IBM System Storage)

News of Storage and Tape Technologies

ReThink your business

ReStore your data

Josef Weingand

+49 171 5526783 - weingand@de.ibm.com

Infos / Find me on:

http://sepp4backup.blogspot.de/ https://www.xing.com/profile/Josef_Weingand

http://www.linkedin.com/pub/josef-weingand/2/788/300

http://www.facebook.com/josef.weingand http://de.slideshare.net/JosefWeingand

https://www.xing.com/net/ibmdataprotection

IBM Smarter Storage is an approach for the design and deployment of storage

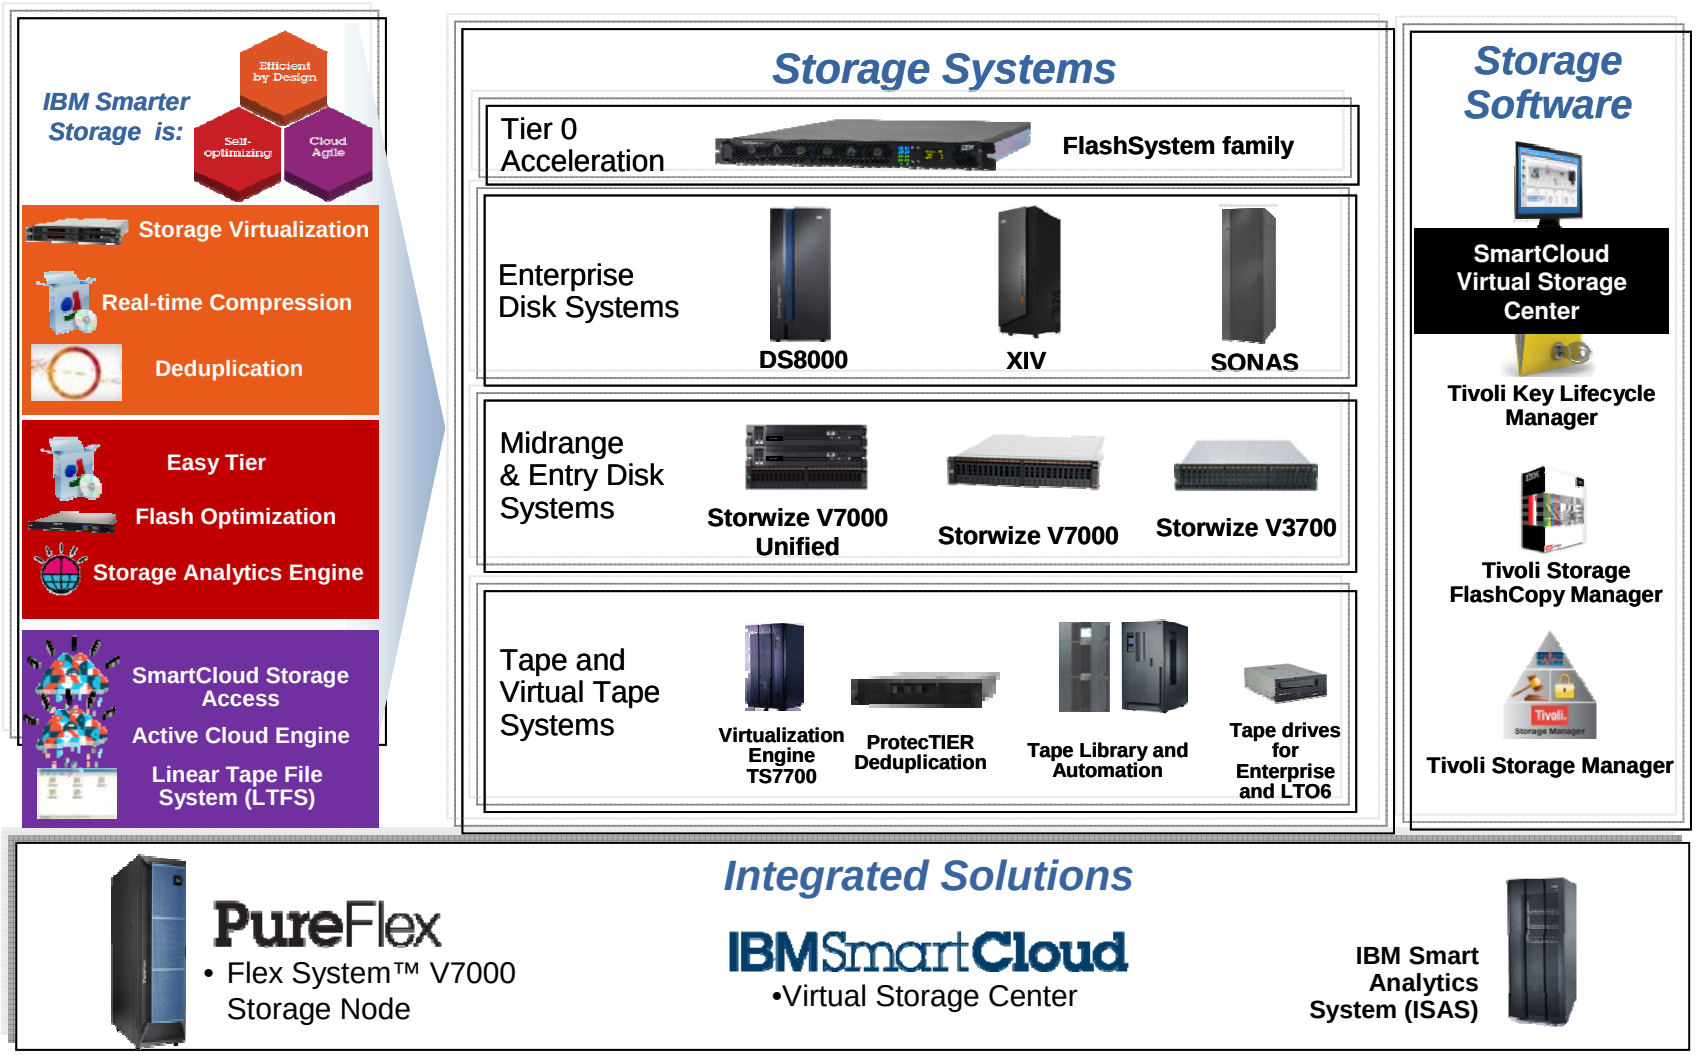

What's New?

Release 7.1 delivers more flash optimization

- Easy Tier 5th generation enhancements for more flash optimization
- New all-flash feature code
- More integration with System z servers
- More integration with IBM Power servers
- More VMware and Windows support
- Support for Symantec space reclamation
- SNIA Emerald™ Power Efficiency Measurement results
- More…

Flash storage removes the historical bottleneck in the data path

In the last 10 years, most elements in the data path have boosted performance substantially... except storage

CPU Performance **8 - 10x**

DRAM Speed **7- 9x**

Network Speed **100x**

Bus Speed **20x**

Disk Speed **1.2X**

…and everything waits

Source: IBM and Industry Estimates

Flash storage delivers more than just faster performance

Improve Performance

Reduce Costs

Enable New Opportunities

- ☑ Accelerate Application Performance
- ☑ Save Energy / Cooling / Floor Space
- ☑ Gain Greater System Utilization
- ☑ Lower Software & Hardware Cost

DS8000 flash evolution

Combining flash optimization plus high-end capabilities

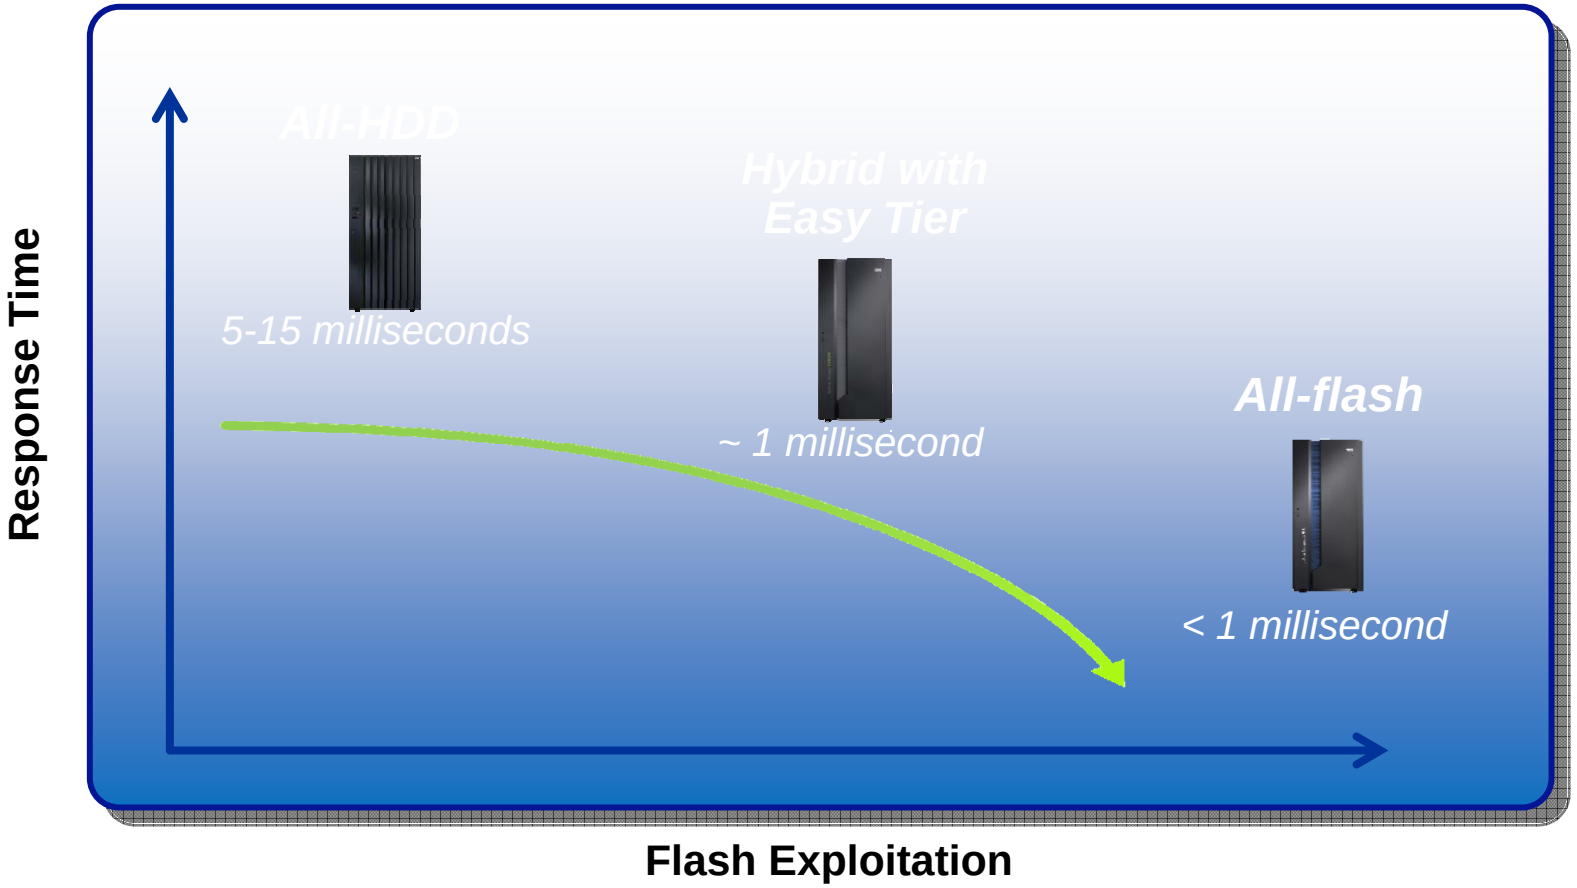

IBM Easy Tier optimizes flash

Smarter storage over 5 generations

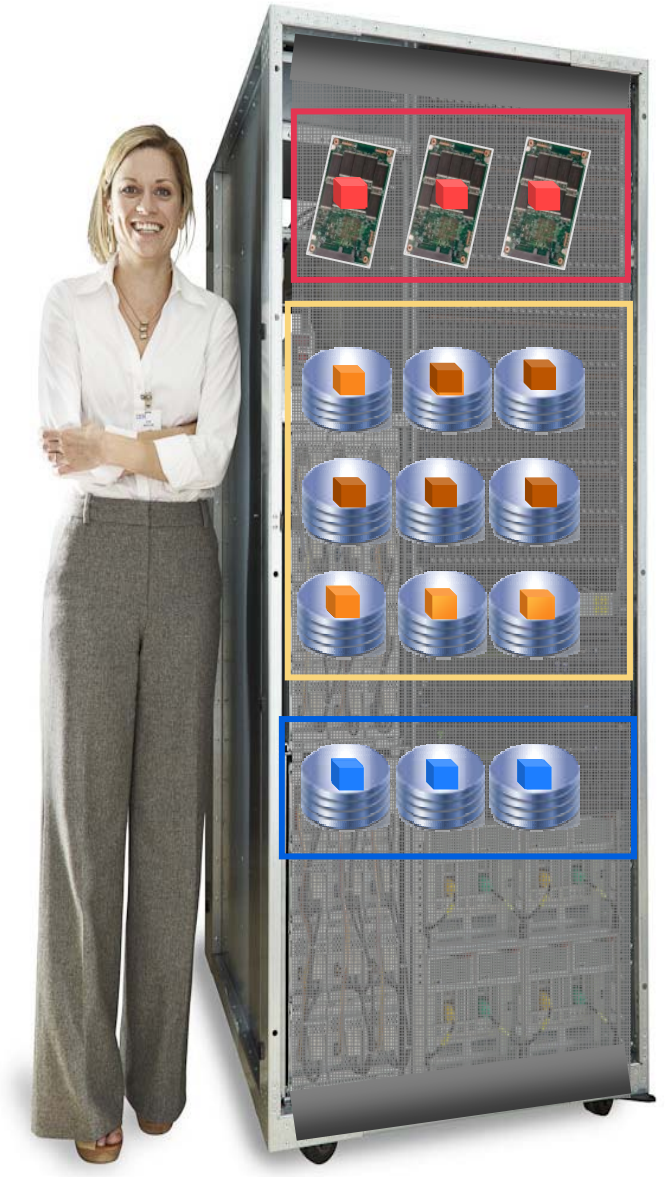

Easy Tier Heat Map Transfer

- **Coordinates heat maps across mirrored systems**
- **Client benefits**
  - Aligns data placement on remote system with that of primary system, ensuring optimized performance for mirrored systems
  - Another business continuity advantage
- **Scheduling and monitoring done through Heat Map Transfer Utility**
  - Runs on Windows, Linux (SLES11+, RHEL 4+)
  - Can also be done with TPC for Replication

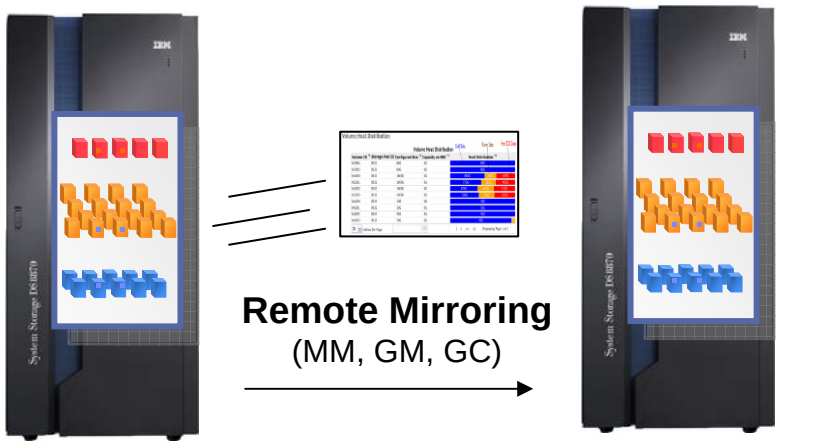

All-flash benefits for transactional (OLTP) workload

Same usable capacity but with…

70%
reduction in response time

80%
reduction in drive count

41%
reduction in raw capacity

Comparing all-flash DS8870 with all-HDD system boost performance and reduces costs with equivalent $/GB

62%
reduction in energy usage

33%
reduction in floor space

Additional flash news

- 40% list price reduction on all DS8870 SSDs
- Feature Code indicator: FC 0600
- Can add HDDs later via RPQ
- Additional performance data will be available after GA

IBM DP&R Storage / DataProtection Overview

IBM Tivoli Storage Manager IBM FlashCopy Manager TSM for VE IBM FastBack

IBM Tape File System/LTFS LTFS Storage Manager LTFS Enterprise Edition – GPFS&LTFS

TSLM

LTO 4
- 800 GB
- 120 MB/sec

LTO 5
- 1500 GB
- 140 MB/sec

New!

