

Replicated De-Duplication Versus Physical Copies – You DON'T Have to Make a Choice: Ultimate Flexibility with CA Vtape

Russell Witt
CA Technologies

Friday, August 12th, 8:00
Session Number 10059

What type of Tape Data Sets are used?



- Data to be shared with another company/vendor
- Backups
 - Disaster Recovery
 - Historical
- Archival
- Log's
 - For recovery
 - For Historical reporting
- Primary application data

Data to be shared with another company/vendor



- Historically one of the primary reasons for tape
- Seldom used with the acceptance of secure electronic transmission
- Product installations via electronic delivery
- Cross platform
 - No longer compatible – too many types of tapes
 - Easier to setup electronic transfer

Backups



- Disaster Recovery
 - Shorter retention
 - Obsolete within 2-4 weeks for DR purposes
 - One copy at least is required at the offsite location
 - Seldom used after creation
- Historical
 - Governmental regulations
 - Financial regulations
 - Retention often measured in years and decades
 - Seldom used after creation

Archival



- HSM L2
- Retention is normally years
- Read very often initially, less often as the data ages
- Often sent offsite for use at DR as well

Log's



- Recovery
 - DB2, IMS
 - Short Retention
 - Seldom used
- For Historical Reporting
 - Long retention
 - SMF Data
 - Used often initially

Primary Application Data



- Not used as often, replaced with larger capacity disk systems
- Retention varies considerably
- Usage Varies

Options available



- Standard Tape Cartridges
- Virtual Tape backed up to physical tape(s)
- Virtual Tape without backup
- Virtual Tape with replication
- Virtual Tape with de-duplication replication

Standard Tape Cartridges



- Capacity
 - Oracle/STK 5Tb/cartridge with the latest
 - IBM is 4Tb/cartridge
- Shelf life now measured in decades, but still not “forever”
- Cost per Megabyte is lowest of any storage product, but when underutilized can be expensive

Virtual Tape backed up to Physical Tape(s)



- Virtual Tape gives faster performance for original writing
- Faster access initially after creation
- Allows for more efficient usage of higher capacity physical cartridges
- Capacity with larger capacity cartridges is virtually unlimited (5,000 high capacity cartridges equals 100Pb of compressed data)

Virtual Tape without backup



- Capacity is also measured in PetaBytes (8 Pb would be a very large Cache box)
- On a megabyte basis, more expensive than cartridges
- Single point of failure
- No DR capability
- Fast access at all times, even for the oldest files

Virtual Tape with replication



- Capacity is also measured in PetaBytes
- Double the cost of Virtual Tape without replication PLUS the additional cost of bandwidth
- No single point of failure
- DR capability
- Fast access at all times, even for the oldest files

Virtual Tape with de-duplication replication

- Capacity is also measured in PetaBytes; possibly another 4-5 times more capacity than without de-duplication which also means less bandwidth
- No single point of failure
- DR capability
- Fast access at all times, even for the oldest files

What type of Tape Data Sets are used?

