

High-Performance Backups

Solutions Symposium, 2002

San Jose, CA

April 2002

Denys Beauchemin



Agenda

- Current bottlenecks
- Splits and multiplex
- DLT Libraries
- Backup Server
- DLT Library sharing between systems
- Conclusion



Current Bottlenecks

- File system backup
- Slow disk drives
- Slow SCSI channel
- On MPE (pre PCI), throughput is about 4 MB/sec. On NT/UNIX w/PCI: 15+ MB/sec
- Operator time to mount proper tapes
- Network speed for remote backups



Device Transfer Rates

	Native	Compressed
DDS-2	.5	1
DDS-3	1	2
DLT4000	1.5	3
DLT7000	5	10
DLT8000	6	12
LTO	15	30

Rates in MB/s



We watch your data

Copyright by HICOMP (August 2001) - www.hicomp.com

Page 4

HI • COMP
“The Innovators...”™

Current Bottlenecks

- If proper speed is not reached, no streaming! Shoe Shine mode
- If reply is not given, indefinite wait
- For remote backups can only go as fast as the network




Increasing Throughput

- If slow devices, use a “split” technique.
- As many splits as you have devices.
- Maybe use remote devices also.
- Overall throughput is based on aggregate speed of the splits.



“Split” Backup




Increasing Throughput

- If fast device, use “multiplex” technique.
- Multiple splits writing to the same tape concurrently.
- Splits can come from same or various systems.



“Multiplexing”




Increasing Throughput

- Requires multiprocessor system
- Perform a split backup to the multiplexer on same system



Split & Multiplex



DLT Libraries

- Wonderful devices, can be expensive.
- Range from 1 DLT drive with a few slots to 20 or more DLT drives with hundreds or thousands of slots.



DLT Libraries

HP Sure Store 818

- 1 DLT 8000
- 8 Slots
- 12MB/s or 43GB/hr
- 640 GB



DLT Libraries

HP Sure Store 2/20

- 2 DLT 8000
- 20 Slots
- 24MB/s or 86GB/hr
- 1.6 TB



DLT Libraries

HP Sure Store 4/40

- 4 DLT 8000
- 40 Slots
- 48MB/s or 172GB/hr
- 3.2 TB



DLT Libraries

HP Sure Store 6/60

- 6 DLT 8000
- 60 Slots
- 72MB/s or 258GB/hr
- 4.8 TB



DLT Libraries

HP Sure Store 6/140

- 4 DLT 8000
- 140 Slots
- 48MB/s or 172GB/hr
- 11.2 TB



We watch your data

DLT Libraries

StorageTek 9714

- 6 DLT 7000
- 100 Slots
- 72MB/s or 258GB/hr
- 8.0 TB



DLT Libraries

StorageTek L180

- 10 DLT 8000
- 174 Slots
- 120MB/s or 430GB/hr
- 13.9 TB



DLT Libraries

ATL P6000

- 32 DLT 8000
- 652 Slots
- 384MB/s 1.34TB/hr
- 52.1 TB



DLT Libraries

HP Sure Store 20/700

- 20 DLT 8000
- 700 Slots
- 240MB/s or 860GB/hr
- 56.0 TB



DLT Libraries

ADIC Scalar 1000

- 48 DLT 8000
- 788 Slots
- 576MB/s or 2.1 TB/hr
- 63.0 TB



We watch your data


Backup Server

Definition of a backup server

- Receives backups from other clients
- Can multiplex backups from other clients.
- Controls a DLT robotic library



Backup Server



We watch your data

Copyright by HICOMP (August 2001) - www.hicomp.com

Page 24

HI • COMP
“The Innovators...”™

Backup Server

Pros

- Allows systems to use robotic devices
- May provide for good throughput
- Central control

Cons

- Central point of failure!
- Requires multiple 100Base-Tx or Gigabit Ethernet
- May require extra network or segmented network with multiple cards
- May require extra server
- May require large server



Library Sharing between Systems

- Each drive in a library is independently connectable.
- Each drive can be connected to a system.
- 4 drives can be connected from 1 to 4 systems.




Library Sharing between Systems

- The library robot or “picker” also has a SCSI connection.
- One system connects to the library robot, the “robotics server.”
- The other systems send library control messages to the robotic server to issue to the library on behalf of the other systems.



Library Sharing between Systems



Library Sharing between Systems

Pros


- No central point of Failure
(Library control easily moved)
- No large central server
- Minimal Network traffic
- Highly flexible, multiple servers
- Very high enterprise throughput
- Local SCSI connection for multiple systems
- Uses Fibre Channel

Cons

- More complex to setup
- Expensive SAN hardware
- Requires Fibre Channel
- More costly in licenses



Storage Area Network (SAN)



HP e3000

HP e3000

HP 9000

Linux

NT



We watch your data

Copyright by HICOMP (August 2001) - www.hicomp.com

Page 30

HI • COMP
"The Innovators..."™

Storage Area Network (SAN)

- Same setup as Library Sharing
- Drives are shared through the SAN architecture
- One system keeps the connection to the Picker



Storage Area Network (SAN)

Pros

- No central point of Failure except for library control
- No large central server
- Minimal Network traffic
- Highly flexible/Scalable
- Very high enterprise throughput
- Local SCSI connection for multiple systems
- Maximize device use

Cons

- More complex to initially setup
- Maybe more complex to administer
- Costly SAN hardware required
- More costly in licenses



Conclusion

- Multiple ways to increase backup throughput.
- Judicious use of hardware, software and network can pay huge dividends.
- Plan wisely.

Denys Beauchemin