LTO 6
- **2500 GB**
- **160 MB/sec**

TS1130
- 1 000 GB
- 160 MB/sec

TS1140
- 4 TB
- 250 MB/sec

Media

Tape Drives

1 HH

2/4 HH
1/2 FH

18 FH

192

TS2900
- 9 Slots

TS31/3200
- 22 / 44 Slots

TS3310
- 30 - 408 Slots

TS3500
- 58 – 20 087 Slots
- Shuttle: 15x20 087 = 300 000 Slots

IBM ReadVerify Appliance (RVA)

Tape Libraries

Neu!

Neu!

TS7620
- SMB Appliance
- 300 MB/s
- 6 / 12 / 25 / 36 TB native

Neu!

TS7650G
- DeDup Gateway
- Cluster
- 3200 MB/sec

D2D
- DS3500, V3700, V7000, XIV, DS8000

TS7700
- VTS für zOS

Tape Virtualization & D2D

LTO – more than 10 years…

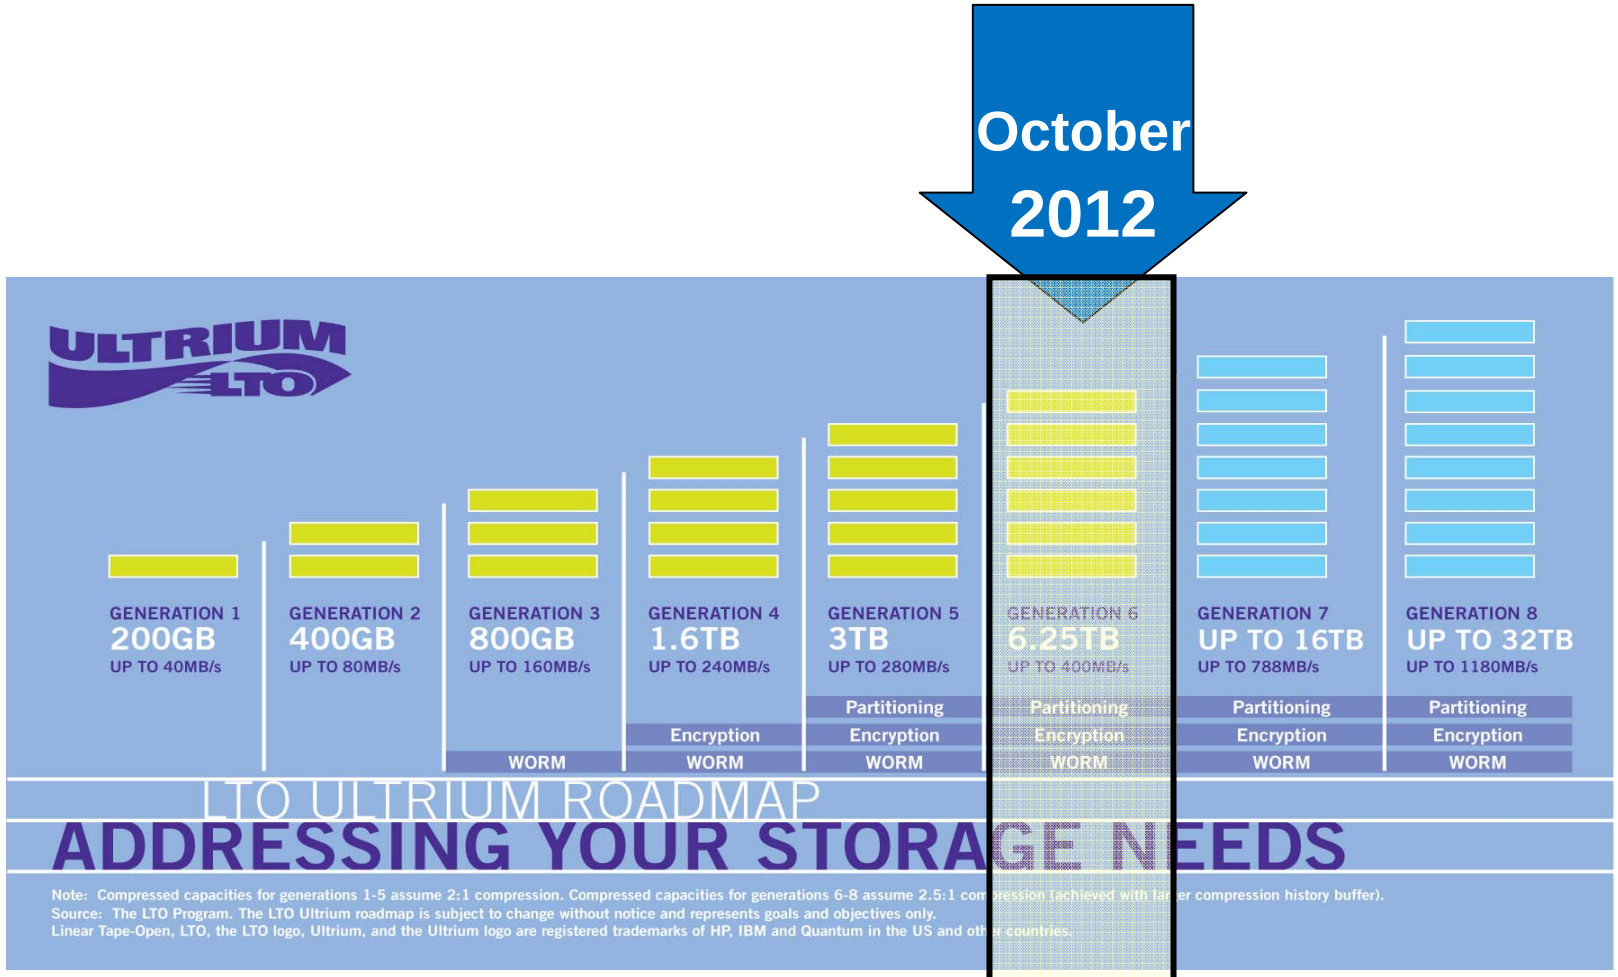

IBM LTO Generation 6

- **New LTO 6 Drive - faster and bigger!**
- Available since October 2012
- 2,5 TB native Capacity
  - 6,25 TB with 2,5:1 Compression
    - Compression improved!
  - 160 MB/s native transfer rate
    - With 2,5:1 Compression 400 MB/sec – **max 645 MB/sec**
    - Variable Speed **40 – 160 MB/sec**
      - 14 Speed Matching steps
      - reduce Start/Stop problem
    - Neu: **1024 MB Data Buffer**
  - Compression improved 2,5 : 1
  - Dual Ported 8 Gbit FC or 6 Gbit SAS
    - Failover & Loadbalancing with IBM Tape Device Driver
  - Read/write of LTO5 cartridges, read LTO4 cartridge
  - Less power consumption
  - Tape File System (LTFS), WORM, Encryption

IBM Barium Ferrite Media

- Higher density of bits require:
  - Smaller magnetic particles
  - Smoother media (so you can get the head closer to the magnetic particles)
  - Better, higher performance heads
- **Barium Ferrite particles (BaFe)**
  - LTO 6 first LTO media to use BaFe
  - Ultra fine, small particles
  - Improved head life
  - Higher signal-to-noise ratios
  - Not all LTO 6 brands use BaFe

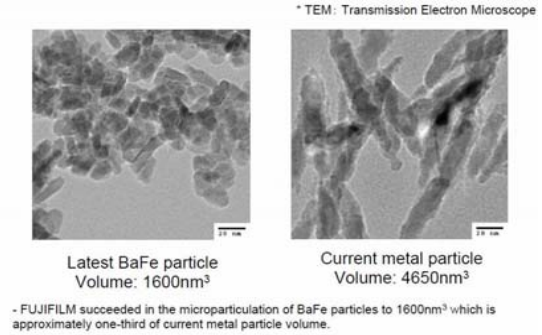

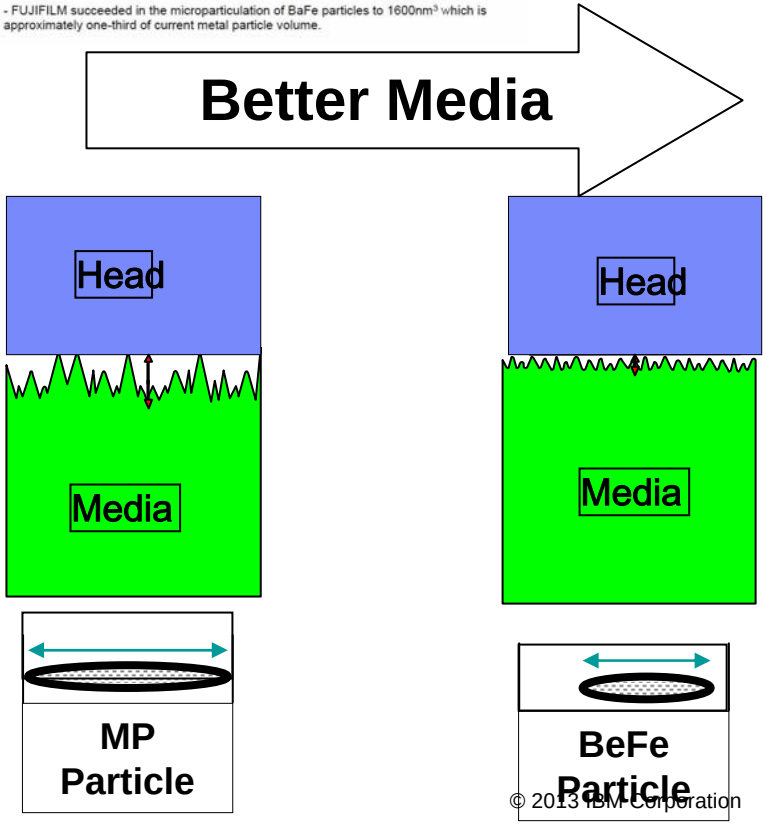

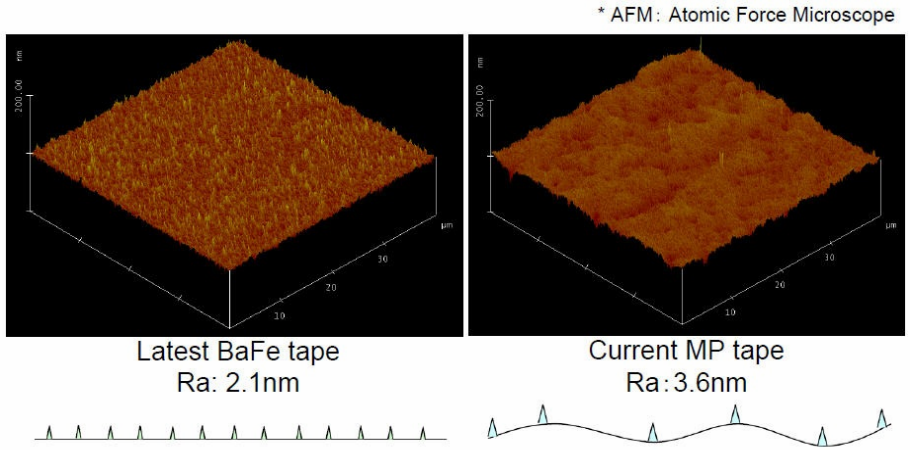

Advantages of IBM LTO6 Tape Drives

- dual-ported 8 Gbit native FC Connection with Failover and Loadbalancing
- GMR Head (Giant Magneto Resistence) with aluminum surface
  - Longer life time
- Speed Matching (40 - 160 MB/s) and 1024 MB Data Buffer
- Fast locate speed
  - 10 % faster than other LTO Drives
  - Faster read/restor, reclaimation

IBM Tape Drive Roadmaps

LTO Generations	LTO-3	LTO-4	LTO-5	LTO-6*	LTO-7	LTO-8
Max Native Capacity	400 GB (L3)	800 GB (L4)	1.5 TB (L5)	2.5 TB (L6)	Up to 6.4 TB (L7)	Up to 12.8 TB (L8)
Other Native Capacities	200 GB L2 100 GB L1 R/O	400 GB L3 200 GB L2 R/O	800 GB L4 400 GB L3 R/O	1.5 TB L5 800 GB L4 R/O	2.5 TB L6 1.5 TB L5 R/O	6.4 TB L7 2.5 TB L6 R/O
Native Data Rate	80 MB/s	120 MB/s	140 MB/s	160 MB/s	Up to 315 MB/s	Up to 472 MB/s

* Data Compression engine enhancement from 2:1 to 2.5:1

TS1100 Generations	Gen-1 3592 J1A	Gen-2 TS1120	Gen-3 TS1130	Gen-4 TS1140	Gen-5	Gen-6
Max Native Capacity	300 GB (JA)	700 GB (JB)	1.0 TB (JB)	4.0 TB (JC)	8-10 TB (JD)	14-20 TB (JD)
Other Native Capacities with Media Reuse		500 GB JA	640 GB JA	1.6 TB JB 640 GB JA R/O	6-8 TB JC	6-8 TB JC
Native Data Rate	40 MB/s	100 MB/s	160 MB/s	250 MB/s	Up to 360 MB/s	Up to 540 MB/s

TS1140 Tape Drive Overview

- **4rd Generation of 3592 enterprise tape drive**
  - Barium Ferrite media (JD) with **4 TB native capacity**
  - **Media Re-Use with JB**
    - Read/Write JB: 700 GB -> 1 TB -> **1,6 TB**
    - Read JA
  - **250 MBps** native drive data rate
  - **Dual 8Gb fiber channel** interfaces with 650 MB/s max compressed data rate
- GMR R/W Head with aluminum surface
  - 32 parallel Tracks (Read/Write)
- Improved Performance
  - 1024 MB Buffer
  - "Virtuelles Back Hitch„
  - Speed Matching with 14 steps
- Fastest Data Access
  - High Resolution Tape Directory
  - Fastest locate Speed with 12,4 m/sec
  - **Read Ahead Feature**
  - **JC Media length: 880m**

- **LTFS Support: Single Drive, Library and with GPFS**
- MES upgrade for TS1130 available (Model Conversion)

TS1140 Drive Format

- TS1140 can read JA Media with Generation 1
- TS1140 read/writes JB & JC Media

- Easy and smooth migration!

