

# Memory Population Rules for 3rd Generation Intel® Xeon® Scalable Processors on PowerEdge Servers

Optimizing Memory Performance

Revision: **1.2**  
Issue Date: **8/16/2021**

## Abstract

Properly configuring a server with balanced memory is critical to ensure memory bandwidth is maximized and latency is minimized. When server memory is configured incorrectly, unwanted variables are introduced into the memory controllers' algorithm, which inadvertently slows down overall system performance. To mitigate this risk of reducing or even bottlenecking system performance, it is important to understand what constitutes balanced, near balanced and unbalanced memory configurations.

Dell Technologies has published this brief to educate PowerEdge customers on what balanced memory means, why it is important and how to properly populate memory to 3rd Generation Intel® Xeon® Scalable Processors for a balanced configuration.

## Revisions

Date	Description
16 August 2021	Updated for 3 <sup>rd</sup> Generation Intel® Xeon® Scalable Processors

## Acknowledgements

This paper was produced by the following people:

Name	Role
Matt Ogle	Technical Product Marketing, Dell EMC
Bruce Wagner	Senior Principal Systems Development Engineer, Dell EMC
Chris Cote	Senior Principal Systems Development Engineer, Dell EMC
Swati Mahajan	Senior Advisor, Product Management, Dell EMC

## Table of Contents

1. Introduction .....	4
2. Memory Topology and Terminology .....	5
3. Memory Interleaving .....	6
4. Guidelines for Balancing Memory.....	7
4.1 Overview .....	7
4.2 Memory Channel Population.....	7
4.3 Identical CPU and DIMM Parts .....	8
4.4 Identical Memory Configurations for Each CPU .....	8
4.5 Summary .....	8
5. Balanced Configurations (Recommended) .....	9
6. Near Balanced Configurations .....	10
7. Unbalanced Configurations .....	12
8. Traditional DIMMs with DCPMMs.....	16
7. Conclusion.....	17

# 1. Introduction

Understanding the relationship between a server processor (CPU) and its memory subsystem is critical when optimizing overall server performance. Every processor generation has a unique architecture, with volatile controllers, lanes and slot population guidelines, that must be satisfied to attain high memory bandwidth and low memory access latency.

Dell EMC PowerEdge products with 3rd Generation Intel® Xeon® Scalable Processors now offer a total of eight memory channels with up to two memory slots per channel; a total of sixteen memory modules per processor. This presents numerous possible permutations for configuring the memory subsystem with traditional Dual In-Line Memory Modules (DIMMs) and Optane™ DC Persistent Memory Modules (DCPMMs), yet there are only a couple of balanced configurations that will achieve the peak memory performance that Dell EMC PowerEdge servers can deliver.

Memory that has been incorrectly populated is referred to as an unbalanced configuration. From a functionality standpoint, an unbalanced configuration will operate adequately, but introduces significant additional overhead that will slow down data transfer speeds. Similarly, a near balanced configuration does not yield fully optimized data transfer speeds but it is only suboptimal to that of a balanced configuration. Conversely, memory that has been correctly populated is referred to as a balanced configuration and will secure optimal functionality and data transfer speeds.

This white paper explains how to balance both traditional DIMMs and DCPMMs configured for 3rd Generation Intel® Xeon® Scalable Processors with Dell EMC PowerEdge servers.

