



## Oxford University TSM Symposium 2007

# D2D2T Backup The right Storage and the right Architecture for your Backup/Restore Environment


Josef Weingand  
IBM Germany

# Agenda


- **Backup Overview**
- **Tape vs Disk**
  - TCO
- **Backup Tiers and Architecture**
  - D2T, D2D2T, D2D2N
  - LAN-free vs LAN-based
- **Virtual Tape Libraries**

## Why Data Backup

### ■ Backups are created for Restoring Data

- For any kind of failures
    - Disaster
      - Natural Disaster: Fire, Hurrigan, ...
      - Technical Defekt: Harddisk-Crash, ...
    - Human Errors
      - Deletion of Files
      - Failures during editing
      - ...
    - Application Errors
      - Software Failures
      - Virus
  - For keeping different Versions
    - Monthly and annual balance sheets
    - Batch-Processing
    - Testsystems
- 
- Therefore Backups are stored on different places
  - Therefore Backups are stored on different media
    - Don't make Backups of Netapp-Data on Netapp
  - Therefore Backups are made with different Software
    - Don't make Backups of Netapp-Data with Ontap
  - Therefore Backup should be done by a different team/person

© 1999 Randy Glasbergen. www.glasbergen.com




**“You caught a virus from your computer  
and we had to erase your brain. I hope  
you kept a back-up copy.”**

## Restore

- **There are two different kinds of Restore**
  - Disaster Recovery
    - Fast Restore of a big amount of Data
  - Single File Restore
    - Restore of single Files
  
- **Backup should be designed for Restore and not vice versa!**
  - For optimal performance of Disaster Recover and Single File Restore you need different Backup process

## Traditional Backup



- **Clients writes backup data via LAN to Backup Server and Tape**

- Since some time also LAN-free Backup to Tape possible

- **Problems:**

- Current Tape Drives are very fast (LTO3 = 80 MB/sec, TS1120=104 MB/sec), but most of the clients still connected over 100 MBit Ethernet

(~ 8 MB/sec)


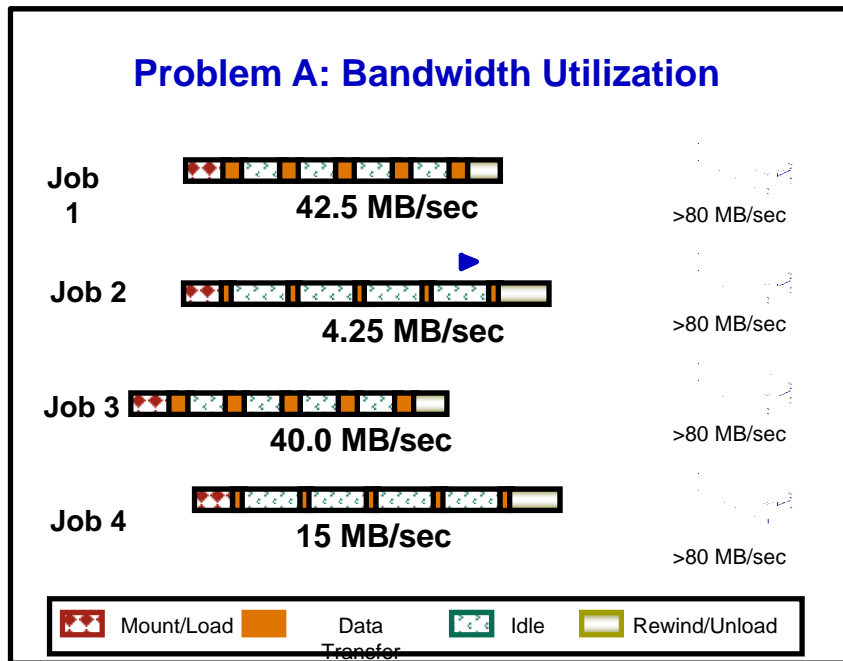
- No advantage from fast Tape Drives
- Limited bandwidth utilization

- LAN-free backups needs dedicated Tape Drives

- Many Tape Drives needed during Backup Window

- **Many Tape Drives needed during Backup Window**


# Problems with traditional Backup



## ■ The Solution: Backup to Disk (D2D)

- Several Backups are able to backup simultaneously
- No tape drive shortness
- Improved bandwidth utilization
- No multiplexing needed

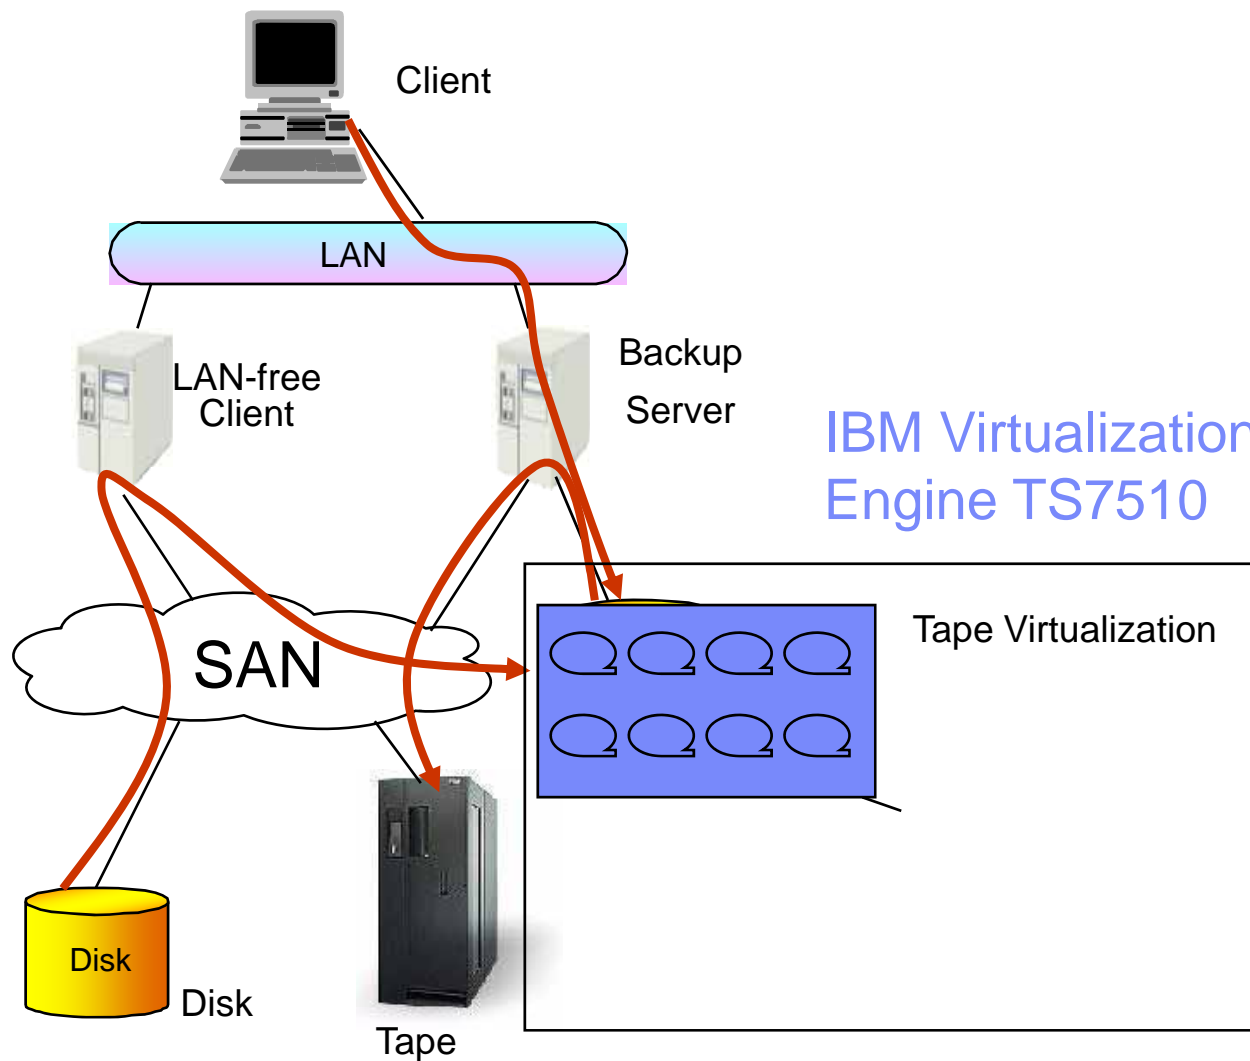
## Backup with Disk Buffer (eg. TSM) – D2D2T



