

Unbalanced Memory Performance

Joseph Stanfield

Sept. 26 2012

It is well understood that a server with a “balanced” memory configuration yields the best performance for your servers (See [Memory Selection Guidelines for HPC and 11G PowerEdge Servers](#)). Balanced implies that all memory channels of the server are populated equally and with identical memory modules (DIMMS). But there are certain situations where an unbalanced configuration might be needed. Cost limitations, capacity requirements, and application needs are all possible factors. This blog will provide a brief overview of how to gain the best performance from an unbalanced memory configuration.

To better understand the demerits of unbalanced configurations and to determine which unbalanced configuration is the best, several tests were conducted in our lab. We have seen many requests for servers configured with 48GB and 96GB of memory. Satisfying these capacity requirements on the latest generation of servers that have four memory channels per socket is only possible with unbalanced configurations. Using the available 2GB, 4GB, 8GB and 16GB DIMMs, we tested the configurations described below.

For the purpose of this study, a Dell PowerEdge M620 was used with the following configuration:



Dual CPU	Intel Xeon E5-2680 @ 2.70GHz
BIOS	1.1.2
CPLD	1.0.2
iDRAC Version	1.06.06
Node Interleaving	Disabled
Memory Mode	Optimized

Memory Used For Testing	
	2GB 1Rx8 @1600 MT/s
	4GB 2Rx8 @1600 MT/s
	8GB 2Rx4 @1600 MT/s
	16GB 2Rx4 @1600 MT/s

Figure 1: PowerEdge M620 configuration and memory used for testing.

Two capacity tests (48GB and 96GB) were performed with eight different memory organization options using the STREAM memory bandwidth benchmark. Due to the similarities in memory channel population and benchmarking results, this blog will focus on the 96GB options. For a comparison of the capacities tested, see the *figure 6* at the end of the blog. All results report the total measured system memory bandwidth.

The first test utilized fully populated memory banks across all four channels (*see figure 2*). Each CPU in this case supports up to 3 DIMMs per channel but, a maximum capacity configuration reduces the speed at which the memory operates significantly, impacting the overall performance as evident by the result.

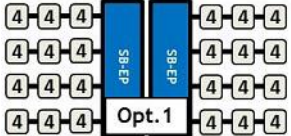
Option	CPU1	CPU2	Triad
	12x4GB	12x4GB	56GB/s

Figure 2

For the remaining seven unbalanced tests, three of the options were unbalanced across processors and four were unbalanced across memory channels.

Balanced Configuration Reference

Before we began the unbalanced testing, we needed a reference point and some actual STREAM results from a balanced configuration. Six valid configurations were benchmarked which gave us an average benchmark outcome to baseline against when testing the unbalanced options.

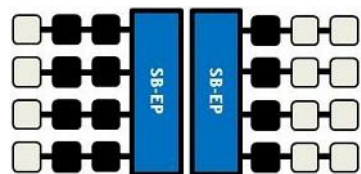
Figure 3 shows an example of a balanced configuration that has an equal amount of memory per channel for each CPU. All of the results are in the same ball park and the maximum performance was achieved when the system board was identically populated with 1600 MT/s DIMMs across four channels per CPU. From the results, it's clear that the DIMM capacity is not a factor for memory bandwidth. The balanced configurations consisted of 1 or 2 DIMMs per channel, with either 8 or 16 identical DIMMs in the server.

Confgs	CPU1	CPU2	Triad Results
Opt 1 32GB	8x 2GB	8x 2GB	74GB/s
Opt 2 32GB	4x 4GB	4x 4GB	77GB/s
Opt 3 64GB	8x 4GB	8x 4GB	78GB/s
Opt 4 64GB	4x 8GB	4x 8GB	77GB/s
Opt 5 128GB	8x 8GB	8x 8GB	71GB/s
Opt 6 128GB	4x 16GB	4x 16GB	76GB/s

Figure 3: An example of two DIMMs per channel populated in a balance configuration.

Unbalanced Across Processors

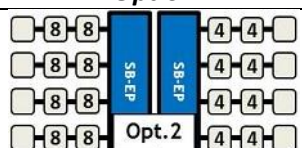
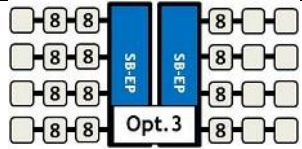
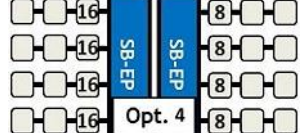
The next three options proposed balanced memory across channels but an unbalanced configuration between the processors. The example in figure 4 has all four memory channels assigned to CPU 1 populated with two DIMMs. This will operate with a larger capacity and will generally have lower latency. CPU 2 also has all four memory channels populated but, with one DIMM each. Depending on which CPU executes the process, there may be a performance reduction due to the higher latency caused by a remote memory request. This happens if the memory required is more than what is assigned to CPU. Interestingly, the performance results from



benchmarking were comparable with the balanced configuration test. This is likely due to the memory benchmark itself; the limits of the memory capacity per CPU are not exercised and the symmetric population across memory channels keeps the memory bandwidth high.

Figure 4: An example of an unbalanced configuration across CPUs configuration.