## 2. Memory Topography and Terminology

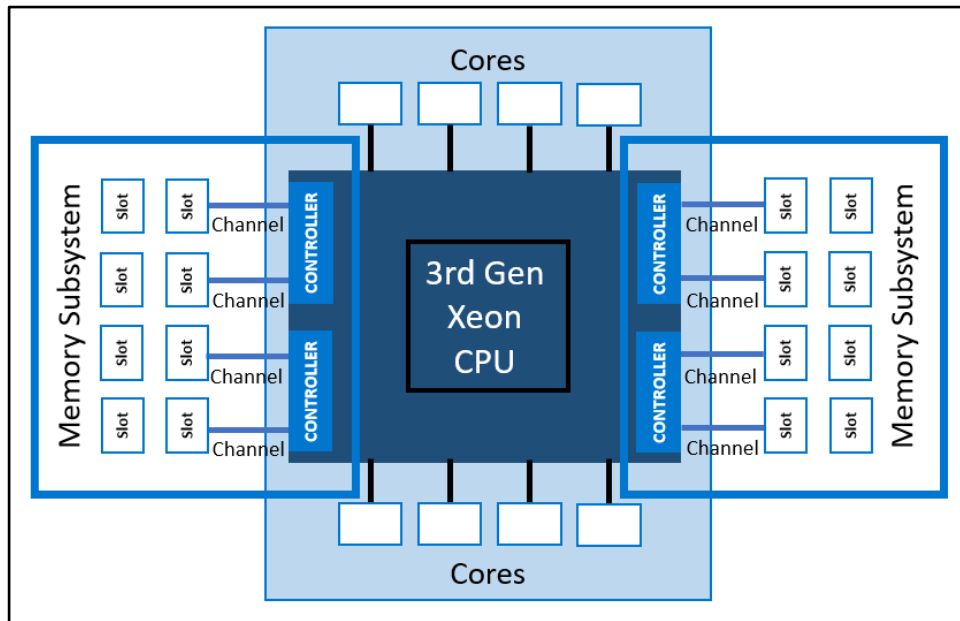


Figure 1: Overview of the CPU-to-memory subsystem connectivity




Figure 2: Decoding DIMM lettering/numbering to Channel lettering/numbering

To understand the relationship between the CPU and memory, terminology illustrated in Figure 1 must first be addressed:

- Memory controllers are digital circuits that manage the flow of data going from the computer's main memory to the corresponding memory channels. Each controller was intended to be populated with at least one DIMM. Xeon processors have four memory controllers.
- Memory channels are the physical layer on which the data travels between the CPU and memory modules. Channels were intended to be populated in a symmetrical fashion, so that when two channels horizontal of one another are populated an interleave set will be created. Xeon processors have eight memory channels.

PROCESSOR	CHANNEL A	CHANNEL B	CHANNEL C	CHANNEL D	CHANNEL E	CHANNEL F	CHANNEL G	CHANNEL H
Processor 1	Slots A1 and A9	Slots A5 and A13	Slots A3 and A11	Slots A7 and A15	Slots A2 and A10	Slots A6 and A14	Slots A4 and A12	Slots A8 and A16
Processor 2	Slots B1 and B9	Slots B5 and B13	Slots B3 and B11	Slots B7 and B15	Slots B2 and B10	Slots B6 and B14	Slots B4 and B12	Slots B8 and B16

*Figure 3: Memory Channels*


- Memory slots host individual memory modules, such as DIMMs or DCPMMs. Xeon processors have two slots per channel, so there are a total of sixteen slots per CPU for memory module population. As seen in [Figure 2](#), for CPU1, DIMM 0 slots A1-A8 (the white slots) are the first eight memory modules to be populated, while DIMM 1 slots A9-A15 (the black ones) are the last eight to be populated. Similarly, for CPU2, DIMM 0 slots B1-B8 (the white slots) are the first eight memory modules to be populated, while DIMM 1 slots B9-B15 (the black ones) are the last eight to be populated. DIMM slot/channel numbering can also be referenced in [Figure 3](#).
- The memory subsystem is the combination of all the independent memory functions listed above.

### 3. Memory Interleaving

Memory interleaving allows a CPU to efficiently spread memory accesses across multiple DIMMs. When memory is put in the same interleave set, contiguous memory accesses go to different memory banks. Memory accesses no longer must wait until the prior access has completed before initiating the next memory operation. For most workloads, performance is maximized when all DIMMs are in one interleave set creating a single uniform memory region that is spread across as many DIMMs as possible. Multiple interleave sets create disjointed memory regions.

