



SPECsfs2008 RESULTS

Results for 1 shelf, pages 2-7

Results for 2 shelves, pages 8-13


SPECsfs2008_nfs.v3 Result

Panasas, Inc. : ActiveStor 14T

SPECsfs2008_nfs.v3 = 20745 Ops/Sec (Overall Response Time = 1.99 msec)

Performance

Throughput (ops/sec)	Response (msec)
2096	0.6
4205	0.7
6303	0.8
8418	0.9
10543	1.8
12693	3.1
14755	1.2
16975	4.8
19084	2.8
20745	7.6



Product and Test Information

Tested By	Panasas, Inc.
Product Name	ActiveStor 14T
Hardware Available	November 2012
Software Available	November 2012
Date Tested	September 2012
SFS License Number	3866
Licensee Locations	California, USA

The Panasas ActiveStor 14T scale-out NAS Storage Appliance features a parallel clustered file system that turns files into smart data objects and then dynamically distributes and load balances data transfer operations across a networked blade architecture. Panasas PanFS distributed file system creates a cluster with a single global namespace. Self-contained shelf chassis include operating and file system, network connectivity, redundant and hot swappable

meta data director and storage blades servers, power supplies and battery backup. Per shelf configuration includes 40 TB storage, 168 GB of system memory, high availability active meta data servers, 10 GigE network connectivity with network failover option, and snapshot license.

Configuration Bill of Materials

Item No	Qty	Type	Vendor	Model/Name	Description
1	2	Panasas ActiveStor 14T Storage Shelf	Panasas, Inc.	ActiveStor 14T	ActiveStor 14T, 40TB, 2+9 (DB+SB), 24+144GB Cache, Switch Module x2 (10GbE x2 CX4/SFP+, 1GbE x8)

Server Software

OS Name and Version	PanFS 5.0
Other Software	none
Filesystem Software	PanFS 5.0

Server Tuning

Name	Value	Description
iopath.iscsi_rcv_buf_size	131072	Receive Buffer Size
intr_queue_maxlen	1024	Maximum size of the IP input queue

Server Tuning Notes

ActiveStor 14T Server tuning only impacts communication between the storage and directorblades. It does not impact clients (load generators) in any way.

Disks and Filesystems

Description	Number of Disks	Usable Size
Each director blade has one 250 GB SATA 7200 RPM HDD that holds the PanFS OS, but not file system metadata or user data.	2	500.0 GB
Each storage blade has two 2000 GB SATA 7200 RPM HDD that contains the PanFS OS and user data.	18	35.2 TB
Each storage blade has one 480 GB SSD that contains PanFS OS, file system metadata and user data	9	4.2 TB

Total	29	39.9 TB
Number of Filesystems	1	
Total Exported Capacity	33320	
Filesystem Type	NFSv3	
Filesystem Creation Options	Default	
Filesystem Config	The AS14T shelf (comprising of 2 director blades and 9 storage blades) in the SUT was configured in one bladeset (containing 2 volumes). A bladeset is a fault-tolerant shared pool of disks. Each volume stripes data uniformly across all disks in the pool using Object RAID 1/5, which dynamically selects RAID -1 (mirroring) or RAID-5 (XOR parity) on a per-file basis. The bladeset was configured with vertical parity enabled. Vertical parity corrects media errors at the storage blade level, before they are exposed to Object RAID, and provides protection equivalent to RAID-6 against a sector error during a RAID rebuild.	
Fileset Size	2456.4 GB	

Network Configuration

Item No	Network Type	Number of Ports Used	Notes
1	10GigE	8	None.

Network Configuration Notes

Each Panasas shelf used two 10 Gigabit Ethernet interfaces configured with MTU=9000 (jumbo frames). These links were connected to the Dell Force10 S4810 switch. Additionally, all (2) director blades were connected with 10 GigE uplinks to the Force10 S4810 switch.

