



ETI-NET, Inc.

## ***ETI-NET Enterprise Integration Solutions***

### ***Virtual Tape Integration Solutions***

***Presented at***

# ***VNUG***

*September 2006*



# OUTLINE

- § **ETI-NET Background**
- § **ETI-NET Product Family Evolution**
- § **Two solutions:**
  - 4 **EZX/Gateway**
  - 4 **EZX/BackBox**
- § **ETI-Net NEWS**
- § **Summary**



## ETI-NET Background

- § **Montreal-Based**
- § **Worldwide Customers : 1200+ Licenses**
- § **Enterprise Solutions since 1988**
- § **Focus: “Tandem-to-Other” Integration**
- § **Development Partner with HP**
- § **Development Partner with IBM/Tivoli**



## A few customers

- § Wells Fargo Bank
- § AT&T
- § Royal Bank Canada
- § ING Bank
- § Barclays (UK)
- § Washington Mutual
- § Banque de France
- § Citicorp (Germany)
- § Loto-Quebec
- § Park Nicollet Health
- § West Corporation
- § Credit Agricole (France)
- § CSC/SAS (Sweden)
- § KDDI (Japan)
- § Dubai Bank (UAE)
- § Nordea

## ETI-NET Solutions

- § **BCOM<sup>®</sup>:** File Transfer, NSK «-» MVS
- § **BackHome<sup>®</sup>** (“software virtual tape”): NSK Backup to  
“Classic” Mainframes via SNA, TCP/IP  
“BackHome/TSM” TSM via TCP/IP
- § **EZX/BackBox<sup>®</sup>** (“hardware virtual tape”): NSK Tape I/O
  - 4 “Conduit” direct connection to TSM. No intermediary storage.
  - 4 “LocalStore” to virtual tape on local/remote disk
- § **EZX/Gateway<sup>®</sup>** : High Speed File sharing between NSK &  
**MVS** via StorageWorks XP Disk Array as 3390s;  
**Unix** via File Server technology;



ETI-NET, Inc.