- **Slow Backups are first buffered on Disk and will be migrated later on**
  - Advantage: Improved bandwidth utilization and no tape drive shortness
  - Afterwards migration from Disk to Tape (D2D2T)
  - Most effective solution
- **Some Data could be stored longer on Disk in order to get better Restore performance**
  - For single file restore, file server, etc.
- **Fast Backups with a large amount of Data could be still stored directly to tape**
- **IBM can deliver all necessary components for a D2D2T**
  - Server (System p, System x)
  - Disk Storage (DS4000, DS8000)
  - Tape System
  - Virtual Tape (TS7500)
  - Software (TSM)



# Virtual Tape Library - VTL



# Agenda

- Backup Overview
- **Tape vs Disk**
  - TCO
- Backup Tiers and Architecture
  - D2T, D2D2T, D2D2N
  - LAN-free vs LAN-based
- Virtual Tape Libraries
- De-Duplication

## What's the Problem with Tape ?

- **Tape is a sequential device**
  - Search time is much longer than on disk
  - This will reduce Restore Performance if Data are not stored as a stream on Tape
    - Multiplexing Backups
    - Incremental Backups
  
- **Tape allows just one stream at a time**
  - Parallel backup to one tape drive is not possible
  
  - Not every server/client can't deliver a single stream with enough performance
    - Problem of the client or of the read performance of the disk
    - Therefore multiple session backup
      - Either use multiple tape drive
      - Or backup to disk/VTL

# Tape vs. Disk

- **Fast (Single-) Streaming Device**

- Single Tape stream is not slower than a Disk stream

- **Multiple Streams possible**

- No advantage for large files (full/image Backup)

## Backup to Disk is not faster!

- **Sequential Device**

- A Single File-Read needs a Locate (tape positioning)
    - Restore of Incremental / Differential Backups are very slow if data are on Tape

- **Random Access**


- Advantage at Single File and FileSystem Restore

## Restore of small Files and Filesystem is faster from Disk

- **Tape is OnDemand**

- **Tape has lowest TCO!**

What is the greatest facility problem with your primary data center?



Gartner, Best Practices In Data Center Facilities, Michael Bell, October 2006

- **Any Disksystem has limited Throughput and Capacity**


- **Disk needs always power and has therefore an higher TCO**



Microsoft Excel Worksheet

# Use Disk where it is usefull, where you get benefit from it!

# Electrical Power & Cooling - Cost for 6 Years



## Calculating Electrical Power Cost for Usage and Cooling

- **Power Input = Heat output**
  - All IT devices change electrical power of about 99% into heat
- **Power x Time = Work (kWh)**
- **Cooling needs about 0,5 – 3 times of the electrical power input**
  - Assumption „1“
    - E.g: 100 W heat output needs 100W electrical Power for Cooling
      - Double the input power in order to get power cost for running and cooling

## Example

- **TS7520 (VTL)**

– CV6	510 W
– SV6 (6,5 TB)	450 W
– SX6 (6,5 TB)	450 W

- **Tape Library**

– TS3500	185 W
– One Tape Drive	46W

- **Example: Compare power cost for running and cooling of a VTL with 500,5 TB with a tape library with 10 tape drives for one and for six years. Use power cost of 0,15 € / kW/h**

## Example

- **500,5 TB = 77 SC6/SX6**
  - 77 x 450 W = 34 650 W
  - 1x CV6 = 510 W
  - Sum = 35 160 W
  - With cooling = 70 320 W
  - Per day = 1687,68 kWh
  - Per year = 616 003,2 kWh
  - = 92 400,48 €
  
- **Tape Lib**
  - 185 W
  - 10x46W = 460 W
  - sum 645 W
  - With cooling= 1 290 W
  - Per day = 30,96 kWh (a 2 – 4 person household needs about 10 – 20 kWh per day)
  - Per year = 11 300,4 kWh
  - = 1695,06 €
  
- Difference per year = 90 705,42 €
  - **Over six years = 544 232 €**



## Tape is dead! Really ?

- **IBM shipped in 2006 3116 PB Tape Hardware**
- **If IBM tape is replaced by Disk/VTL (TS7520), then we would need 480 000 SV6/SX6 (a 6,5 TB)**
  - 480 000 \* 0,45 kW
  - =216 000 kW = 216 MW
  - With Cooling = 432 MW
- **One nuclear power plant has about 1 000 MW**
- If tape is dead, then we would need every 2 years a new nuclear power plant for Disk replacing (IBM) tape!

## The complementary Disk/Tape Backup Solution = D2D2T

- **A DiskPool supports the Backup process**
  - Store slow Backupstreams on Disk and migrate data from Disk to Tape afterwards
  - Reduce the amount of needed Tape Drives
    - Replace the 1:1 relationship between Backup and Tape Drive
- **A DiskPool supports the Restore process**
  - Restore of small Files and File system are faster from Disk
    - Random-Access possible therefore no long locates need as on tape

### ***But***

- **Tape will be always the lowest cost Storage Hierarchy**

### ***Therefore***


- **Use DiskPool only for the *right* Data**
  - Disk Residency/Retention Time impacts the total cost
- **Use direct Tape Backup for DBs, Image and Full Backups**
- **Small Files (< 1-2 MB) store on Disk only**
  - Reduce Tape usage for housekeeping (Reclamation) and improves Restore

# Agenda

- Backup Overview
- Tape vs Disk
  - TCO
- **Backup Tiers and Architecture**
  - D2T, D2D2T, D2D2N
  - LAN-free vs LAN-based
- Virtual Tape Libraries
- De-Duplication

## Backup Tiers and Architecture

- Just one Backup Tier / Architecture will not meet the requirements for Restore and Backup
- You need several Backup Tiers / Architectures to meet the requirements
- Several Backup Tiers will reduce cost



## Backup Tiers and Architecture: FlashCopy

- **For very large DBs**
- **Eliminate Backup Window**
  - Creation of Flash just takes a few seconds
- **Very fast restore from Flash (Flashback)**
- **Save FlashCopies LAN-free direct to Tape**
  
- **For TSM use Tivoli Storage Manager for Advanced Copy Services**
  - Data Protection for IBM ESS / DS8000 for mySAP
  - Data Protection for Disk Storage and SAN Volume Controller for mySAP

## Backup Tiers and Architecture: LAN-free

- **Large DBs LAN-free to Tape**
- **Only Clients with backup speed faster than 50 - 80 MB/sec**
  - For LTO3&4 and TS1120
  - Keep it simply and use it just for a few clients
- **Clients who has SLAs for Restore which can be only meet with LAN-free**
- **Ressource Sharing and Tape-Handling can be done with eRMM**
  - To reduce administrative and maintenance tasks
    - Maintenance Device Paths
- **However LAN-free means additional administrative task and additional license cost**
  - Therefore use LAN-free restrictive and only for a few clients
- **But throughput through the backup server will be reduced**

## Backup Tiers and Architecture: D2Tape

- **Large DBs direct to Tape**
- **Only Clients with backup speed faster than 50 MB/sec**
  - For LTO3&4 and TS1120
- **Will reduce the requirements for capacity and throughput for the disk pool**

