

Force10 Networks and NetEffect Performance Report



The test results in this document show MPI latency and throughput performance between Force10 Networks' S-Series S2410 low latency data center switch and NetEffect's NE010 10 Gigabit Ethernet PCI-X adapters.

Table of Contents


- Introduction 2
- Test Diagram..... 2
- Test 1 - OSU “mpiexec” Performance Test 3
- Test 2 - NetEffect “vperf” Message Latency Test 4
- Test 3 - NetEffect “neseval” Performance Test 5
- Appendix 6
 - Hardware Configurations..... 6
 - Test 1 Output 6
 - Test 2 Output 8
 - Test 3 Output 8

Introduction

Performance testing was successfully measured between Force10 Networks’ S-Series S2410 low latency data center switch and NetEffect’s NE010 10 Gigabit Ethernet PCI-X adapters. The testing measured end-to-end MPI latency and throughput results over 10 Gigabit Ethernet CX4 links between NE010 adapters connected by an S2410 switch.

The results show end-to-end MPI latency as low as 9 µsec and switching latency as low as 300ns for 64 byte messages, and throughput of up to 910 MB/s for 65536 byte messages.


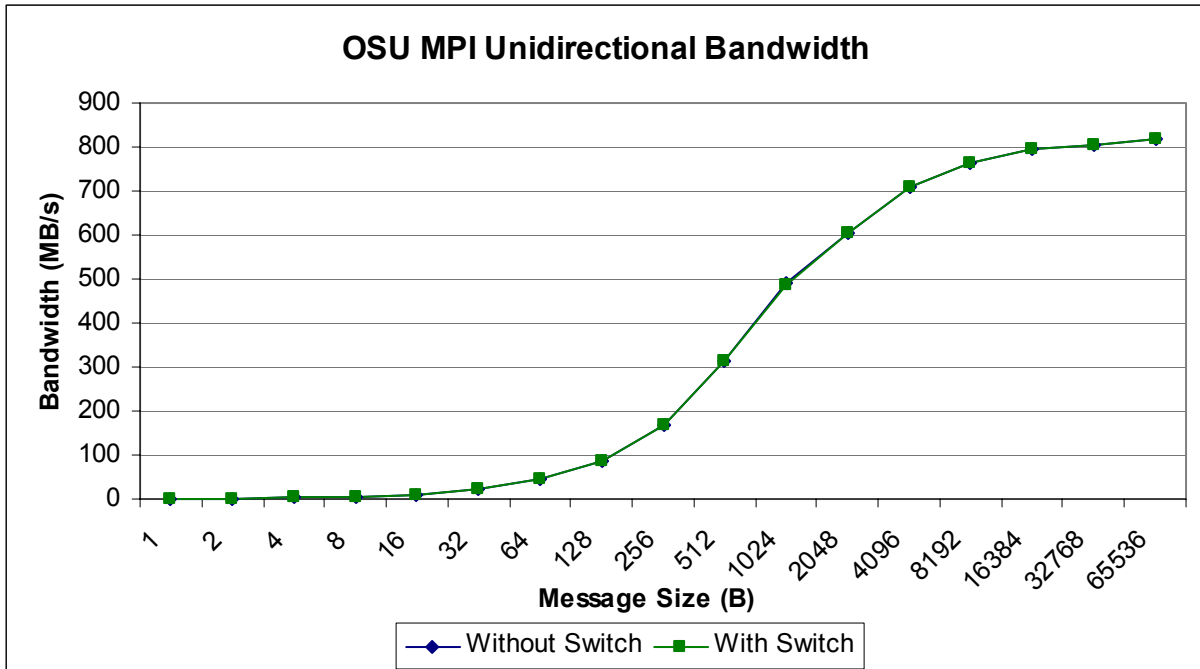
Test Diagram

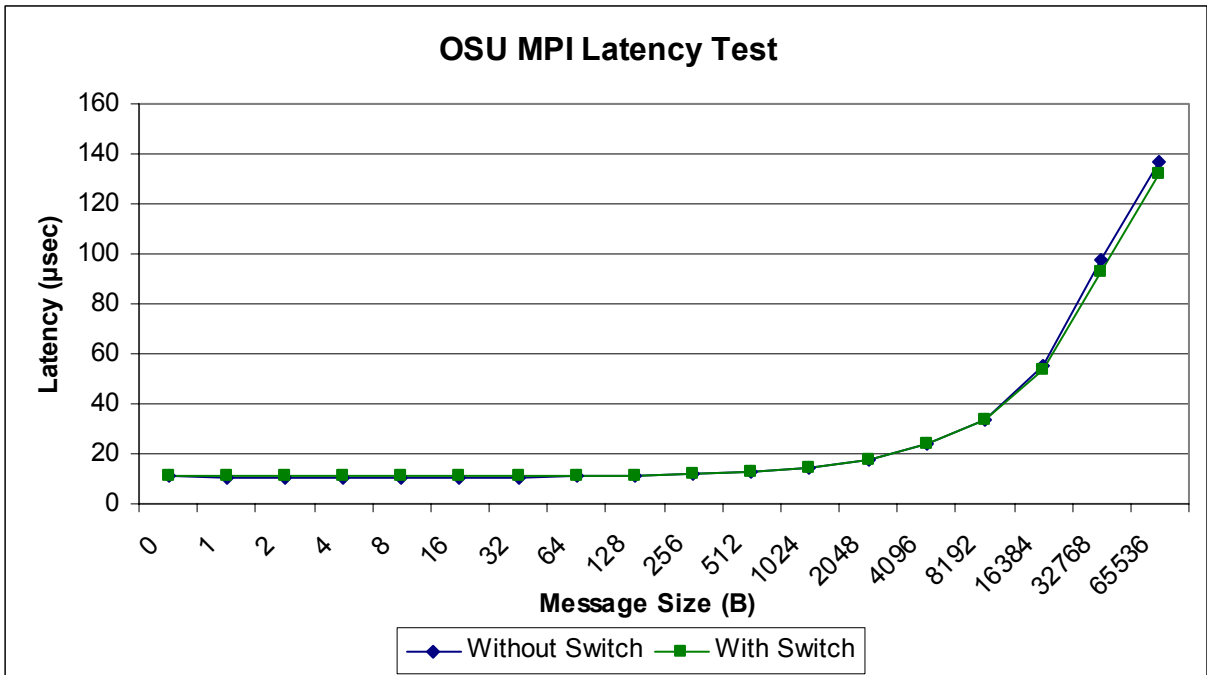


Test 1 - OSU “mpiexec” Performance Test

Objective: To measure MPI unidirectional bandwidth, bidirectional bandwidth and latency between NE010 NICs using the OSU “mpiexec” tool. The test was first run with the NICs connected to each other to establish baseline results. Then the test was run again between NE010 NICs connected to the S2410 switch.

Results: The test results are summarized in a separate graph for unidirectional bandwidth, bidirectional bandwidth and latency. The throughput was almost identical in both tests, and the latency difference when the switch was connected was on the order of hundreds of ns.