# KDDI Japan

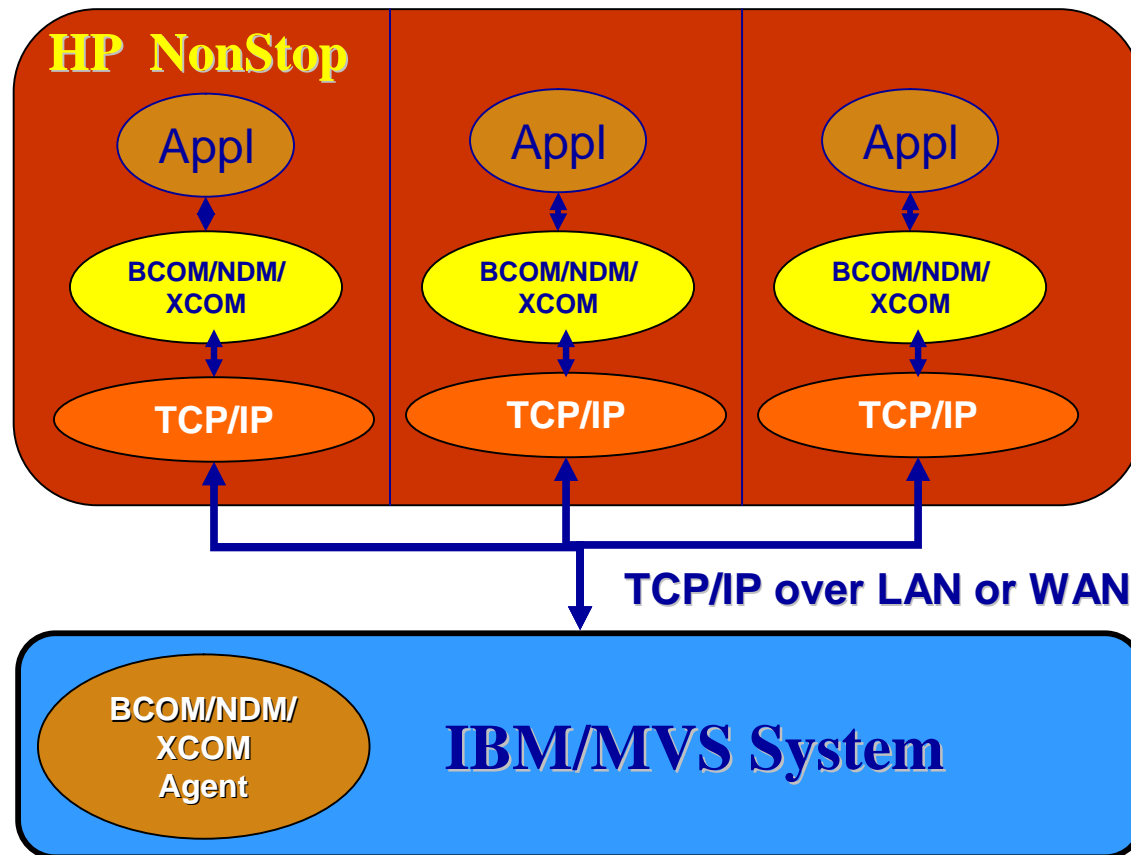
***EZX Gateway File Mover***

## Cross Platform Sharing at KDDI: The Problem

- § Japanese Telecom carrier– end of month processing
- § Move **2.6TB** from NSK to IBM in 2hr window
- § Equals **360 MB/sec**
- § Translate from ACSII to EBCDIC-J en route
- § Assure every file transferred intact
- § Minimize operator actions required on both systems