## LAN-free Backup vs. LAN-based Backup

- **Additional License cost**
- **Additional Setup and Administrative task need**
  - Need to handle all Device Pathes
- **May need to invest in FC/SAN Infrastructure**
  - Only 4 Gbit FC available
  - Separate Disk and Tape traffic on the FC HBAs
- **Reduce IOs on backup server**
  - May reduce the need to invest in new backup server hardware
    - But price for backup server hardware is reducing
- **No additional license cost**
- **Easy setup**
- **May need to invest in additional Gbit Ethernet**
  - 10 Gbit available and might be cheaper than 4 Gbit FC/SAN
- **IOs goes through backup server**
  - High IO bandwidth is needed for the backup server
    - Additional Gbit Ethernet and FC Card might be needed
    - Only certain server gives the possibility to increase the IO bandwidth (like pSeries with additional IO Drawers)


**Recommendation: Use LAN-free only for large amount of data and for clients which can deliver data fast enough!**



## Backup Tiers and Architecture: D2D2T

- **For the majority of the backup clients**
- **Consider separate Backup LAN**
  - GigE / 10 GigE
- **DiskBuffer on DS4000/DS8000 with FC Disk**
  - Capacity Size: at least the size of the daily backup data
  - Will free up tape drives during backup window for backup direct to tape
    - Flasch-Copy, LAN-free and D2Tape
  - A VTL may not replace the primary disk pool
    - Or can you handle 200 (virtual) tape drives
- **First create CopyPool after that start migration from disk to tape**
  - Need less tape drives for creation of CopyPool

## Backup Tiers and Architecture: D2D2T



- First create CopyPool
- Then start Migration

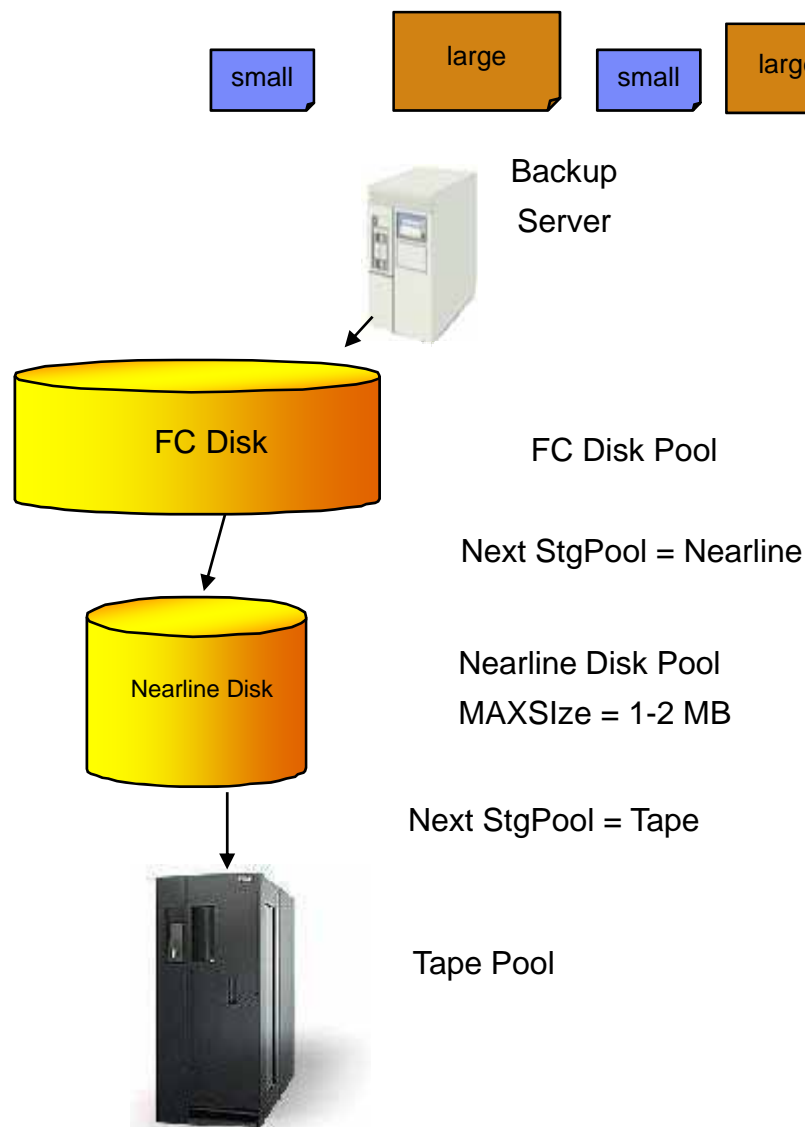
## DS4000/DS8000 as Disk Buffer

- **Use always a different storage system as a Disk Buffer**
  - A shared system would get 3 Workload types
    - Normal Production
    - Backup (reading the production data)
    - Disk Buffer (writing the backup data)
  - Backup and Production Data on the same box!
- **Raid5 is normally sufficient**
  - Backup means Streaming with 256k Blocksizes
- **S-ATA might be sufficient**
  - But only if Backup (write) and Migrating to Tape (read) is not in parallel
    - Migrating should start after the backup
      - Use large enough Capacity
        - > *Consider Data Growth*
  - Gets performance impact during RAID reconfiguration
- **FC Disk with Raid5 is my recommendation**

## Backup Tiers and Architecture: D2D2Nearline

- **For Backup clients with a majority of small Files**
  - Use it for Files smaller than 1 - 2 MB
  - Better Restore performance
  - Reduce Tape usage
    - Reduce Tape Reclamation
- **For Fileserver**
  - Meets Restore SLAs
  - Use Collocate ActiveData (TSM 5.4)
- **Nearline Storage with S-ATA (DS4000, DS8000, TS7500)**
  - Primary Diskbuffer with FC Disk
  - Secondary Diskbuffer with S-ATA
  - CopyPool with S-ATA or with Tape