	Generation 1 Format	Generation 2 Format	Generation 3 Format	Generation 4 Format
JA Media	Read only 300 GB	Read only 500 GB	Read only 640 GB	
JB		Read only 700 GB	Read/Write 1 TB	Read/Write 1,6 TB
JC				Read/Write 4 TB

IBM Linear Tape File System (LTFS)

Self-describing tape format for archiving data to tape

Run applications designed for disk files on tape

Improves efficiency, simplifies direct access and management of files on tape

Four offerings:

Standalone Drive Edition (SDE)

No charge software allows tape cartridges to managed in the same way as USB storage devices

Library Edition (LE)

Manages tape drive mounts and unmounts for access to data in IBM tape libraries

Storage Manager (SM)

Attaches to multiple systems for file accessibility

Instead of being presented as a tree structure, objects are arranged in a logical order irrespective of their location on physical tape.

Enterprise Edition (EE)

Allows tape to be defined as a Tier 2 or 3 storage

GPFS™ automates the movement of data between tiers as per user defined policies

Supports LTO 5 and 6 and TS1140 tape drives

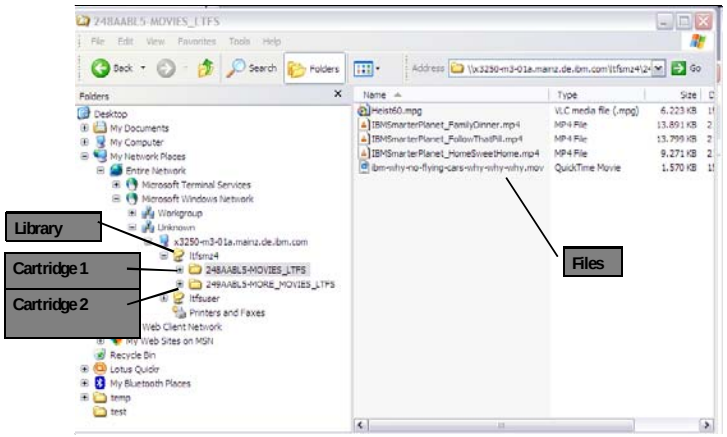

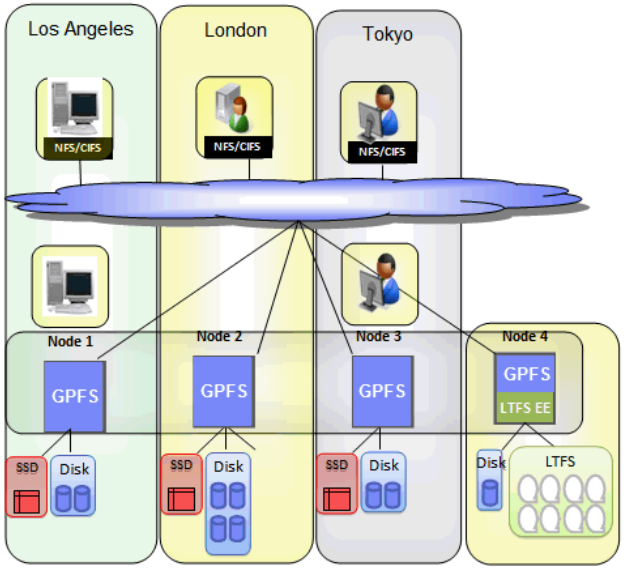

LTFS Enterprise Edition introduction

- LTFS Enterprise Edition integrates LTFS with GPFS
  - LTFS represents external tape pool to GPFS
  - Files can be migrated using on GPFS policies or LTFS EE commands
  - Similar implementation as with TSM HSM
- LTFS EE can be configured on multiple nodes
  - Multiple instances of LTFS EE share the same tape library

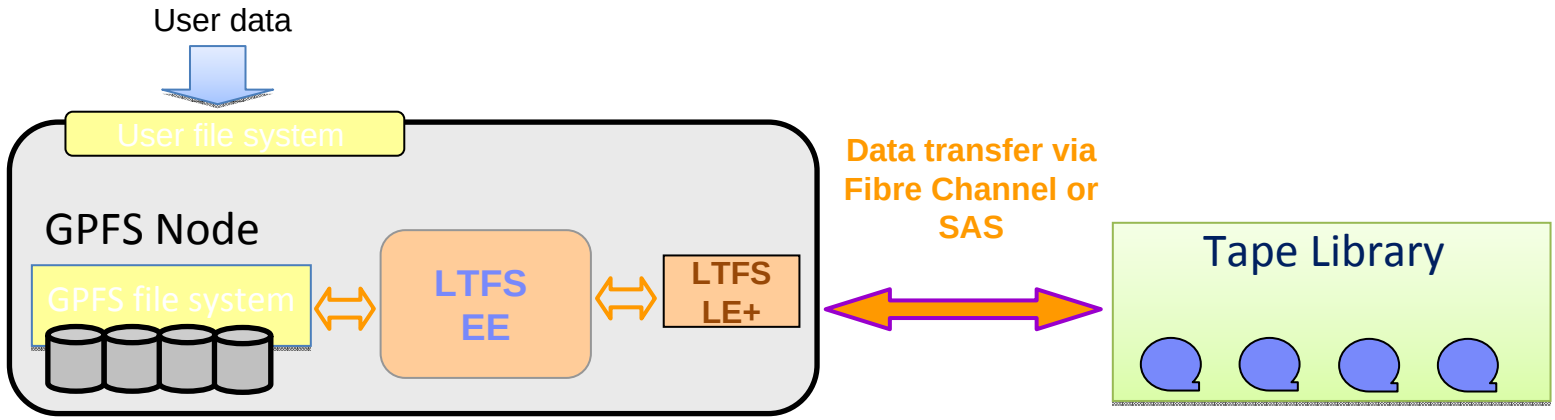

LTFS EE system view

- LTFS Enterprise Edition integrates LTFS with GPFS
  - LTFS represents external tape pool to GPFS
  - Files can be migrated using on GPFS policies or LTFS EE commands
  - Similar implementation as with TSM HSM
- LTFS EE can be configured on multiple nodes
  - Multiple instances of LTFS EE share the same tape library, LTFS EE is installed on one or more GPFS nodes
- GPFS provides global name space
- Each LTFS EE instance has dedicated tape drives attached
  - In the same TS3500 library partition
- LTFS EE instances share the same tape cartridges and LTFS index
- Workload is distributed over all LTFS EE nodes and tape drives

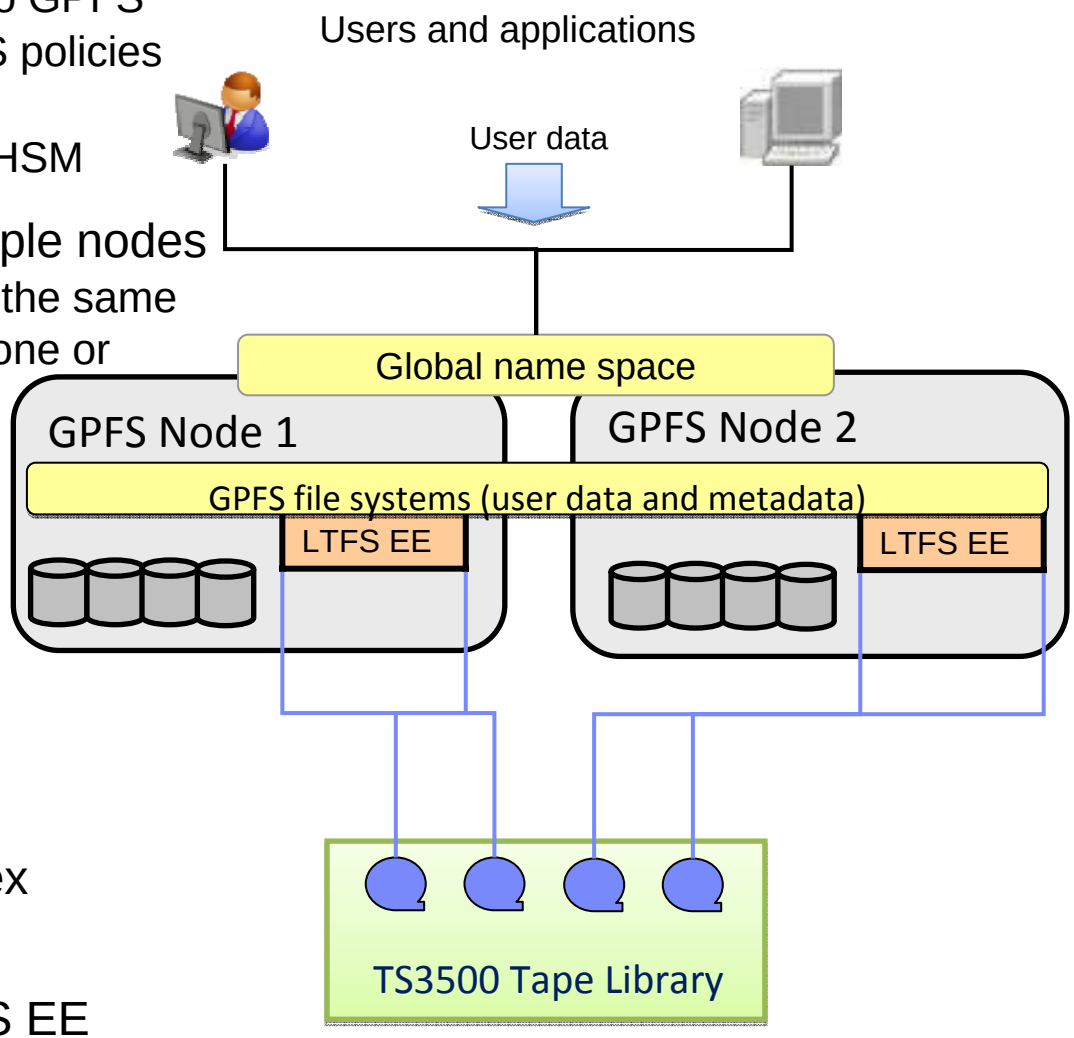

IBM System Storage TS7650G ProtecTIER® Deduplication Gateway

- **IBM HyperFactor® deduplication technology**
  - Offers patented deduplication algorithm (not hash-based) designed to reduce storage needs by up to 25 times or more, radically reducing costs while maintaining enterprise-class data integrity
- **Backup performance**
  - Delivers sustained inline deduplication throughput rates up to 2,500 MBps or more
  - Offers throughput that can scale to meet the most demanding data center requirements
- **Restore performance**
  - Delivers sustained data recovery throughput rates up to 3,200 MBps or more
  - Restores data faster to get mission-critical systems back online quicker
- **Capacity**
  - Provides scalability to 1 petabyte (PB) of physical storage, which provides up to 25 PB or more backup storage capacity
  - Enables easy scaling of capacity to meet ever-increasing requirements over time
- **Clustering**
  - Allows two nodes to be clustered within a single system for higher performance and availability

ProtecTIER Values!!

Performance

Sustainable 1500 MB/s **1800 Restore** per node (2800 MB/s two node cluster **3500 Restore**), performing ***inline*** de-duplication

Scalability

Up to 1 PB physical capacity per node

Highest Level of Data Integrity

Non-hash & Binary diff process during de-dupe designed for the highest data integrity

Non-Disruption

Daily Operations

Inline de-duplication eliminates need for significant secondary processing

Implementation

Integrates well with existing backup environment and infrastructure

Clustering

Two (2) node, high availability

Global Deduplication

Locates & Deduplicates ALL data

IBM System Storage TS7620 ProtecTIER Deduplication Appliance

- **Capacity**
  - 6 TB, 12 TB, 23TB oder 35 TB native Capacity
- **Performance**
  - Up to 300 MBps or more inline deduplication performance
  - Recovery is even faster, enabling systems to get back online quicker
- **2U Server, 12x 2TB SAS-NL, 8GB FC, 48GB Memory**

16 years IBM Enterprise Tape Virtualization

- **1997: First Virtual Tape Server (VTS)**
  - IBM 3494-B16
    - Capacity TVC: 72 GB
    - Throughput: 6 MB/s
- **2006: Generation 4 – new Grid based architecture**
  - IBM TS7700 Virtualization Engine R1.0
    - Capacity TVC 6.000 GB (6 TB)
    - Up to 600 MB/s throughput
    - Available as two models (Tape and “Disk only”)
- **2011: Generation 5 – Power7 based architecture**
  - IBM TS7700 Virtualization Engine R2.0
    - Capacity TVC 28 TB
    - Up to 1.000 MB/s throughput

