SPECsfs97 Result

© Copyright 1997-2002 Standard Performance Evaluation Corporation

EMC Corp.
Celerra NS600 cluster (2FS, primary/primary, 4GB BE)

SPECsfs97_R1.v3 =
38459 Ops/Sec (Overall Response Time = 3.94 msec)

Tested:	January 2003	By:	EMC Corp.	License #:	47	Hardware Avail:	January 2003	Software Avail:	March 2003

Throughput (ops/sec)	Average Response (msec)
3906	1.1
7914	1.4
11848	2.4
15709	2.8
19765	3.4
23739	4.3
27764	5.3
31769	6.4
35892	8.9
38459	**10.3**

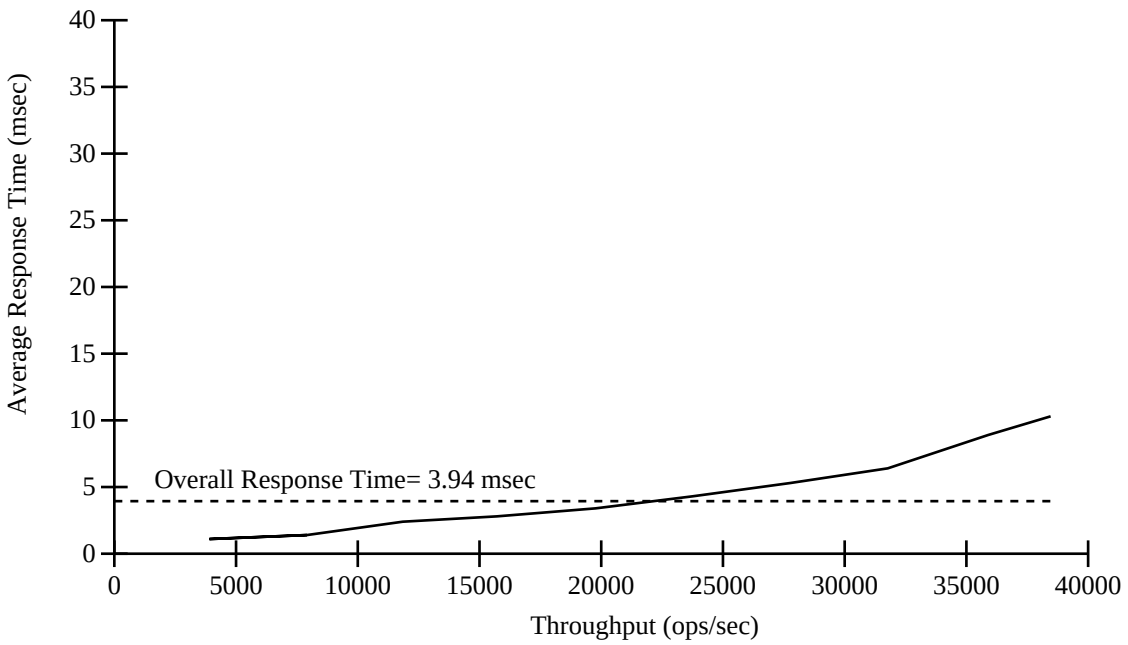

CPU, Memory and Power

Model Name:	Celerra NS600 cluster (2FS, primary/primary, 4GB BE)
Processor:	2GHz Pentium 4
# of Processors:	4 [2 per datamover]
Primary Cache:	16KBI+16KBD on chip
Secondary Cache:	512KB(I+D)
Other Cache:	N/A
UPS:	N/A
Other Hardware:	N/A
Memory Size:	8 GB (4 GB per datamover)
NVRAM Size:	N/A
NVRAM Type:	N/A
NVRAM Description:	N/A

Server Software

OS Name and Version:	Dart 5.1.11.3007
Other Software:	Redhat Linux 7.2-1 on Control Station
File System:	N/A
NFS version:	3