# Backup Tiers and Architecture: D2D2N

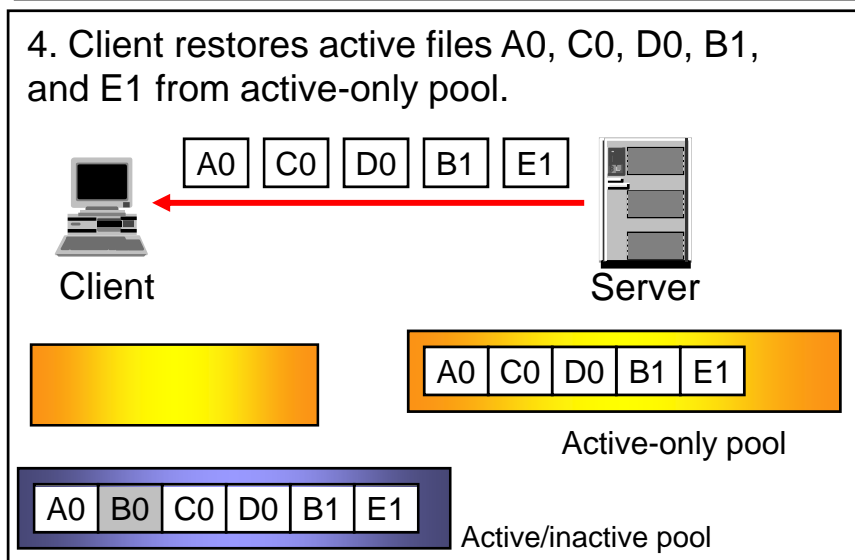
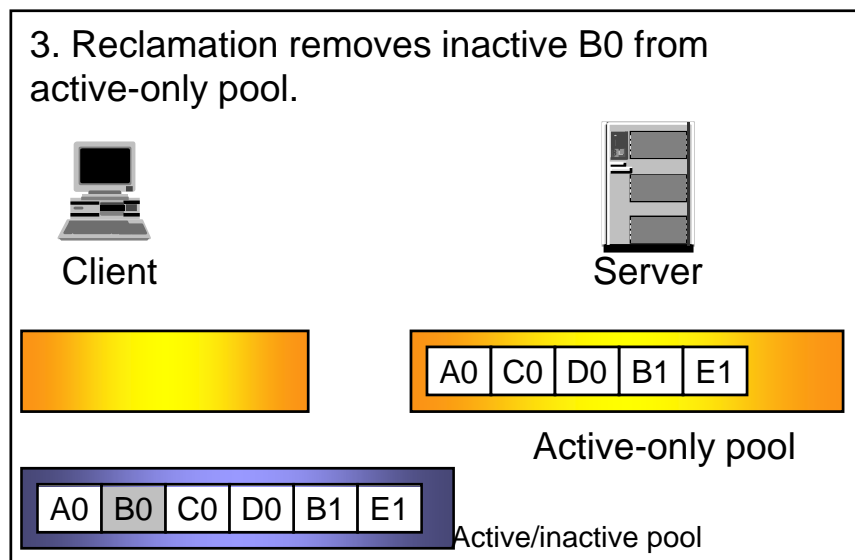
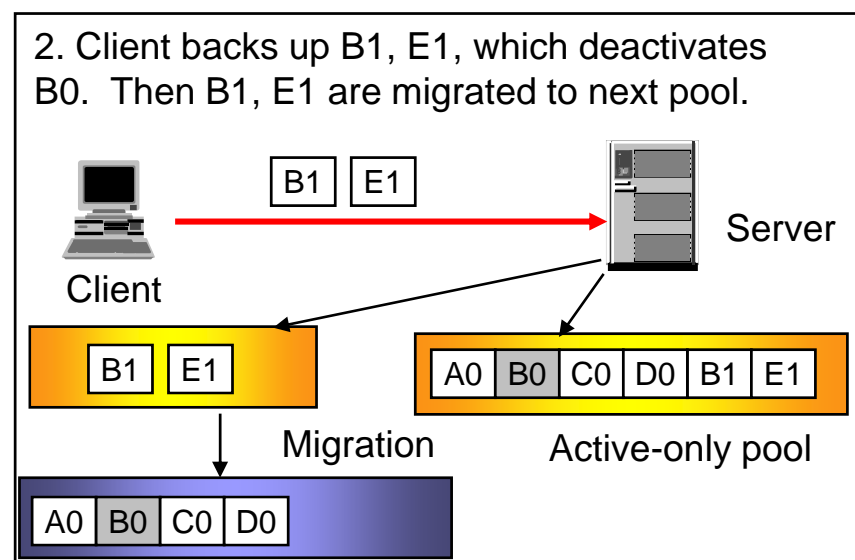
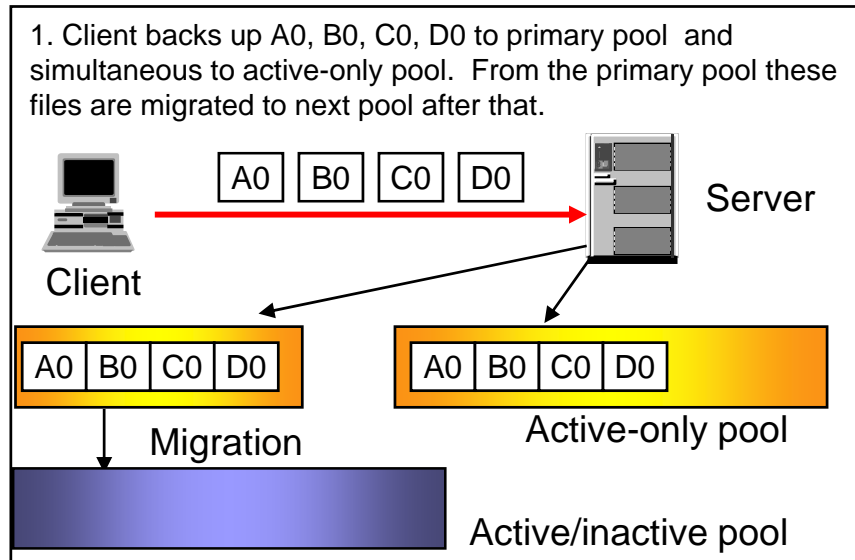


