Disk-to-Disk backup customer experience

Harald Burose
Architect
Hewlett-Packard

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice

"I believe between 85 and 100% of all companies will implement disk-to-disk backup within the next 18 months"

Arun Taneja (Dec 2003)

Founder, Taneja Group.

Agenda

- Backup Challenges & Dimension
- Today's Backup-To-Disk solutions
  - solution offering
  - pros & cons
  - solution positioning
- Customer Scenarios
- Emerging Backup-To-Disk solutions
  - solution offering
  - pros & cons
- Conclusion
- Questions

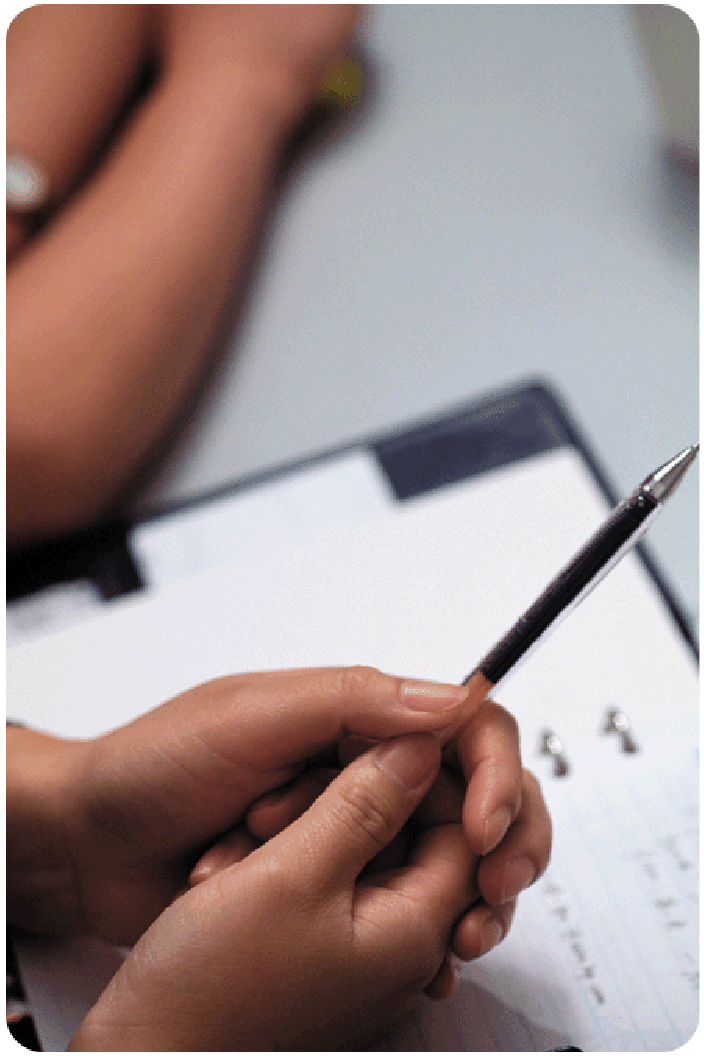

Backup/restore challenges

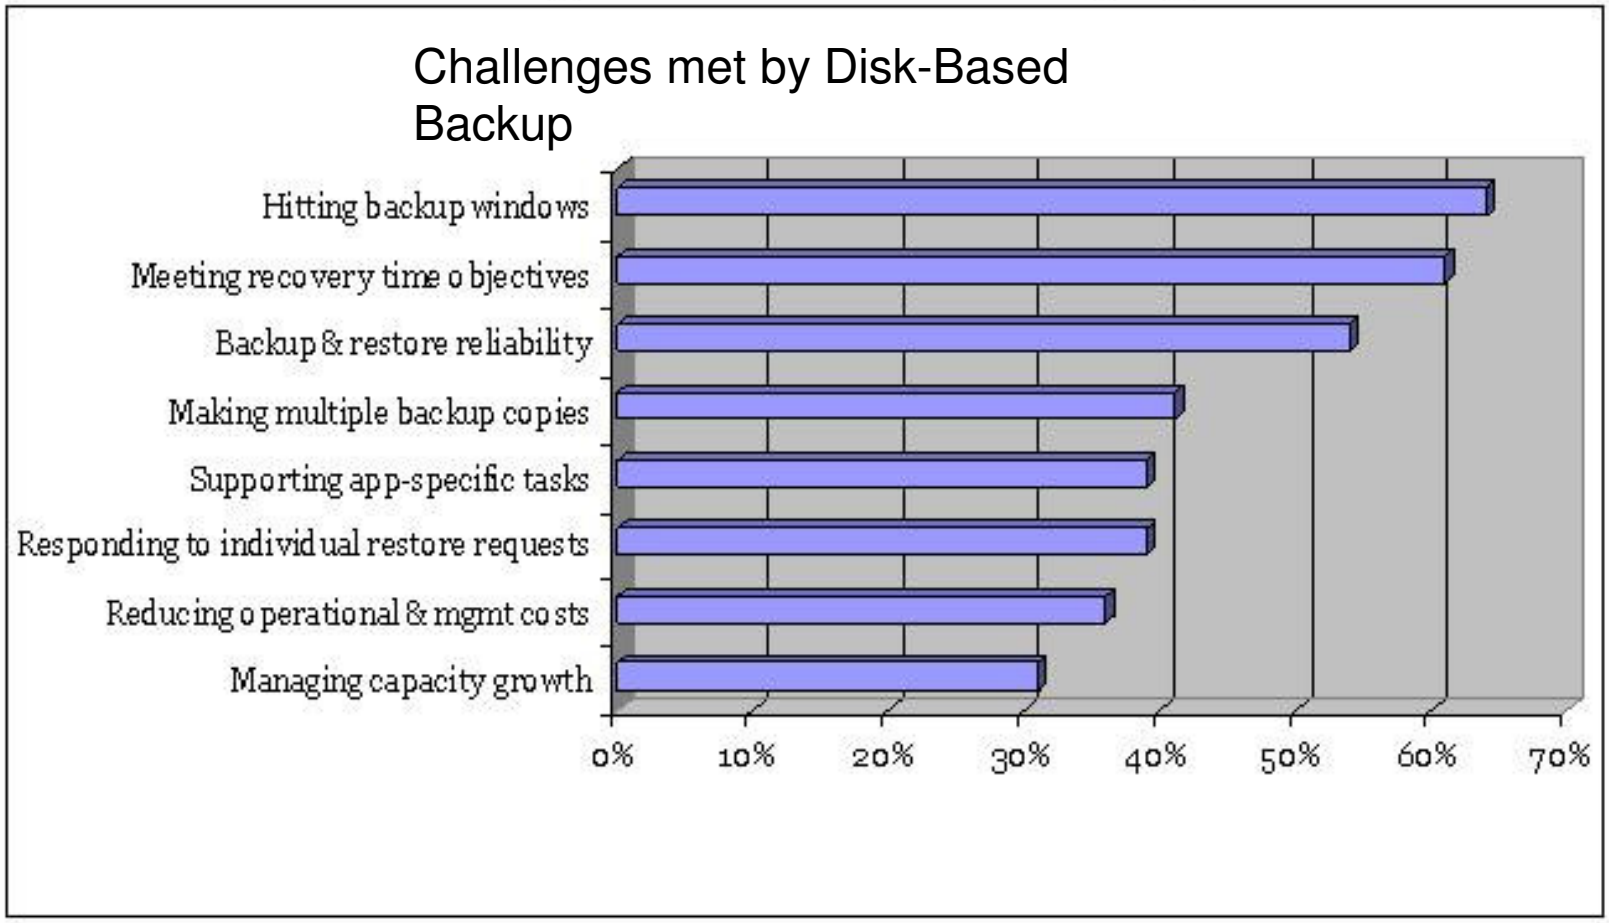

Source: Taneja Group, 2004 (300 customers surveyed)