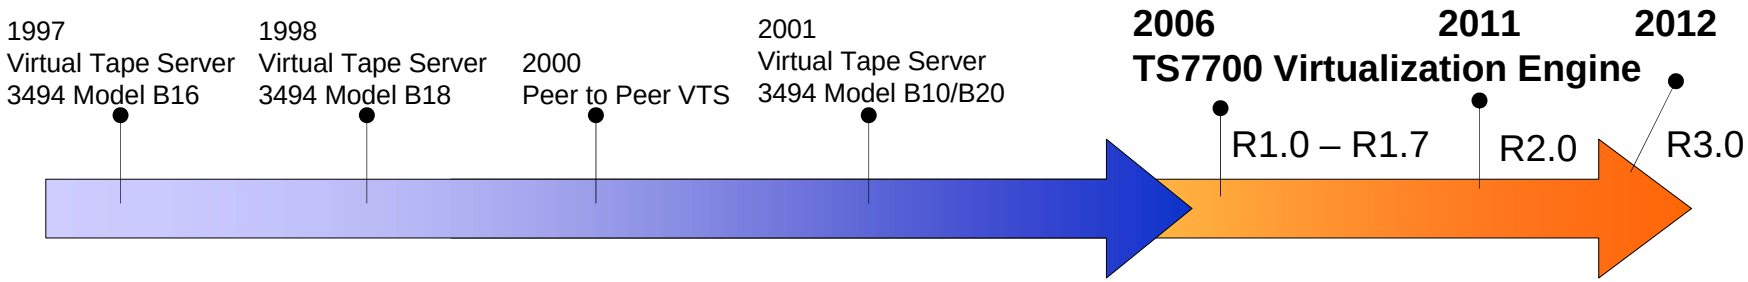

TS7700 Functional Summary by Release

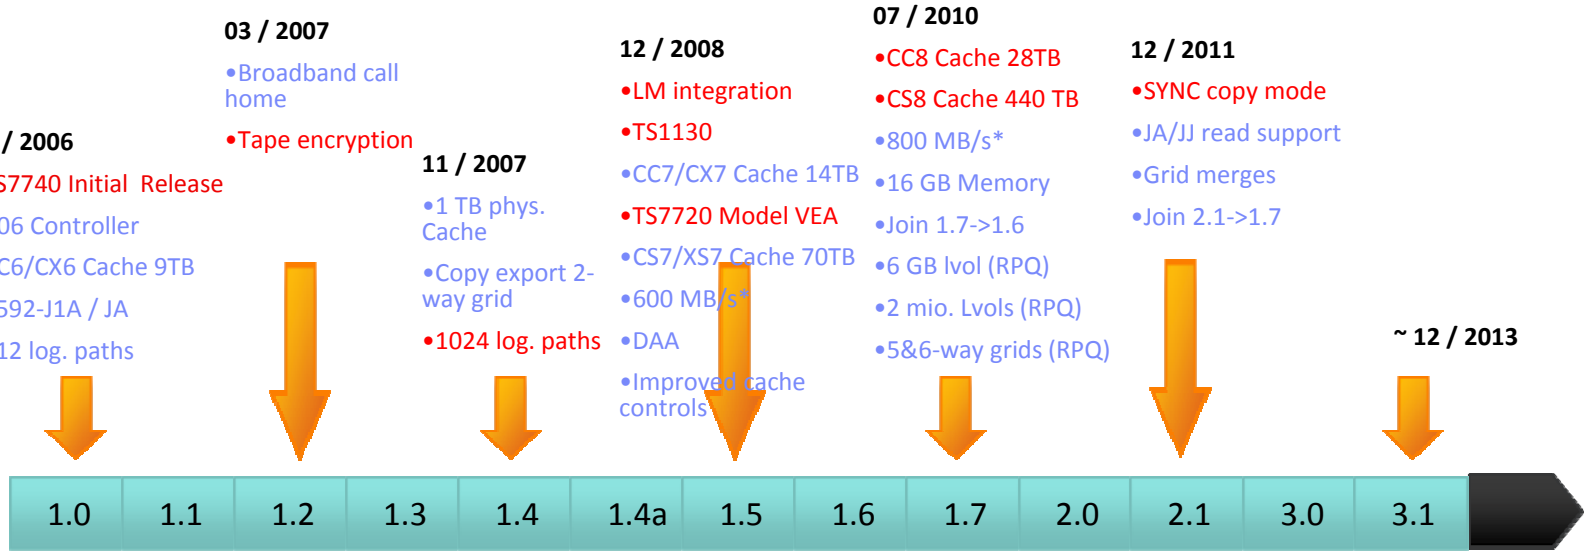

09 / 2006
- TS7740 Initial Release
- V06 Controller
- CC6/CX6 Cache 9TB
- 3592-J1A / JA
- 512 log. paths

01 / 2007
- AOTM
- TS1120 / JB
- 2 cache drawer configs
- Cache increments
- Optical gridlink (RPQ)

03 / 2007
- Broadband call home
- Tape encryption

08 / 2007
- 3-way grid
- 1 mio. lvols
- 2->4 FICON MES
- 1->3 CX6 MES
- Copy export standalone
- Perf. increments
- HCR

11 / 2007
- 1 TB phys. Cache
- Copy export 2-way grid
- 1024 log. paths

04 / 2008
- Link load balancing
- 1->2 drw. Cache MES
- Host copy control
- Copy audit
- Cluster removal
- Cluster cleanup

12 / 2008
- LM integration
- TS1130
- CC7/CX7 Cache 14TB
- TS7720 Model VEA
- CS7/XS7 Cache 70TB
- 600 MB/s*
- DAA
- Improved cache controls

12 / 2009
- 4-way grid
- Hybrid grid
- LWORM
- Families
- SSPC
- Retain copy mode
- HCR additions

07 / 2010
- CC8 Cache 28TB
- CS8 Cache 440 TB
- 800 MB/s*
- 16 GB Memory
- Join 1.7->1.6
- 6 GB lvol (RPQ)
- 2 mio. Lvols (RPQ)
- 5&6-way grids (RPQ)

06 / 2011
- P7 models V07/VEB
- 1200 MB/s*
- TS1140 / JC
- 2x 10GE gridlinks
- 4x 1GE gridlinks
- SAA
- SDAC
- 6GB lvol (GA)
- 2 mio. lvols (GA)
- 8 Gb FC switches
- 5&6-way grids

12 / 2011
- SYNC copy mode
- JA/JJ read support
- Grid merges
- Join 2.1->1.7

12 / 2012
- CC9 Cache 28TB
- CS9 Cache 624TB
- 2100 MB/s*
- Disk encryption
- 4 mio. lvols
- Unified GUI
- Join 3.0->1.7 (RPQ)

~ 12 / 2013

* Peak write, 2-way dual-active clusters, immediate copy mode, 2.66:1 compression

IBM TS7700 R3.0 VE Overview

- **4Gb/s FICON (2 / 4) with HW Compression**
  - Up to 4.000.000 virtual Volumes, WORM Support
  - 256 virtual Devices per TS7700 Cluster (3490E)
  - Ext. log. Volume-Size 1, 2, 4, 6GB
  - 256 Path per Adapter = up to 1.024 per TS7700
- **Fully integrated in in z/OS, DFSMS (and HSM, OAM, RMM)**
- **Front-End Throughput per Cluster up to 1.000MB/s**
  - Scalable in 100MB/s steps
- **Based on 100% IBM components (P7 Server, Disk-System, uCode**
- **Available in 2 versions**
  - “tapeless” Virtualization Engine (TS7720)
  - With Back-End Tape (TS7740)
- **TS7740 Disk Cache**
  - Scalable up to 28TB per Cluster (native)
    - 1TB steps
- **TS7740 with Backend-Tape-Support:**

  IBM 3592 Laufwerke/TS3500 Library
  - 4 to 16 Tape Drives per System (Cluster)
  - Tape Encryption Support
  - Copy Export (Snapshot of aTS7740)
  - Supports alle Enterprise Tape Drives and Tape Media
    - 3592J1A, TS1120, TS1130, TS1140
    - 3592 Medien JA, JB, JC

IBM TS7700 R3.0 VE Overview cont.

- **TS7720 tape less with large Disk-Cache**
  - From 20TB up to 620TB per Cluster (usable without compression)
  - Scalable in 20TB steps
- **Multi-Cluster-Grid**
  - Up to 6 TS7700 Cluster Nodes
  - “hybride” Grid configuration possible (TS7720 and TS7740)
  - Many different Grid Configuration possible for HA, DR and Performance
  - Synchronous or Asynchronous Replication based on policy possible
  - IP-based Replication

Summary of key specifications R3.0

Specifications per cluster	TS7740	TS7720
Number of Virtual Devices	256	256
Usable Cache	1TB to 28TB	Up to 624TB
Compressed Cache Capacity (3:1)	3TB to 84TB	Up to 1.872 PB
Number of Virtual Volumes	4,000,000	4,000,000
TS1140/1130/1120/3592 Tape Drives	4 – 16	NA
FICON channels	2 or 4	2 or 4
Logical paths per FICON channel	256	256

What is a Grid?

- Two to 6 TS7700s
- worldwide access; as little as one physical instance of the data
- automated failover
- Read / written at disk speeds, sent at network speeds
- For Transparent Business Continuance

Multiple System z Hosts have access to all data at all times, regardless of physical location

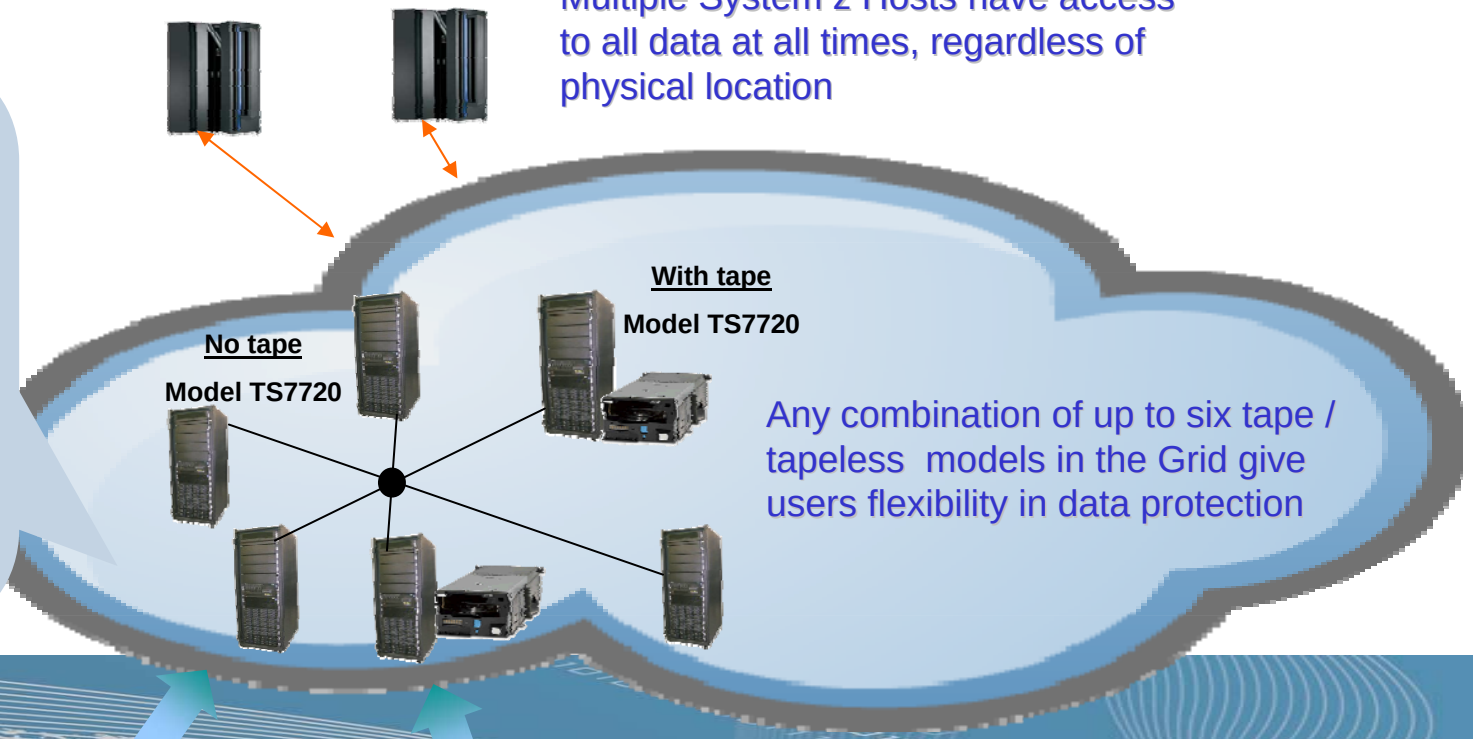

Any combination of up to six tape / tapeless models in the Grid give users flexibility in data protection

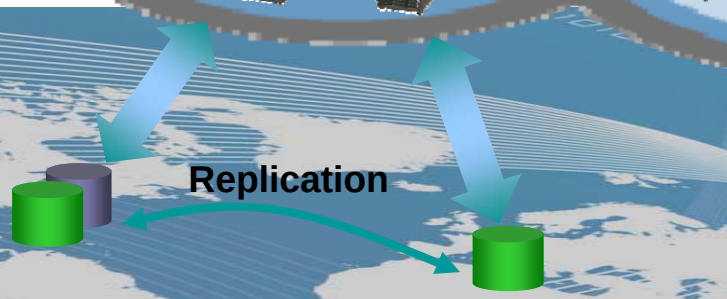

Grid: access to all data independent of where it exists on planet earth + automated failover!

30

What is a TS7700 Grid?

WAN interconnected TS7700s clusters form a Grid configuration

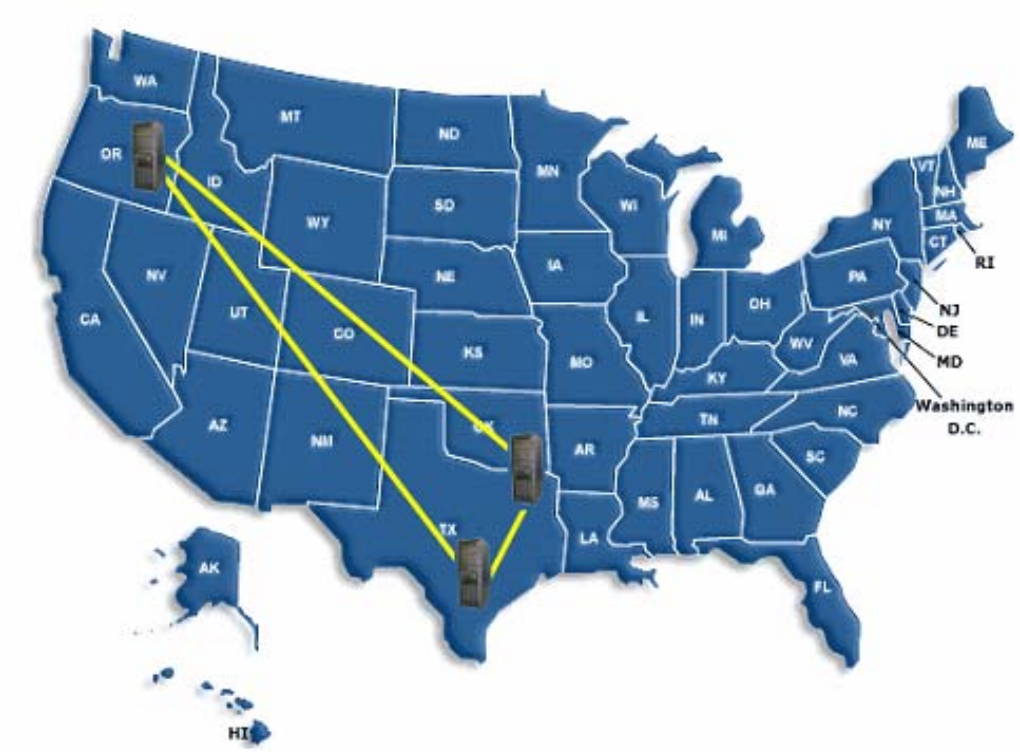

Optimized for Recoverability and Automatic Failback

- All clusters are peers to each other
  - The Grid forms one large composite library
  - Interconnect is standard TCP/IP using dual 1Gb / 10 Gbit links
    - Supporting 1000s of miles separation
  - Data written to one is transparently replicated to the other
    - Policy control for location of copies and how copies are made
  - A volume's data can be accessed through any TS7700
  - Can be configured for disaster recovery and/or higher availability environments
    - One to six site configurations
- Each TS7700 Cluster provides
  - 256 virtual devices
  - 4 - 4Gb FICON Channels
  - Up to 84TB of Cache (3:1 C/R) TS7740
  - Up to 1.87PB of Cache (3:1 C/R) TS7720
  - 4,000,000 logical volumes

Example: Six-Way Grid Configuration for High Availability and High Performance

- 6-way provides 2 HA production sites with two cluster at the DR site.
- All scratch mounts are directed to the desired clusters (ie. Scratch Allocation Assist is available).
- Retain copy modes is enabled to prevent extra copies during outages when the mounts are not routed to the expected cluster.
- Large caches on the TS7740 will allow premigration activity to be done outside of the batch window by setting the premigration thresholds much higher.

