Memory Population Requirements for Lenovo ThinkSystem Servers

Describes the rules for memory DIMM population for Lenovo ThinkSystem

Describes the available memory modes

Gives examples of supported and unsupported DIMM population

Introduces the memory wizard to help select the best memory configuration

Sumanta K Bahali

LENOVO PRESS

Abstract

This paper examines the memory population requirements of the Lenovo® ThinkSystem™ platforms based on the Intel Xeon Processor Scalable family. We describe both Intel and Lenovo specific requirements complete with various examples that show supported and unsupported memory population. This paper explains the value of using the Memory Wizard in the Lenovo Enterprise Server Configurator (LESC) as an easy way to determine the best memory configuration for ThinkSystem servers.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

http://lenovopress.com

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction

Lenovo ThinkSystem servers are based on Intel Xeon Processor Scalable family CPUs. These CPUs support six memory channels. Each memory channel can support a maximum of 2 DIMMs. Thus one CPU can support a maximum of 12 DIMMs.

Most ThinkSystem servers support the full 12 DIMMs per processor (2 DIMMs per channel), however, some models support fewer DIMMs:

- ThinkSystem SR530, SR550, ST550 support a maximum of 6 DIMMs per processor which means they only support 1 DIMM for each memory channel.
- ThinkSystem SD530, SR570 and SR590 support 8 DIMMs per processor, which means that one memory channel per memory controller supports 2 DIMMs and the other memory channels only have one DIMM per channel.

Previous Intel E7 v4 and E5 v4-based systems generally supported any combination of memory DIMMs. Systems based on the Intel Xeon Processor Scalable family have more complex rules for validated memory configurations.

There are specific rules on memory populations for the Intel Xeon Processor Scalable family. Some of the rules are based on Intel validation and support. Other rules are based on Lenovo DIMM population requirements. The goal of these rules is to provide the best memory solution for any specific memory capacity.

Intel Xeon Scalable Family Processors

The Intel Xeon Scalable Family series of processors provides a common building block across the entire Lenovo ThinkSystem platforms.

Figure 1 and Figure 2 on page 4 illustrate CPU and memory architecture block diagrams for ThinkSystem two- and four-socket platforms. Some system architectures may differ from what is shown, so always refer to platform-specific documentation for precise system details.

The 2-socket (2S) block diagram shows two Intel UltraPath Interconnect (UPI) links between processors, and six memory channels per processor with two DIMMs maximum per memory channel for a total of 24 DIMM slots.

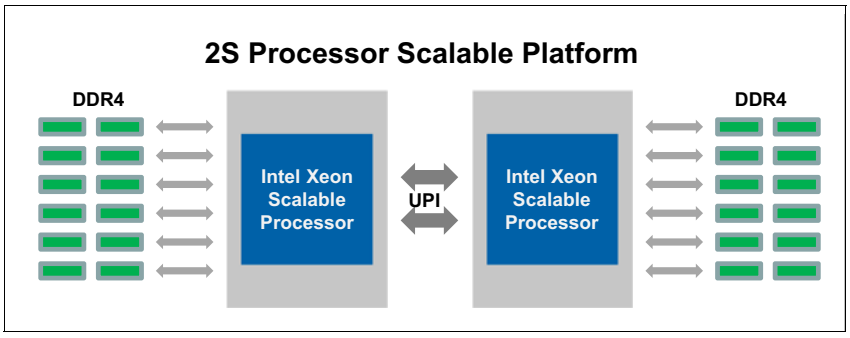

Figure 1 2-socket CPU and memory architecture block diagram

The 4-socket (4S) block diagram, Figure 2 on page 4, shows a fully-connected UPI topology between all four processors, also called a *mesh topology*. Each processor has six memory channels with two DIMMs maximum per memory channel for a total of 48 DIMM slots.