Dimensions of the data protection problem

- Recovery time – how quickly do the systems need to be back online?
- Recovery point – how recent does the data have to be?
- Data capacity – how much data has to be managed
- Application impact – backup must be as transparent and quick as possible
- Cost of ownership – including all the cost factors
- Fit into Process – backup & recovery must fit into the overall customer process

Backup-To-Disk categorization

There are two classes of disk-assisted backup solutions today

Virtual tape

- tape hardware emulation
- simple backup to files on disk
- staged backup to disk then migrate/archive to tape

Snapshots, mirrors and clones

- maintaining point-in-time copies or full image copies on disk
- off host backup from split disk mirrors to tape

Virtual Tape

Software

- backup is written to a file by backup application
- backup objects can be subsequently migrated to tape
- format on disk is the same as on tape
- restore can be from disk or tape and must be managed by backup application

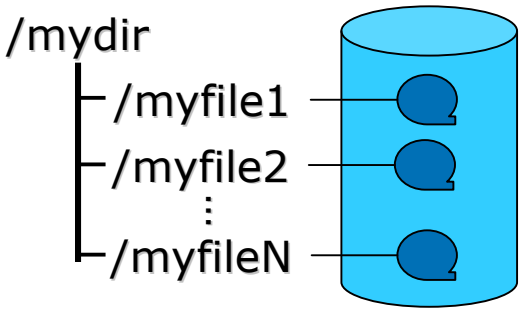

Hardware

- a controller presents disk storage as a set of scsi tape devices to hosts
- data can be migrated to tape either transparently, or through control of the backup application

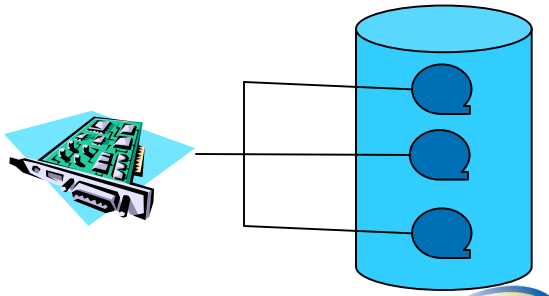

Snapshots, Mirrors and Clones

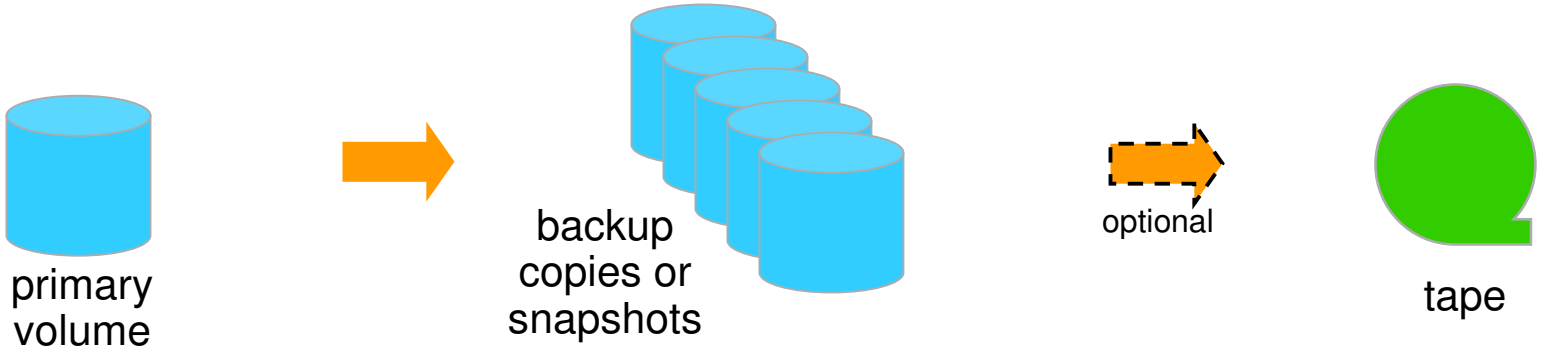

- Multiple replicas are kept on disk for fast recovery through the OS file system structure
- Replicated objects can be migrated to tape for long-term retention via a backup application

Hardware

- utilizing disk array controller, within the same frame or using remote mirrored copies between homogeneous disk arrays

Software

- implemented within the I/O stack on the host or in the SAN allowing the use of different disk array types including low-cost disk

Disk-assisted backup HP Solutions

Virtual Tape	• HP OpenView Storage Data Protector – File Device – Advanced backup to disk – Disk staging feature • HP MSA disk array with low cost SATA • HP EVA disk array with low cost FATA
Snapshots, Mirrors & Clones	• HP OpenView Storage Data Protector –Windows 2003 Volume Shadow Copy (VSS) support – HP Zero Down-time Backup (ZDB) – HP Instant Recovery • HP OpenView Storage Mirroring • HP XP disk array with Business Copy & Continuous Access • HP EVA with snapshot and clones

Solution pros & cons

Virtual tape

- ✓ no media problems
- ✓ improved backup reliability
- ✓ fast single object access

- ✗ sequential full bkp/rest slower than tape
- ✗ adjust backup process

HW (only)

- ✓ address shared SAN's

SW (only)

- ✓ re-use existing HW (storage)

Snapshots, Mirrors and Clones

- ✓ zero impact HA backup
- ✓ instant full image recovery
- ✓ leverage high availability configurations

- ✗ using low-cost disk will limit performance of mirror and use as redundant copy
- ✗ increased number of copies typically influence performance of original

HW (only)

- ✗ today requires use of expensive disk space

Will tape go away

- Tape is still - the medium for long term storage
  - a good choice for disaster recovery
    - it provides a possibility to move data offsite
  - a valid medium for large sequential data transfer
  - the most flexible medium
    - move data between locations & company borders
- Disk shines in immediate access of small granularity
- The amount of data kept on tape will be reduced but not eliminated
- Disk basically just adds a layer of protection for faster short term recovery

Customer scenario A
SAP DB backup

Environment:

- SAP DB based on Oracle
- capacity 2 – 3 TB (growing)
- batch process takes 3 – 6 hrs
- DB used between 6 a.m. & 10 p.m.
- 10 LTO I drives used for backup
- data is 1:2 compressible

Backup:

- runs every night after batch process
- data streamed directly from DB server
- takes ~ 3 hrs

Restore:

- takes ~ 3 hrs for restore
- takes ~ 3 hrs for recovery (apply logs)

