced ECC	1.35
1600	RDIMM	4	2	x8	Advanced ECC	1.5
1333	RDIMM	4	2	x8	Advanced ECC	1.35
1333	RDIMM	4	1	x4	All Modes	1.35
1333	UDIMM	4	2	x8	Advanced ECC	1.35
1600	RDIMM	8	2	x4	All Modes	1.5
1333	RDIMM	8	2	x4	All Modes	1.35
1600	RDIMM	16	2	x4	All Modes	1.5
1333	RDIMM	16	2	x4	All Modes	1.35
1333	LRDIMM	32	4	x4	All Modes	1.35
1333	RDIMM	32	4	x4	All Modes	1.35

Not all servers offer every option. Visit [Dell.com/PowerEdge](http://Dell.com/PowerEdge) for details on memory options.

## Summary

Dell is committed to bringing the latest technologies in memory to our server customers. The addition of 32GB RDIMMs to the portfolio provides greater memory capacity on our Dell PowerEdge 12<sup>th</sup> generation servers. With the addition of 1600MT/s DIMMs to the portfolio, you will see higher memory bandwidth and lower memory latency, resulting in higher application performance. Dell's memory strategy for the PowerEdge 12<sup>th</sup> generation servers provides the latest technologies and advanced features for you to meet your workload needs and solve real-world applications.