3rd Generation Intel® Xeon® Scalable Processors create interleave sets to efficiently spread memory accesses across memory controllers and channels, in order to increase memory data transfer speeds. The number of interleave sets created corresponds to how the memory modules populating each slot are configured. Preserving only one interleave set will improve memory bandwidth by utilizing all memory channels while any memory is accessed. As a result, the distribution of information is divided across several channels instead of just one, and the total memory bandwidth is increased. Generating any more than one interleave set will create additional work for the memory controllers, which will reduce memory bandwidth. Therefore, the optimal memory subsystem design centers around generating only one interleave set.

Configuring memory to be balanced ensures that the number of required interleave sets is minimized; a total of one set. Anything beyond one interleave set is not fully balanced and will be referred to as near balanced (partial fulfillment of balanced conditions) or as unbalanced (no fulfillment of balanced conditions) within this brief. As seen in [Figure 3](#), balanced memory configurations are heavily dependent upon properly configuring DIMMs/DCPMMs within the subsystem. Below are reflective of 1 CPU configurations; for CPU2 configurations, populate respective DIMMs on CPU2 (B1-B16).



*Figure 4: Side by side comparison of an unbalanced and balanced memory configuration. The 7 DIMM unbalanced configuration has created two interleave sets due to non-symmetrical memory channel population, represented by blue and green colors. The IMC has essentially broken up memory into disjointed regions that degrade performance and create unpredictable process patterns. The 8 DIMM balanced configuration has symmetrical memory channel population across all memory channels, therefore creating only one interleave set and thus increasing memory bandwidth*

## 4. Guidelines for Balancing Memory

### 4.1 Overview

DIMMs must be populated into a balanced configuration to yield the highest and most consistent memory bandwidth with the lowest memory access latency. Various factors will dictate whether a configuration is balanced or not. Please follow the guidelines below for best results:

- Memory Channel Population
  - Balanced Configuration
    - Each memory channel must be fully populated with one or two DIMMs for best performance; a total of **eight** or **sixteen** DIMMs per CPU
  - Near Balanced Configuration
    - Each memory controller must be populated with at least one DIMM in a symmetrical fashion; a total of **four**, **six**, **twelve** or **fourteen** DIMMs per CPU; DIMMs must be populated in sequential order starting from A1-A16
    - Symmetrical refers to two memory channels that are horizontally flipped
- CPU and DIMM parts should be identical
- Each CPU must be identically configured with memory



## 4.2 Memory Channel Population

To achieve a balanced configuration, populate either eight or sixteen DIMMs per CPU. By loading each channel with one or two DIMMs, the configuration is balanced and has data traveling across channels most efficiently on one interleave set. Following this guideline will yield the highest memory bandwidth and the lowest memory latency.

If a balanced configuration of sixteen or eight DIMMs per CPU is not possible, the next best option is a near balanced configuration. To obtain a near balanced population, populate four, six, twelve or fourteen DIMMs per CPU in a symmetrical fashion. When any number of DIMMs other than four, six, twelve or fourteen is populated, disjointed memory regions are created which introduce more interleave sets.

