


NERSC System Architecture

April 2004

Visualization Server – “escher”
SGI Onyx 3400 – 12 Processors
2 Infinite Reality 4 graphics pipes
24 Gigabyte Memory
5 Terabyte Disk

HPSS
12 IBM SP servers
30 TB of cache disk
8 STK robots, 44,000 tape slots,
24 200 GB drives, 60 20 GB drives
max capacity 9 PB

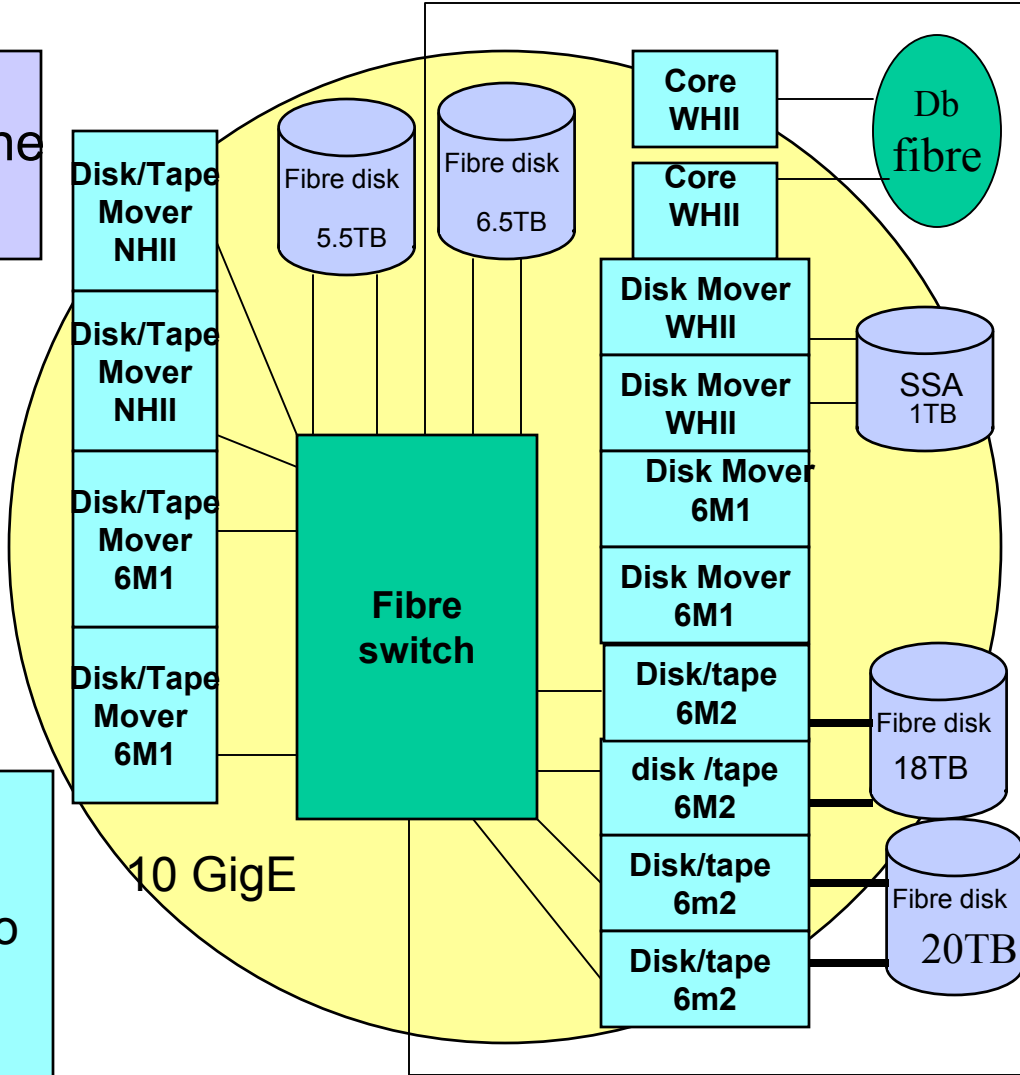


Ratio = (RAM Bytes per Flop, Disk Bytes per Flop)