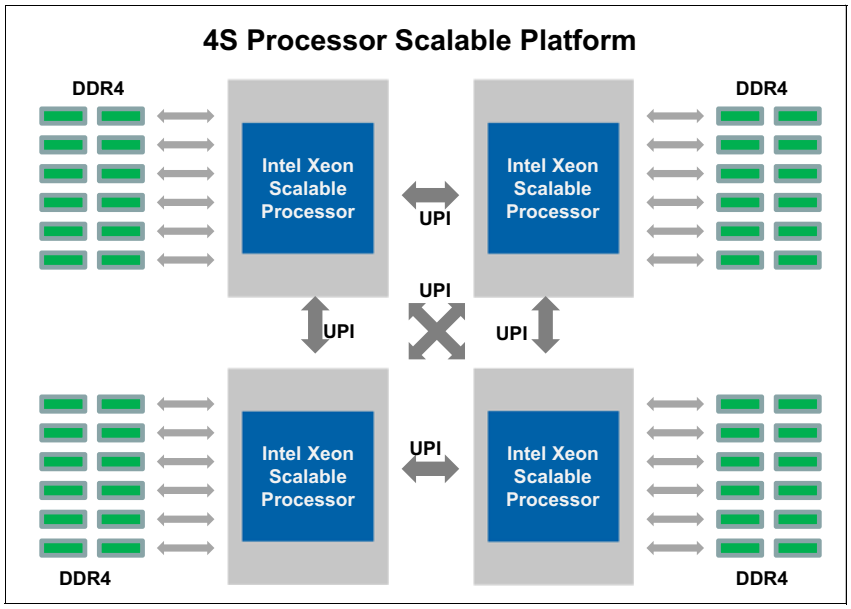

Figure 2 2-socket and 4-socket CPU and memory architecture block diagrams

The processor has two independent memory controllers. As shown in Figure 3, each memory controller supports three memory channel and each memory channel can support up to two DIMMs. The two memory controllers are referred to as iMC1 and iMC2. The memory channels 1 to 3 are connected to iMC1 and memory channels 4 to 6 are connected to iMC2.

These two memory controllers work independently to fulfill the memory requirement from the processor cores. System firmware will attempt to interleave memory address range across the memory controllers when possible e.g. for symmetric memory population. Interleaving memory addresses across controllers improves memory performance.

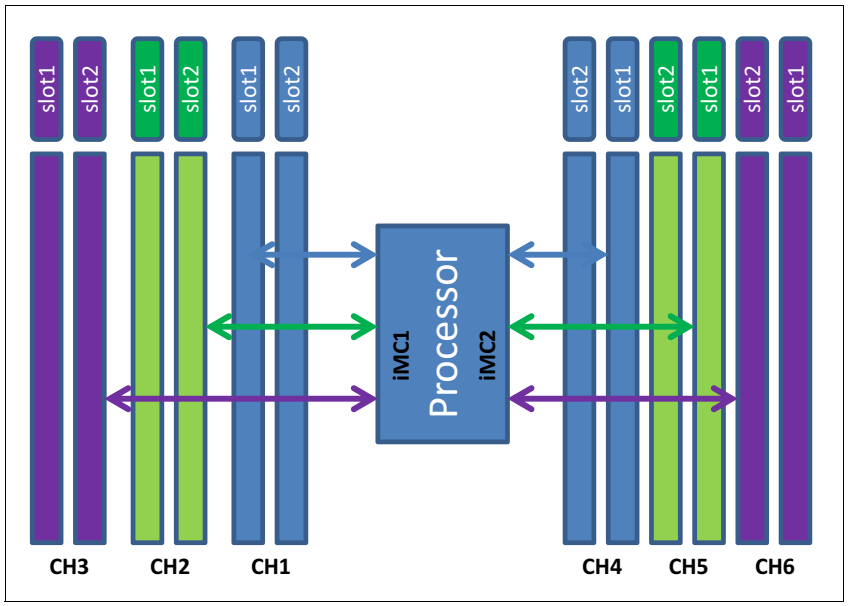

Figure 3 Processor memory controller and memory channels

Types of DIMMs

The Lenovo ThinkSystem server family supports a variety of memory DIMM types and capacities. Refer to platform-specific documentation for supported memory DIMM types for the most up to date information.