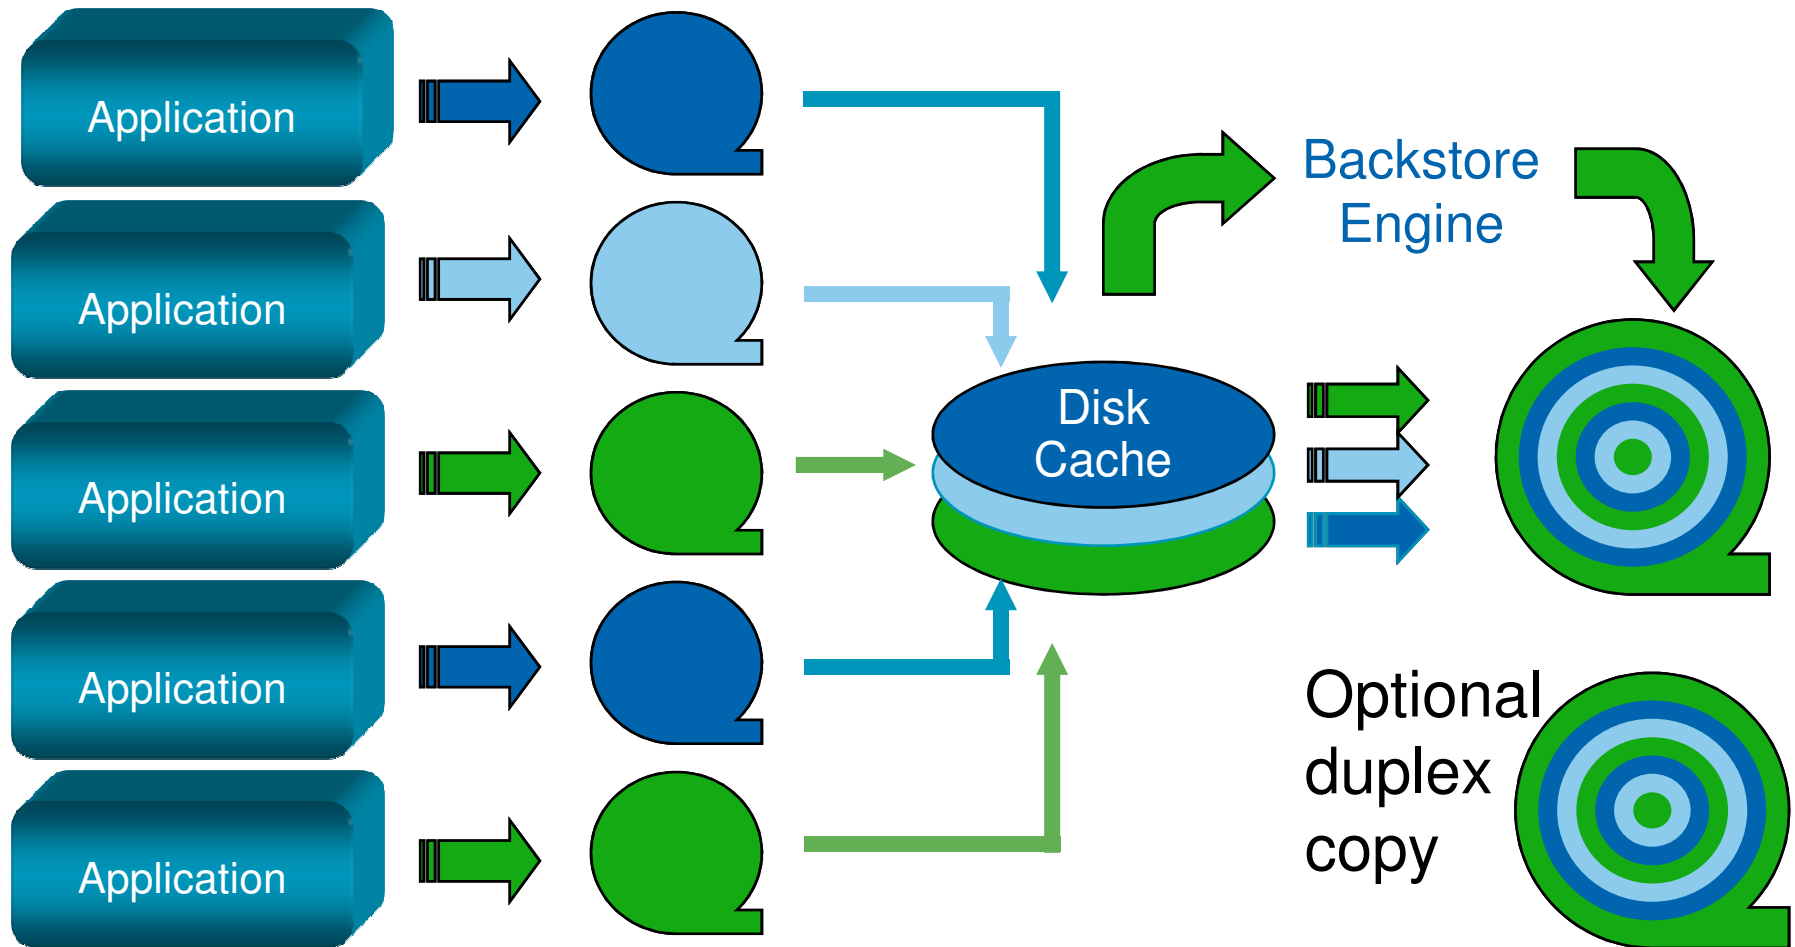


- Data to be shared with another company/vendor
- Backups
 - Disaster Recovery
 - Historical
- Archival
- Log's
 - For recovery
 - For Historical reporting
- Primary application data

CA Vtape Basic Operation

- Application data from any z/OS application is written to standard DASD cache
- Data can be compressed (zIIP exploitation) when written to cache to reduce DASD usage
- Data can be grouped together and each group handled differently (how long to remain resident in cache, how long after dismount before backing up to physical tape, should one or two physical copies be made)
- Encryption can be done with CA Tape Encryption (also zIIP exploitation) when written to CACHE or when written to physical cartridges – or device encryption can be used
- CACHE can be erased when released


CA Vtape basic overview



Peer to Peer Option

- Employing TCP/IP instead of expensive channel extenders, a peer-to-peer option provides real-time tape mirroring and immediate sharing of Virtual Volumes to disaster recovery or remote sites.
- The size and number of TCP/IP links is handled by z/OS
- Transmission can be deferred to times when links are less utilized
- Lets IT organizations centralize physical tape processing and support remote sites that don't use physical tape


Peer-to-Peer option




New NFS options

- Can be done with or without physical copies being created (no physical tape, primary only, both primary and duplex)
- Can be replicated only (with NetApp as an example)
- Can be done with replicated de-duplication (Falconstor or EMC's Data Domain)
- Any NFS attached file server you choose supported by z/OS
- Because the backstore operation is done much later than when the Virtual Volume was written, the bandwidth requirements can be much more balanced.


New NFS option



Primary
copy -
optional



Duplex
copy -
optional



NFS Copy replicated to cold-site

