


# Texas Memory RAM-SAN 520

## Test Team

Dan Million  
Harvard Holmes  
Nancy Johnston  
Michael Gleicher

## RAM-SAN 520



- no internal disk
- backup options
  - mirror
  - dedicated RAID (TMS)

- 28 usec latency
- 3,000 MB/sec transfer rate
- 8-64 Gigabytes
- 750,000 IOPs
- Memory Upgrades: 4GB increments

## RAM-SAN 210 (not tested)

20 usec latency    200,000 IOPs  
700 MB/sec



8-32 GB, 2 or 4 FC interfaces per chassis  
dual disk drives, dual redundant power supplies

Memory Upgrades: 4GB increments

## Areas of Interest

- Raw I/O performance
- HPSS performance
- SFS performance


## Test Configuration

- RAM-SAN 520, 4 FC interfaces, 16GB
- ORNL:
  - S80 (stingray), H70 Emulex LP 7000, AIX driver
  - F50 - Emulex LP 8000, Emulex driver
  - Sun T3 (1 FC interface)
  - SSA RAID (32MB cache)
- NERSC:
  - H70 (swift) Emulex LP 8000 (3 - 2TMS,1 DDN)
  - F50 (eagle) Emulex LP 8000 (IBM driver)
  - F50 (raven) Emulex LP 8000 (IBM driver)

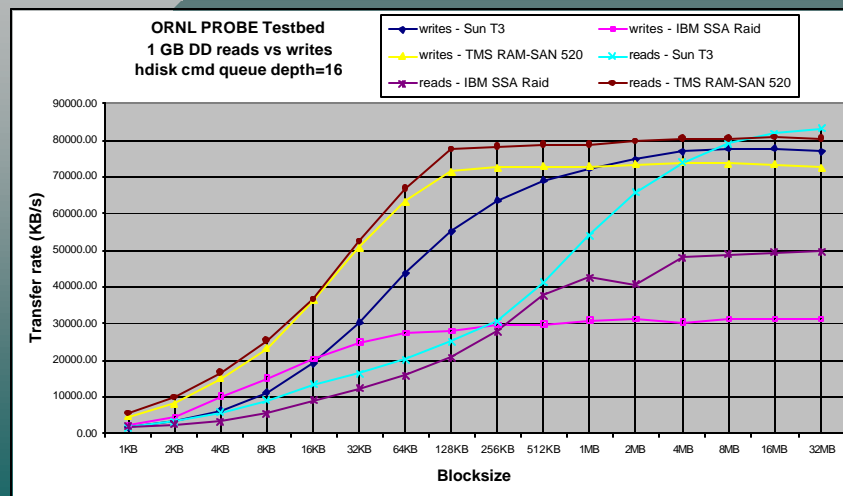
## Testing Summary - ORNL

- Raw DD, iozone tests
- Mirroring tests
- HSI read/write tests
- HPSS random access tests

## iozone results - TMS



## Raw I/O Test - (S80) sample




## TMS Tests - Mirroring tests

Notes:  
 1. queue depth=8 for both TMS and T3  
 2. "typical" non-mirrored rate used from initial tests  
 3. "best" rate (for reads) used from all mirror tests; all mirrored write tests had approximately the same results

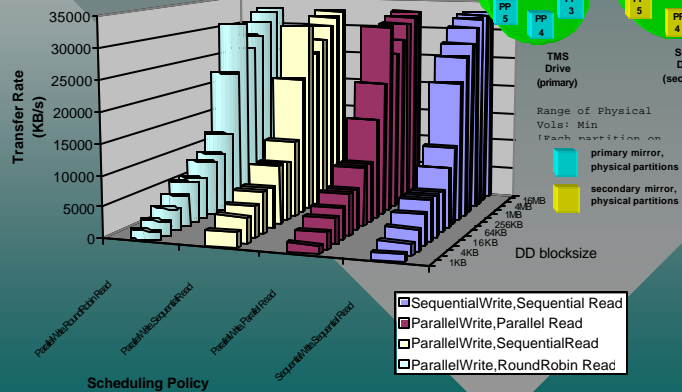

ORNL PROBE Testbed  
 Texas Memory RAM-SAN  
 Mirrored vs Nonmirrored  
 DD reads/writes

Mirror Info  
 TMS=primary  
 Sun T3=secondary  
 LVM Scheduling policy  
 parallel writes, round-robin reads