- **Nearline DiskPool for small files**

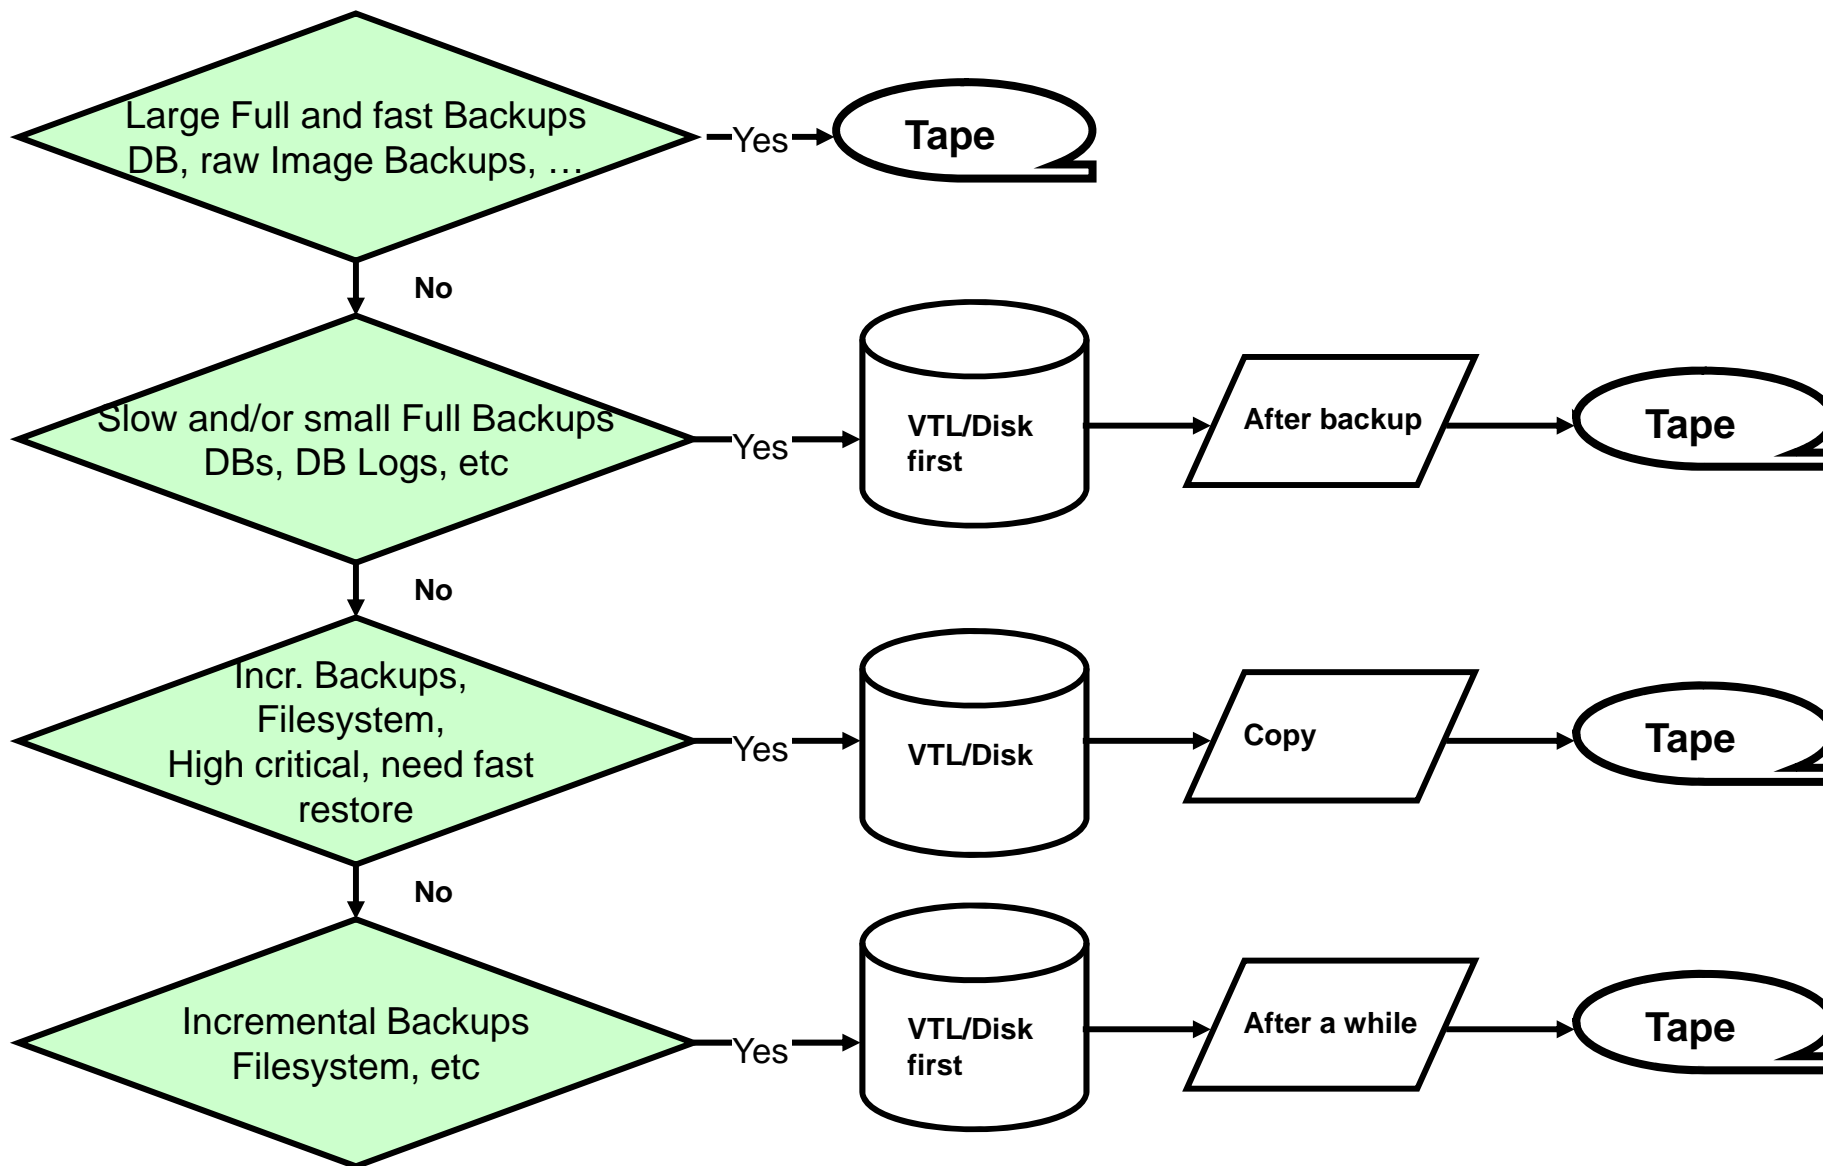
- Nearline Pool stores only Files smaller than 'MAXSize'
- Nearline Pool has as NextPool the Tape Pool

- **During Migration StgPool definition selects file size and store data either on Nearline or on Tape**

# Collocation of Active Data: Example



# Which Data should go to real Tape versus VTL/Disk?



# Agenda

- Backup Overview
- Tape vs Disk
  - TCO
- Backup Tiers and Architecture
  - D2T, D2D2T, D2D2N
  - LAN-free vs LAN-based
- **Virtual Tape Libraries**
- De-Duplication



## Why virtual Tape? What is virtual Tape?

- **Virtualization make sense if physical hardware resource is not fully utilize**

- Tape Cartridge does not filled to 100%
- Tape Drives bandwidth not fully utilize
  - If no D2D2T Backup used
  - Therefore too less tape drives availalbe
- Tape Library Resource are bad utilized
  - If no (or bad) sharing possibilities of the Library

Virtualization does not make sense everywhere!

- **IBM offers different Tape Virtualization Layers:**

- ALMS virtualize Library Slots
- eRMM virtualize Libraries
- B2D and B2D2T (TSM Disk Pool) elimnates bad tape drive utilization and a lack of tape drives
- Virtualization Engines with Tape Drives Emulation
  - VTS – TS7700
  - VTL – TS7500


Virtualization layers does not exclude each other.

Different virtualization layers may complementary

## Why virtual Tape? What is virtual Tape?


- **Virtual Tape Libraries (VTLs) are disk systems that emulate tape drives and libraries**
  - Backup application thinks that it is writing to tape
  - No change in backup application needed
- **VTLs “may” replace native disk pool**
  - VTL solve the problem where backup application does not support native disk pool
  - VTLs may not replace primary TSM Disk Pool
    - At least in larger environments
      - Because of the needed mount points
        - > *E.g. 200 session needs 200 virtual tape drives (= 200 device paths)*
- **VTLs (also called DiskLibraries) are pushed from Disk Vendors, like EMC (CDL) and NetApp to replace tape**
- **There is just one VTL Vendor (FSC CentricStor) which does Tape Stacking for OpenSystems**

# Virtual Tape Library - VTL



## When to use VTLs?

- **Use VTLs instead of native secondary Disk buffer**
  - Don't use VTLs as a replacement of physical Tape
  - If it is a replacement of the primary pool
    - Consider Mountpoints
      - Use small virtual volume size
    - Consider mixed performance (write&read = backup&migration)
- **Use VTLs for LAN-free backup streams which can not utilize native tape**
  - For streams slower than 50-80 MB/sec
- **Still consider different Backup Tiers:**



# VTL (TS7500) vs native Disk Pool

- **LAN-free Backup possible**
- **Compression possible**
  - Use it only for compress able files
  - TS7520 will offer hardware compression
- **Black-Box**
  - No Disk tuning (Raid, Stripping) needed
  - Maintenance done complete by IBM ITS
- **Share Disk Capacity with several Backup Server**
- **Capacity expansion transparent to the ISV host**
  - No LUN Configuration needed by the customer
- **Tape Connection**
  - No additional FC HBAs
  - No Device Driver problem
  - Need to handle tape device paths
    - May not replace primary disk pool
- **Asynchrony Mirroring**
  - With VTL Replication
- **LAN-free Backup only possible with additional software**
- **No Compression**
- **Openness**
  - Disk might be shared with other Applications
    - Not recommended
  - Tuning might be required
  - Maintenance might be done by customer
- **Disk Capacity can not be shared over several Backup Servers**
  - Disk need to be partitioned
  - Or use GPFS
- **Capacity expansion needs additional administrative task**
- **Disk connection**
  - Several FC HBAs are needed
  - Use of different Device Driver
- **Synchrony Mirroring**
  - With LVM or Disk Mirroring