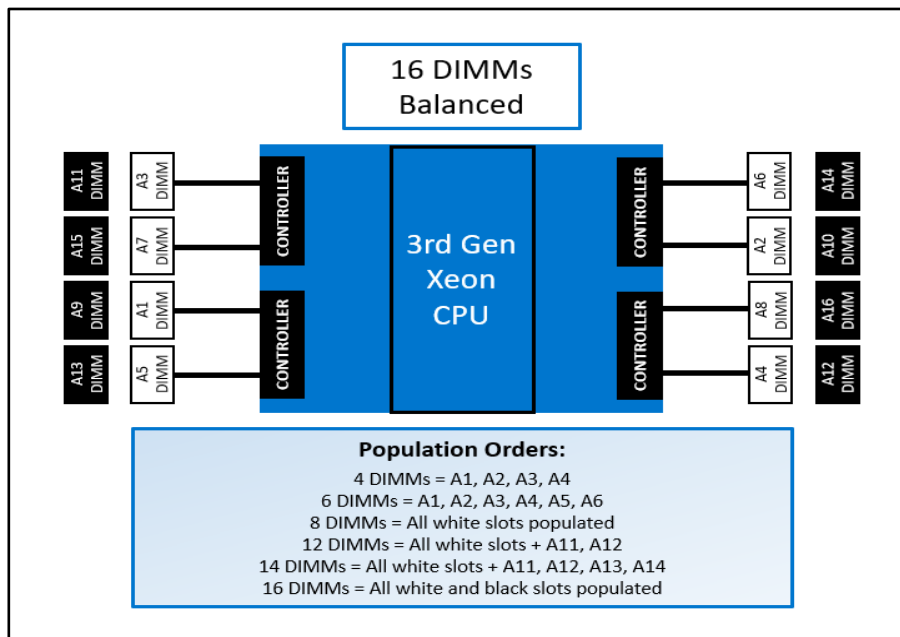


Figure 5: DIMM population order for balanced and near balanced configurations

The last guideline is that DIMMs must be populated in an assembly order because Xeon processors have a symmetrical architecture for each type of CPU core count. Figure 4 illustrates the assembly order in which individual DIMMs should be populated, starting with A1 and ending with A16. Figure 4 is reflective of 1 CPU configurations; for CPU2 configurations, populate respective DIMMs on CPU2 (B1-B16).

## 4.3 Identical CPU and DIMM Parts

Although Dell Technologies does allow for DIMM mixing (up to two different DIMM PNs in a system), and this would still be considered a balanced configuration, it is good practice to have identical DIMMs used across all DIMM slots. If memory DIMMs are installed with different speeds (they must have the same rank, capacity and DIMM type) then they will operate at the speed of the slowest DIMM. This principle applies to the processors as well; multi-socket Intel systems shall be populated with identical CPUs.

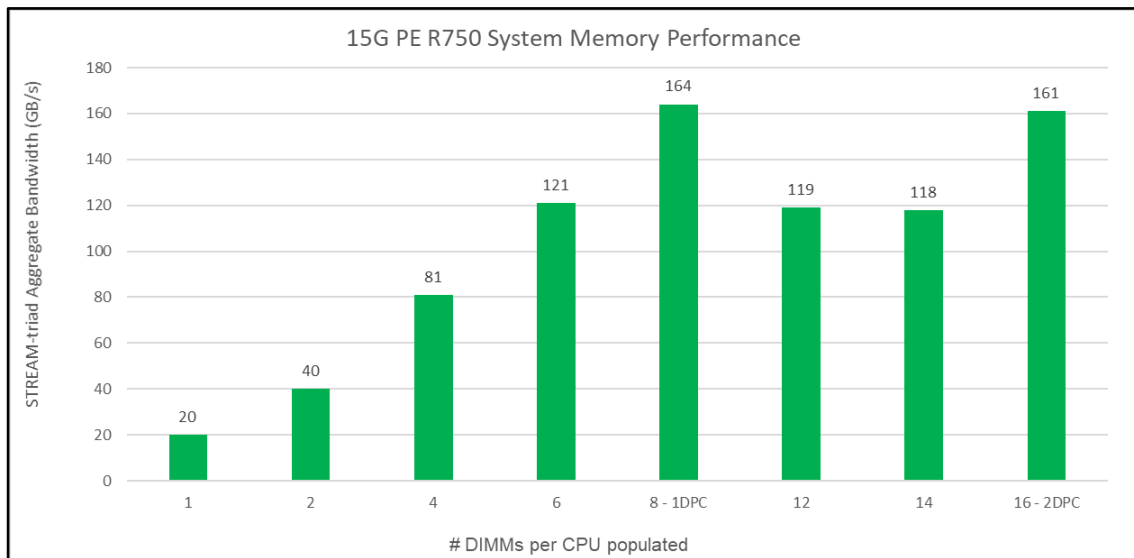
## 4.4 Identical Memory Configurations for Each CPU

Every CPU socket within a server must have identical memory configurations. When only one unique memory configuration exists across all the CPU sockets within a server, memory access is further optimized.