The memory DIMMs supported across the various ThinkSystem servers are:

- ThinkSystem 8GB TruDDR4™ 2666 MHz (1Rx8 1.2V) RDIMM
- ThinkSystem 16GB TruDDR4 2666 MHz (1Rx4 1.2V) RDIMM
- ThinkSystem 16GB TruDDR4 2666 MHz (2Rx8 1.2V) RDIMM
- ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM
- ThinkSystem 64GB TruDDR4 2666 MHz (4Rx4 1.2V) LRDIMM
- ThinkSystem 128GB TruDDR4 2666 MHz (2S4Rx4 1.2V) 3DS RDIMM

Many end users configure server memory based on a capacity requirement. Historically, RDIMMs have been favored for lower capacity requirements and LRDIMMs have been selected for higher capacity requirements. A relatively new technology is three-dimension stacking (3DS) which stacks DRAM components using Through Silicon Vias (TSV) and allows for greater capacity DIMMs without increasing the load on the memory bus. The 128GB RDIMM listed above is a 3DS RDIMM.

Memory population rules

In this section we document the memory population rules from Intel and Lenovo. The memory in ThinkSystem servers can be configured in three different memory modes:

- Independent Channel Mode

  All six channels can be populated independently as long as the memory population rules are met. All channels run at the same interface frequency, however individual channels may run at different DIMM timings (RAS latency, CAS latency, and so forth) if the installed DIMMs require it. All the memory DIMMs and ranks within DIMMs are available for system use. Most users will use Independent mode.

- Rank Sparing Mode

  A DIMM may contain one or more rank. A rank refers to a set of DRAM in the DIMM that fulfils a memory operation. For example a 32GB 2Rx4 DIMM has 36 DRAM devices. But only 18 devices (one rank) fulfill a memory transaction such as memory read. Different ranks are selected using DIMM Chip Select signal.

  In Rank Sparing Mode, one or more ranks are spares for the other ranks on the same channel. The spare rank(s) are held in reserve and are not available as system memory. The spare rank(s) must have identical or larger memory capacity than any of the other ranks (sparing source ranks) on the same channel. After sparing, the sparing source rank will no longer be available as system memory.

- Mirror Mode between Channels

  In Mirrored Channel Mode, the memory contents are mirrored between Channels 1, 2, and 3 or Channels 4, 5, and 6. Mirrored mode provides resiliency even when one of the DIMM encounters a critical uncorrectable error. The available memory capacity is only half of the total DIMM capacity.

Memory population rules for Independent Mode

In this section we discuss the memory population rules when system memory is configured in independent mode.

Intel rules

The memory rules for ThinkSystem servers consist of two sections, Intel rules and additional Lenovo rules.

The Intel rules are based on Intel and Lenovo validation and are as follows:

- Within a server, RDIMMs, LRDIMMs or 3DS DIMMs cannot be mixed. Each DIMM type is mutually exclusive within a server
- Within a CPU populate Integrated Memory Controller 1 (iMC1) first (The iMCs are being referred to as iMC1 and iMC2)
- Within an Integrated Memory Controller:
  - Populate Memory Channel 1 (CH1) first
  - CH2 will be either empty or must be identically populated as CH1
  - CH3 will be either empty or must be identically populated as CH2
  - For an system that supports memory channels with 2 slots each, 5 DIMMs is a supported exception to these population rules.
  - 5 DIMMs would be populated 2-2-1 where channels 1 and 2 each have 2 DIMMs and channel 3 has 1 DIMM
  - For 2-1-1 memory topology, e.g. as in SD530, CH1 can support 2 DIMMs while CH2 and CH3 have one DIMM each. This is a special case covered by Intel and Lenovo validation.
  - Within a memory channel populate SLOT1 first.
  - If a memory channel has two DIMMs populated and the DIMMs have different numbers of ranks, populate the DIMM with the higher number of ranks in SLOT1. Doing so improves signal integrity
  - If two DIMMs on a channel have identical ranks, then populate the one with the higher capacity in SLOT1
- Intel Xeon Processor Scalable family may limit the memory capacity depending of SKU. Please select appropriate SKU to address the memory capacity. For example “M” SKU CPU is required for a CPU socket supporting more than 768 GB of memory

Additional rules from Lenovo

At Lenovo we strive to get the maximum performance out of memory subsystem. These additional rules help improve the memory performance.

Please also refer to the Lenovo Press paper on memory performance, *Intel Xeon Scalable Family Balanced Memory Configurations*, available from:

https://lenovopress.com/lp0742

Lenovo configurations rules are as follows:

- All CPUs within a server will have identical memory population
- Each individual system specification provides the DIMM locations to be populated given a number of DIMMs. In most cases, we will alternate the DIMM population order between

CPUs first and then between iMCs on each CPU. There are some exceptions as noted in system specifications.