HPSS: 2004/2005

Disk Cache
50 TBs


Movers:
(2X) GigE
(2X) Jumbo
10GigE
400MB/s



44,000 Tape Capacity
16 PBs



Bytes in Storage



Yearly I/O transfers to storage


Network Traffic to Storage
for 1999 (145.7 TB)

□ write ▨ read




Network Traffic to Storage
for 2001 (302.7 TB)

□ write ▨ read




Network Traffic to Storage
for 2002 (689.8 TB)

□ write ▨ read



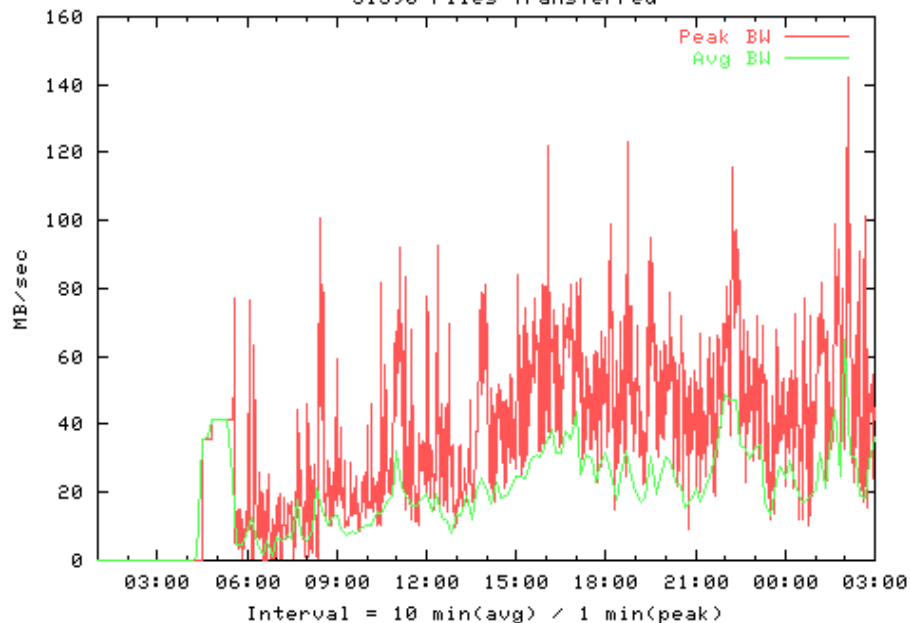
Network Traffic to Storage
for 2003 (914.4 TB)

□ write ▨ read




NERSC Aggregate Bandwidth for 12/17/2001

31590 Files Transferred




NERSC Aggregate Bandwidth for 09/12/2002

51599 Files Transferred




NERSC Aggregate Bandwidth for 10/17/2003

131,515 Files Transferred




NERSC Aggregate Bandwidth for 02/04/2004

125,929 Files Transferred




NERSC File Transfers: ALL-HOSTS 02-04-2004

Fetches: 86,956 (2,389,299 MB) Stores: 39,331 (3,325,684 MB)