## 4.5 Summary

- Populate DIMMs in a balanced configuration to yield best performance
  - Near Balanced configurations are second best
  - Unbalanced configurations are not recommended
- DIMMs must be populated symmetrically & in sequential order starting from A1-A16, as discussed in 4.2
- Each memory slot should have identical DIMMs
- Each server must contain identically configured CPUs

If these principles are applied across Dell EMC PowerEdge servers with 3rd Generation Intel® Xeon® Scalable Processors, socket-level bandwidth and memory-level performance will be optimized. [Figure 6](#) below illustrates the expected memory bandwidth curve when these rules are followed. Although not shown, please note that unbalanced configurations will see significant performance degradation compared to both the balanced and near balanced configurations displayed.



*Figure 6: Bar graph illustrating expected performance variation as # of dimms increases*

## 5. Balanced Configurations (Recommended)

These recommended configurations satisfy balanced conditions by requiring each memory channel to be populated with one or two identical DIMMs. By doing this, one interleave set can optimally distribute memory access requests across all the available DIMM slots; therefore, maximizing performance. Memory controller logic was designed around fully populated memory channels, so it should come as no surprise that **eight or sixteen populated DIMMs are recommended**. Having eight DIMMs will reap the highest memory bandwidth while having sixteen DIMMs will yield the highest memory capacity.

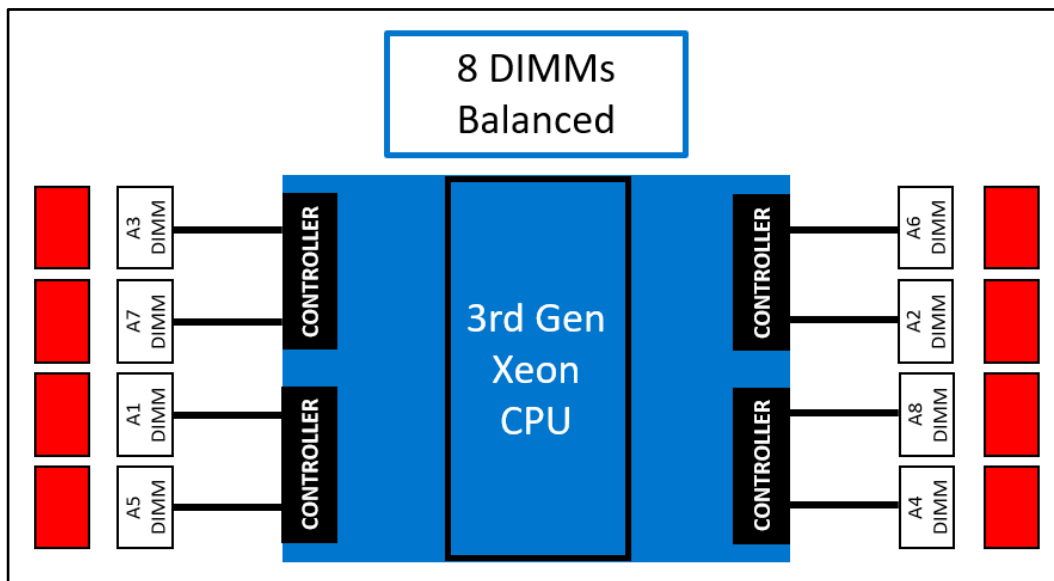


Figure 7: Eight DIMMs are populated in a balanced configuration, producing the highest memory bandwidth while at a lower capacity than sixteen