- If a CPU has exactly three identical DIMMs (same Lenovo part number), populate all three DIMMs on iMC1.
- If a CPU has exactly 10 identical DIMMs (same Lenovo part number), populate these DIMMs with five on iMC1 and five on iMC2. If the DIMMs are not identical then populate six DIMMs on iMC1 and four DIMMs on iMC2

Examples of supported memory populations

This section provides example configurations that use valid and supported memory configurations.

In Figure 4, CH1 on iMC1 has an 8GB DIMM and CH4 on iMC2 has a 16GB DIMM.

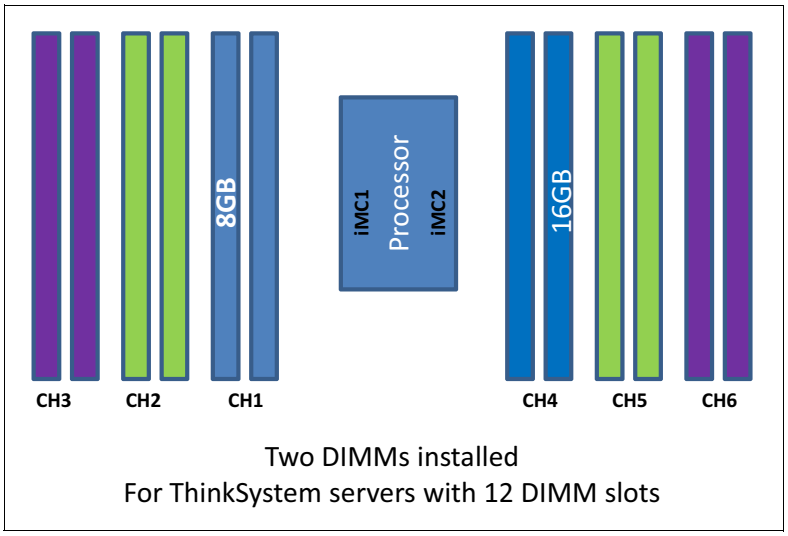

Figure 4 Configuration of two DIMMs: 1x8GB and 1x16GB

In Figure 5 on page 8, DIMM population on iMC1 and iMC2 are different but on each iMC the channels are populated the same. This is supported.

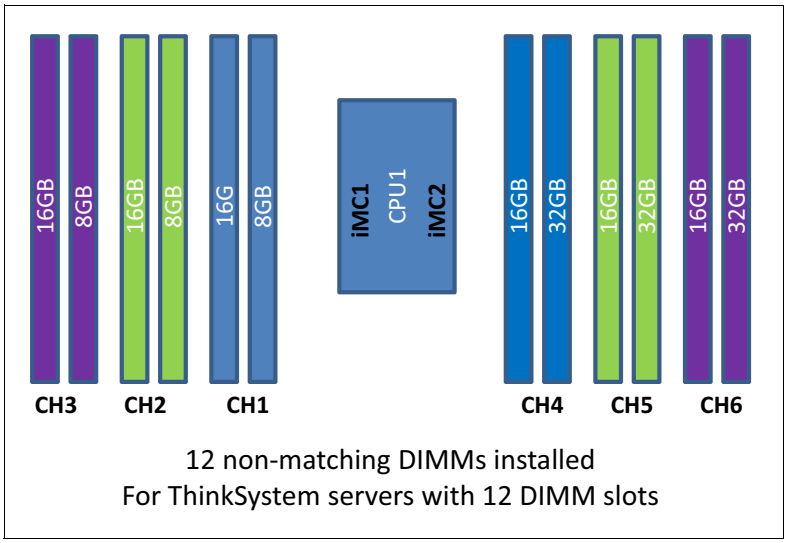

Figure 5 Configuration of 12 non-matching DIMMs

In Figure 6, Slot 1 of CH1, CH2 and CH3 on each iMC populate identical DIMMs as are Slot 1 on CH4, CH5 and CH6, however Slot 2 of CH1 and CH4 are different. This is a special case and is supported for SD530.

