Lenovo ThinkSystem SD530 Performance Considerations with 12 DIMMs and 16 DIMMs

Describes how to optimize the memory configuration of the SD530 for peak performance or maximum capacity

Explains why more DIMMs isn't always better for performance

Provides details about the lab measurements performed

Guidance by experts in the Lenovo Data Center Group Performance Lab

Joe Jakubowski

Dan Colglazier

LENOVO PRESS

Abstract

The Lenovo® ThinkSystem™ SD530 server supports memory configurations that are optimized for performance and for memory capacity but with reduced performance. This performance brief explains the performance considerations for both memory configurations and when to choose one over the other. This brief is intended for IT architects, IT administrators, customer support specialists and engineering staff who are responsible for optimizing server performance with the SD530.

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

Many factors can affect the performance of servers based on the Intel architecture. The memory subsystem plays a key role in overall server performance. Optimal memory performance is achieved when applications running on the server have adequate memory capacity, bandwidth, and low latency.

Memory performance is influenced by factors such as the number of DIMMs populated, the type of DIMM (for example LRDIMM or RDIMM), and whether the DIMMs are installed such that they are evenly balanced across a processor's memory channels and its memory controllers and also across processor sockets.

Populating too few DIMMs or populating the DIMMs in an unbalanced manner across a processor's memory channels or memory controllers can reduce memory bandwidth and increase latency.

The Lenovo ThinkSystem SD530 supports balanced memory configurations with up to 12 DIMMs that are optimized for performance. For increased capacity, the SD530 also supports an unbalanced 16 DIMM memory configuration but with reduced performance.

After reading this performance brief you will understand when to choose 12 DIMM and 16 DIMM memory configurations.

For information about the ThinkSystem SD530, see the Lenovo Press product guide, available from:

http://lenovopress.com/lp0635-thinksystem-sd530-server

Quick summary: Applications that require peak performance and whose memory capacity can be met with 12 DIMMs should be deployed on the SD530 configured with 12 DIMMs.

Only populate 16 DIMMs on the SD530 when the application's memory capacity cannot be met with 12 DIMMs and when peak performance is not required.

SD530 memory topology

The SD530 memory topology has an extra memory DIMM slot on one memory channel for each of the processor memory controllers. Since there are three memory channels on each memory controller, the topology is referred to as 2-1-1 on each memory controller as shown in Figure 1 on page 4.

A 2-1-1 topology means that one memory channel has two DIMM slots and the other two memory channels each have one slot.

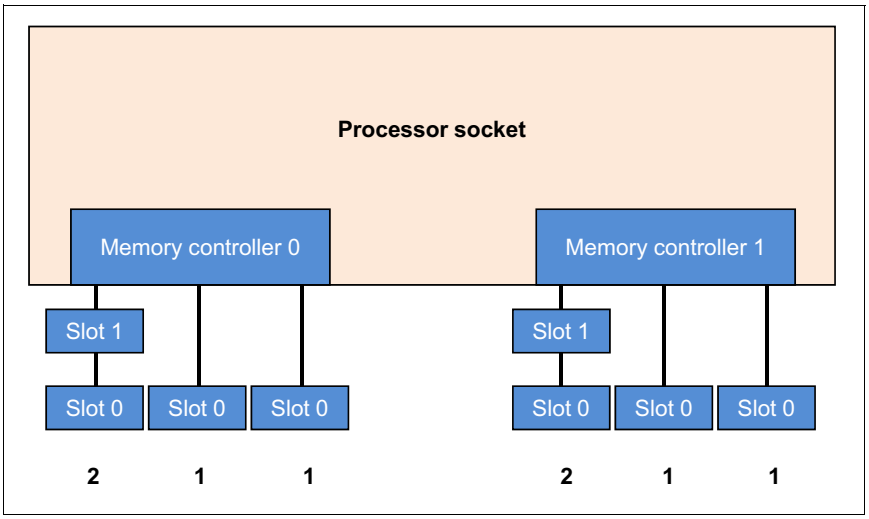

Figure 1 SD530 2-1-1 memory topology

To achieve optimal memory performance, all memory channels should be populated with exactly the same number and type of memory DIMMs. The memory type refers to the memory technology (for example RDIMM or LRDIMM), memory capacity and number of ranks for the DIMM.