Benchmark Network

Eight Supermicro X8DTT-HIBQF servers were used as load generators and connected via 10 GbE NICs to the Dell Force10 S4810 switch. The AS14T shelf was also connected to the same Dell Force10 S4810. There were no additional settings on the clients or the shelves.

Processing Elements

Item No	Qty	Type	Description	Processing Function

1	2	CPU	Intel(R) Xeon(R) 2.13GHz	NFS Gateway, Metadata Server, Cluster Management, iSCSI, TCP/IP
2	9	CPU	Intel(R) Xeon(R) 1.73GHz	Object Storage File System, iSCSI, TCP/IP

Processing Element Notes

The SUT includes 2 director blades, each with a single Intel Xeon 2.13 GHz CPU, and 9 storage blades, each with a single Intel Xeon 1.73 GHz CPU. The director blades manage file system metadata and provide NFS gateway services. The storage blades store user data and metadata, and provide access to it through the OSD (Object Storage Device) protocol over iSCSI and TCP/IP. One director blade in the SUT was also running the Panasas realm manager and management user interface.

Memory

Description	Size in GB	Number of Instances	Total GB	Nonvolatile
Each director blade has 12GB ECC RAM	12	2	24	NV
Each storage blade has 16GB ECC RAM	16	9	144	NV
Grand Total Memory Gigabytes			168	

Memory Notes

Stable Storage

AS14T storage shelf has two (redundant) power supplies as well as a battery, which powers the entire shelf for about five minutes in the event of AC power loss. In the event of a power failure, each blade saves cached writes from main memory to its local HDD before shutting down. The data saved in this way is maintained indefinitely while the system is powered down, and is automatically recovered when power is restored.

System Under Test Configuration Notes

The SUT is comprised of 1 Panasas AS14T shelf, connected via 10 GbE SFP+ to a Dell Force10 S4810 10GbE switch. Each of the 8 load generators were connected to the Dell Force10 S4810 switch with 10 GbE SFP+.

Other System Notes

Test Environment Bill of Materials

Item No	Qty	Vendor	Model/Name	Description
1	8	Supermicro	X8DTT-HIBQF	X8DTT-HIBQF, 827H-R1400B
2	1	Dell Force10	S4810	Dell Force10 S4810 High-Performance 10/40 GbE Switch

Load Generators

LG Type Name	1
BOM Item #	1
Processor Name	Intel(R) Xeon(R) X5650
Processor Speed	2.67 GHz
Number of Processors (chips)	2
Number of Cores/Chip	6
Memory Size	24 GB
Operating System	Linux 2.6.18-128.el5
Network Type	1 Intel Corporation 82599EB 10-Gigabit SFI/SFP+ NIC

Load Generator (LG) Configuration

Benchmark Parameters

Network Attached Storage Type	NFS V3
Number of Load Generators	8
Number of Processes per LG	16
Biod Max Read Setting	2
Biod Max Write Setting	2
Block Size	AUTO

Testbed Configuration

LG No	LG Type	Network	Target Filesystems	Notes
1..8	Supermicro	N1	F1 (GW1:/vol0 GW2:/vol0 GW1:/vol1 GW2:/vol1 GW1:/vol0 GW2:/vol0 GW1:/vol1 GW2:/vol1	

			GW1:/vol0 GW2:/vol0 GW1:/vol1 GW2:/vol1	
			GW1:/vol0 GW2:/vol0 GW1:/vol1 GW2:/vol1	

Load Generator Configuration Notes

Each of the 8 clients accessed all (2) gateways to all (2) volumes in the file system. In this way, all the volumes were accessed through all gateways from these 8 clients.

Uniform Access Rule Compliance

Each Panasas director blade provides both NFS gateway services for the entire file system, and metadata management services for a portion of the file system (one or more virtual volumes). The system as tested contained 2 director blades and was configured with 2 volumes, one managed by each director blade. To comply with the Uniform Access Rule (UAR) for a single namespace, all the volumes were accessed through all gateways from the 8 load-generating clients and no client was afforded any advantage by colocating the NFS gateway and metadata server for the volume being accessed.