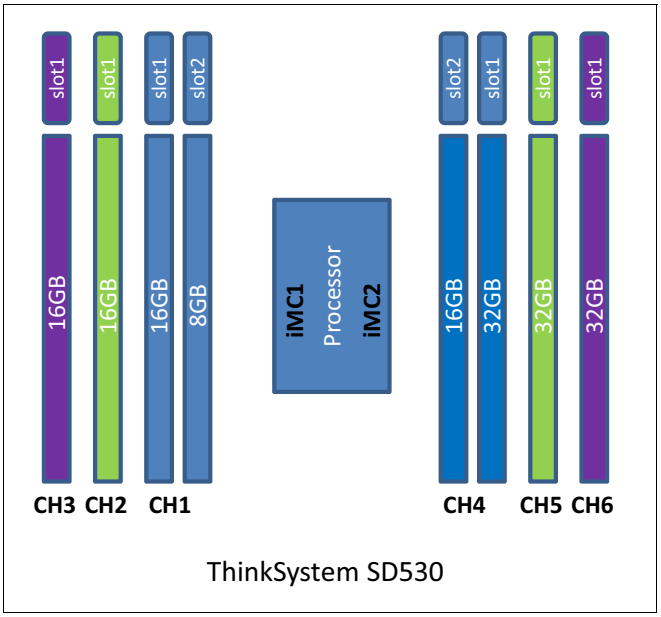

Figure 6 SD530 memory configuration

Examples of unsupported and not recommended memory populations

This section provides example configurations that use memory configurations that are either unsupported or are not recommended.

Figure 7 on page 9 meets Intel's rule but does not meet Lenovo rule of alternately populating between memory channels. Four DIMMs on a CPU will be populated with one DIMM on CH1 (or CH4) and one DIMM on CH2 (or CH5) on each iMC.

If 48GB per CPU is needed then the then a configuration of 6x8GB is the best solution.

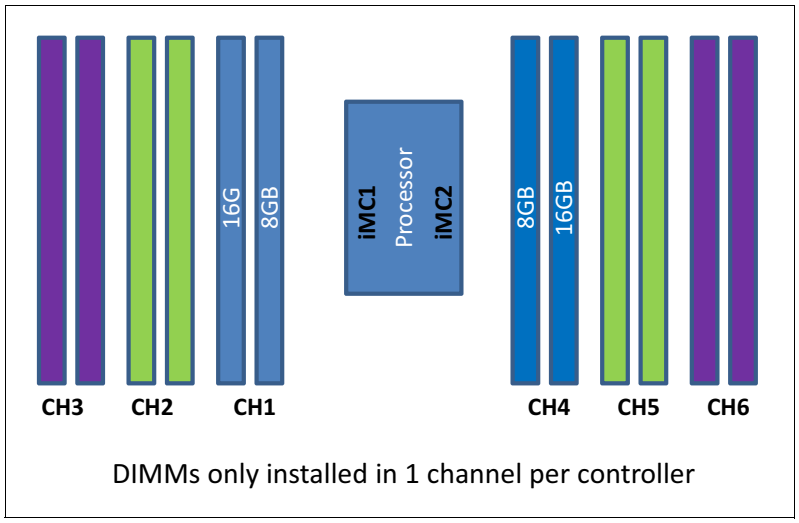

Figure 7 DIMMs only installed in one memory channel

In Figure 8, DIMM population on CH2 of iMC1 is different from CH1. This does not meet Intel rules.