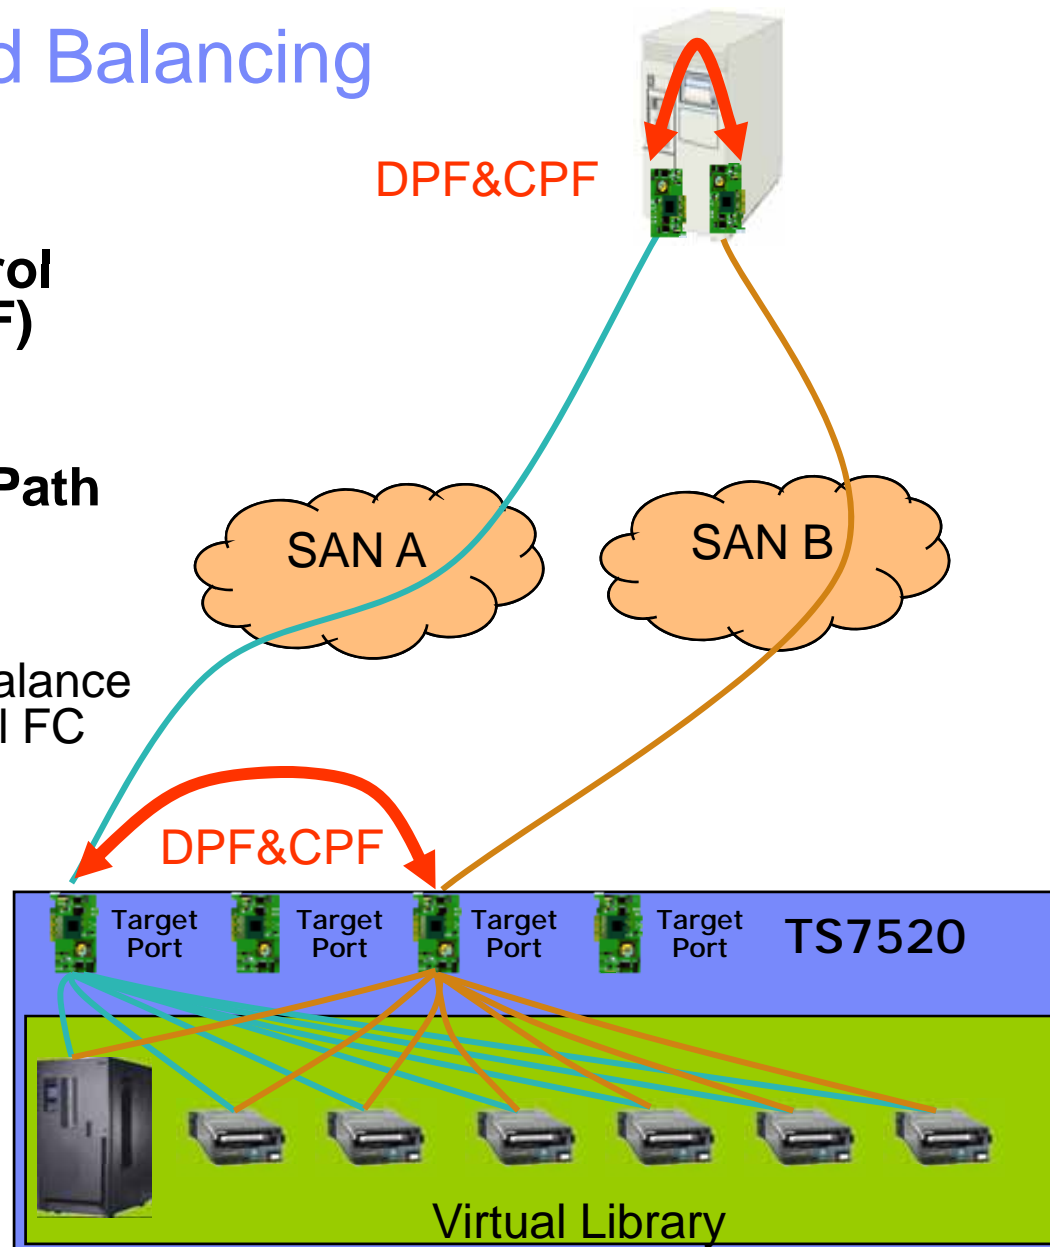
# IBM Virtualization Engine TS7520

- **High Capacity**
  - Up to 884 TB Disk Cache Capacity
  - Up to 1326 TB in 02H07
- **High Performance – Best on the Market!**
  - Up to 4 800 MB/sec Throughput
  - 4 Gbit FC Connections - **Front-end and Back-end!**
    - Up to 32 x 4 Gbit FC for Front-end
    - Up to 16 x 4 Gbit FC for Tape Back-end
    - Up to 16 x 4 Gbit for Disk Cache
  - Load balancing for **FC Paths (DPF)**
- **Functions**
  - Up to 4096 virtual tape drives, up to 512 virtual libraries
  - Up to 256 000 virtual Volumes (Volume Size 1 – 700 GB)
  - Encryption
  - Supports all IBM Tape Drives and Libraries in the Back-end
    - **Policy based migration**
      - **Pre-migration** (data on disk and tape)
        - several policies possible (age, time, end-of backup/demount, cache usage)
      - **Direct access to migrated data** (no recall is needed)
  - Hardware assist Compression
- **High Availability**
  - Up to 4 Cluster Servers
    - Cluster Failover
  - **FC Path Failover**
    - For Library Controll (CPF)
    - For Tape Drive Pathes (DPF)
  - **Fully redundant Power**
  - **Call Home**
- **Scalability**
  - **In Capacity**
    - From 6,5 TB to 884 TB
    - 6,5 TB steps
  - **In Performance**
    - From 500 MB/sec to 4 800 MB/sec
    - 1 - 4 Cluster Server
    - 1 - 8 FC Ports
    - Additional CPUs and Memory



## Path Failover and Load Balancing

- **Unique IBM feature**
- **Failover for Library Control (CPF) and Data Path (DPF)**
  
- **Load balancing for Data Path (DPF)**
  - Is a must have function!
  - There is no other way to balance your workload over several FC HBAs.



# TS7520 Functions

- **Enhanced Tape Caching**
  - Provides the backup application direct access to data whether in cache or on physical tape and a more intelligent policy based data migration
    - **Policy based migration**
      - **Pre-migration** (data on disk and tape)
      - several policies possible (age, time, end-of backup/demount, cache usage)
      - Creates a 1 to 1 copy
      - Recommendation: Migrate only if cartridge is full!
        - > *Because append writes to migrated cartridges goes to physical tape*
  - **Direct access to physical tape for migrated data** (no recall to cache is needed)
  - Scratch mount always mounts the physical tape cartridge first (in order to do the label verification), then writes to disk cache
  - Use frontend compression
    - Consider performance impact
    - Without compression, consider more physical tape cartridges
- **Encryption for Backend Tape Drives**
  - Use software Encryption
  - Or use Hardware Encryption from IBM Tape Drives
    - IBM TS1120 or LTO 4
    - No Performance Impacts!
- **Path Failover (CPF & DPF)**
- **iSCSI**
- **Call Home to IBM Support Center (post GA)**



## Capacity Sizing for Tape Caching

- **Number of parallel jobs**
- **Number of filling cartridges**
- **Keep data in disk for restore**
  