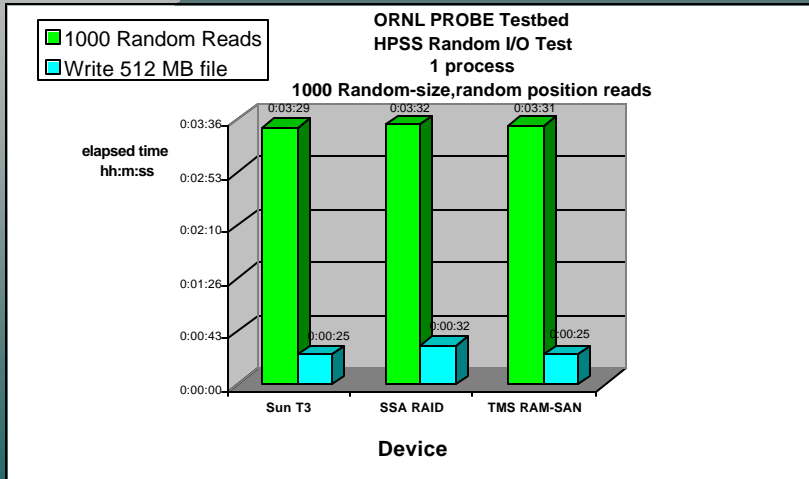


## Mirroring Tests - LVM Scheduling Policy

ORNL/PROBE  
 Texas Memory Mirrored LV Tests  
 DD writes - 1 GB from /dev/zero  
 LVM Scheduling Policy Comparison




## HPSS Random I/O Tests




Conclusion: bottleneck on issuing I/Os (?)

## HPSS Random I/O (9840)



## HPSS Sequential I/O - 1-way stripe



## Testing Summary - NERSC

- Raw I/O - dd
- HPSS Speed tests - PDSF (linux) cluster
  - 300+ MB/s HPSS aggregate rate achieved
- SFS Database tests

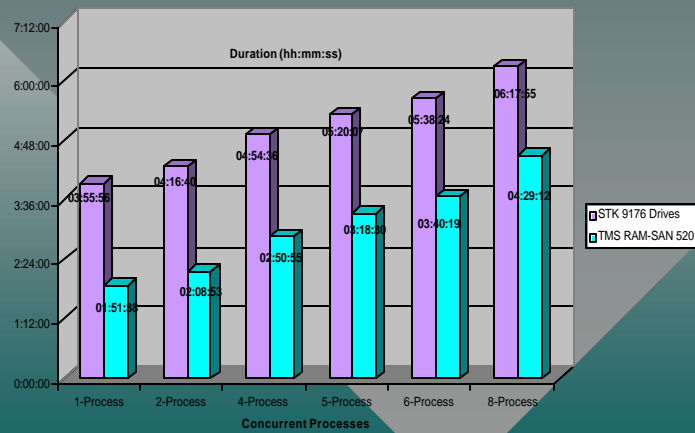
## NERSC Tests - Database


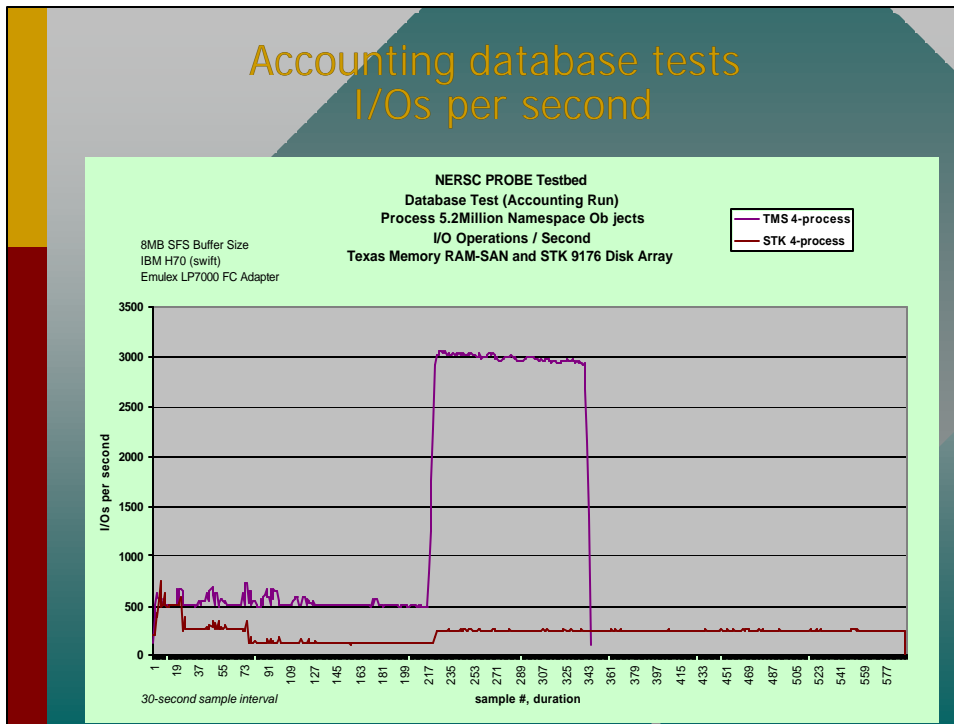
- Application: Accounting Program
- Sequentially read CNSOBJECTS
- Randomly read associated BITFILE data
- Test File: "Archive" metadata
  - restored into probe DCE cell
  - tests run on TMS, STK 9176 RAID