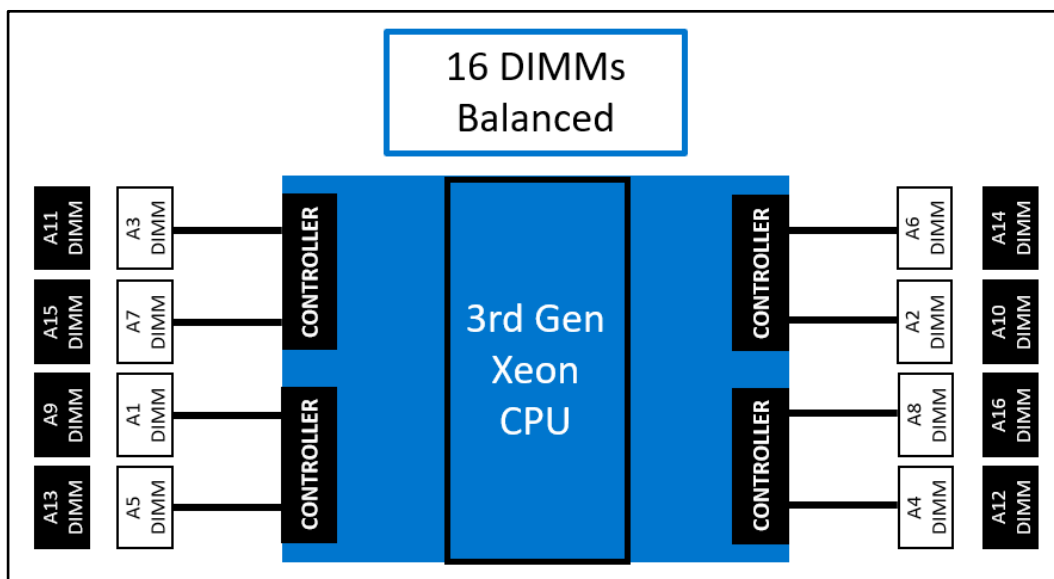


Figure 8: Sixteen DIMMs are populated in a balanced configuration, producing the highest memory capacity while at a lower bandwidth than eight

## 6. Near Balanced Configurations

These configurations satisfy near balanced conditions by populating four, six, twelve or fourteen identical DIMMs per CPU. Although these configurations do have at least two DIMMs populated per controller, they are not optimized because the channels are only partially populated, which creates disjointed memory regions that reduce performance. Memory performance for near balanced configurations can undergo degradation when compared to balanced configurations, but they are still adequate for implementation.