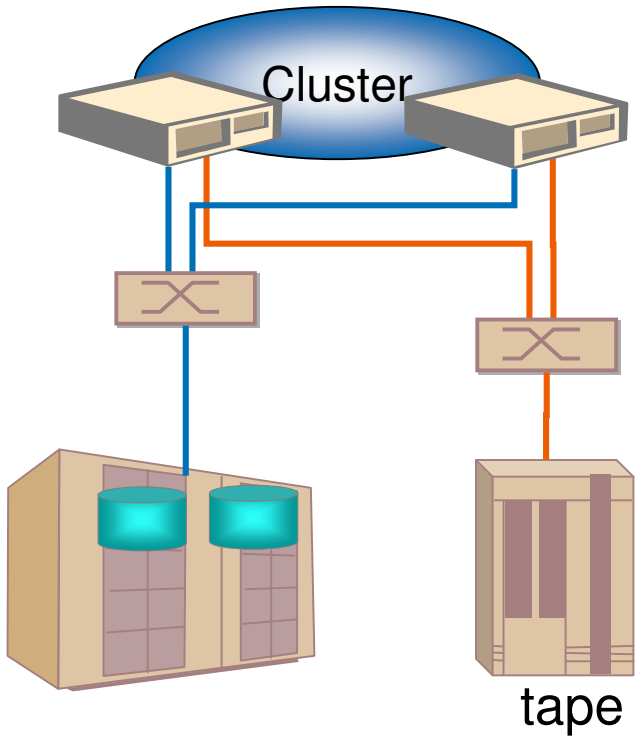

Customer scenario A current issues

Backup:

- backup overlaps with user operation
- backup takes oracle server execution power away
- backup time & batch times are at their limit
  - A single error causes skipped backups

Restore:

- time needed to open the database and verify consistency is too long

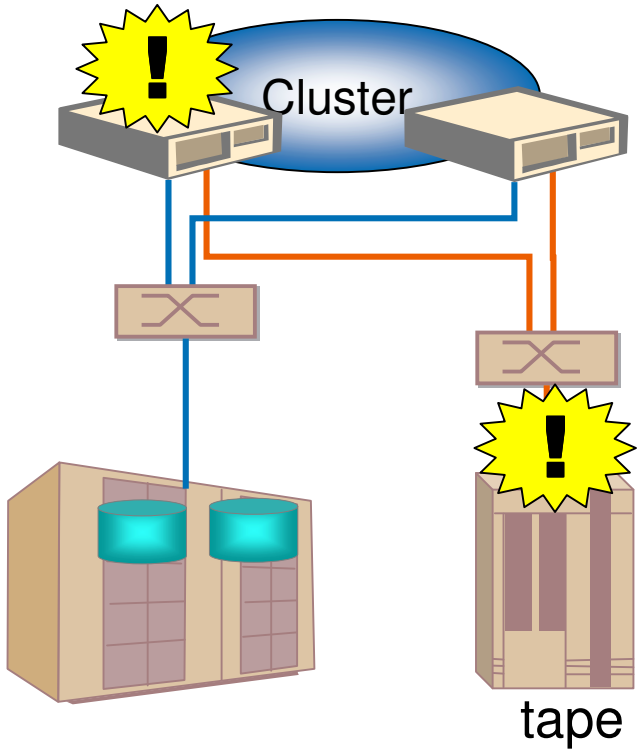

disk FC (SAN)

tape FC (SAN)

Customer scenario A
SAP DB backup using Mirror

Environment changes:

- added BC copy in XP
- adding dedicated backup server

Backup:

- runs every night after batch process
- ✓ new backup window is 15 min
  - batch & bkp fits into quiet time
  - room for data growth
- ✓ Tape backup can be restarted in case of problems without impact to application
- tape backup takes ~ 3 hrs

Restore:

- takes ~ 3 hrs for restore
- ✓ takes ~ 15 min for recovery
  - less logs need to be applied to open the DB
- total RTO improved

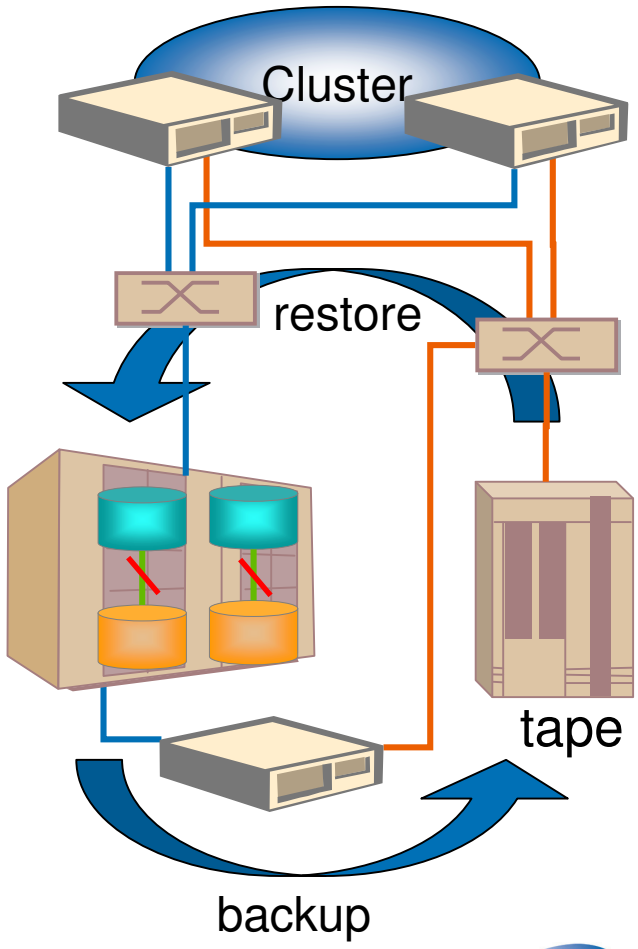

Customer scenario A
SAP DB backup additional value

Environment:

- using 2 BC copies in XP

Backup:

- add a disk only backup
- run it after the batch job
- run ZDB tape backup parallel to batch

Restore:

- improved RPO & RTO
  - restore to time before & after batch using instant recovery
  - can fix batch problems if detected during operation within a day
  - reset DB in case of batch runs into problems