NERSC Active File Transfers for 12/17/2001

31590 Files Transferred



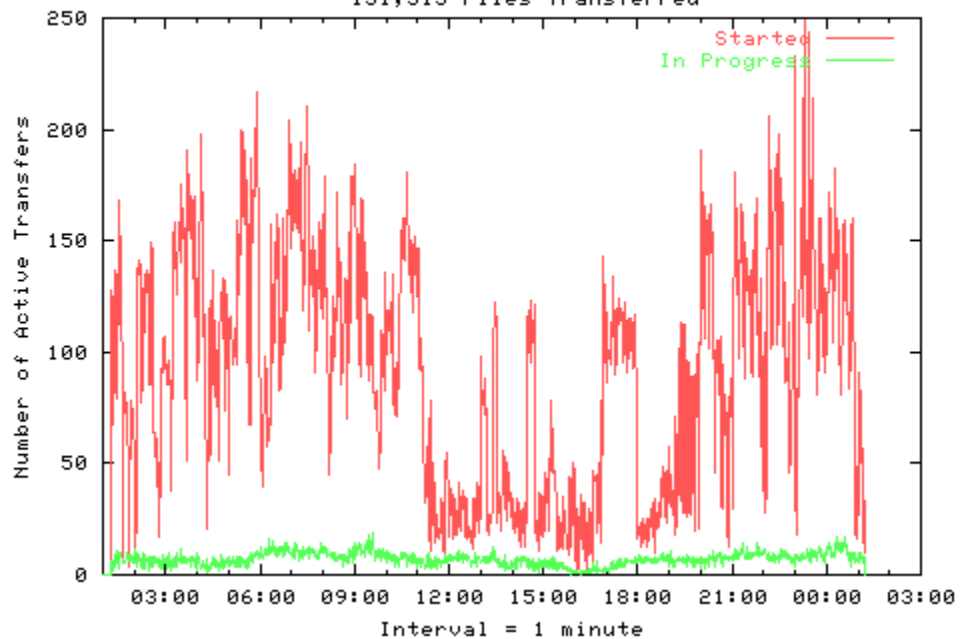
NERSC Active File Transfers for 09/12/2002

51599 Files Transferred




NERSC Active File Transfers for 10/17/2003

131,515 Files Transferred



NERSC Active File Transfers for 02/04/2004

125,929 Files Transferred



Storage Projections

Year	TBs/day	Netspeed	Disk	Data	Cap
2001	1.5TB/day	17MB/sec	7.5TB	298TB	1.6PB
2002	4TB/day	80 MB/sec	20TB	530TB	3.1PB
2003	6TB/day	200MB/sec	30TB	1PB	4.3PB
2004	10TB/day	400MB/sec	50TB	1.6PB	8.8PB
2005	20TB/day	800MB/sec	?100TB	3.1PB	16.8PB

Problems

DDN 8000 Disk Array doesn't perform with IBM 6228 FC driver commands are not queued. Implemented a 4 way stripe to get performance.

9940B Cartridges: manufacturing defect: uneven media: worked with imation scanned all 3000 carts, then repacked and replaced the bad batches 130 cartridges 4 bad files