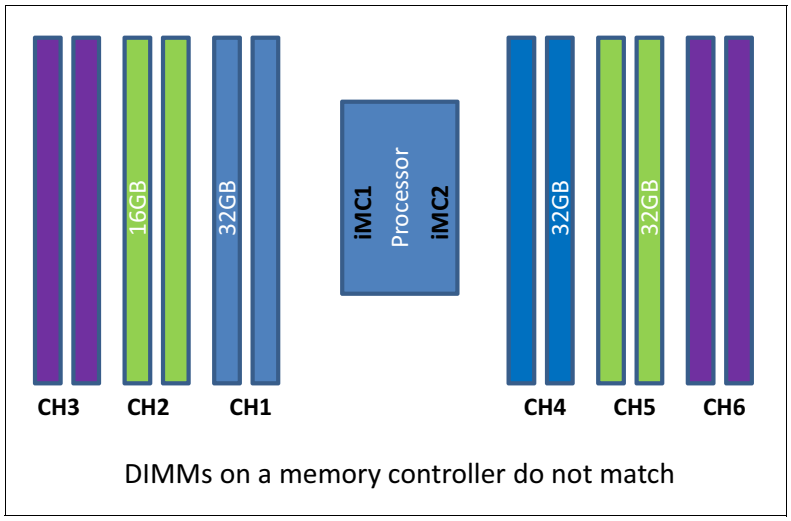

Figure 8 DIMMs on a memory controller do not match

In Figure 9 on page 10, DIMM population on each CPU meets Intel requirement, but the configuration does not meet the Lenovo requirement that all CPUs be populated identically. Populating DIMMs equally among all CPUs helps with better memory subsystem performance and improves customer satisfaction.

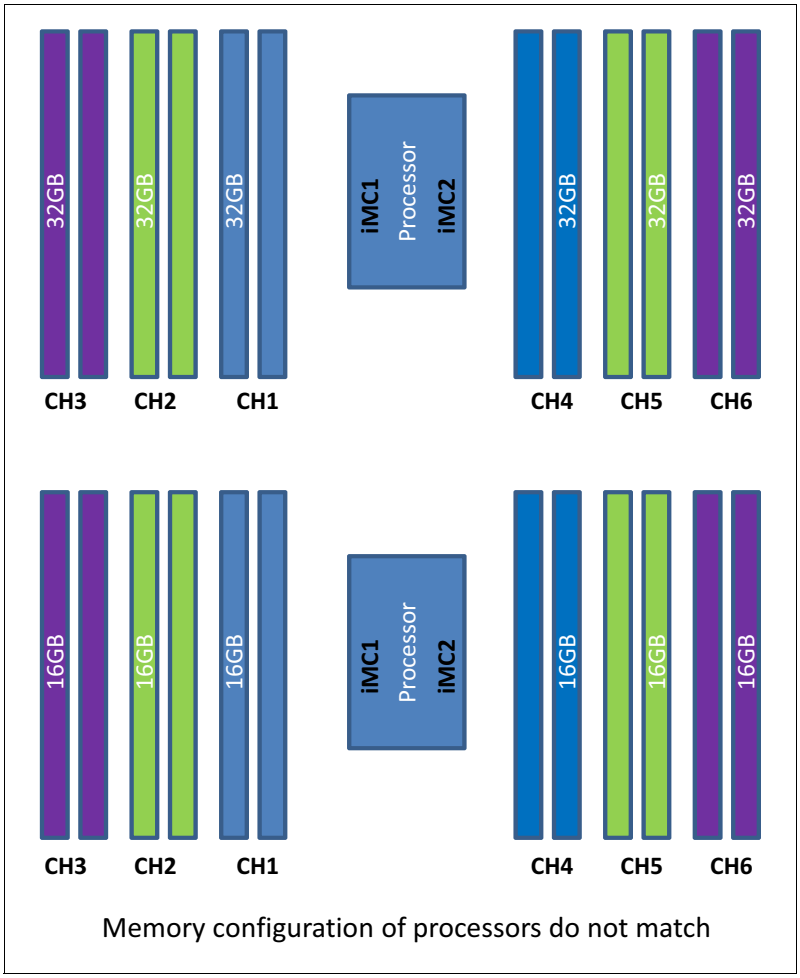

Figure 9 Memory configuration of processors do not match

Lenovo Memory Configuration Wizard

From the above examples it is apparent that memory population support on ThinkSystem Servers is complex. Lenovo developed a Memory Wizard addition to the Lenovo Enterprise Solutions Configurator (LESC) to address this complexity.

The Memory Wizard is located at:

http://lesc.lenovo.com/ss/#/memory_configuration

At the time of writing, the wizard is a standalone tool, however, there are plans to also integrate it into configurations built in LESC.

The Memory Wizard will help end users arrive at a supportable memory configuration for a given capacity for any of our ThinkSystem platforms. In addition, memory configurations can selectively be optimized for performance, price or a combination of performance and price.