## Accounting Database test results Duration

IBM 4-CPU H70  
Fiberchannel-Attached  
Drives (STK & RAM-SAN)

NERSC PROBE Testbed  
Database Test (Accounting Run)  
Time to Process 5.2Million Namespace Ob jects







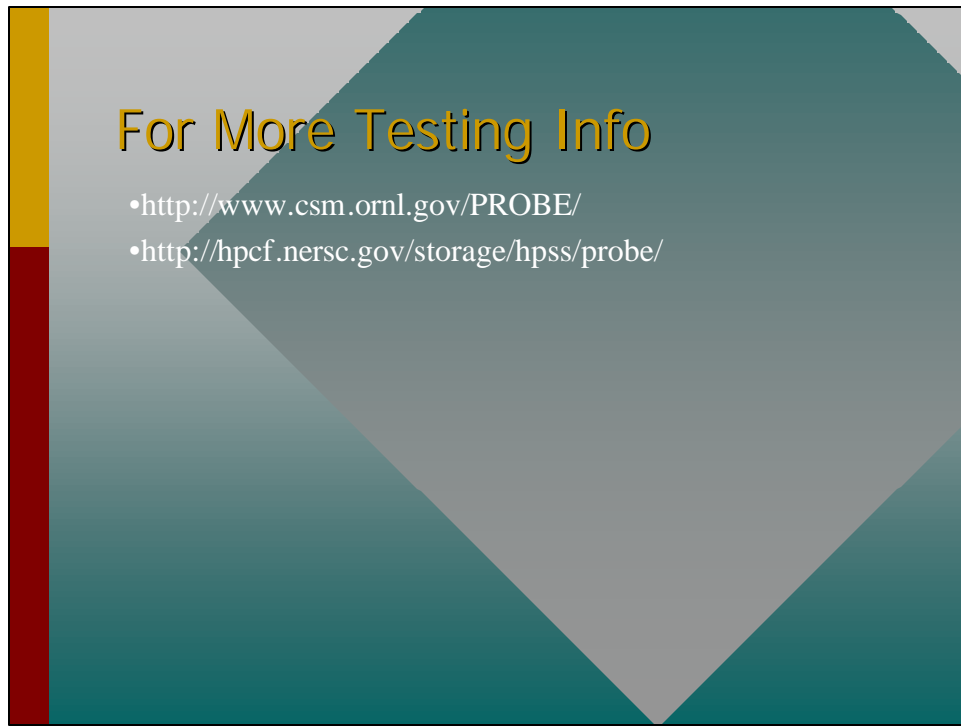
## Conclusions

- Stability was excellent - 0 problems with disk emulation software
- It easily handled everything we tried with it. Bottlenecks were always in AIX, HPSS (DCE) and SFS
- Would be an excellent device to use for metadata consistency checking, or other database-type applications

## Texas Memory Contact Information

[www.texmemsys.com](http://www.texmemsys.com)

Doug Keller



For More Testing Info

- <http://www.csm.ornl.gov/PROBE/>
- <http://hpcf.nersc.gov/storage/hpss/probe/>