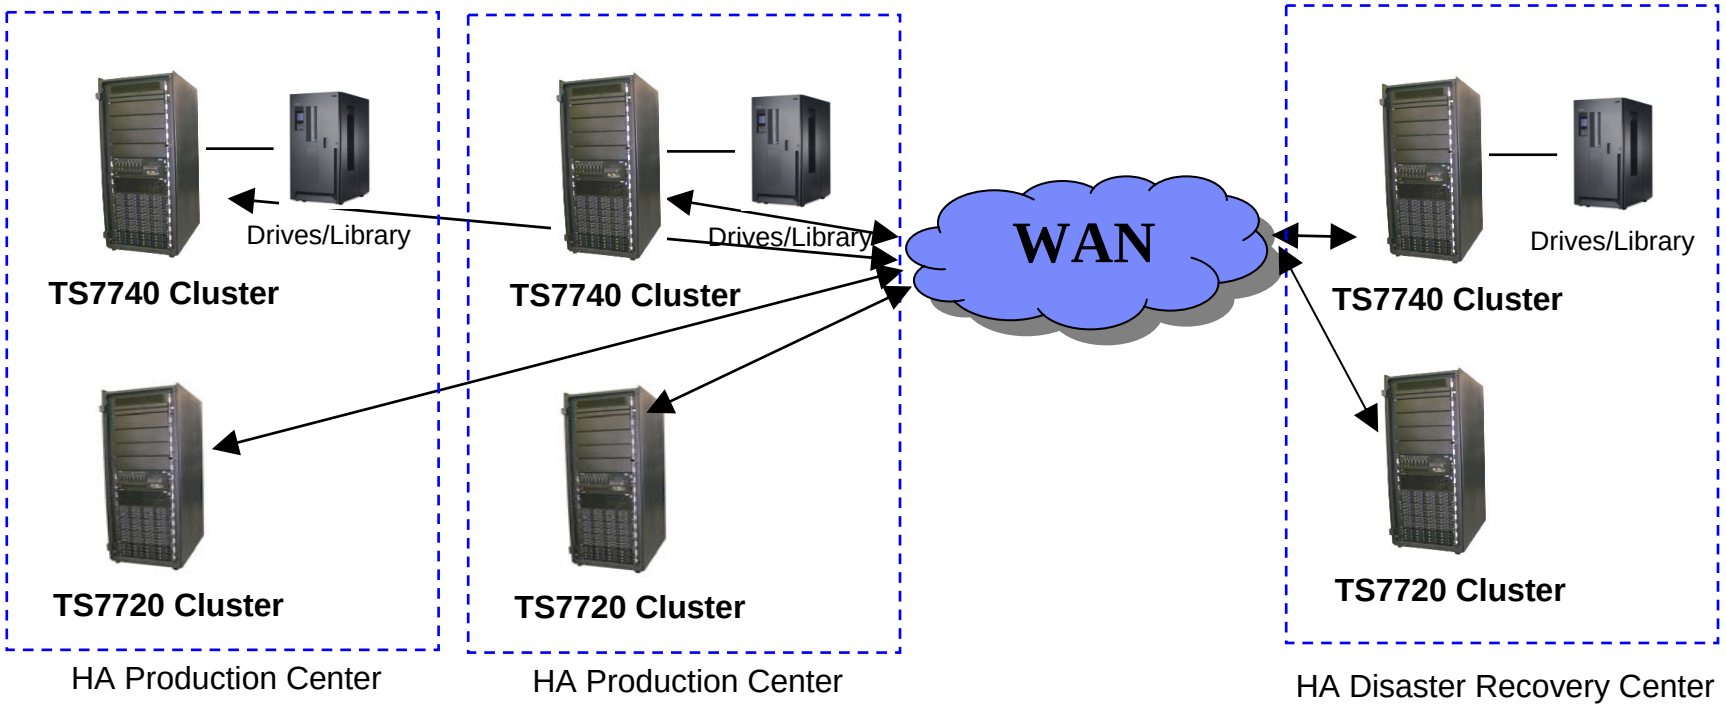

Why Tape is not dead! - Agenda

- Cheap, low TCO
  - Long lifetime
- Green & Cool
  - Less Power and cooling
- High capacity
  - Needs less floorspace
  - Compression included
- Fast
- Secure
  - Offline
  - Bit error rate are better than disk
- Portable
- Roadmap – Future growth

- ***Enterprise Strategy Group***: - Most Enterprises utilize a combination of Physical Tape & Disk for Back-up & Restore
  - More large enterprise users have Tape & Disc installations than Midmarket
  - Tape-only clients: One-third more midmarket users than large enterprise
- ***Gartner*** – Confirms that most clients use disk & tape for backup and restore

Exhibit 8 — Average TCO/Three-Year Cycle (for PBs Required in Each Cycle) for Long-Term Archived Data (Using Clipper's TCO Study Model)

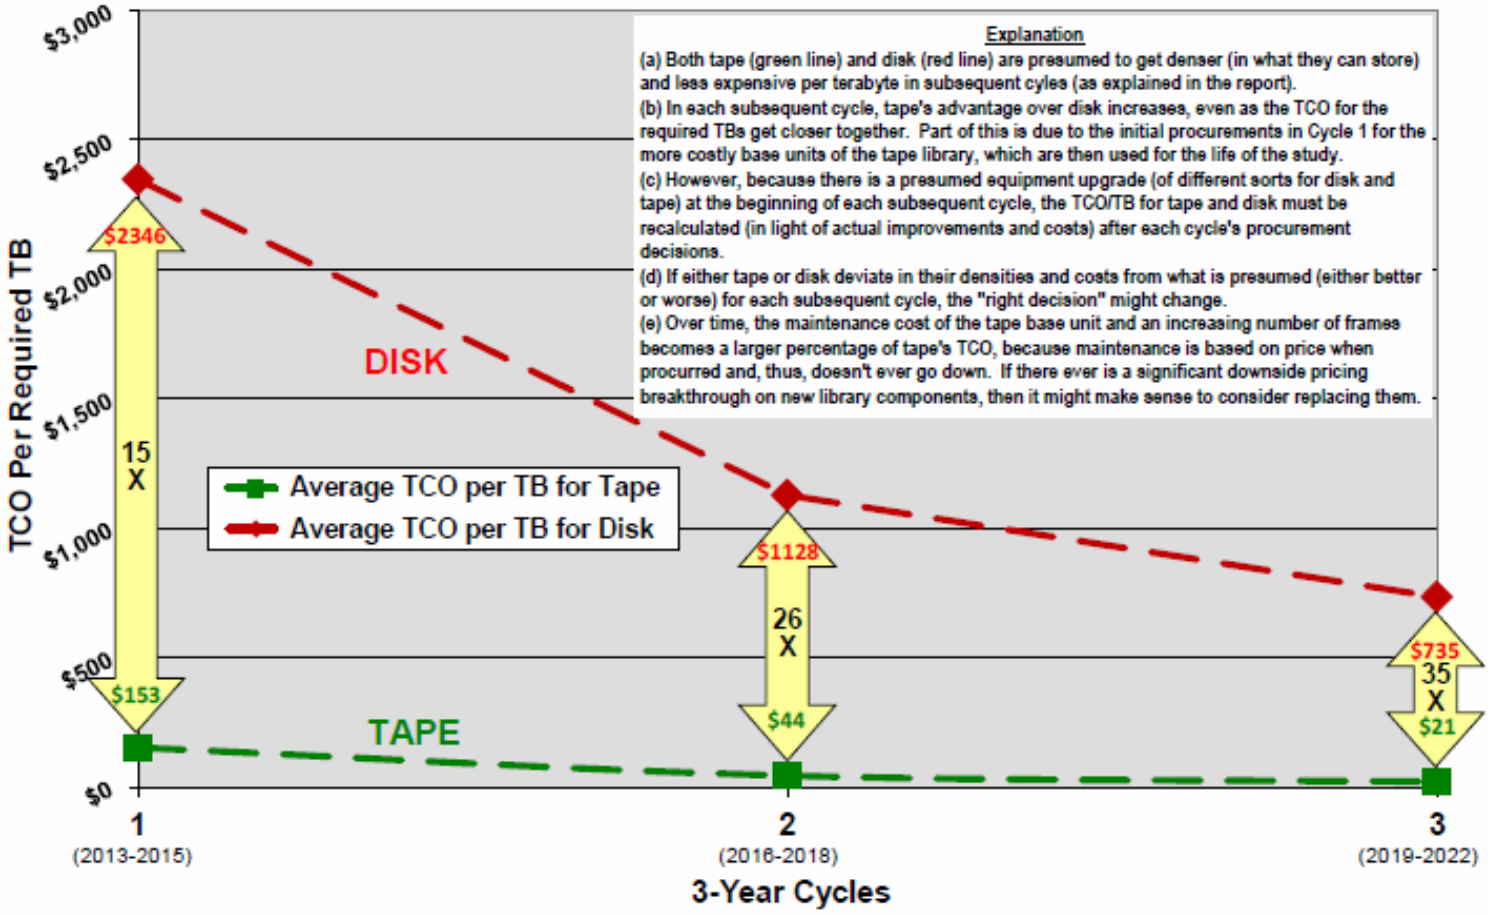

Source: The Clipper Group

- This report is not about whether disk costs more than tape, or not; it is about having the right mix of disk and tape …, taking **advantage of the strengths of both**, i.e., the low-cost, high-capacity of tape and the rapid response of disk.
- Regardless, for large quantities of data, **tape always is much less expensive than disk** and always uses much less energy and floor space, when measured on a per-petabyte ba-sis. Tape should be used whenever its some-what slower retrieval times are acceptable.
- Source: http://www.lto.org/pdf/Clipper_Group-Long-Term_Storage-TCO_Analysis_of_Tape_and_Disk-May_13_2013.pdf

Tape Data Format: 2-Level ECC – Very Strong Data Protection

- **User data (optionally compressed/encrypted) is partitioned into fixed size Data Sets**
- **Data Set = 32 sub data sets**
- **Sub Data Set = 96 × 960-byte array**

ROWS: Protected by C1 ECC
COLUMNS: Protected by C2 ECC

Failure Rate Comparison
HDD (SATA): $P_{bit} < 10^{-15}$ to 10^{-14}
Tape: $P_{byte} < 10^{-19}$

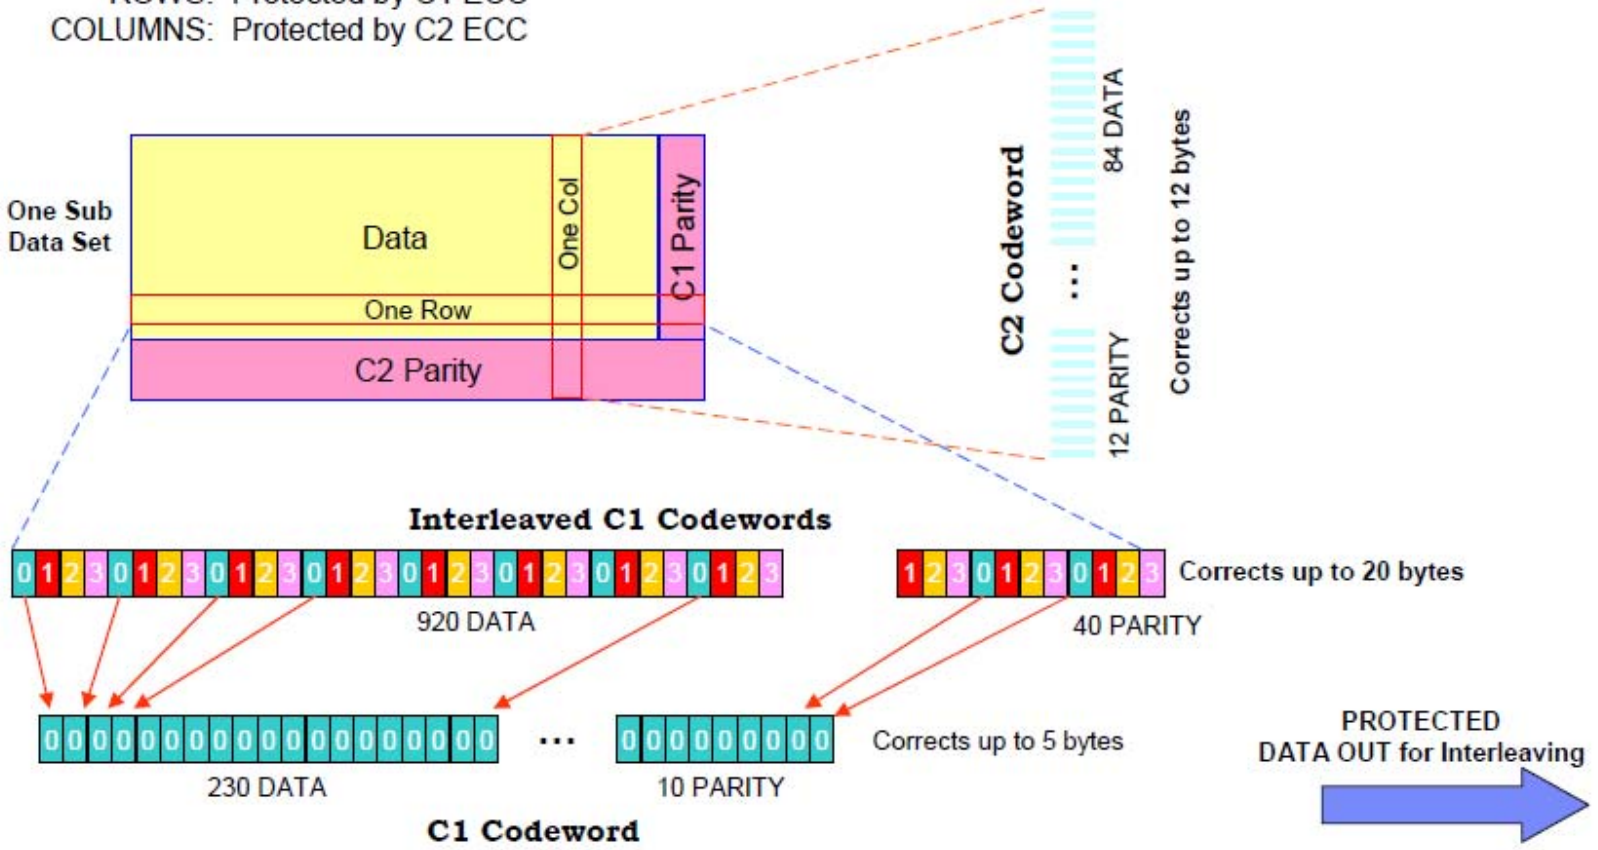

Data layout on tape with deep interleaving

- The drive can tolerate *strip errors of 11 mm length across tape*
- Bit error rate on tape is better than on disk
  - You can write 10PB – 1 EB more data on tape than disk
- The drive can read Data Sets with *2 entirely DEAD tracks out of 16*