Network Subsystem

Network Type:	Jumbo Gigabit Ethernet
Network Controller Desc.:	Broadcom Gigabit Ethernet Controller
Number Networks:	1
Number Network Controllers:	4 (2 per datamover)
Protocol Type:	TCP
Switch Type:	Cisco 6500 GBit Switch (Jumbo)
Bridge Type:	N/A
Hub Type:	N/A
Other Network Hardware:	N/A

A line in your configuration is too long
or you have too many lines.
Please fix the problem and try again

SPECsfs97 Result

© Copyright 1997-2002 Standard Performance Evaluation Corporation

EMC Corp.
Celerra NS600 cluster (2FS, primary/primary, 4GB BE)

SPECsfs97_R1.v3 =
38459 Ops/Sec (Overall Response Time = 3.94 msec)

Load Generator (LG) Configuration

Number of Load Generators:	6
Number of Processes per LG:	56
Biod Max Read Setting:	3
Biod Max Write Setting:	3
LG Type:	LG1
LG Model:	Dell 2550
Number and Type Processors:	2 x 2.1 GHz Pentium 4
Memory Size:	2 GB
Operating System:	Linux 7.2-1
Compiler:	gcc 2.96
Compiler Options:	-O
Network Type:	Broadcom Gbit NIC, MTU=9000

Testbed Configuration

LG #	LG Type	Network	Target File Systems	Notes
1	LG1	1	/s2_stripe, /s3_stripe, ..., /s2_stripe, /s3_stripe	N/A
2	LG1	1	/s2_stripe, /s3_stripe, ..., /s2_stripe, /s3_stripe	N/A
3	LG1	1	/s2_stripe, /s3_stripe, ..., /s2_stripe, /s3_stripe	N/A
4	LG1	1	/s2_stripe, /s3_stripe, ..., /s2_stripe, /s3_stripe	N/A
5	LG1	1	/s2_stripe, /s3_stripe, ..., /s2_stripe, /s3_stripe	N/A
6	LG1	1	/s2_stripe, /s3_stripe, ..., /s2_stripe, /s3_stripe	N/A

Notes/Tuning Information

Other Notes
-param ufs inoBlkHashSize (inode block hash size)
-param ufs inoHashTableSize (inode hash table size)
-param ufs updateAccTime (disable access-time updates)
-param nfs withoutCollector (enable NFS-to-CPU thread affinity)
-param fcTach device_q_length (Fibre Channel HBA queue depth)
-param fcTach per_target_q_length (Fiber Channel HBA per device queue depth)
-param file initialize nodes (Number of inodes)
-param dnlc (Number of dynamic name lookup cache entries)
-param nfs start openfiles (Number of open files for NFS)
-param nfsd (Number of NFS Daemons)
-param bcm cge0 offload (offload IP checksum computation to the NIC)
-param bcm cge3 offload (offload IP checksum computation to the NIC)
-The ufs log was 2GB per datamover. The log was striped over 2 R1 LUNs.
-Each filesystem is built on a stripe across all filesystem LUNs.
-Storage array notes:
-The storage array exposes luns across multiple fibre ports. Each
-data lun is 146GB in size. Each physical data disk is configured
-with 2 logical LUNs. Each filesystem is built with a 64K stripe
-element size that spans across 118 Raid 1 LUNs and all 118 physical
-disks. The net result is that each client can access each LUN.
-The storage array has dual storage processor units that work as
-an active-active failover pair. The mirrored write cache is backed up
-with a battery unit capable of saving the write cache to disk prior to
-a power failure. In the event of a storage array failure, the second
-storage processor unit is capable of saving all state that was managed
-by the first (and vise-versa), even with a simultaneous power failure.
-When one of the storage processors or battery units are off-line, the
-system turns off the write cache and writes directly to disk before
-acknowledging any write operations.
-The length of time the battery can support the retention of data is 2
-minutes - that is sufficient to write all necessary data twice.
-Storage processor A could have written 99% of its memory to disk and
-then fail. In that case storage processor B has enough battery to
-store its copy of A's data as well as its own.