Figure 9: Four DIMMs are populated in a near balanced configuration

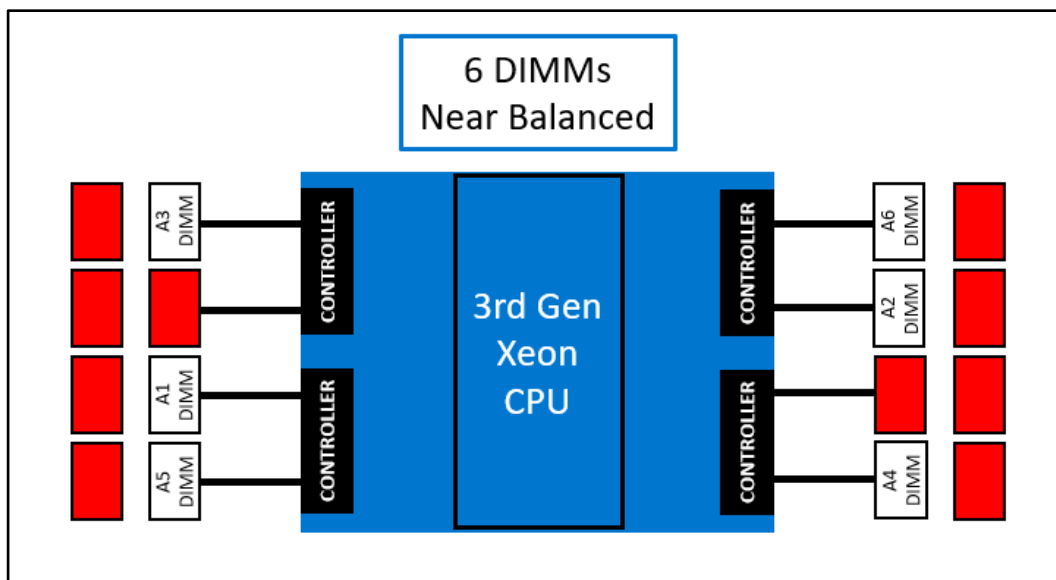


Figure 10: Six DIMMs are populated in a near balanced configuration




Figure 11: Twelve DIMMs are populated in a near balanced configuration




Figure 12: Fourteen DIMMs are populated in a near balanced configuration

## 7. Unbalanced Configurations

Because the below configurations do not have at least one DIMM for each controller, we do not recommend using them. Two or more interleave sets can now be introduced to the memory controller algorithm which causes very disjointed regions. Memory performance for these configurations are *significantly* less than configurations meeting balanced or unbalanced conditions. **We do not support unbalanced configurations except for 1 DIMM & 2 DIMM configurations.**