# Classic Data Transfer



## Problems:

- Slow Speed
- Network Congestion
- NSK CPU Load (TCP/IP)
- User Implemented API

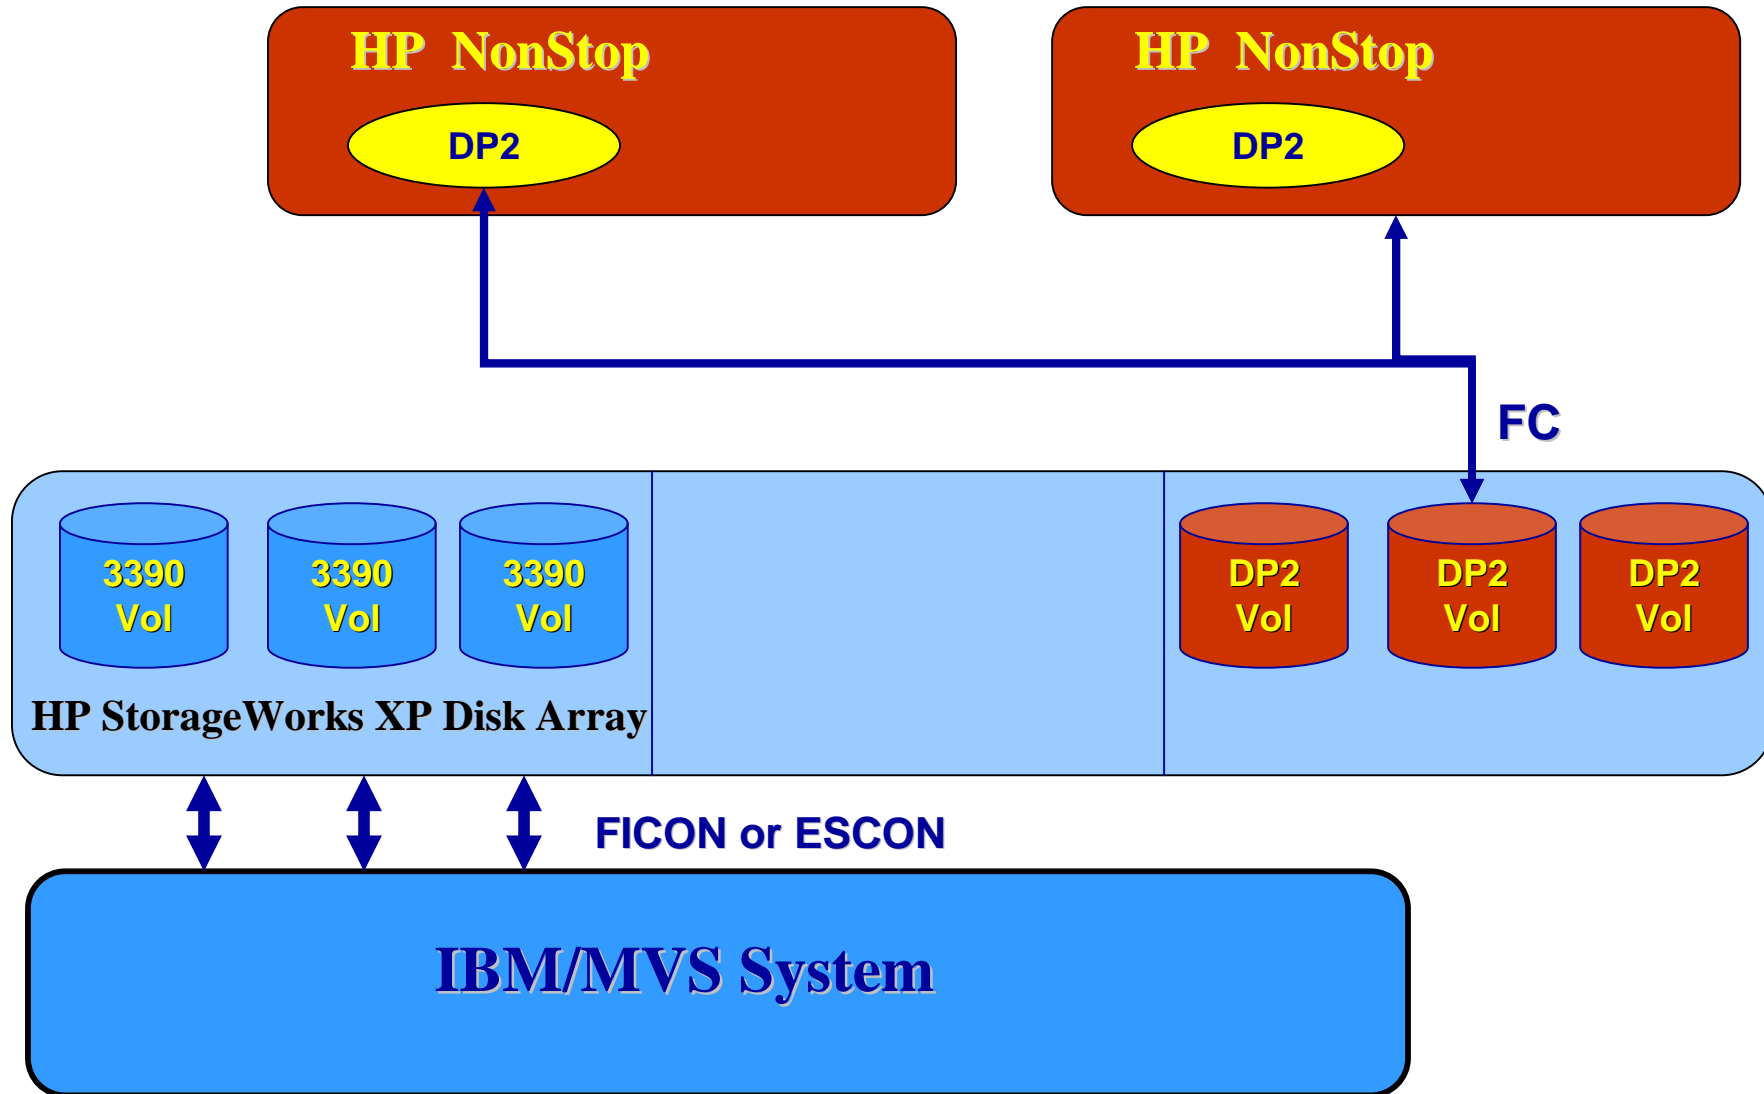


## Disk on NonStop – The Challenge

- § **“DP2”**: Proprietary NSK disk driver – not “open” format
  - 4 integrates database & TMF (data integrity protection) functions
- § **No ability to connect disks other than NSK native**
- § **NSK support for HP StorageWorks XP arrays**
- § **But still no disk sharing with other systems!**
  - 4 **Common use of array does not mean volume sharing!!!**