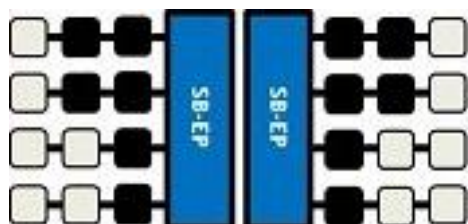
The table below shows the exact configurations used for each option and the corresponding STREAM Triad results.

Option	CPU1	CPU2	Triad
	8x8GB	8x4GB	75GB/s
	8x8GB	4x8GB	77GB/s
	4x16GB	4x8GB	76GB/s

Unbalanced Across Channels

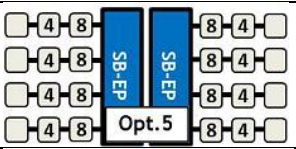

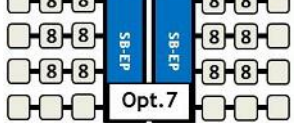
The unbalanced channel configurations that were tested had partially populated memory channels resulting in bandwidth bottlenecks. *Figure 5* shows an example of unequally populated memory channels.


With the exception of option 5, the tests executed with unbalanced channel configurations completed with



significantly lower results than testing executed against unbalanced memory per CPU structure. The actual configurations that were tested are illustrated in the table below.

Figure 5: An example of an unbalanced across memory channels configuration.

Option	CPU1	CPU2	Triad
	4x4GB 4x8GB	4x4GB 4x8GB	75GB/s
	6x8GB	6x8GB	44GB/s
	6x8GB	6x8GB	63GB/s



As mentioned earlier, 48GB capacity configurations were also tested with the eight organization options populated similarly to the 96GB tests. The DIMM channels were populated similarly between both sets of tests with smaller capacity DIMMs used for the 48GB tests. The results followed similar trends, with the exception of Option 5 on the 48GB test. The 48GB benchmark mixed 2GB and 4GB DIMM capacities with mixed ranks resulting in a lower performance. Details of the DIMMs used can be reviewed in *figure 1*.

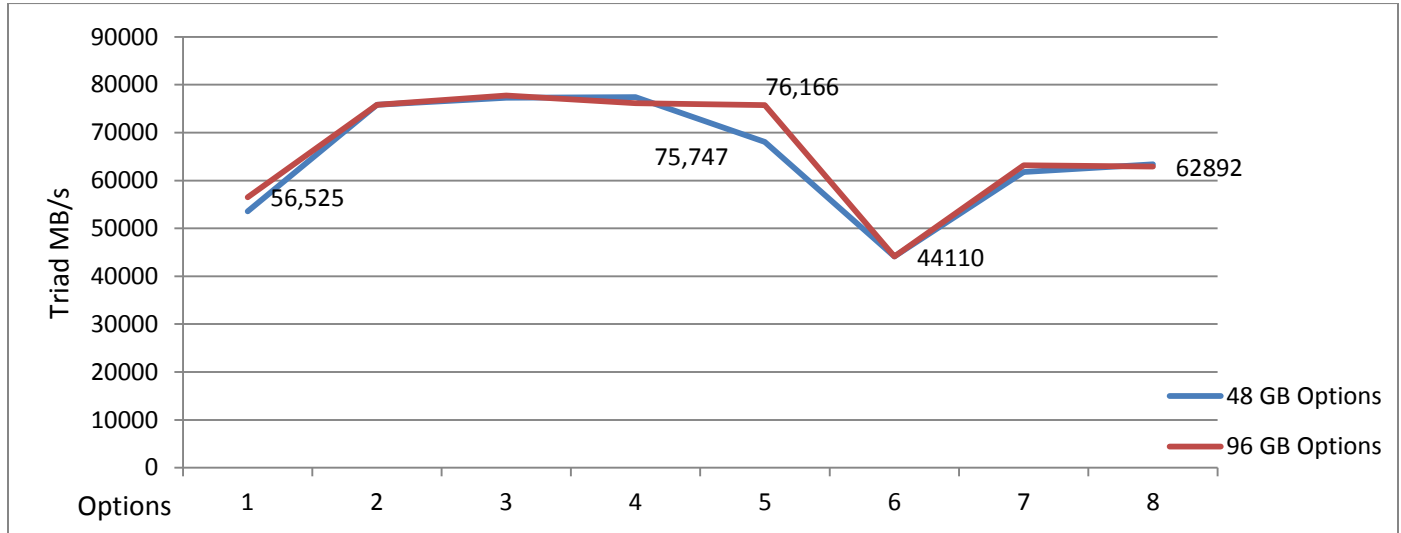


Figure 6: Unbalanced 48GB and 96GB STREAM Triad Comparison Results.

Our conclusion resulting from this exercise is that it is possible to achieve desirable results with an unbalanced memory configuration, provided that the memory assigned to the CPU is identical and does not exceed more than two DIMMs per channel.

References

1. Memory Performance Guidelines for Dell PowerEdge 12th Generation Servers.

<http://en.community.dell.com/techcenter/b/techcenter/archive/2012/07/26/memory-performance-guidelines-for-dell-powerededge-12th-generation-servers.aspx>

2. Nehalem and Memory Configurations.

<http://en.community.dell.com/techcenter/b/techcenter/archive/2009/04/08/nehalem-and-memory-configurations.aspx>

3. Memory Selection Guidelines for HPC and 11G PowerEdge Servers

<http://content.dell.com/us/en/enterprise/d/business~solutions~whitepapers~en/Documents~11g-memory-selection-guidelines.pdf.aspx>