WildFire Storage Server
Network Performance Report

September 6, 2017

A series of simple tests were run on the production release of WildFire Storage Server over a small sample of network configurations likely of interest for WSS users.

Server Configuration:

Open-E: DSSv7 build 21661
WildFire Storage: Enterprise Storage Stack release 5.0.0
WSS: Text utility version 2.0.0
WSS: Scripts dated 9/1/2017

SuperMicro 24-bay 2U server system
X10-SRL-F motherboard, E5-1650 v3 CPU, 64 GB DDR4 ECC RAM
24 x Samsung 850 Pro 256 GB SSDs, 3 x SuperMicro 3008 SAS HBAs

ESS array: RAID-5 with default configuration

Client Configuration

Centos 7.0 – 64-bit
SuperMicro X9-SRL-F, E5-1650 CPU, 32 GB DDR3 ECC RAM

Network Configuration:

Single Card, Single Link
- Single Mellanox ConnectX-3 PCI-e v3 x8 card – CX314A
- Link speed set to 10GigE, 40GigE, 56 GigE
  - 56 GigE only tested with a single link due to limited 56 GigE cables

Single Card, Dual Links
- Single Mellanox ConnectX-3 PCI-e v3 x8 card – CX314A
- Link speeds set to 10GigE, 40GigE

Dual Card, Single Link
- Two Mellanox ConnectX-3 PCI-e v3 x8 cards – CX314A
- Link speeds set to 40GigE

Each network link was configured on a separate IPv4 subnet (10.0.1.0, 10.0.2.0, 10.0.3.0, 10.0.4.0).
All cards were configured with MTU=9000 (Jumbo Frames)
The network connections were all copper “point to point” connections. No ethernet switch was used.

Storage Configuration

The host WSS system was configured with 8 x 100 GB iSCSI targets.
With multiple cards, and/or multiple links, the iSCSI mount were spread across the available links.

Benchmark Configuration and Tools

‘fio’ for Linux was used for all tests.
- direct IO, and async IO were selected

block size set to 4K, 16k, 64k, 256k, and 1M
100% read, 50% read and 50% write, 100% write
note that 50/50 read/write can appear to run faster than the interface speed
queue depth varied from q=1 to q=32 depending on block size