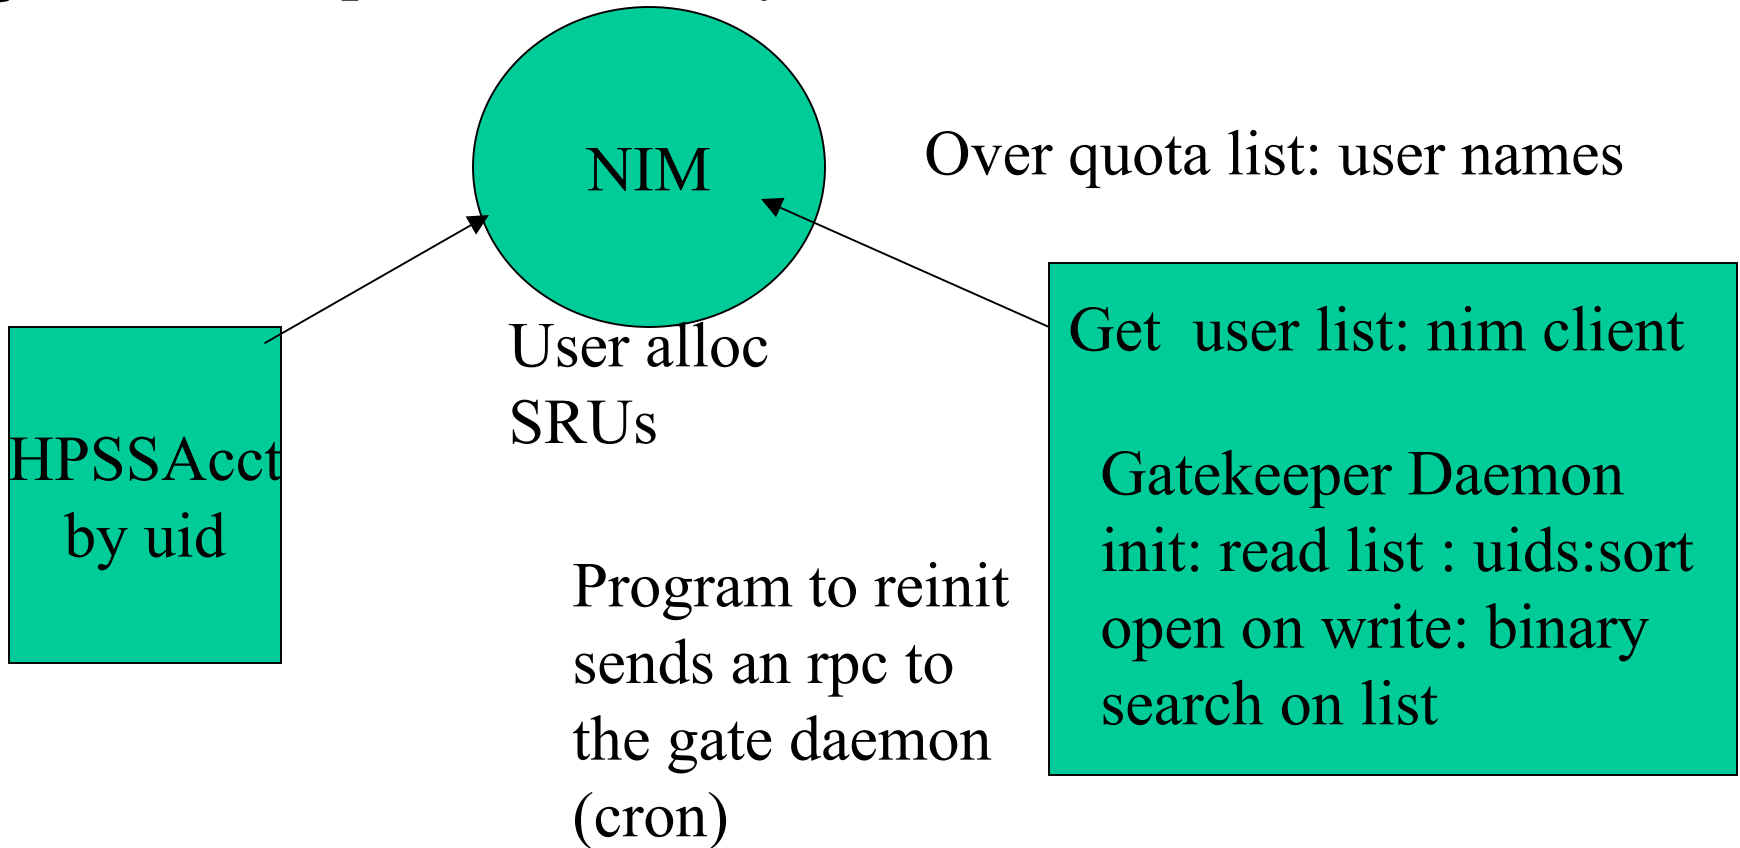
Security vulnerabilities fixed: (nessus internal testing)

We fixed buffer overflows in pftp, logc and logd

Applied IBM fix dced it cored with the scan

Unique site items

Quotas: Implemented this year



Grid/ANL testbed running HPSS 6.1

LSTape utilities (what's on tape added #tapesused/month)

NERSC Wish List

1. Aggregation of small files on disk and on tape one read brings a bunch online. Using Htar with collections
2. Secure remote management of storage systems with acceptable performance.
3. Integration of hpss with center wide filesystem solutions storage tank, luster, gpfs, etc. A DMF type of capability
4. Full grid integration with HPSS with performance
5. Mover affinity...choose mover closest to the requestor.
6. File lifetimes.