For the SD530, the performance is optimized when all memory channels are populated with one of the same type DIMM. This is referred to as a *balanced* 1-1-1 memory topology since each memory channel has exactly one DIMM installed and the DIMM type is identical for all DIMMs. The SD530 with two processor sockets can support 12 total DIMMs when optimized for performance.

The SD530 supports 16 total DIMMs when both slots of the memory channels with two DIMM slots are populated. Again, this is referred to as a 2-1-1 memory topology. A 2-1-1 memory topology is optimized for memory capacity and not for performance. A 2-1-1 memory topology is unbalanced because the memory channels are not identically populated.

Memory regions and interleave sets

The Intel processor optimizes memory accesses by creating memory regions on which interleave sets are formed across the memory controllers and/or the memory channels. For example, if identical DIMMs are populated in Slot 0 on all three memory channels attached to each memory controller, the memory controllers create a single memory region and a 6-way interleave set across the six DIMMs as shown in Figure 2 on page 5.

Interleaving enables higher memory bandwidth by spreading contiguous memory accesses across all the memory channels rather than sending all memory accesses to one memory channel or another.

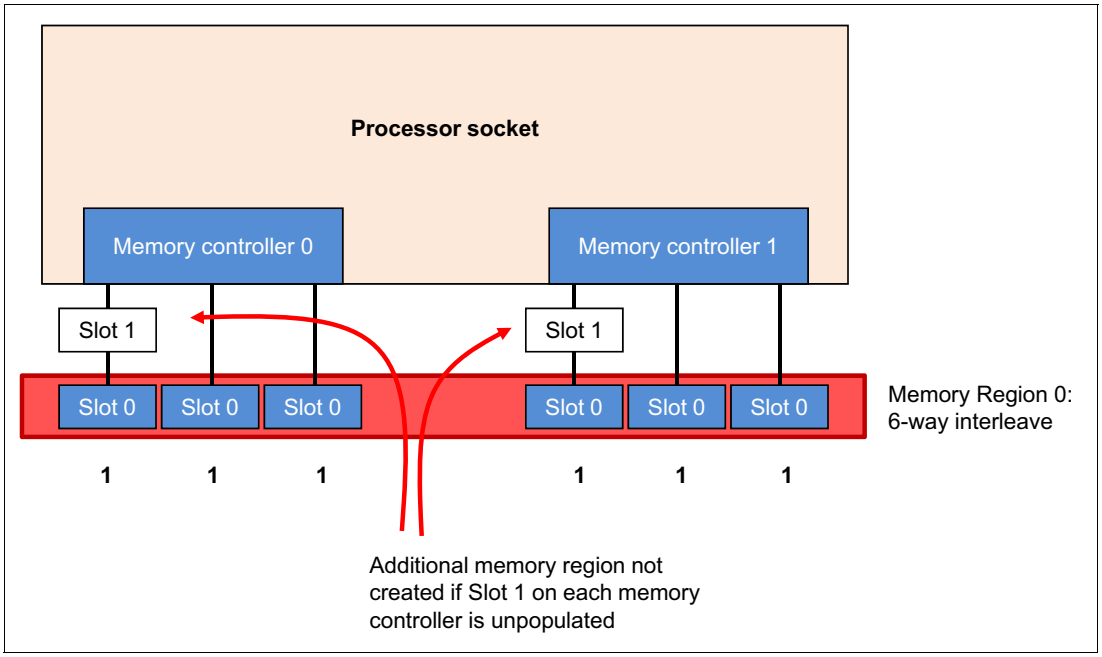

Figure 2 1-1-1 memory topology forming a 6-way interleave set

If all DIMM slots are populated on the SD530 in an unbalanced 2-1-1 configuration, as shown in Figure 3, the memory controller creates multiple memory regions and interleave sets. If the DIMMs are all identical, two memory regions and two interleave sets are created. Some interleave sets will have fewer DIMMs. Managing multiple interleave sets creates overhead for the memory controller which can reduce memory bandwidth.