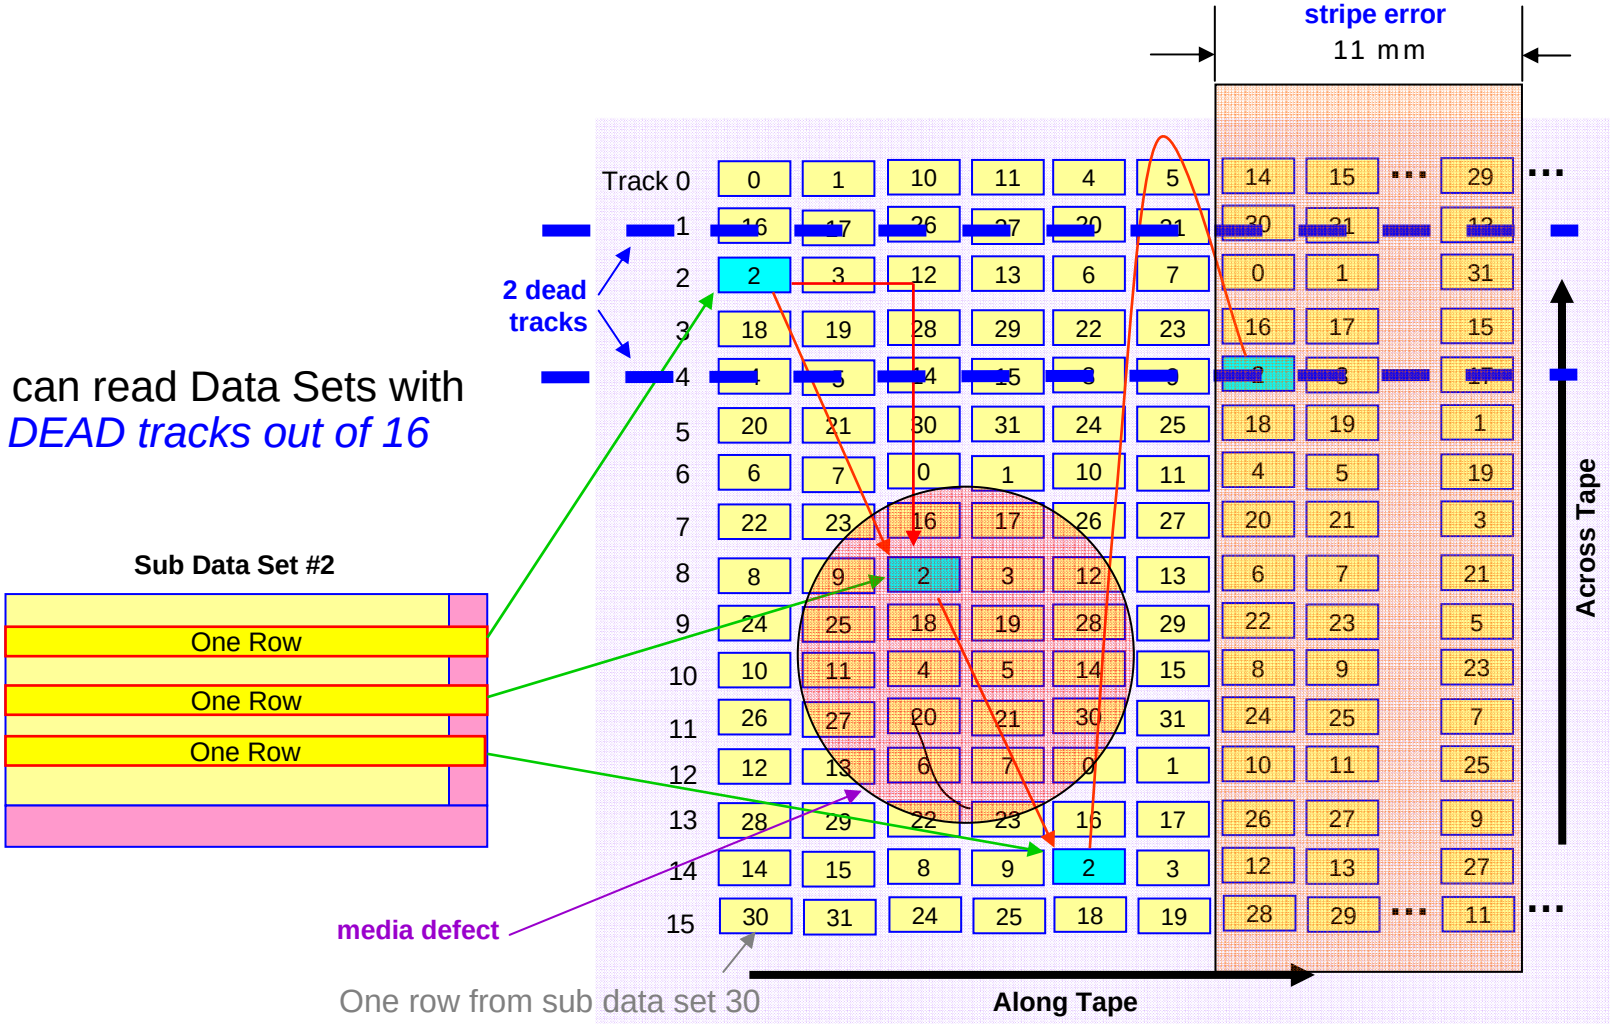

Tape Technology Development

- 35 TB demonstration done in Jan 2011
- Working on 125 TB demonstration

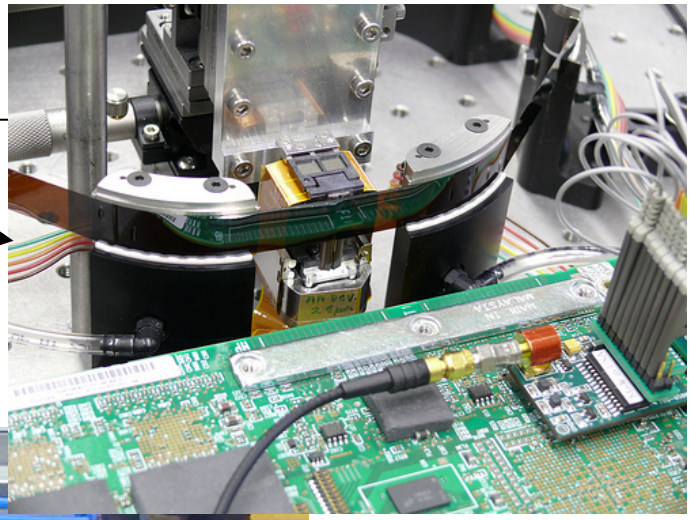

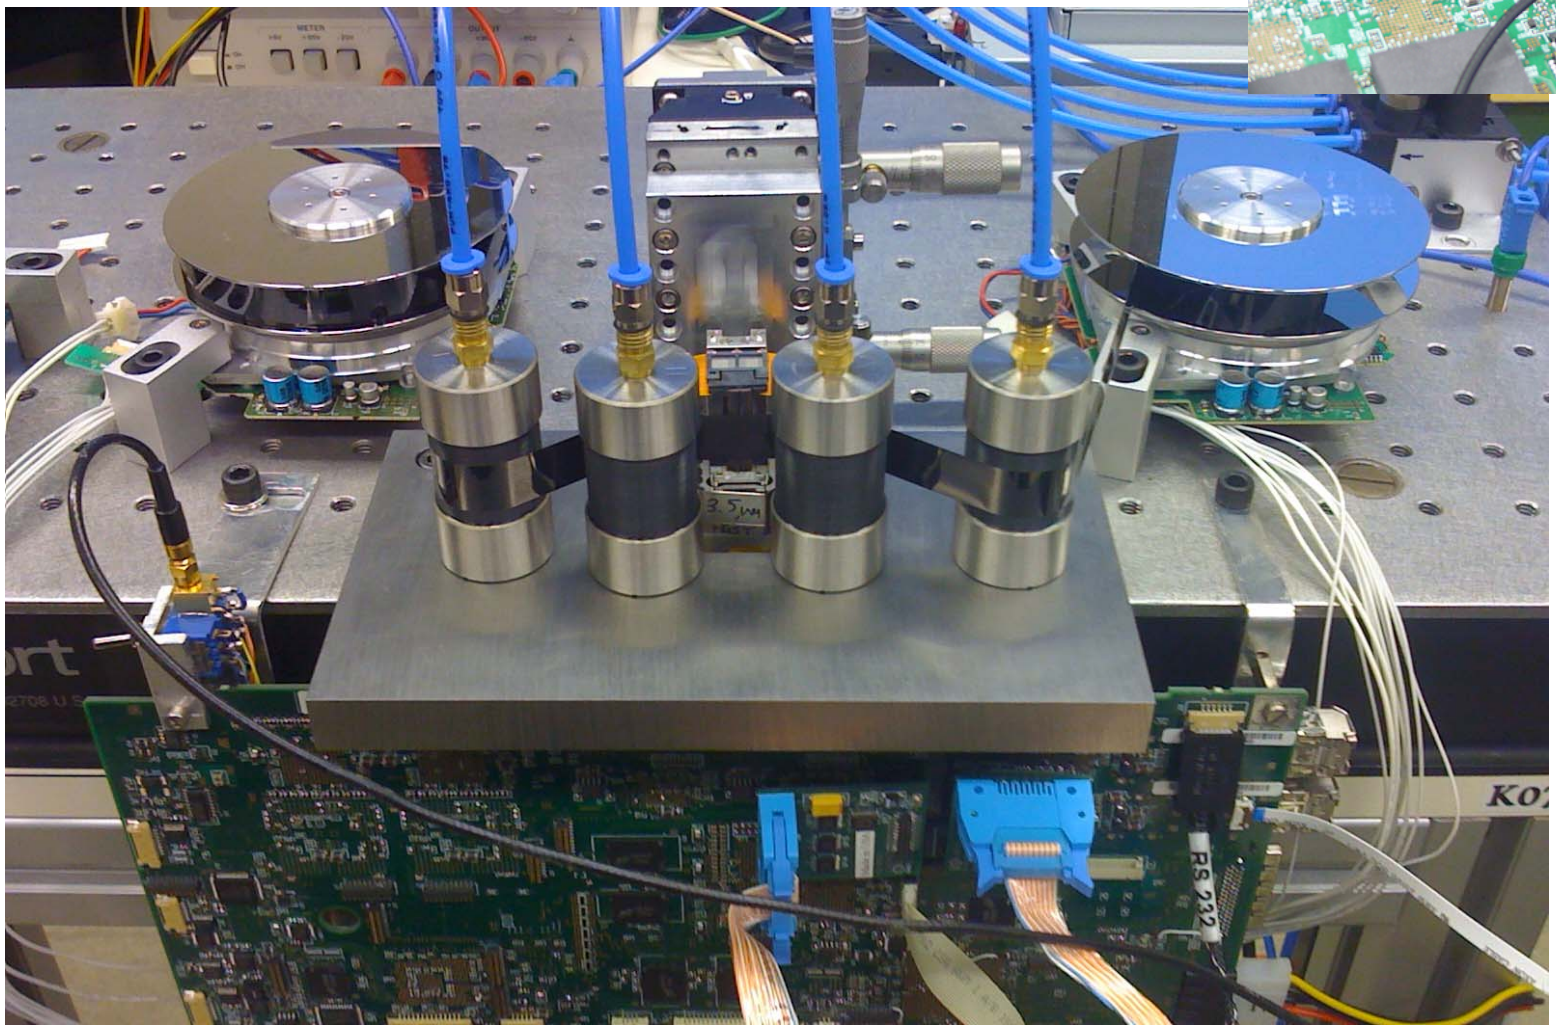

Take a look at the Bits …

- **Magnified View of Scaled Bit Cells**

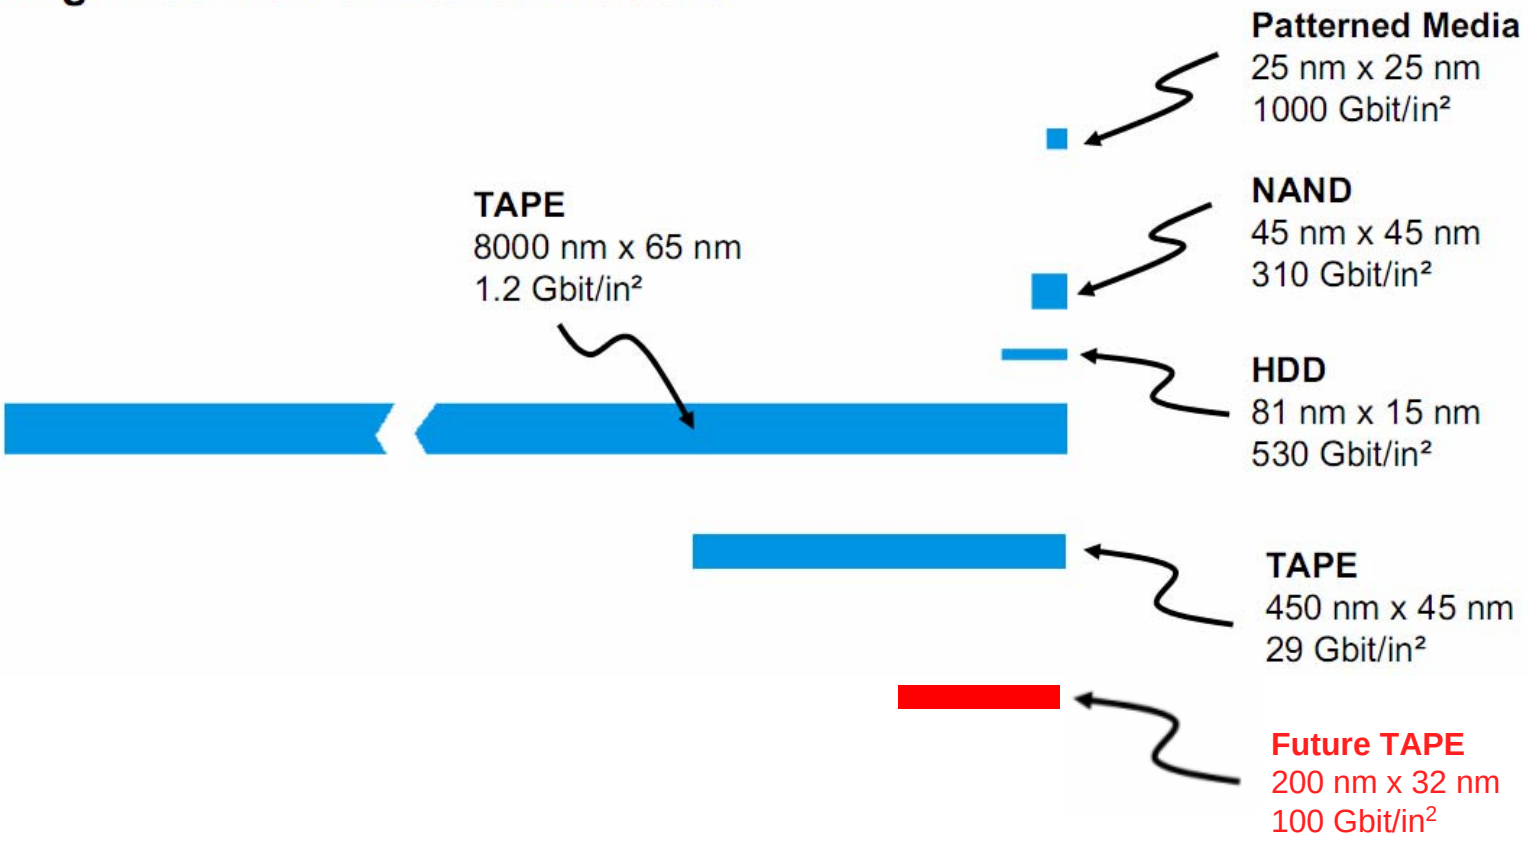

Tape has most potential to scale out in capacity

* assuming current tape technology

Tape – yes, but how…?