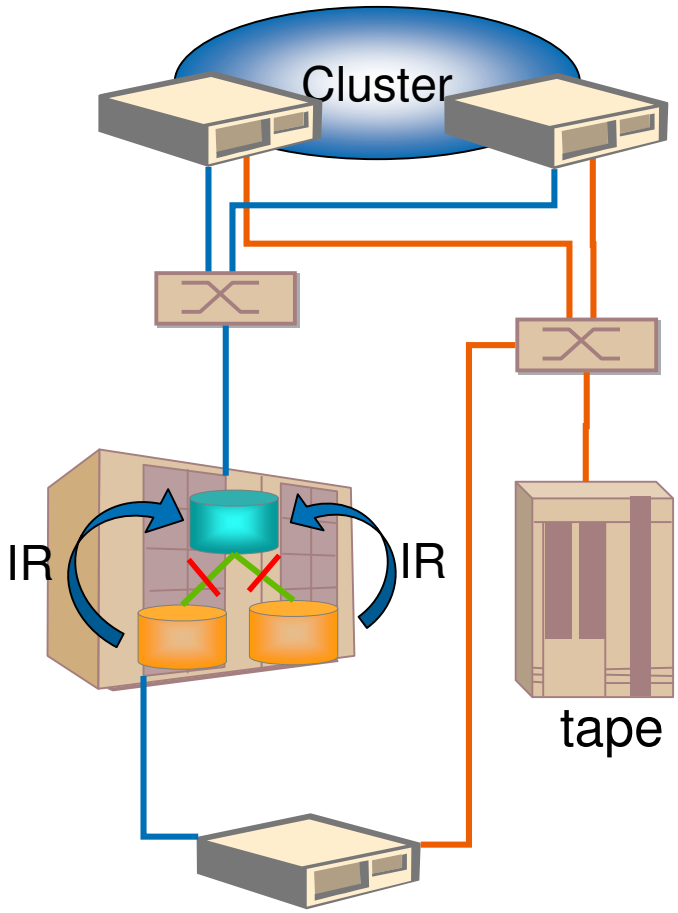

Customer scenario B
Server farm backup

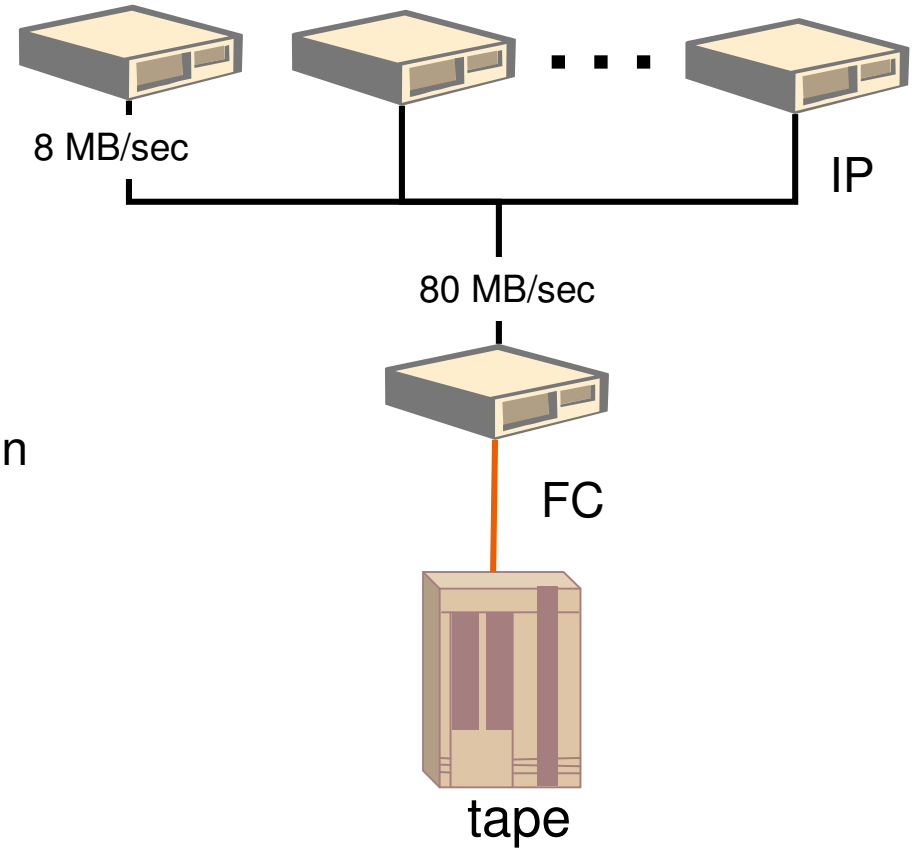

Environment:

- Client machines backup
  - capacities: 100 – 200 GB/server
- 1 Gib Ethernet into Bkp Server
- 100 baseT lines from the clients
- LTO II drives in the library
- data 2:1 compressible

Backup:

- Require 10 servers to be protected in parallel to stream 1 drive
  - 8 MB/sec from the client
  - 80 MB/sec into the backup server
- Backup window of 5 – 10 hrs
  - tape blocked for 5 – 10 hrs

Restore:

- Restore ~ 3 – 5 hrs
  - Even restore directly to host takes too long due to high multiplexing

Customer scenario B
Server Farm with SW virtual tape

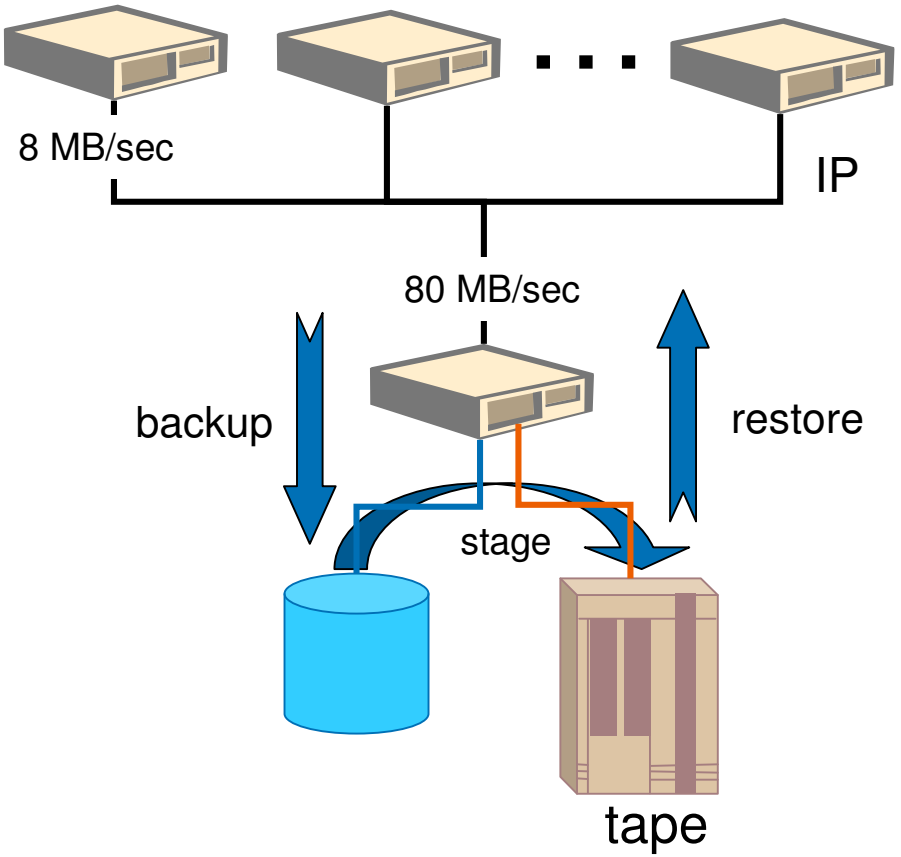

Environment change:

- add a low cost disk array to the backup server (e.g. MSA)
- utilize the backup SW (i.e. DP) backup to disk & staging features

Backup:

- Protect every server when time allows
  - improved management
  - shorter backup window for individual server
- Create tape copy as local copy
- Reduce multiplexing to tape

Restore:

- Faster single system restore due to non-multiplexed data
- Improved RTO

Additional scenarios

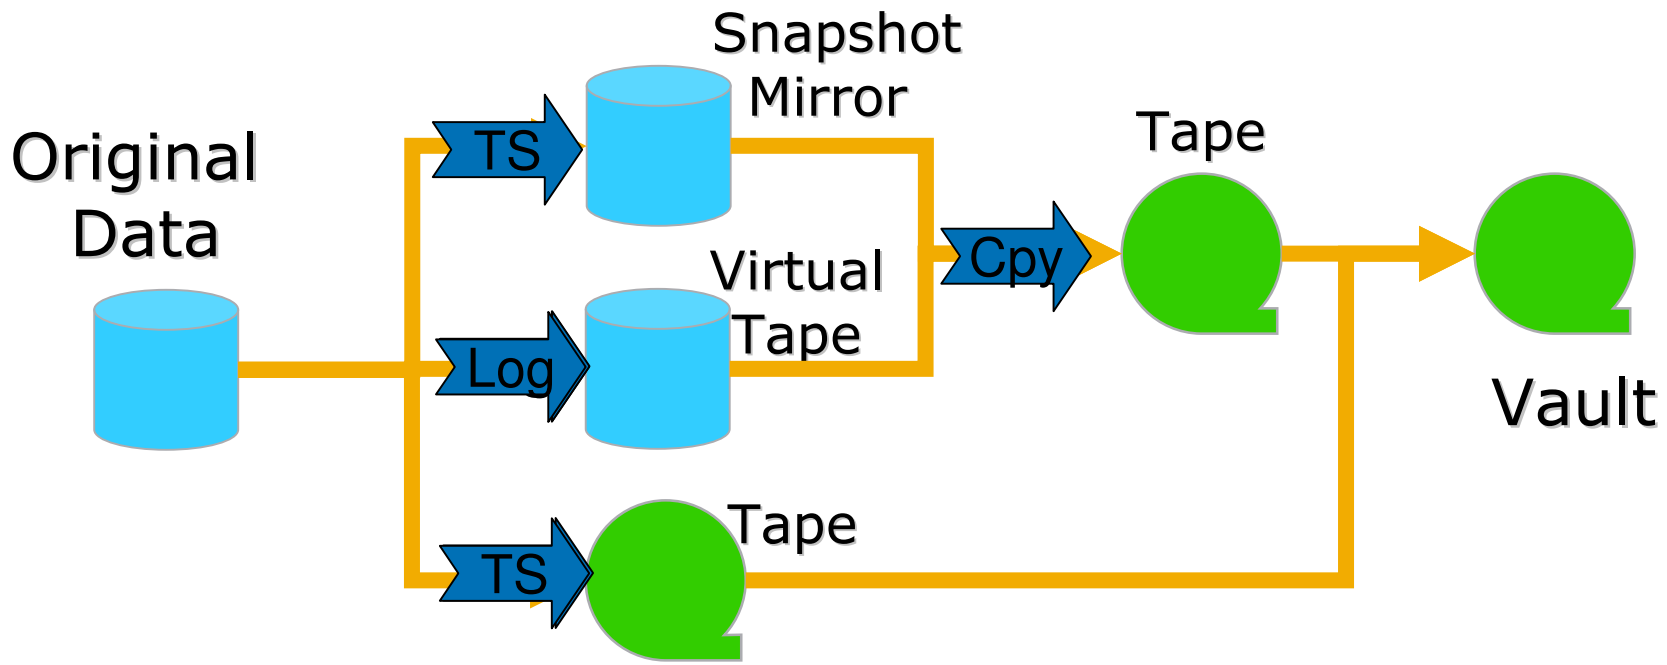

Combination use cases

- Stage data on disk to make it worth streaming it to tape
  - accumulate DB log files on disk before moved to tape
  - send data via the IP network to disk before moved to tape
- Keep short lived data on disk
  - maintain incremental backups on disk
  - take full backups to the vault

Solution positioning

- Snapshots and Mirrors are excellent for databases & applications
  - major restore case is the full DB
- Virtual Tape is excellent for File servers, consolidating data streams & log file backup
  - major restore case is single objects where it outperforms tape
- Tape works for both use cases

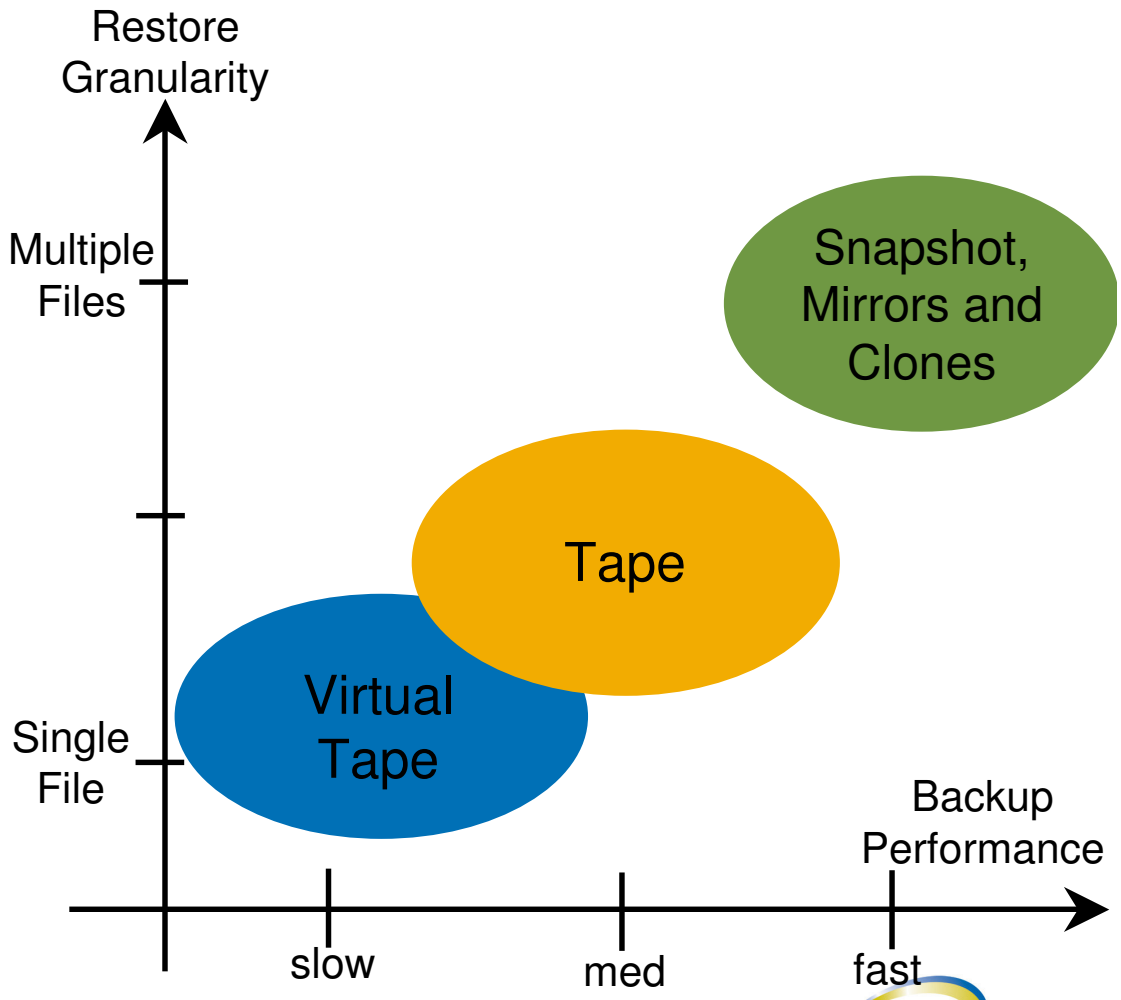

Dimensions addressed by today's solution

- ✓ Recovery time – how quickly do the systems need to be back online
- ○ Recovery point – how recent does the data have to be
- × Data capacity – how much data has to be managed
- ✓ Application impact – backup must be as transparent and quick as possible
- ○ Cost of ownership – including all the cost factors
- ✓ Fit into Process – backup & recovery must fit into the overall customer process