Other Notes

None

Config Diagrams

- [AS14T 1-shelf Setup](#)


SPECsfs2008_nfs.v3 Result

Panasas, Inc. : ActiveStor 14T

SPECsfs2008_nfs.v3 = 41116 Ops/Sec (Overall Response Time = 1.39 msec)

Performance

Throughput (ops/sec)	Response (msec)
4098	1.1
8217	0.7
12350	0.7
16569	1.8
20646	1.0
24738	1.2
28910	1.6
33202	1.6
37823	2.5
41116	4.0



Product and Test Information

Tested By	Panasas, Inc.
Product Name	ActiveStor 14T
Hardware Available	November 2012
Software Available	November 2012
Date Tested	September 2012
SFS License Number	3866
Licensee Locations	California, USA

The Panasas ActiveStor 14T scale-out NAS storage appliance features a parallel, clustered file system that turns files into smart data objects and then dynamically distributes and load balances data transfer operations across a networked blade architecture. Panasas PanFS distributed file system creates a cluster with a single file system and single global namespace. Self-contained shelf chassis include operating and file system, network connectivity,

redundant and hot swappable meta data director and storage blades servers, power supplies and battery backup. Per shelf configuration includes 40 TB storage, 168 GB of system memory, high availability active meta data servers, 10 GigE network connectivity with network failover option, and snapshot license.

Configuration Bill of Materials

Item No	Qty	Type	Vendor	Model/Name	Description
1	2	Panasas ActiveStor Storage Shelf	Panasas, Inc.	ActiveStor 14T	ActiveStor 14T, 40TB, 2+9 (DB+SB), 24+144GB Cache, Switch Module x2 (10GbE x2 CX4/SFP+, 1GbE x8)

Server Software

OS Name and Version	PanFS 5.0
Other Software	none
Filesystem Software	PanFS 5.0

Server Tuning

Name	Value	Description
ioath.iscsi_rcv_buf_size	131072	Receive Buffer Size
intr_queue_maxlen	1024	Maximum size of the IP input queue

Server Tuning Notes

ActiveStor 14T Server tuning only impacts communication between the storage and directorblades. It does not impact clients (load generators) in any way.

Disks and Filesystems

Description	Number of Disks	Usable Size
Each director blade has one 250 GB SATA 7200 RPM HDD that holds the PanFS OS, but not file system metadata or user data.	4	1000.0 GB
Each storage blade has two 2000 GB SATA 7200 RPM HDD that contains the PanFS OS and user data.	36	70.3 TB
Each storage blade has one 480 GB SSD that contains PanFS OS, file system metadata and user data	18	480.0 GB

Total	58	71.8 TB
Number of Filesystems	1	
Total Exported Capacity	66400	
Filesystem Type	NFSv3	
Filesystem Creation Options	Default	
Filesystem Config	The 2 ActiveStor 14T shelves (18 storage blades) in the SUT were configured in two bladesets (2 volumes per bladeset). A bladeset is a fault-tolerant shared pool of disks. Each volume stripes data uniformly across all disks in the pool using Object RAID 1/5, which dynamically selects RAID-1 (mirroring) or RAID-5 (XOR parity) on a per-file basis. The bladesets were configured with vertical parity enabled. Vertical parity corrects media errors at the storage blade level, before they are exposed to Object RAID, and provides protection equivalent to RAID-6 against a sector error during a RAID rebuild.	
Fileset Size	4800.7 GB	

Network Configuration

Item No	Network Type	Number of Ports Used	Notes
1	10GigE	8	None.

Network Configuration Notes

Each Panasas shelf used four 10 Gigabit Ethernet interfaces configured with MTU=9000 (jumbo frames). These links were connected to the Dell Force10 S4810 switch.

Benchmark Network

Eight Supermicro X8DTT-HIBQF servers were used as load generators and connected via 10 GbE NICs to the Dell Force10 S4810 switch. The 2 AS14T shelves were also connected to the same Dell Force10 S4810. There were no additional settings on the clients or the shelves.