# Isolated Islands in the Array

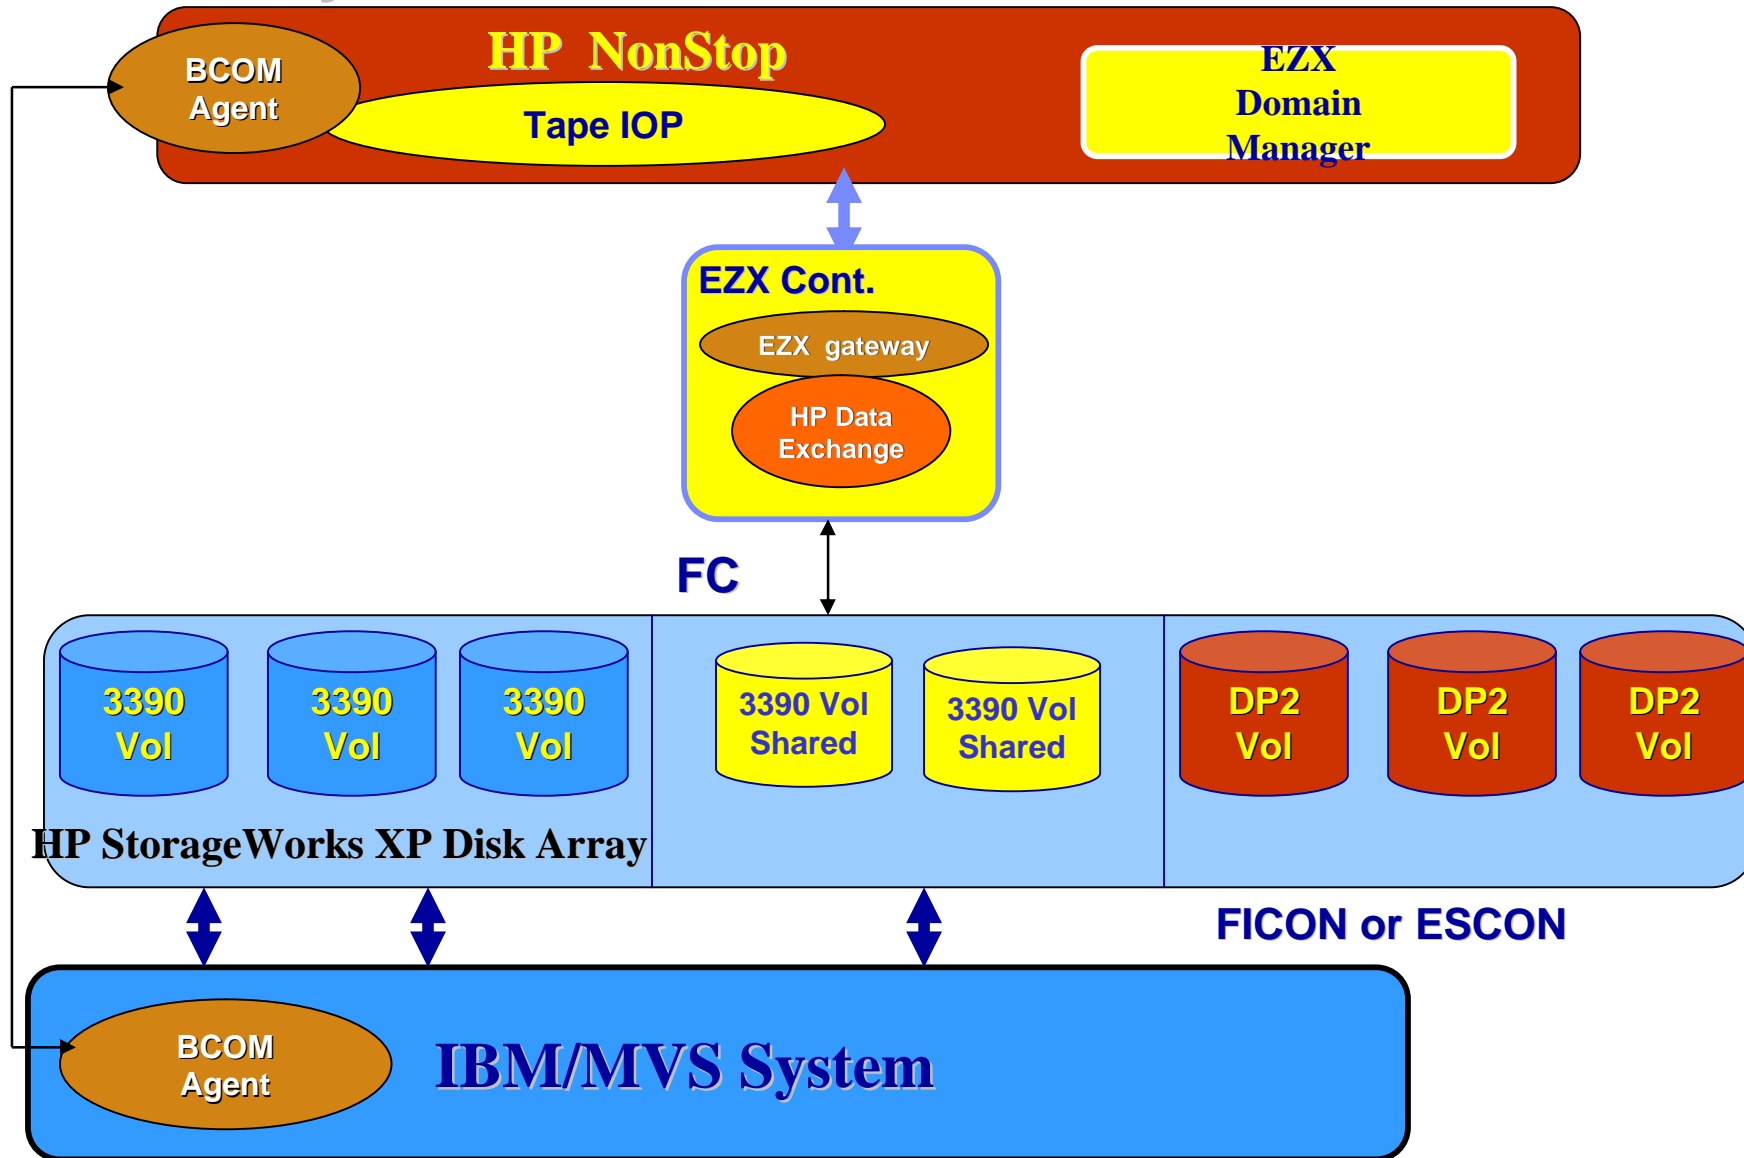


## Elements of the Solution

- § **EZX Virtual tape interface to NSK**
- § **HP StorageWorks XP array as common storage**
- § **HP DataExchange™ API for sharing 3390 volumes**
- § **Fibre Channel & multiple controllers for speed**
- § **EZX Domain Management & GUI**
- § **BCOM for end-to-end management & agents**



# Gateway: Share NSK «-» MVS




## EZX/Gateway® Performance

- § **High throughput per EZX Controller (FC or multiple SCSI)**
- § **Maximize performance using multiple streams/devices**
- § **Avoid bottlenecks: distribute Tape IOPs across NSK CPUs**
- § **Application: Blocksize and NSK disk I/O are key**
- § **Proliant server CPUs: power for Data Exchange processing**
- § **Use Multiple EZX Controllers for linear growth**



## Comparing Transfer Rates – NSK Write



**Writing to XP is 7.1x faster than to JBOD**

**Writing via EZX to XP is 4.8x faster than to JBOD**



## EZX/Gateway Summary

- § **Transfers to/from shared IBM 3390 volumes**
- § **Field proven performance: 60-95 MB/sec per Controller**
- § **Job submission & locking via BCOM**
- § **Assured file delivery & integrity**
- § **Linearly expandable**
- § **Fault-tolerant: Re-routes around failed controllers or FC**
- § **Helps justify cost of XP array for use on NSK & MVS**
- § **Chosen by HP for their largest NS Integrity installation**
- § **In production at KDDI, Japan.**





ETI-NET, Inc.