- **Example with LTO3 in the backend**
  - Virtual volume size is 400 GB

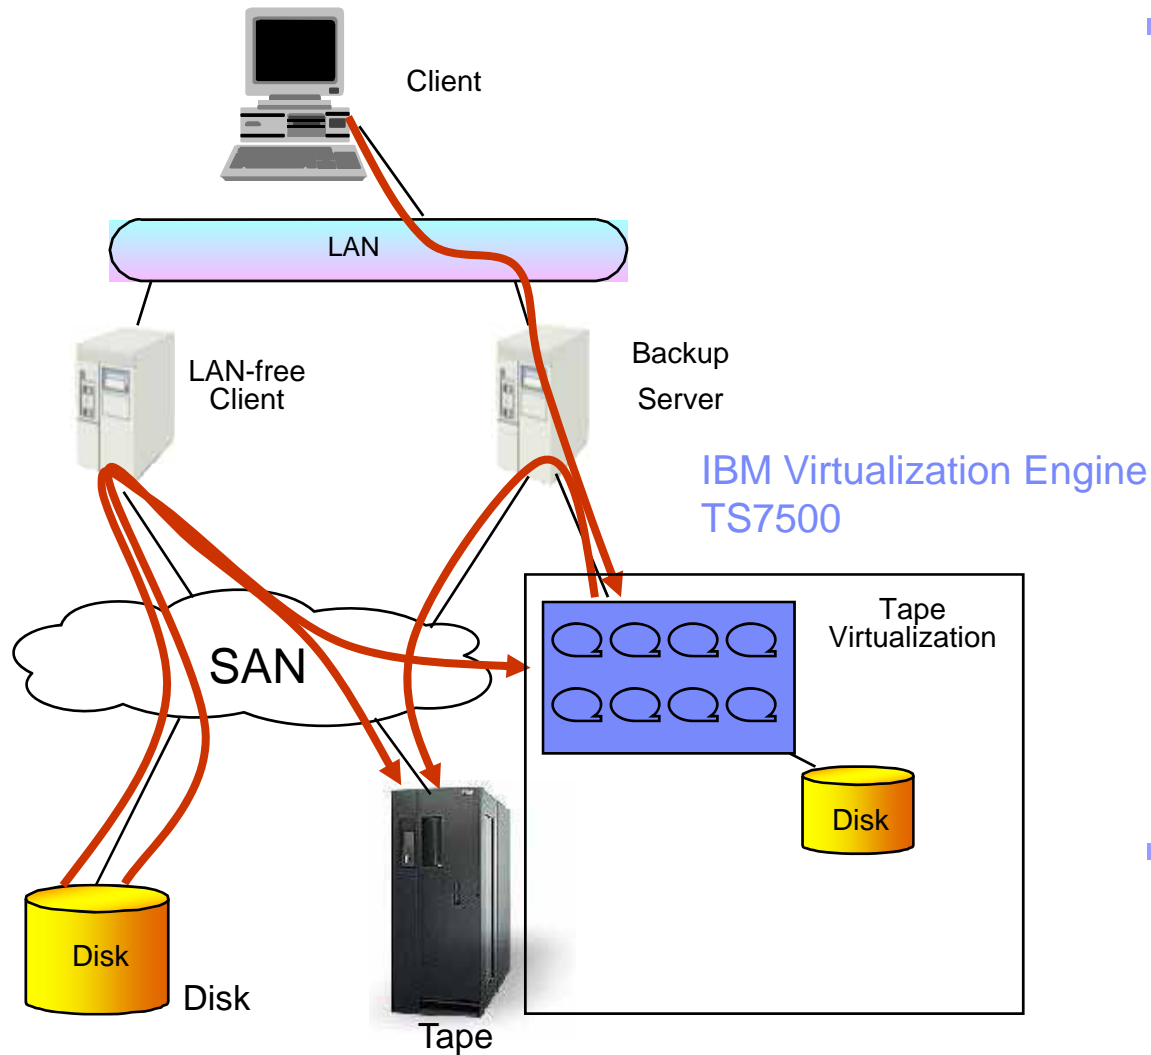
	Amount	Volume size	
Number of parallel jobs	100	400 GB	40 TB
Number of filling cartridges	40	400 GB	16 TB
Restore data			20 TB
Sum			76 TB

## How to Implement VTL and VTS?

- Migration done by the backup server
  - Backup software must handle the migration
    - No big task for TSM
    - Symantec(Veritas NBU 6 can now also “stage” from VTL
    - EMC/Legato Networker needs to be scripted
  - Virtual and physical volumes can be on different size and type
  - Physical library and tape drives must be only supported on the Backup software
- 
- VTL migrates data
    - Creates a 1:1 copy from the virtual volume to the physical volume
    - Virtual volume must be the same size and type as the physical volume
      - If you want to run 100 jobs in parallel, then the cache must large enough to hold at least 100 virtual volumes
        - > *E.g. using TS1120 with JB Media (700GB), and want to run 100 jobs in parallel, then the cache must be at least 100 x 700 GB = 70 TB*
    - Backend Library and Drives must be supported by the VTL

# D2D2T Backup with VTL / TS7500

## Migration done by the Backup Server



### ■ Advantage:

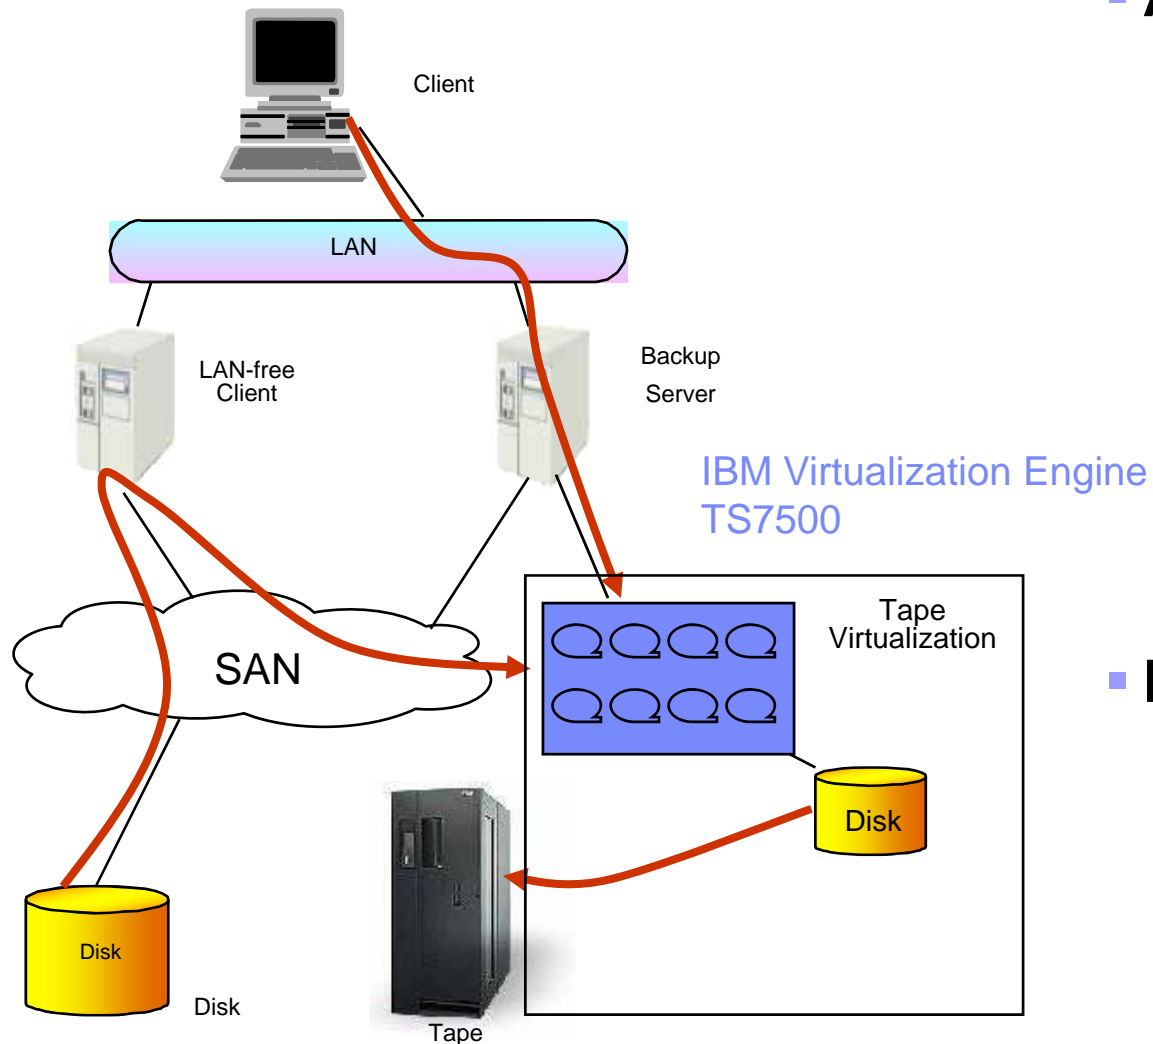
- LAN-free Backup to Disk possible
- Volume Size can be different from Physical Volume size
  - Migration can be handled better
- Physical Library & Drives still direct useable
  - Restore of migrated data is done directly from Tape
  - Backup/Restore direct to Tape