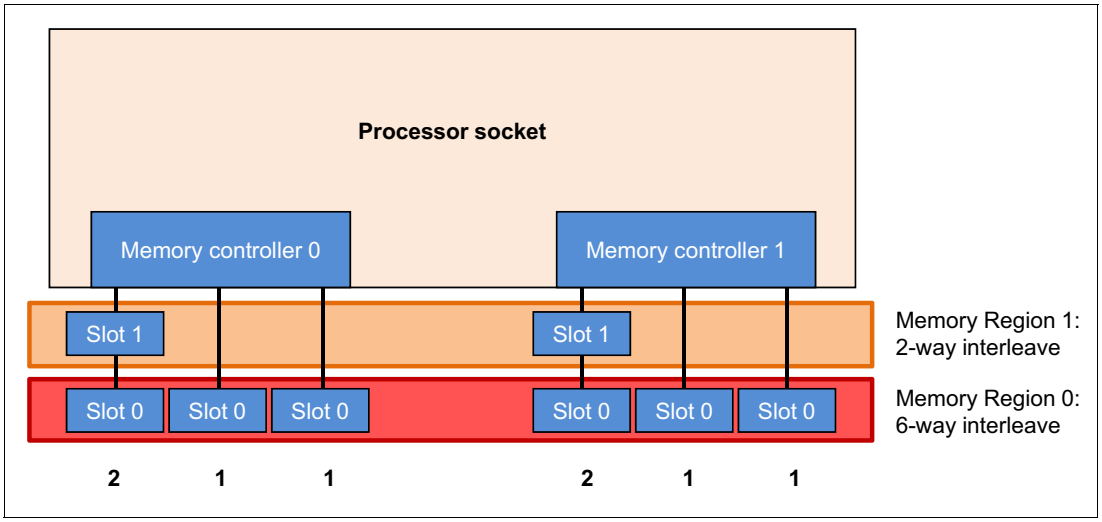

Figure 3 2-1-1 memory topology forming multiple interleave sets

In addition, the performance of a specific memory access depends on which memory region is being accessed and how many memory DIMMs comprise the interleave set. Contiguous memory accesses to a memory region with fewer DIMMs in the interleave set will have lower performance compared to accesses to a memory region with more DIMMs in the interleave set.

12 DIMM and 16 DIMM performance comparison

Figure 4 compares the memory bandwidth and latency between the SD530 configured with 12 DIMMs (1-1-1) and with 16 DIMMs (2-1-1).