New emerging solutions

Block-level incremental backup

- incremental backups forever
- virtual full image restore

Time Addressable Storage

- continuous backup of all changes
- restore to any point in time

Block-level incremental backup

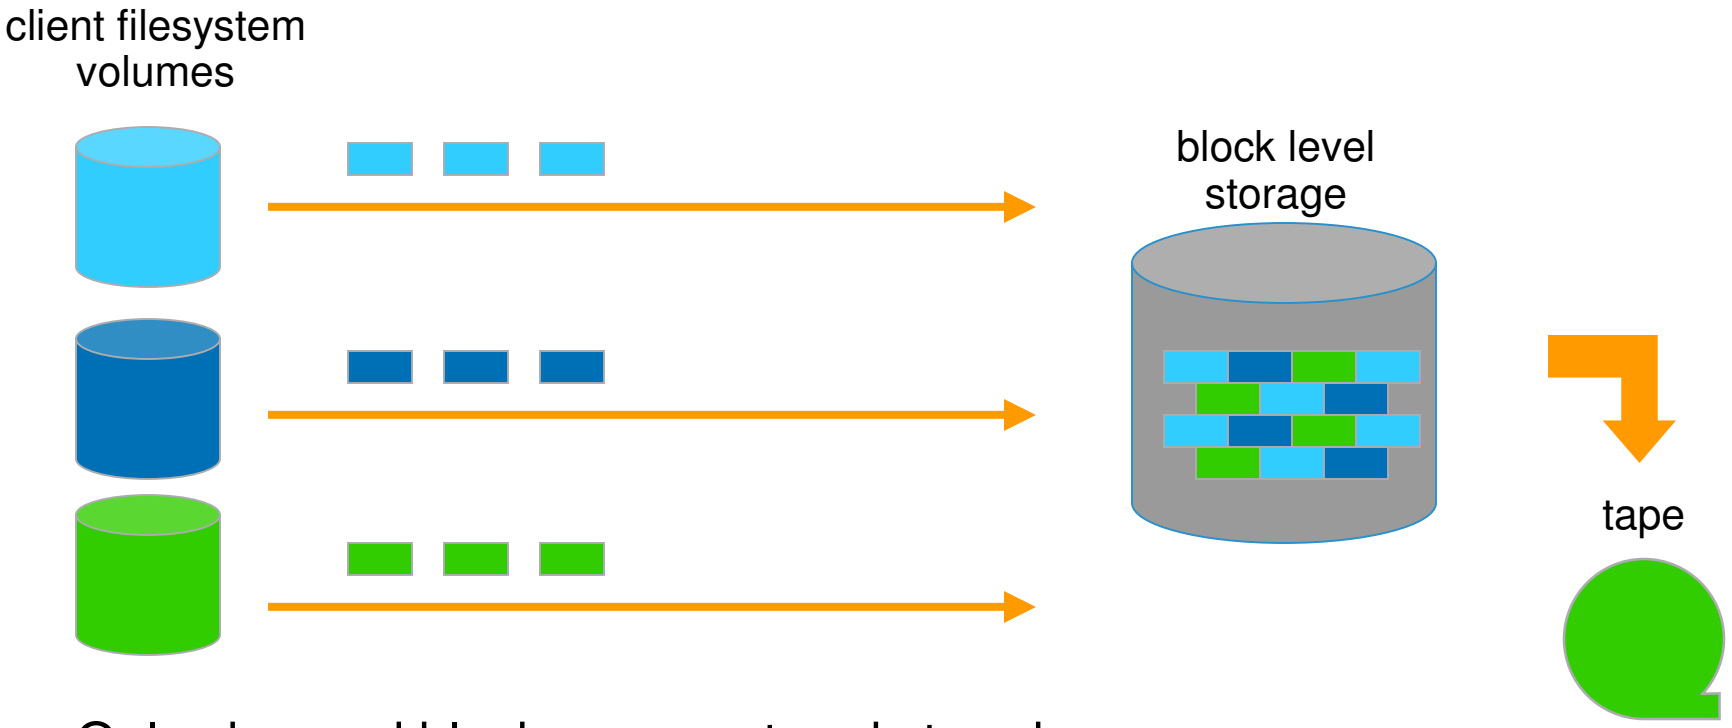

- Only changed blocks are sent and stored.
- Virtual full restore image can be created by assembling correct blocks, and archived to tape for vaulting and fast image restore.
- Identical blocks are stored only once, reducing storage requirements.

Time Addressable Storage

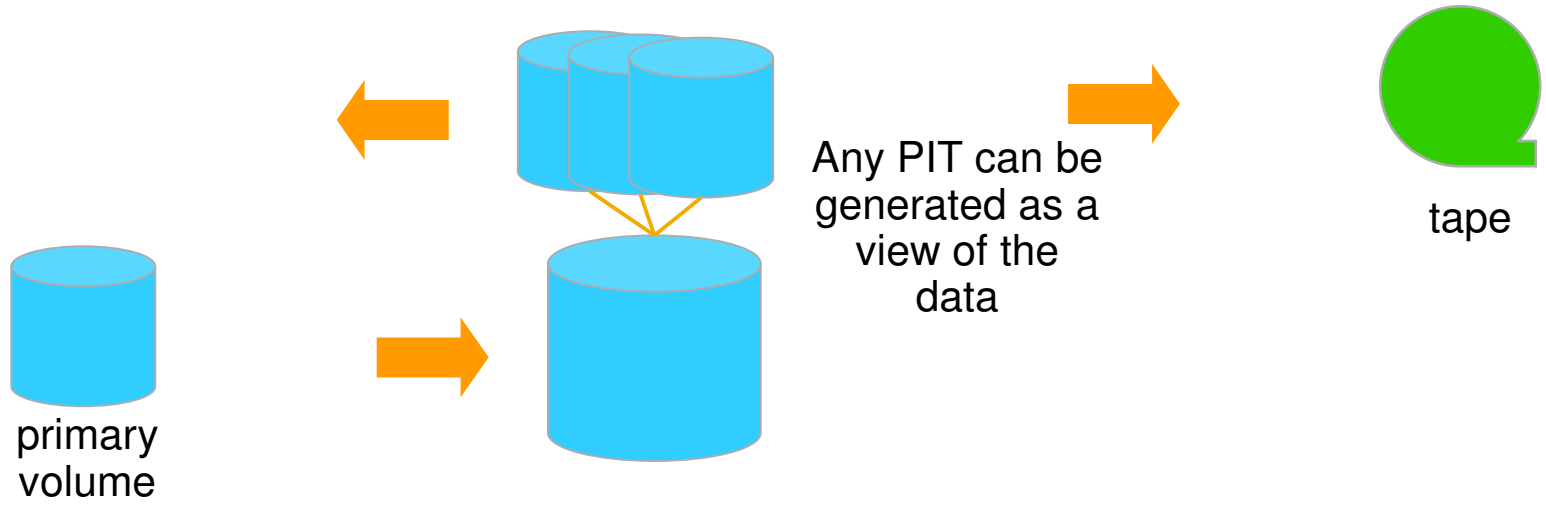

- Every write is replicated to a remote copy and journalled on that copy
- Any point in time (PIT) can be created as a view of the data after the fact
- The PIT can be used for tape copies, restore or simple data investigation