Over 80%

consider tape an integral part of their backup process

Source: Enterprise Strategy Group Research Report, 2010 Data Protection Trends, April 2010

- „One fits all“ can not full fill all Requirements
- Establish backup & restore policies by type of data
  - Critical data
  - Important application data
  - Data with business value
- Use combination of disk & tape to balance RPO / RTO, costs
- Protect data against logical corruption, physical destruction, unauthorized access
  - Maintain multiple copies, geographically separated
  - Use different types of media

- **Only the combination of different Backup-Method guarantees Security, Performance and Efficiency!**

Trend: SnapBackups

- Snap-Backups reduce Backup and Restore/Recovery time
- Snap‘s within Storage Systems does not influence Workload on Server/Applications
- **But:**
  - Take Care about Application integration – consistence!
    - Complex Integration
    - Skill for different layers needed (Application, Disk-System, Backup)
  - **Snap-Backups does not protect fully**
    - **See Amzon Cloud, Google Mail,….**
- **Only the combination of SnapBackups and classic Backup fully protected your data and full fill future requirements!**

IBM DP Blueprint

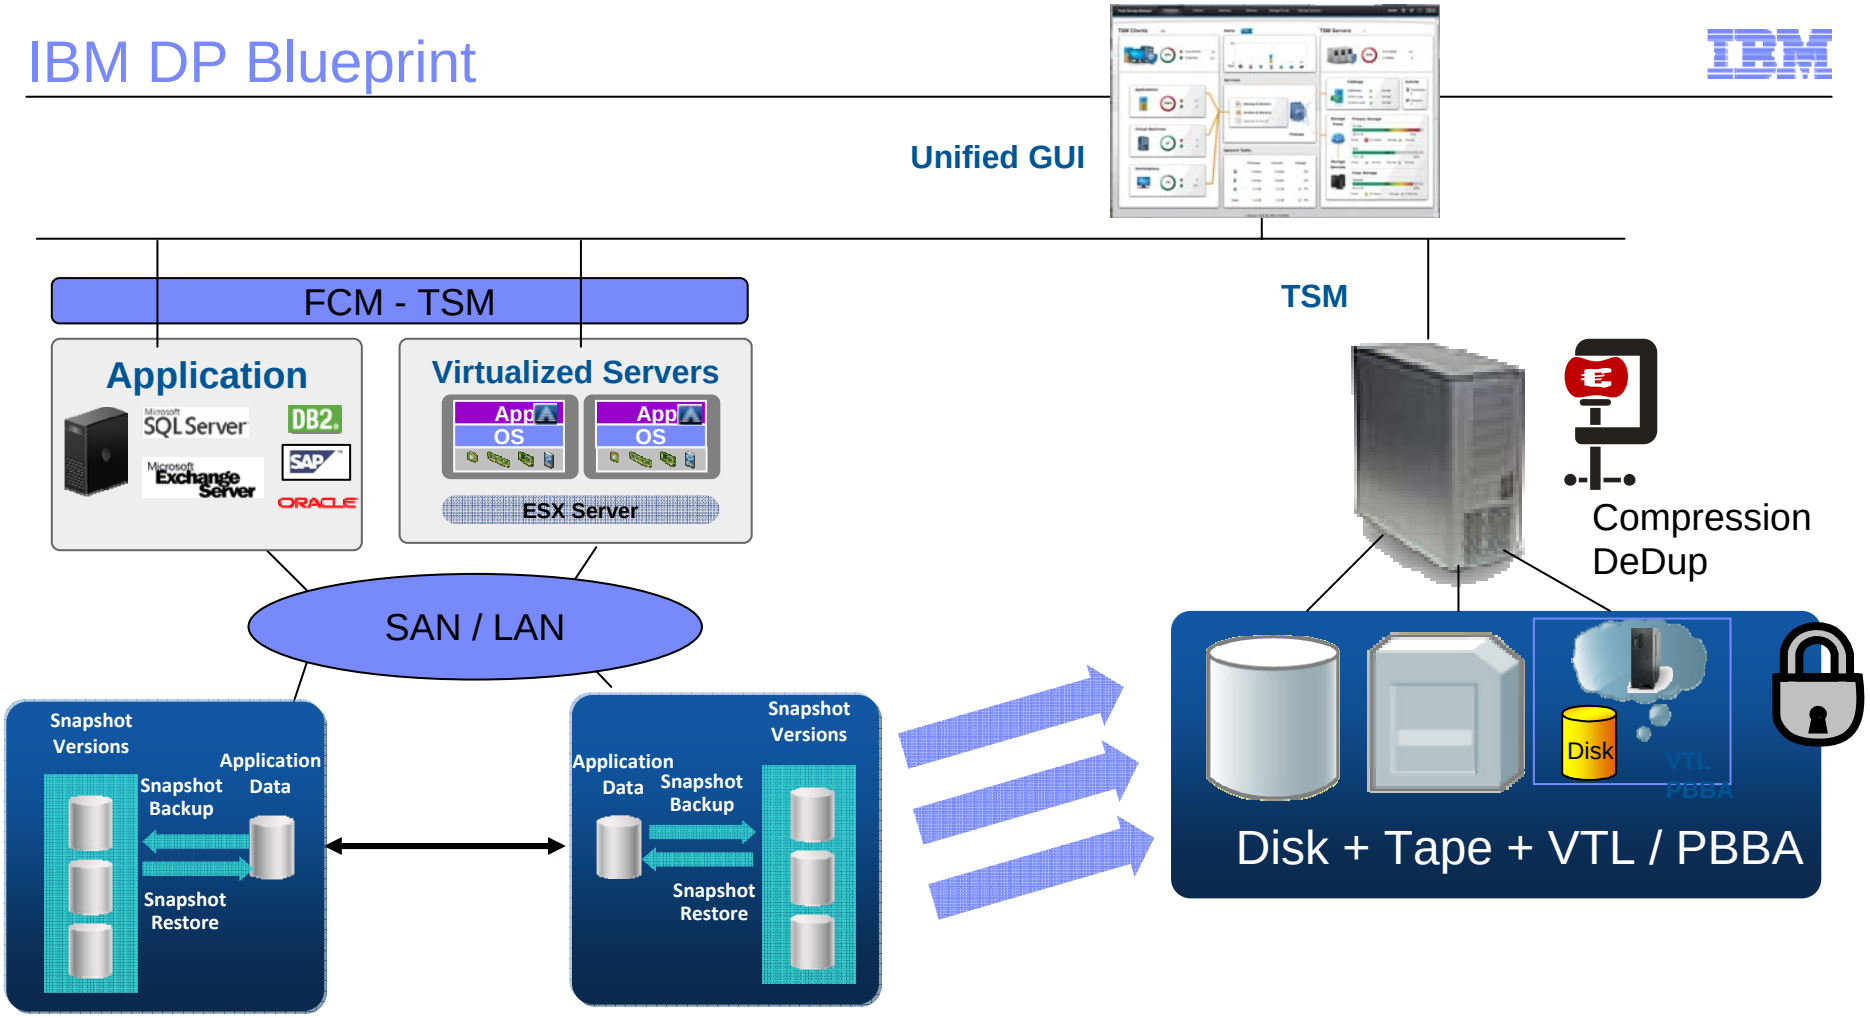

- **Integration:** Applications- consistence Snap-Backups
- **Security:** offline-Media, Chance of media and technology
- **Cost:** Tiering, Compression, DeDup, Tape > 80%
- **Scalability :**
- **Efficiency :** Unified GUI für Server / Infrastructur / Storage / Backup
- **Compliance:** WORM und Encryption
- **IBM: Fit For Purpose**

Large Databases

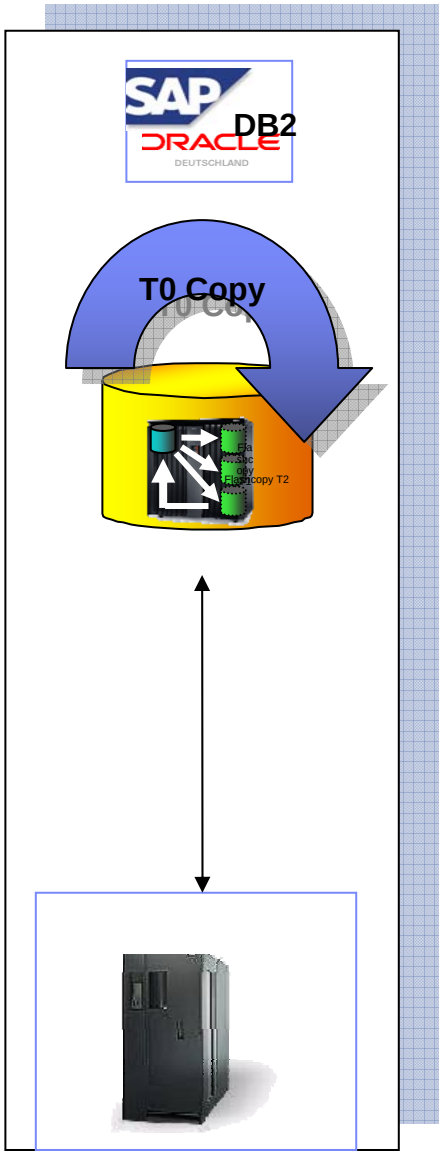

- Tape backup
  - LAN-free
- For (large) databases
- For backup streams faster than 50-80 MB/sec
  - Cheaper and faster than disk/VTL
- Consider FlashCopy Manager

General Data/ Universal using TSM Pool and Migration to Tape

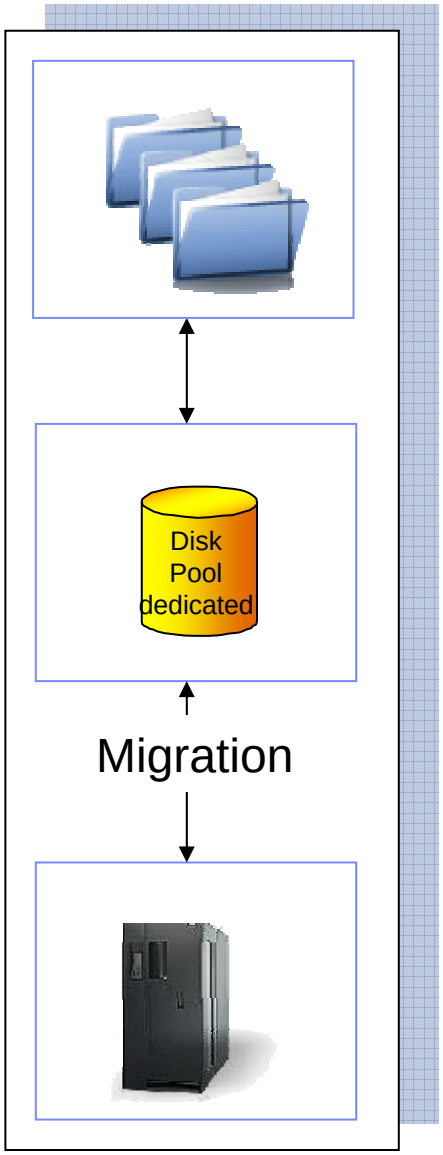

- Use TSM Disk Pool for “caching” the data
  - Some time later migrate data to Tape
- Size Disk Pool for at least 1-2 days of daily backup data
  - May increase the disk pool size to 5-7 days
- Take care about performance – most important single stream performance
  - Use dedicated disk system for TSM
    - Does not use disk system which is shared with productive data/workload
      - Performance and data protection issues
  - NL SAS/SATA may not full fill the performance requirements
  - DeDuplication appliances will not full fill performance requirements
  - Recommendation: use SAS/FC Disk for primary disk pool with FC/SAN connection
  - Rule of Thumb“
    - SAS HDDs: Raid5 8+1 -> ~260 MB/sec
    - NL SAS HDDs: Raid6 8+2 -> ~85 MB/sec

Small Files / **Restore optimized**

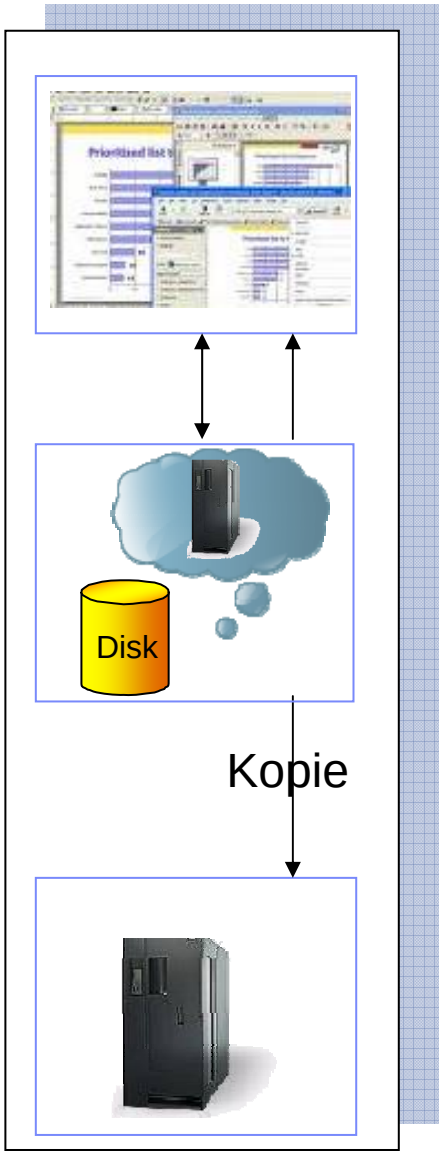

- Disk Backup
  - Copy Pool on Tape
- For (small) Files to improve restore
- Consider DeDuplication
  - TSM or ProtecTier
- Save license cost with ProtecTier