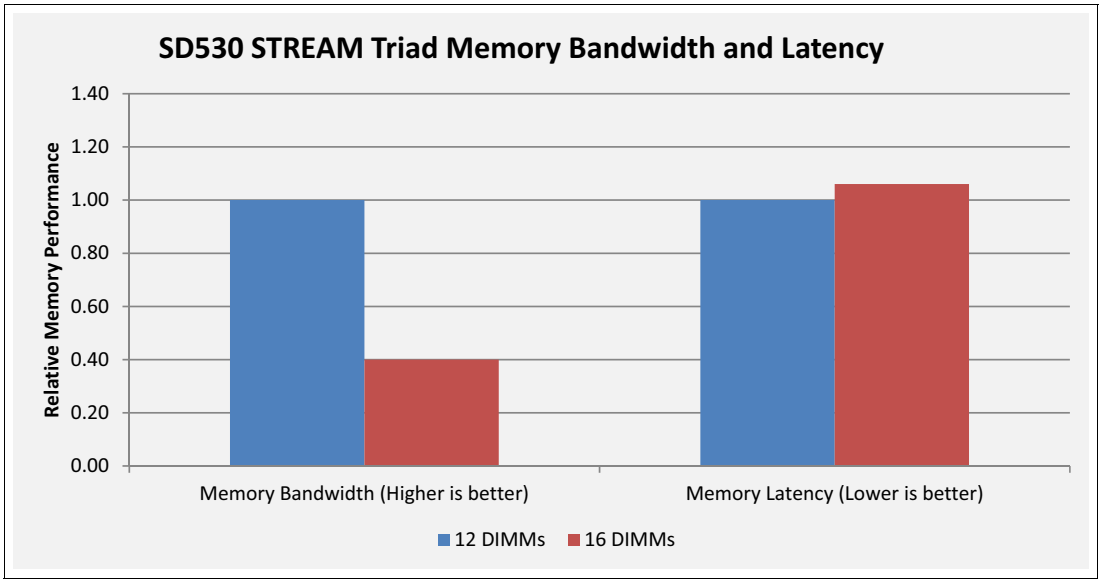

Figure 4 SD530 STREAM Triad Memory Bandwidth and Latency

Using the standard STREAM Triad workload which measures memory bandwidth, the 16-DIMM configuration had only 40% of the peak memory bandwidth compared to the 12-DIMM configuration. Memory latency was separately measured to be 6% higher with 16 DIMMs. The lower memory bandwidth and higher memory latency are both due to the unbalanced 2-1-1 topology used in the 16-DIMM configuration, the resulting creation of multiple memory regions and the negative effect that has on memory interleaving.

The effect that this reduced memory bandwidth from using 16 DIMMs has on application performance is based on two primary factors:

- The application's memory bandwidth demand – Applications that require moderate to high memory bandwidth have a higher likelihood of experiencing degraded performance.
- The processor's memory bandwidth demand – Lower performing processors will have lower memory bandwidth requirements. Under some circumstances when using these processors, the relative system performance degradation between 16 DIMMs and 12 DIMMs can be lower.

Many enterprise applications with moderate memory bandwidth requirements such as databases, enterprise resource management and virtualized hosts can experience degraded performance with a 16 DIMM memory topology when compared to a 12 DIMM memory topology. Applications with high memory bandwidth requirements like some high performance computing applications can experience severely degraded performance with a 16 DIMM memory topology when compared to a 12 DIMM memory topology.

Summary

The SD530 supports memory configuration options for maximum performance and maximum memory capacity:

- A memory configuration with 12 DIMMs is optimized for performance.
- A memory configuration with 16 DIMMs is optimized for memory capacity with reduced performance.

Applications that require peak performance and whose memory requirements can be met with 12 DIMMs should be deployed on the SD530 configured with 12 DIMMs. Only populate 16 DIMMs on the SD530 when the application's memory capacity cannot be met with 12 DIMMs and when peak performance is not required.

Authors

This paper was produced by the following team of specialists:

Joe Jakubowski is a Senior Engineering Staff Member in the Lenovo Data Center Group Performance Laboratory in Morrisville, NC. Previously, he spent 30 years at IBM. He started his career in the IBM Networking Hardware Division test organization and worked on various Token Ring adapters, switches and test tool development projects. He has spent the past 23 years in the server performance organization focusing primarily on database, virtualization and new technology performance. His current role includes all aspects of x86 server architecture and performance. Joe holds Bachelor of Science degrees in Electrical Engineering and Engineering Operations from North Carolina State University and a Master of Science degree in Telecommunications from Pace University.

Dan Colglazier is a Senior Engineer in the Lenovo Data Center Group Performance Laboratory in Morrisville, NC. He had previously spent over 30 years at IBM, where he started out in the IBM Data Processing Division performance team helping improve the designs of future mainframe storage hierarchies. Assessing the performance of future System x® servers in both IBM and Lenovo has been his main focus. He is currently the leader of the Performance Design Guidance team. Dan holds a Bachelor of Science degree in Computer Engineering and a Master of Science degree in Electrical Engineering both from the University of Illinois.

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

- Jamal Ayoubi
- Sumanta Bahali
- Larry Cook
- David Watts

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.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp0659

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®	System x®
Lenovo(logo)®	ThinkSystem™

The following terms are trademarks of other companies:

Intel, 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.