Hardware

- utilizes controller that acts as a target for replication

Software

- implemented within the I/O stack on the host

Emerging disk-assisted backup HP's position

- HP views this as an emerging market
  - Few players in the market
  - Needs proof of concept
  - Missing interface & integration into tape backup schema
- These concepts derive real value from random access characteristics of the disk
- A combination of both technologies would address the problems of the backup market
  - while utilizing the backup resources most efficiently
- HP will mature this technology to enable it for broad usage
- By combining archive & backup concepts, truly new backup paradigms can be achieved

Solution pros & cons

Blocklevel Incremental

- ✓ optimize backup resource utilization
- ✓ virtual full restores (compose image when required
- ✓ Same as virtual tape solution
- ✘ Sequential full backup/restore slower than tape
- ✘ Adjust backup process

Time Addressable Storage

- ✓ capture every state of the environment
- ✘ Synchronization with the application is mandatory to guarantee consistency at restore time (avoid trial and error)
- ✘ Additional load to the system
- ✘ Might capture unnecessary versions

Conclusion

- Backup to disk is not a single solution
- Each of the solutions is justified for specific market segment / business applications
- Tape will not disappear as technology in the backup & restore market
- The integration of the different layers of protection (Disk & Tape) is required to provide the lowest TCO
- Just changing the storage medium will not solve all the problems

For More Information

- For More Data Protector Info:
  - http://www.hp.com/go/dataprotector
    - Product Info
    - Support Matrices
    - Data Protector Manuals
    - Evaluation Software (60-day Trial)
- For more ILM info:
  - http://www.hp.com/go/ilm
- See you at the Technology Showcase
- Speak to your HP Storage Specialist