### ■ Disadvantage:

- Transfer is done by the Backup Server

# D2D2T Backup with VTL / TS7500

## Migration done by VTL



### Advantage:

- Transfer is done by the VTL / TS7500
  - Migration does not need any Backup Server resources (DB, etc.)
    - Advantage for many small files
- LAN-free Backup to Disk possible
- 1:1 Copy of virtual volume to physical volume
  - Data native readable

### Disadvantage:

- Virtual volume size and type must be the same as the physical volume size
- No direct Backup to tape possible
  - Or physical Library must be partitioned
- Needs maybe more disk capacity

## Can a primary disk pool replaced with TS7520?

- **It depends!**
- **Consider:**
  - Number of needed mountpoints
    - If you currently run for instance 200 parallel backup jobs (session), then you need with a VTL 200 virtual tape drives.
      - It is not a prbl of the VTL
      - But may you able to handle 200 tape drives on your server?
      - If you use x lan-free clients, then you need to handle x times 200 tape device pathes on your server. Lets say if you have 5 lan-free clients, then you get 1 000 tape device pathes on TSM. Can you handle this?
        - > *Consider the use of eRMM*
  - If you want to use Tape Caching, then consider the VTL cache capacity
    - For tape caching virtual and physical volume size must be same
      - E.g. for LTO4 you need 800 GB virtual volume size
      - If you need to run 200 session in parallel, then your VTL cache size must be  $200 \times 800 \text{ GB} = \mathbf{160 \text{ TB!!!}}$
  - If you want to use Tape Caching, then consider Scratch Mounts
    - Scratch Mounts on a VTL means, that the migrated cartridge must be mount a physical tape drive first!
    - If you want to start 200 jobs on the same time, then only the first x jobs (x stands for number of physical tape drives connected to the VTL) may start immediately
- **On a medium to large TSM environment I don't recommend to replace the primary disk buffer with a VTL!**


## Why not using IBM VTS (or other Tape Stacking VTLs) for Open Systems

### ■ IBM VTS offers for zOS

- Tape Stacking
  - Because of one DataSet per Tape Volume
- Virtual Tape Drives with Disk Buffer
- Automatic Migration from Disk to Tape and back from Tape to Disk

### ■ Open Systems needs virtual Tape Drives with Disk Buffer, but no Tape Stacking

- All Open Systems ISV fully utilize tape cartridges
- Disadvantages of Tape Stacking:
  - With Tape Stacking only the VTL can read the physical cartridge
  - If data are needed for restore (reading) from volume which are already migrated to physical tape, then the total logical volume must be recalled to Tape Volume Cache before any data transfer to host can take place
    - No big problem with small volume (800 MB)
    - But it takes hours for large volumes (e.g LTO1 1-2h)



## Disadvantage of Tape Stacking for TSM

- **Virtual Volume size between 0,8 – 5 GB**
  - Many more volumes must be managed with TSM (up to 500 times more Volumes as with TS1120 native)
  - Restore performance will be degraded, because virtual volume must be recalled into the cache before application/server can restore any data
  - Reclamation needs to be turned off, because of the heavily workload on the backend drives
    - More physical tape capacity is needed if reclamation needs to be turned off
  - Collocation does not give any advantages if volume size is small and virtual volumes are stacked on different physical cartridges
- **Fragmentation on the backend (physical) tape cartridges reduce the usable tape capacity**
  - **Average loss about 50 %**
    - More physical tape capacity is needed
  - Reorganization / Reclamation reduce the loss, but needs additional physical tape drive resources
- **Data can be read only with VTL**
  - HA solution necessary to guarantee the access to the data
- **No direct backup to tape possible**
  - Or Tape Library needs to be partitioned

# Agenda

- **Backup Overview**
- **Tape vs Disk**
  - TCO
- **Backup Tiers and Architecture**
  - D2T, D2D2T, D2D2N
  - LAN-free vs LAN-based
- **Virtual Tape Libraries**
- **De-Duplication**



# Data Reduction

- **Also called De-Duplication, de-dup, single instance store**
  - De-Dup only for Disk!
- **Eliminates redundant attachments in emails**
- **Eliminate redundant unchanged files**
- **Different implementation**
  - Byte-level delta reduction
  - File-level reduction
- **When reduction is performed**
  - Immediately
    - Need sufficient processing power
  - Post processing
    - Need more capacity
  - How much capacity can be managed?
    - Dilligent may manage up to 30 TB if server runs with 8 GB RAM
  - How is data protected
    - How are the pointers protected

## Why using De-duplication

- **On Mail- and Fileserver there could be many duplicated files and mails**
  - If some send out a mail with an attachment to x persons then this mail is stored x times and also the attachment is stored x times
  - Space will be wasted because of duplicated files and mails
- **Also during backup space will be wasted if backup software does not recognize duplicated files**
- **Backup process will save the same file several times, even the file was not modified, depending on the backup software**
- **But Mail Archive Solution (like IBM Common Store) already reduce the space and duplicated files (Singel Instance Store)**

## IBM's view of Current De-duplication offerings in the Marketplace

### ■ Appliance-based Products

#### – Common characteristics

- Generally affords storage efficiency only for disk – exported physical tapes are not de-duped
- Duplicate recognition may be meta-data based or purely algorithmic, or both
- De-duplication may be inline or run as 2<sup>nd</sup> pass, independent process
- Some products assume equal tokens always imply identical data – used as basis for reducing network traffic;
  - Some products do byte-wise compare of data on token collision
- Inter-appliance replication can leverage shared token information to reduce replication network demands
- Storage allocation may be token driven & highly distributed or sequentially optimized

#### – Products with unique client software (Avamar, Centera)

- Typically requires proprietary API between client and appliance
- Client can run token generating algorithm (often crypto hash) or optionally the appliance

#### – Products without unique client software (Diligent, Data Domain, Sepaton, FalconStor)

- Usual environment is VTL or NAS backup appliance
- Appliance generates the token

### ■ Application-based Products

- Best at recognizing duplicate objects created by the application
- Generally limited to the application
- Generally limited to single server's data stream
- Examples: Exchange & Domino Single Instance Repository

## User experience with current de-duplication

- **Current de-duplication products are limited in**
  - Capacity
    - Just a few TB possible
  - Performance
    - Maybe up to 200 MB/sec depending on system and data
  - Availability
    - Just one server
    - Or no cluster failover
  - Scalability
- **De-duplication may not work if**
  - Data are already compressed by the client
  - Data are encrypted by the client or backup server
  - If multiplexing is used
  - If zero compression is used (for DBs)
  - Mail archiving is used
- **De-duplication should not be used for data which needs high restore performance, such as DBs**

## Question about De-dup

- **What happens if the „Hash“ DB gets corrupted?**
  - All your Backup Data are lost!
- **The “Hash” DB is normally stored on the memory of the server because of performance**
  - Not a large amount of data can be managed
  - DB can be gets corrupted
- **What happens if one block of the very first version gets corrupted?**
  - Then all other version can be used any more
  - Normal backups stores normally different version

# Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml): AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/30, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation  
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries  
LINUX is a registered trademark of Linux Torvalds  
UNIX is a registered trademark of The Open Group in the United States and other countries.  
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.  
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.  
Intel is a registered trademark of Intel Corporation  
\* All other products may be trademarks or registered trademarks of their respective companies.

## NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.