Processing Elements

Item No	Qty	Type	Description	Processing Function
1	4	CPU	Intel(R) Xeon(R) 2.13GHz	NFS Gateway, Metadata Server, Cluster Management, iSCSI, TCP/IP

2	18	CPU	Intel(R) Xeon(R) 1.73GHz	Object Storage File System, iSCSI, TCP/IP
---	----	-----	-----------------------------	---

Processing Element Notes

The SUT includes 4 director blades, each with a single Intel Xeon 2.13 GHz CPU, and 18 storage blades, each with a single Intel Xeon 1.73 GHz CPU. The director blades manage file system metadata and provide NFS gateway services. The storage blades store user data and metadata, and provide access to it through the OSD (Object Storage Device) protocol over iSCSI and TCP/IP. One director blade in the SUT was also running the Panasas realm manager and management user interface.

Memory

Description	Size in GB	Number of Instances	Total GB	Nonvolatile
Each director blade has 12GB ECC RAM	12	4	48	NV
Each storage blade has 16GB ECC RAM	16	18	288	NV
Grand Total Memory Gigabytes			336	

Memory Notes

Stable Storage

Each shelf has two (redundant) power supplies as well as a battery, which powers the entire shelf for about five minutes in the event of AC power loss. In the event of a power failure, each blade saves cached writes from main memory to its local HDD before shutting down. The data saved in this way is maintained indefinitely while the system is powered down, and is automatically recovered when power is restored.

System Under Test Configuration Notes

The SUT is comprised of 2 Panasas AS14T shelves, with each shelf connected via 10 GbE SFP+ to a Dell Force10 S4810 10GbE switch. Each of the 8 load generators were connected to the Dell Force10 S4810 switch with 10 GbE SFP+.

Other System Notes

Test Environment Bill of Materials

Qty	Vendor	Model/Name	Description
-----	--------	------------	-------------

Item No				
1	8	Supermicro	X8DTT-HIBQF	X8DTT-HIBQF, 827H-R1400B
2	1	Dell Force10	S4810	Dell Force10 S4810 High-Performance 10/40 GbE Switch

Load Generators

LG Type Name	1
BOM Item #	1
Processor Name	Intel(R) Xeon(R) X5650
Processor Speed	2.67 GHz
Number of Processors (chips)	2
Number of Cores/Chip	6
Memory Size	24 GB
Operating System	Linux 2.6.18-128.el5
Network Type	1 Intel Corporation 82599EB 10-Gigabit SFI/SFP+ NIC

Load Generator (LG) Configuration

Benchmark Parameters

Network Attached Storage Type	NFS V3
Number of Load Generators	8
Number of Processes per LG	16
Biod Max Read Setting	2
Biod Max Write Setting	2
Block Size	AUTO

Testbed Configuration

LG No	LG Type	Network	Target Filesystems	Notes
1..8	Supermicro	N1	F1 (GW1:/vol0 GW2:/vol0 GW4:/vol0 GW3:/vol0 GW1:/vol1 GW2:/vol1 GW4:/vol1 GW3:/vol1 GW1:/vol2 GW2:/vol2 GW4:/vol2 GW3:/vol2 GW1:/vol3 GW2:/vol3 GW4:/vol3 GW3:/vol3)	

Load Generator Configuration Notes

All (4) volumes were accessed through all (4) gateways from each of the 8 load generating clients.

Uniform Access Rule Compliance

Each Panasas director blade provides both NFS gateway services for the entire file system, and metadata management services for a portion of the file system (one or more virtual volumes). The system as tested contained 4 director blades and was configured with 4 volumes, one managed by each director blade. To comply with the Uniform Access Rule (UAR) for a single namespace, all the volumes were accessed through all gateways from the 8 load-generating clients and no client was afforded any advantage by colocating the NFS gateway and metadata server for the volume being accessed

Other Notes

None.

Config Diagrams

- [AS14T 2-shelf Setup](#)