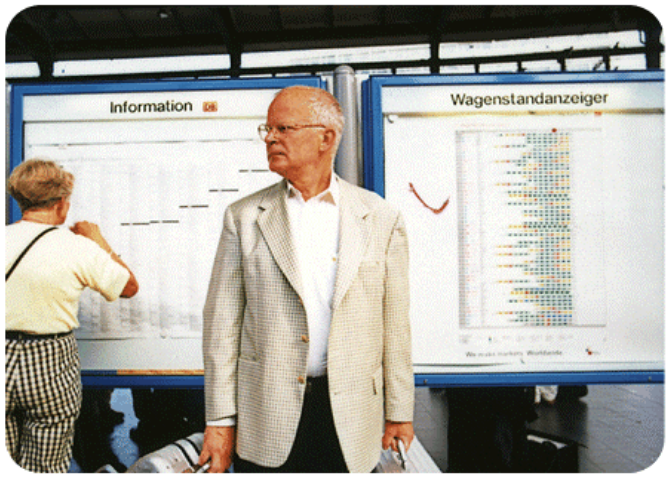

HP WORLD 2004
Solutions and Technology Conference & Expo

Questions

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE

HP Certified Professional

Backup slides

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice

Configure Copy Sessions

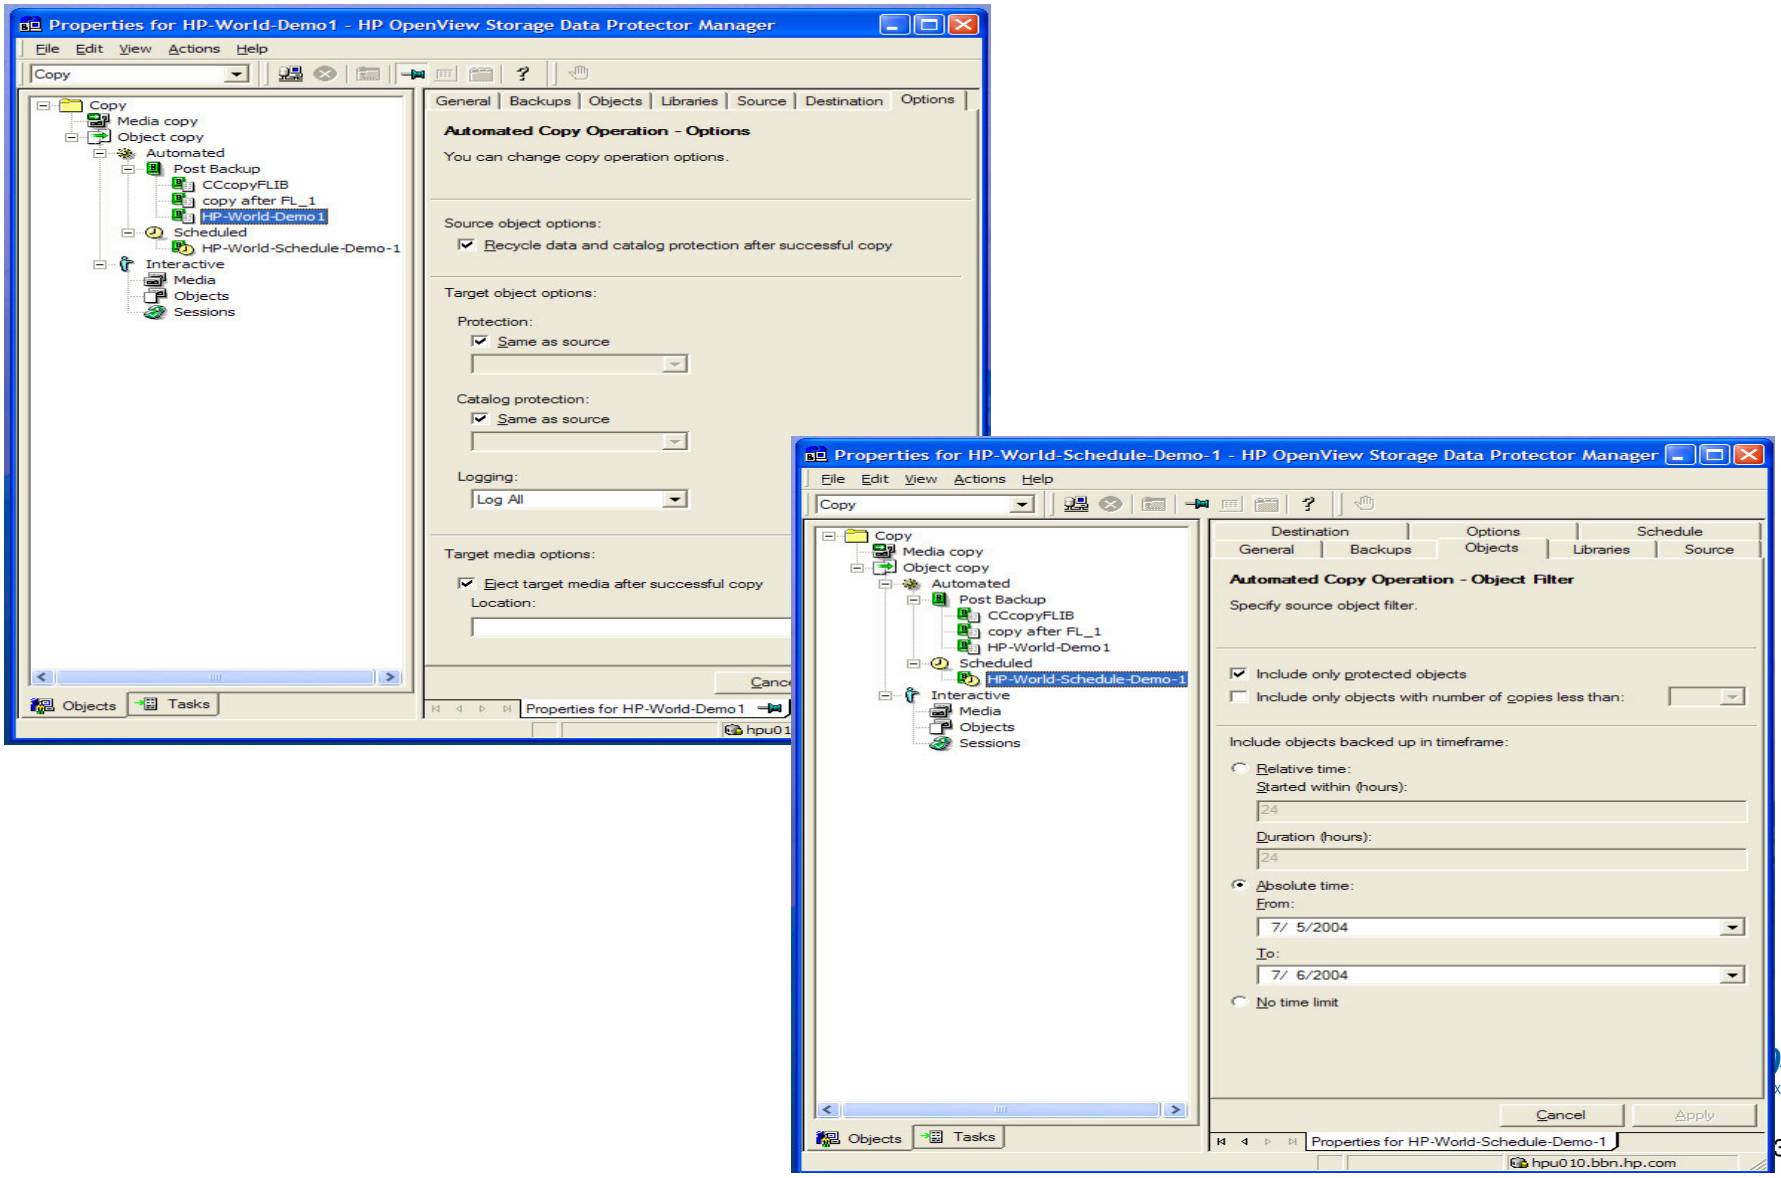

HP openview storage data protector advanced backup to disk

- Disk staging functionality
  - restore data transparent from tape or disk
- Automatic space management
  - delete data that expired
    - before every backup
    - in case a new file (media) must be created
- Simultaneous read
  - allow restore while backup is performed
- GUI wizard for configuration

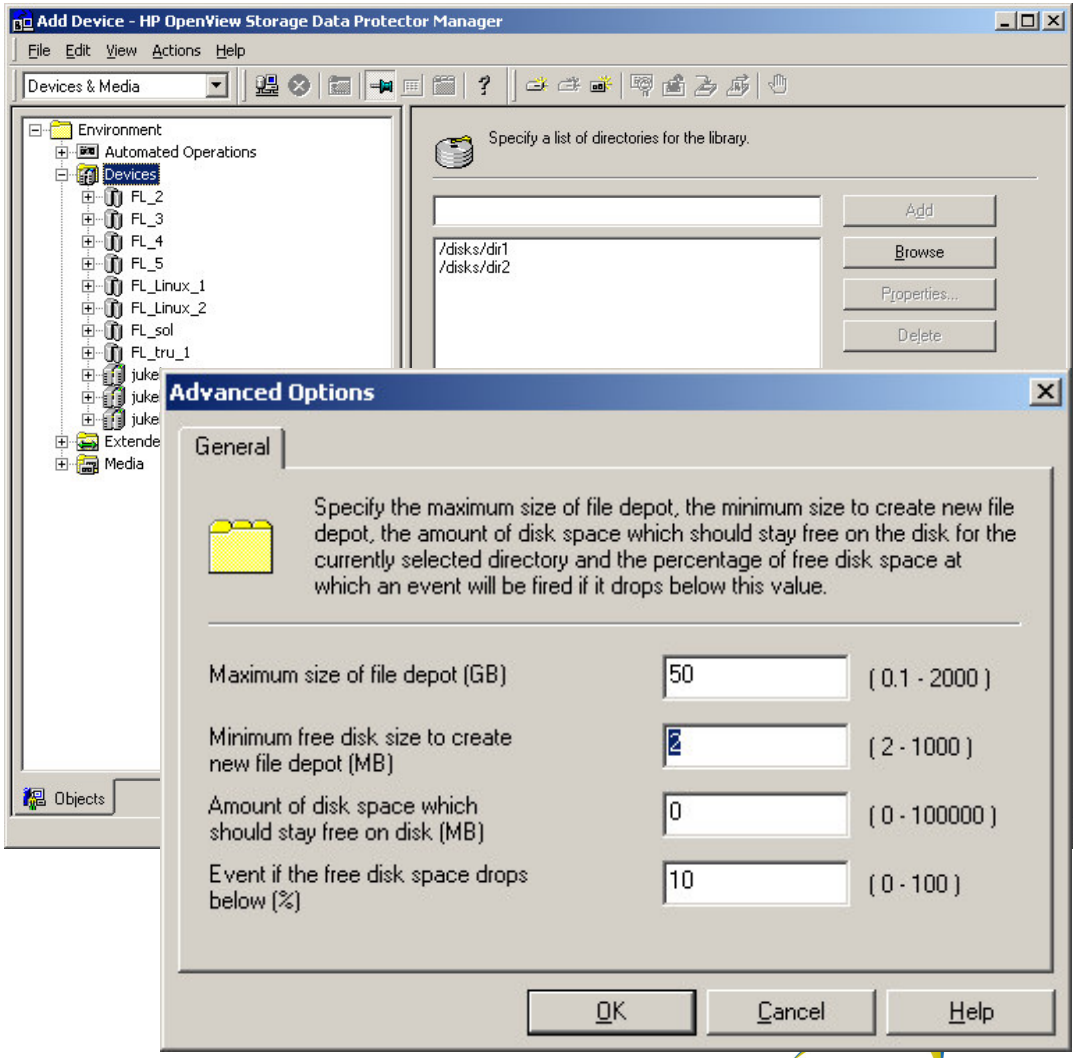

For more details on DP 5.5 see presentation

Advanced backup to disk details