Figure 13: One DIMM is populated in an unbalanced configuration

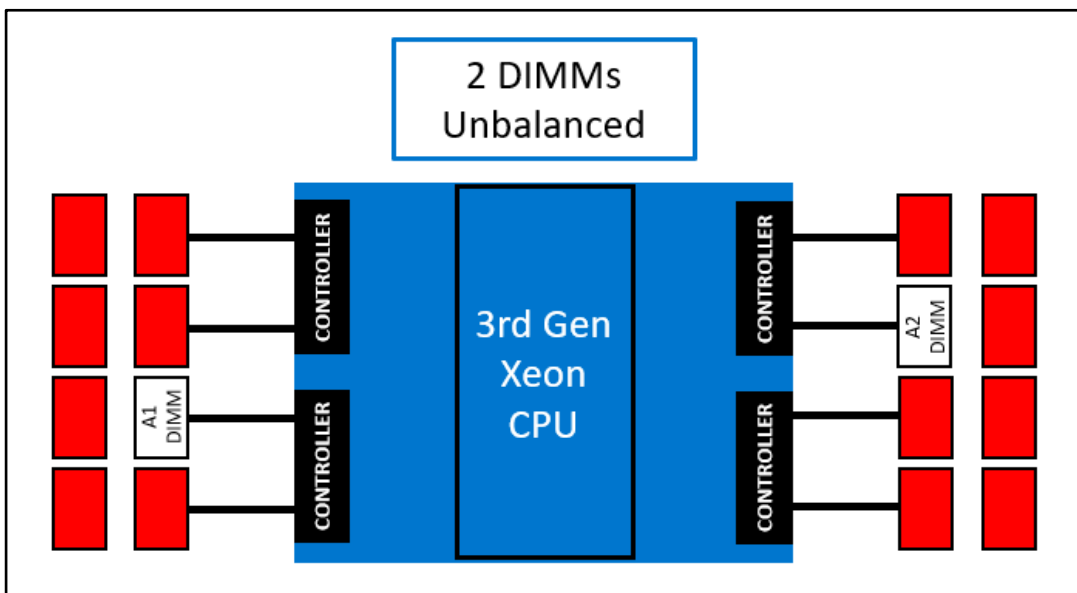
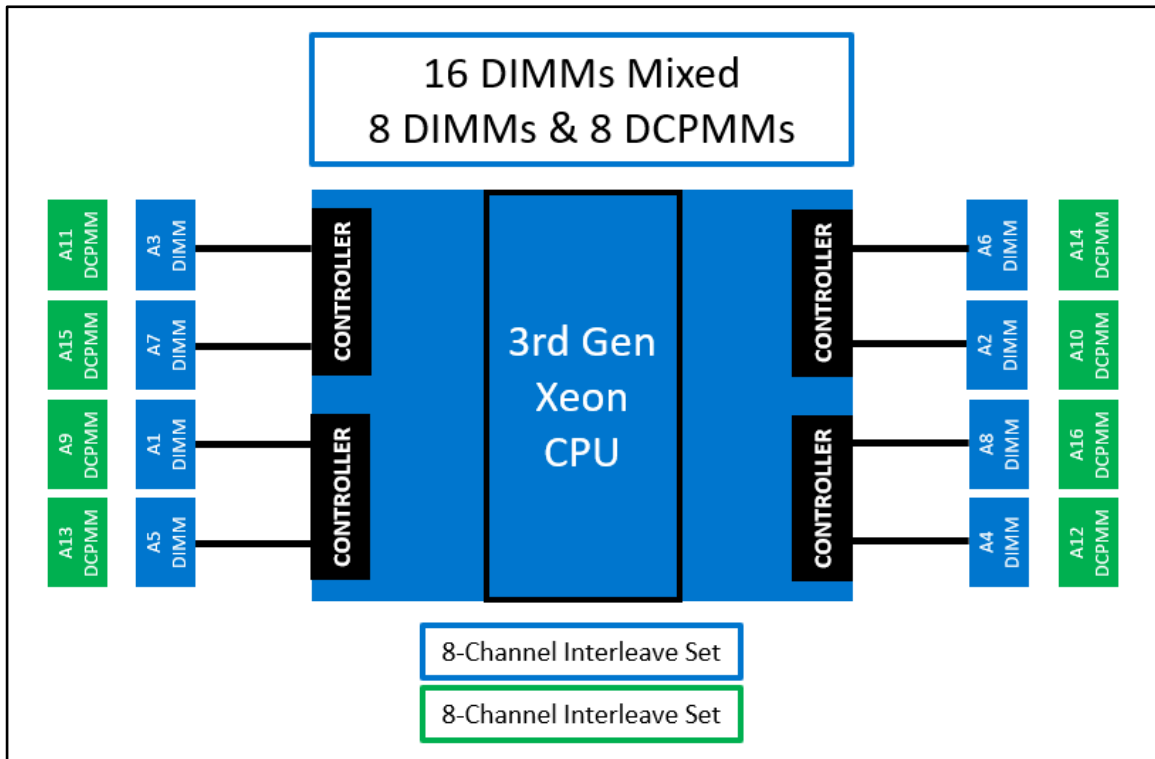


Figure 14: Two DIMMs are populated in an unbalanced configuration

## 8. Traditional DIMMs with DCPMMs

For mixed memory configurations, populating eight traditional DIMMs in the inner eight slots and eight DCPMMs in the outer eight slots will reap the highest memory bandwidth. The benefits gained from DCPMMs persistence and increased capacity will largely offset the memory bandwidth degradation from having two interleave sets.



*Figure 15: When using both DIMMs and DCPMMs, populating the inner eight slots with same-capacity traditional DIMMs and the outer eight slots with same-capacity DCPMMs is recommended. Two interleave sets are created because DIMMs and DCPMMs are unique drives with unique capacities, making them unbalanced. Despite the memory bandwidth degradation that occurs from having two interleave sets, the persistence and increased capacity gained from DCPMMs will still enable various potential benefits*

## 9. Conclusion

Balancing memory with 3rd Generation Intel® Xeon® Scalable Processors increases memory bandwidth and reduces memory access latency. When memory modules are configured in such a way that the memory subsystem capacities are identical, interleave sets will sync together to perform most efficiently at both a socket-level and server-level.

Applying the balanced memory guidelines demonstrated in this brief will ensure that both memory bandwidth and memory access latency are optimized for the best system performance possible.

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell Technologies and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.