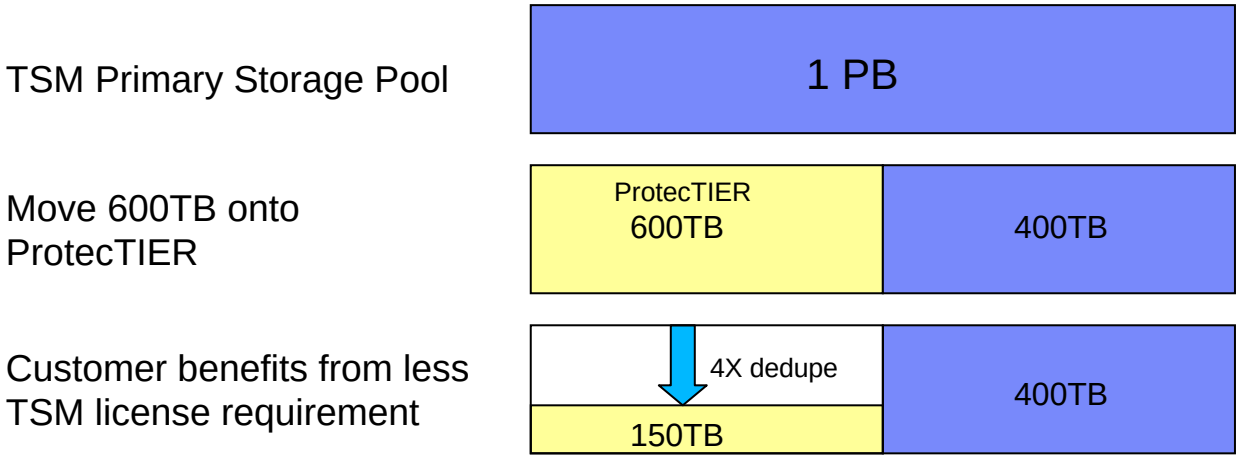

Data DeDuplication (DDD)

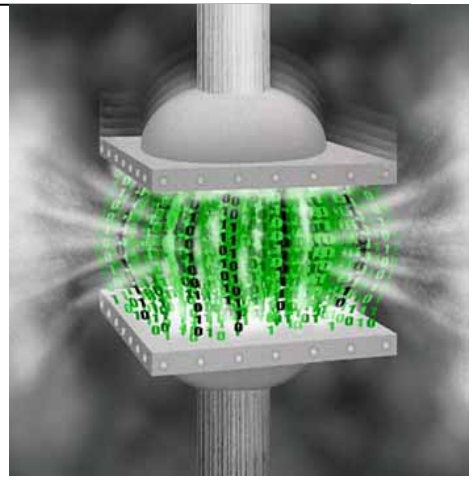

- Goal: cheaper Disk-Backup
- When do use DDD
  - Data with high redundancy
  - Data which needs to be stored for a long time on disk
  - Replication over WAN
    - Saves bandwidth requirements
- However
  - DDD needs CPU “power”
  - Performance might be slow
  - Restore might be much slower
    - Example backup of database
      - with LTO3 -> 166 MB/sec
      - With DDD Appliance 180 MB/sec
    - Restore
      - With LTO3 -> 160 MB/sec
      - With DDD Appliance 110 MB/sec
- DDD ratio depends on data, retention time, number of versions
  - DDD ration with TSM 2:1 and 15:1
  - Good for DDD
    - Full Backups, database, OS
  - Challenging for DDD:
    - Encrypted data, compressed data
    - Databases with many reorgs, like Exchange
    - Databases with compression, like Oracle, DB2
    - Logs
- Keep one copy on tape

VTL for LAN-free

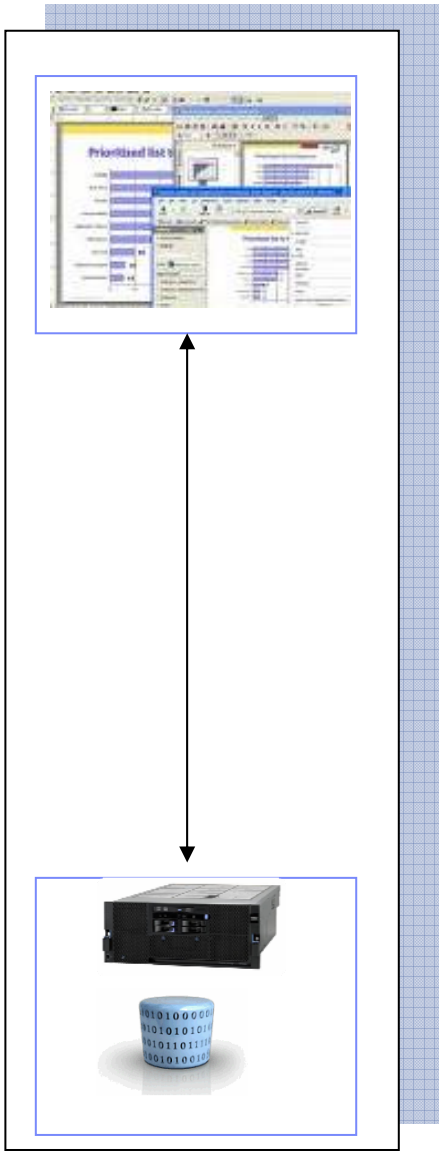

- Use it for “slow” and small LAN-free Backups
- Save license cost with ProtecTier

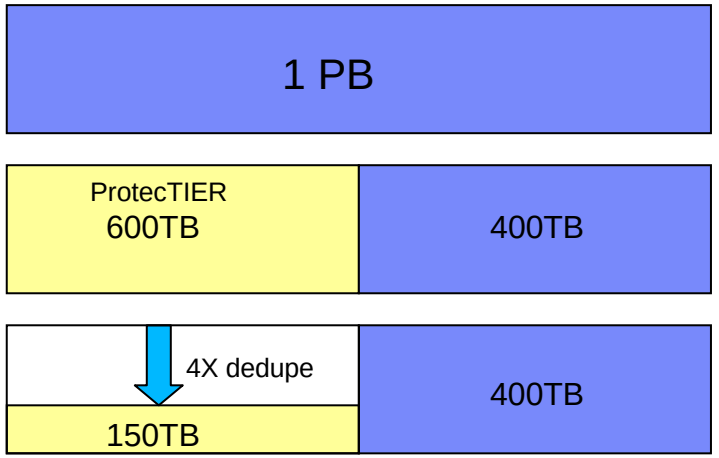

IBM DP&A: Smarter Backup Solution

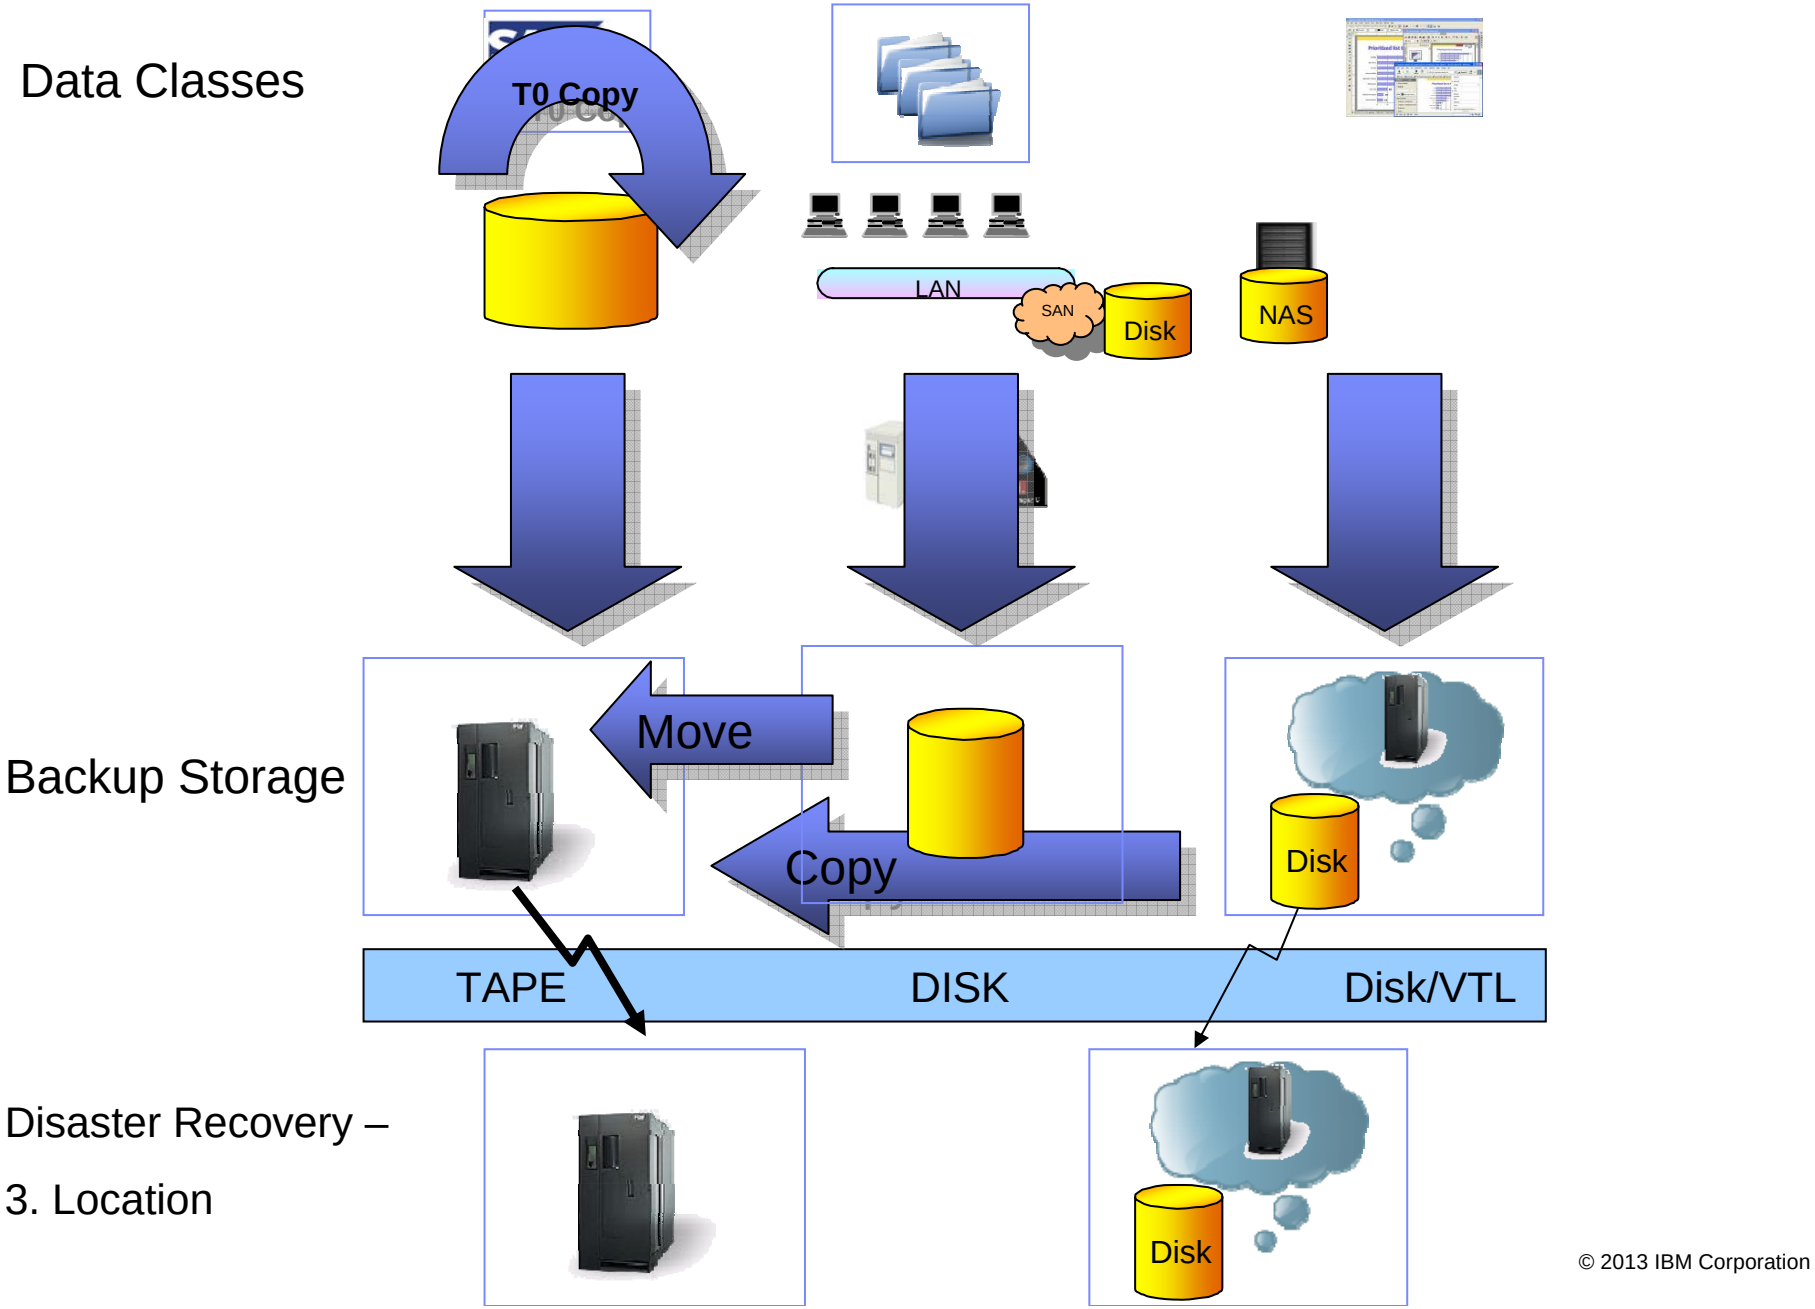

Questions?

Infos / Find me on:

http://sepp4backup.blogspot.de/
https://www.xing.com/profile/Josef_Weingand
http://www.linkedin.com/pub/josef-weingand/2/788/300
http://www.facebook.com/josef.weingand

Thank You!

Disclaimers

The performance data contained herein was obtained in a controlled environment based on the use of specific data. Actual results that may be obtained in other operating environments may vary significantly. These values do not constitute a guarantee of performance.

Product data is accurate as of initial publication and is subject to change without notice.

References in this document to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM program product in this document is not intended to state or imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

The information provided in this document has not been submitted to any formal IBM test and is distributed "As Is" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into their operating environment.

While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in either the United States, other countries or both.

IBM, S/390, ES/3090, ES/9000, AS/400, RS/6000, MVS/ESA, OS/390, VM/ESA, VSE, TPF, OS/2, OS/400, AIX, DFSMS/MVS, DFSMS/VM, ADSTAR Distributed Storage Manager, DFSMSdfp, DFSMSdss, DFSMShsm, DFSMSrmm,
FICON, ESCON, Magstar, Seascape

Other company, product, and service names mentioned may be trademarks or registered trademarks of their respective companies.

Windows NT is a registered trademark of Microsoft Corporation.