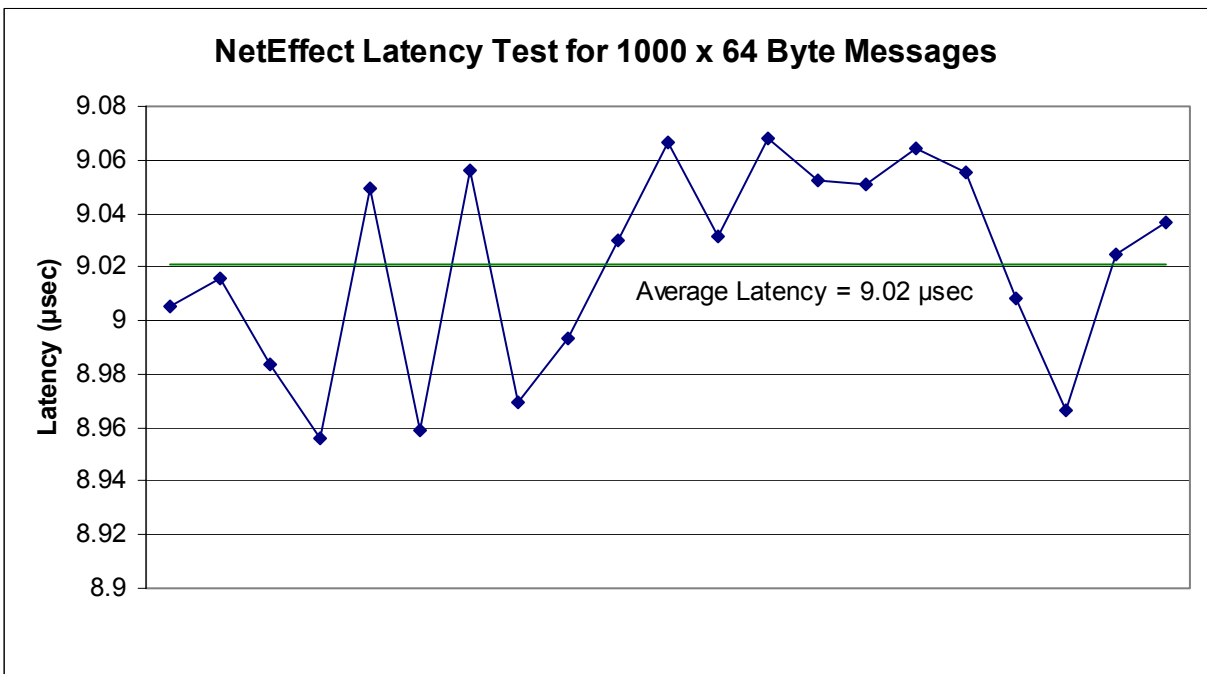




Test 2 - NetEffect “vperf” Message Latency Test

Objective: To measure MPI latency between NE010 NICs connected to the S2410 switch using NetEffect’s “vperf” tool.