For more information on populating balanced memory configurations for optimal memory performance, see the Lenovo Press paper *Intel Xeon Scalable Family Balanced Memory Configurations* available from:

https://lenovopress.com/lp0742-intel-xeon-scalable-balanced-memory-configurations

Memory mirroring rules

Memory mirroring mode provides greater memory resiliency. The available memory is one half of the total memory.

The following are additional Intel rules for mirroring support:

- ThinkSystem servers can support only 2, 3, 4 or 6 DIMMs per iMC when memory mirroring is to be enabled. The servers cannot support either 1 DIMM or 5 DIMMs per iMC with memory mirroring
- As with independent mode, equal DIMM sizes must be installed on the populated memory channels. DIMM slot population within a channel does not have to be identical, but the same DIMM slot location across Channel 0, Channel 1, and Channel 2 OR between Channel 3, Channel 4 or Channel 5 must be populated identically. See Figure 10.

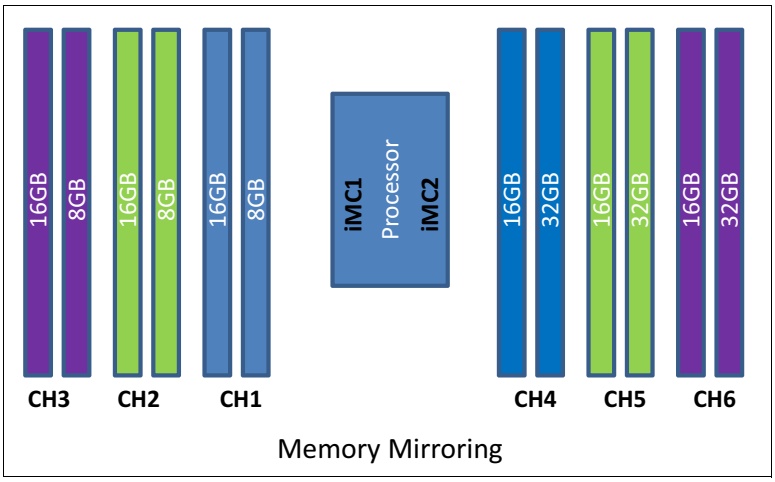

Figure 10 Mirroring example

- If DIMMs are installed in all three channels, then mirroring occurs across all three DIMM channels. In such a situation, the memory data is cloned between Channels 0, 1 and Channel 1, 2 and channels 2, 0.If DIMMs are installed in only two channels, then mirroring occurs across two DIMMs. That means memory data is cloned between Memory DIMM on CH0 and Memory DIMM on CH1. You cannot mix 2-channel mirroring and 3-channel mirroring in the same IMC.

Table 1 shows the DIMM population for memory mirroring.

Table 1 DIMM population for mirroring

Number of DIMMs to mirror	CH0		CH1		CH2	
	DIMM1	DIMM0	DIMM1	DIMM0	DIMM1	DIMM0
2	No	Yes	No	Yes	No	No
3	No	Yes	No	Yes	No	Yes
4	Yes	Yes	Yes	Yes	No	No
6	Yes	Yes	Yes	Yes	Yes	Yes

Memory Sparing rules

Sparing mode is another mode that provides greater memory resiliency for higher RAS. One rank of memory is kept as a spare of other ranks of memory in the channel and the spare rank is used in case UEFI detects too many correctable memory errors in the other memory ranks.

- Every populated memory channel must have at least two ranks. So if a channel has only one DIMM and it is a 1R DIMM then that channel can't support Sparing

Further reading

For more information, consult these resources:

- Intel Xeon Scalable Family Balanced Memory Configurations

  https://lenovopress.com/lp0742
- Lenovo ThinkSystem SD530 Performance Considerations with 12 DIMMs and 16 DIMMs

  https://lenovopress.com/lp0659

Author

Sumanta K Bahali is memory architect at the Lenovo Data Center Group (DCG) campus in Morrisville, NC.

Thanks to the following people for their contributions to this project:

- Charles Stephan, System Performance Team lead
- Joseph Jakubowski, System Performance Architect
- David Watts, Lenovo Press

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on May 22, 2018.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp0774

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®	ThinkSystem™
Lenovo(logo)®	TruDDR4™

The following terms are trademarks of other companies:

Intel, Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.