Performance Charts

Single adapter, Single Link

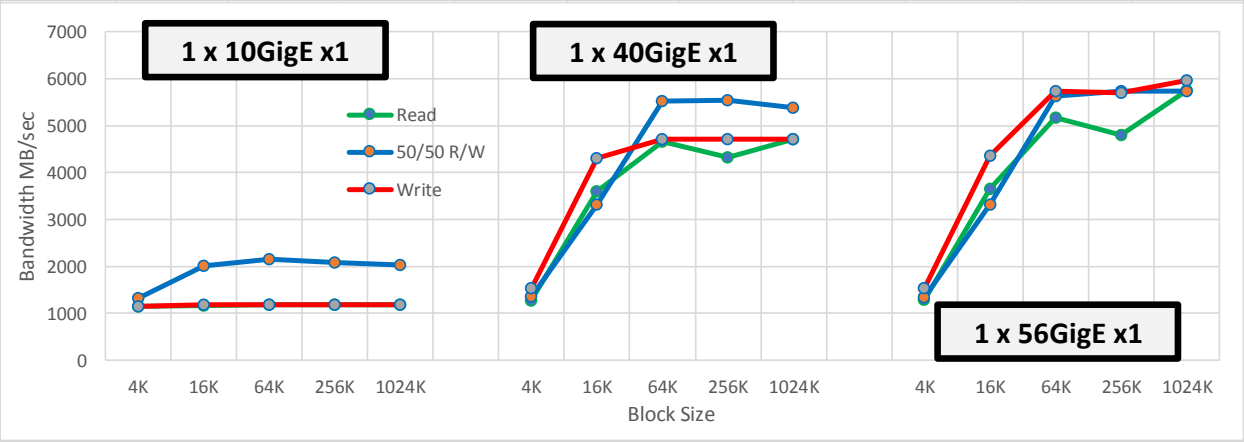

Multiple Links on Single Adapter, plus Multiple Adapters

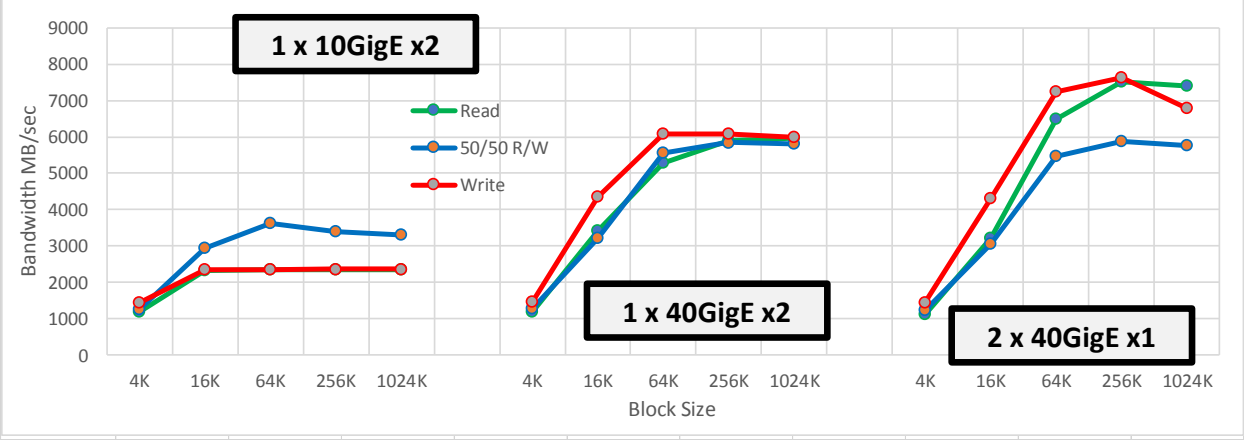

Performance Discussion

As is typical for ESS, writes are faster than reads, although network latency hides a bit of this. 4K random IO clocks about 300K IOPS and larger blocks start to quickly the saturate the network.

For slow links, a read/write mix can take advantage of the full-duplex nature of the network. At faster network speeds, this becomes more difficult to sustain. Regardless, ESS remains amazingly efficient at all read/write mix ratios.

Performance Scalability

Multiple cards in multiple slots on the server, plus multiple clients should scale, at least somewhat.

Other brands or models of network adapters may have better performance or lower overhead. WildFire Storage has not tested other high speed adapters.

In general, this client/server configuration just about runs out of steam at about 60 Gbit/sec. Configurations with faster more more CPUs on the server or more clients are likely to push this number higher.

This level of performance implies support for 2000 VDI seats for a single 10GigE adapter and in excess of 5000 VDI seats with multiple adapters. This is in-line with existing ESS installs.

Local Performance

This version of WSS can run local benchmarks to test the performance of the array without the overhead of the network iSCSI connection. This 24 SSD array benchmarks at over 1.5 Million 4K random writes and reaches 11 GB/sec for large block random writes. This level of performance is reached at very low queue depth (q=10). Read performance is actually limited by the SSDs to much lower values and requires much higher queue depths.

Read Tests:

Block Size	1 thread IOPS	1 thread BW	10 threads IOPS	10 threads BW	40 threads IOPS	40 threads BW
4K	10532	41.1M	111156	434.2M	416294	1626.1M
8K	9099	71.0M	101363	791.9M	352001	2750.0M
16K	8149	127.3M	87599	1368.7M	274527	4289.4M
32K	7157	223.6M	69268	2164.6M	190435	5951.0M
64K	6827	426.6M	57650	3603.1M	128142	8008.8M
128K	6518	814.7M	44862	5607.7M	80097	9.7G
256K	5519	1379.7M	31254	7813.6M	47488	11.5G
512K	4763	2381.6M	20088	9.8G	26244	12.8G
1M	3138	3138.7M	12031	11.7G	13947	13.6G

Write Tests:

Block Size	1 thread IOPS	1 thread BW	10 threads IOPS	10 threads BW	40 threads IOPS	40 threads BW
4K	261819	1022.7M	1653862	6460.3M	1556555	6080.2M
8K	205346	1604.2M	1225132	9.3G	1261607	9.6G
16K	169605	2650.0M	734893	11.2G	732940	11.1G
32K	104718	3272.4M	369849	11.2G	366963	11.1G
64K	75210	4700.6M	184155	11.2G	184545	11.2G
128K	40843	5105.3M	92206	11.2G	92309	11.2G
256K	19225	4806.2M	46088	11.2G	46038	11.2G
512K	9742	4871.0M	23024	11.2G	22263	10.8G
1M	4828	4828.0M	9601	9.3G	8859	8.6G