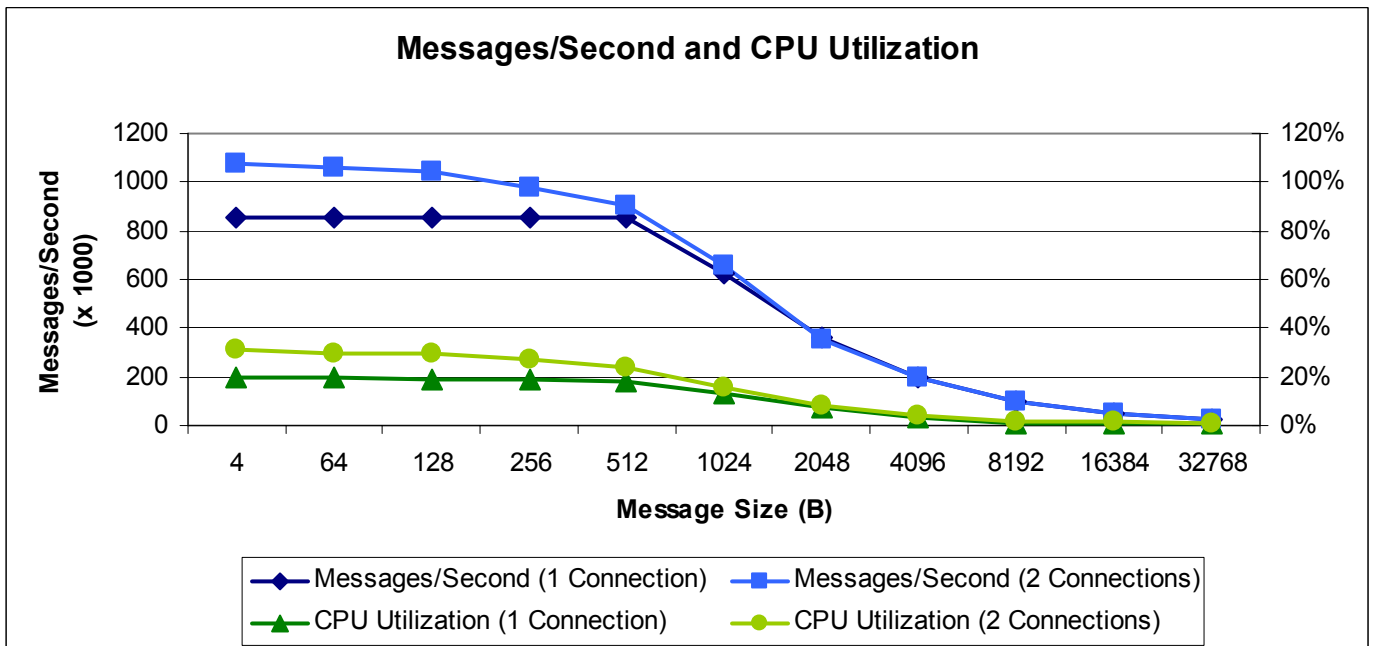

Results: The results demonstrate a sustained average MPI latency of 9.02 µsec for 1000 x 64 byte messages.



Test 3 - NetEffect “neseval” Performance Test

Objective: To measure CPU utilization and bandwidth with a variety of message sizes using the “neseval” tool.

Results: The results demonstrate lower CPU utilization and higher bandwidth as message sizes increase.



Appendix

Hardware Configurations

Force10 Networks	NetEffect
Chassis: S2410	Adapters: NE010
Interfaces: CX4	Interfaces: CX4
	Servers: Dual AMD 2GHz Opteron CPU Super Micro H8DC8 Motherboard 2GB RAM 80GB IDE Disk Drive FC4 32-bit Linux OS

Test 1 Output

Tests were performed using OSU mpiexec Version 2.1 (<http://www.osc.edu/~pw/mpiexec/index.php>).

Results Without S2410 Switch	
<pre>[mpi@mil osu]\$ mpiexec -n 2 ./bw # OSU MPI Bandwidth Test (Version 2.1) # Size Bandwidth (MB/s) 1 0.695970 2 1.392806 4 2.761327 8 5.551923 16 11.026024 32 22.143002 64 44.123354 128 86.955860 256 169.891658 512 311.511681 1024 489.324267 2048 606.620411 4096 707.445392 8192 762.289846 16384 794.401769 32768 802.782972 65536 817.978697</pre>	<pre>[mpi@mil osu]\$ mpiexec -n 2 ./bibw # OSU MPI Bidirectional Bandwidth Test (Version 2.1) # Size Bi-Bandwidth (MB/s) 1 1.136283 2 2.236589 4 4.193649 8 8.959608 16 17.746699 32 35.285633 64 70.506101 128 136.396881 256 253.760000 512 451.500315 1024 666.596922 2048 773.742727 4096 836.435830 8192 877.102543 16384 897.468934 32768 901.292138 65536 910.290649</pre>
<pre>[mpi@mil osu]\$ mpiexec -n 2 ./lat # OSU MPI Latency Test (Version 2.1) # Size Latency (us) 0 11.12 1 10.61 2 10.62 4 10.62 8 10.65 16 10.65 32 10.80 64 10.88 128 11.24 256 11.70 512 12.63 1024 14.52 2048 17.46 4096 23.91 8192 33.90 16384 55.31 32768 97.25 65536 136.79</pre>	

Force10 Networks and NetEffect Performance Report

Results With S2410 Switch

```
[mpi@mil osu]$ mpiexec -n 2 ./bw
# OSU MPI Bandwidth Test (Version 2.1)
# Size      Bandwidth (MB/s)
1           0.690918
2           1.381089
4           2.745581
8           5.500727
16          11.007374
32          21.988825
64          43.798264
128         86.786740
256         169.047445
512         311.448149
1024        488.498146
2048        606.734225
4096        706.990501
8192        761.672008
16384       793.298433
32768       803.307171
65536       817.691647
```

```
[mpi@mil osu]$ mpiexec -n 2 ./bibw
# OSU MPI Bidirectional Bandwidth Test (Version 2.1)
# Size      Bi-Bandwidth (MB/s)
1           1.128757
2           2.263320
4           4.507827
8           8.892640
16          18.002965
32          36.098229
64          72.150977
128         141.569965
256         280.986555
512         496.824167
1024        678.853357
2048        771.035208
4096        836.919702
8192        877.263500
16384       897.002174
32768       901.670936
65536       910.483808
```

```
[mpi@mil osu]$ mpiexec -n 2 ./lat
# OSU MPI Latency Test (Version 2.1)
# Size      Latency (us)
0           11.43
1           10.87
2           10.87
4           10.86
8           10.92
16          10.93
32          11.07
64          11.17
128         11.49
256         11.91
512         12.84
1024        14.70
2048        17.52
4096        23.64
8192        33.38
16384       53.65
32768       92.63
65536       132.31
```

Note: At larger message sizes the latency decreases slightly when NICs are connected to the switch. This is because the buffer in the switch reduces the number of link pauses.

Test 2 Output

```

TOTAL MSG CNT=0x1000 time=0.073774 latency=9.00562
TOTAL MSG CNT=0x1000 time=0.073856 latency=9.01562
TOTAL MSG CNT=0x1000 time=0.073594 latency=8.98364
TOTAL MSG CNT=0x1000 time=0.07337 latency=8.9563
TOTAL MSG CNT=0x1000 time=0.074135 latency=9.04968
TOTAL MSG CNT=0x1000 time=0.073391 latency=8.95886
TOTAL MSG CNT=0x1000 time=0.074185 latency=9.05579
TOTAL MSG CNT=0x1000 time=0.073477 latency=8.96936
TOTAL MSG CNT=0x1000 time=0.073676 latency=8.99365
TOTAL MSG CNT=0x1000 time=0.073974 latency=9.03003
TOTAL MSG CNT=0x1000 time=0.074274 latency=9.06665
TOTAL MSG CNT=0x1000 time=0.073986 latency=9.03149
TOTAL MSG CNT=0x1000 time=0.074284 latency=9.06787
TOTAL MSG CNT=0x1000 time=0.074159 latency=9.05261
TOTAL MSG CNT=0x1000 time=0.074145 latency=9.0509
TOTAL MSG CNT=0x1000 time=0.074253 latency=9.06409
TOTAL MSG CNT=0x1000 time=0.074182 latency=9.05542
TOTAL MSG CNT=0x1000 time=0.073799 latency=9.00867
TOTAL MSG CNT=0x1000 time=0.073456 latency=8.9668
TOTAL MSG CNT=0x1000 time=0.07393 latency=9.02466
TOTAL MSG CNT=0x1000 time=0.074029 latency=9.03674
    
```

Test 3 Output

Message Size (B)	1 Connection			2 Connections		
	Bandwidth (MB/s)	Messages/Second x 1000	CPU Utilization (%)	Bandwidth (MB/s)	Messages/Second x 1000	CPU Utilization (%)
4	80	855	20	101	1074	31
64	132	855	20	162	1060	30
128	187	857	19	227	1042	30
256	297	857	19	340	980	27
512	515	855	18	542	904	24
1024	700	628	13	732	658	16
2048	766	358	7	760	354	8
4096	825	197	3	814	194	4
8192	835	100	1	832	100	2
16384	839	50	1	840	51	2
32768	851	26	1	850	26	1



Force10 Networks, Inc.
 1440 McCarthy Boulevard
 Milpitas, CA 95035
www.force10networks.com

Phone: 408-571-3500
 Fax: 408-571-3550
 Email: info@force10networks.com



NetEffect, Inc.
 9211 Waterford Centre Boulevard
 Suite 100
 Austin, TX 78758
www.neteffect.com

Phone: 512-302-0002
 Fax: 512-493-3399
 Email: sales@neteffect.com