# Australian DOD

## *EZX BackBox Database Protection*





# Australian Dept of Defense: The Issues

- § **Massive high ingest-rate database**
  - 4 Adds 6.7 TB to database (gross disk space) per week
  - 4 Uses separate groups of 48 partitions for each week
  - 4 Current week's set of partitions backed up daily
- § **24x7 Database ingest & query – mission critical**
- § **Large (Terabytes) and growing much larger**
- § **Based on NonStop platforms**
- § **Must minimize impact of backups on normal use**
  - 4 Schedule for “off-hours, limit backup time window
  - 4 Limit performance impact during backup
- § **Bottlenecks in platforms' I/O architectures**



# HP NonStop Database Backups: A Primer

- § **NonStop SQL databases, often multiple partitions**
- § **Protected by TMF (Transaction Monitoring Facility)**
  - 4 **Periodic “online dumps” to backup TMF protected files**
  - 4 **Audit logs for all DB changes**
- § **Restore to point in time by:**
  - 4 **Restore a given “online dump”**
  - 4 **“Roll-forward” to point in time using audit log contents**
- § **DB partitions can be backed up independently**
  - 4 **Allows for parallelism during dumps & restores**



## The Challenge: How to Get Data Off the Host Platform?

- § **“Native” NSK backup application**
  - 4 HP NSK: TMF “Online Dump” to tape, or Backup/Restore for non-TMF-protected files
  
- § **TSM client**
  - 4 ETI-NET’s BackHome/TSM™: implements TSM API with “software virtual tape” (high CPU load penalty for TCP/IP)
  
- § **Disk Array snapshot mechanism**
  - 4 HP NSK: Not feasible due to NSK proprietary disk format



## The Challenge: Host Platform I/O Limitations

- § **NSK S-series (legacy) systems: SCSI ports**
  - 4 Ultra SCSI (40 MB/sec) on recent system models
  - 4 8 MB/sec on older system models
  - 4 More than 2 SCSI devices/port = diminished performance
- § **NSK Integrity (Itanium-based) systems: FC ports**
  - 4 Multiple ports “networked” across processors
  - 4 Theoretical limit: 200 MB/sec per port
  - 4 Multiple devices/port, connectivity configurable in FC fabric
- § **Blocksize limit:**
  - 4 O/S limit: 56 Kbytes (52 or 32 Kbytes for some applications)



## EZX Controller

- § **Windows 2003 Server platform**
- § **Fibre Channel or SCSI to NSK systems**
- § **Key ingredient: ETI-NET tape emulation engine**
- § **Automates tape mounts**
- § **Up to 32 tape devices emulated per controller**
- § **Connect to multiple NSK CPUs or systems**
- § **Controllers can function as (mix-and-match):**
  - 4 **EZX/BackBox - VT on local, SAN or FS storage**



## EZX/BackBox®

- § **Uses EZX Controllers as virtual tape subsystem**
- § **Totally transparent to NSK system operation**
  - 4 **Use Backup/Restore, TMF or any application**
- § **Integrated with DSM/TC & TMF catalogs**
  - 4 **Creating media => automatic cataloging**
  - 4 **Security => checks NSK security before load**
  - 4 **Deletion protection => no delete unless scratch**
- § **Easy, intuitive subsystem management**



## **EZX<sup>®</sup> Fault Tolerance**

- § **All management metadata stored on NSK (TMF prot'd)**
- § **EZX Controllers are context-free**
- § **Multiple paths can be defined to each data store**
  - 4 **Gigabit Ethernet between EZX Controllers**
- § **Traffic automatically re-routed by Domain Manager**
  - 4 **To other EZX Controllers**
- § **Data stores can be on SAN-based storage**
  - 4 **If a controller fails, just replace it & re-connect to storage**



## **EZX/BackBox<sup>®</sup> (“LocalStore”)**

- § **Tape images on local or SAN-based storage**
  - 4 **Cost-effective storage on local disks**
  - 4 **SAN storage for fault tolerance & centralized management**
  
- § **Integrate with TSM “back-end”**
  - 4 **Scriptable**
  - 4 **Asynchronous migration of tape images**
  - 4 **Space can be freed on local disks once image is migrated**
  - 4 **LAN-Free option for high-speed movement to managed tape**





# Massive Online Database Backup: The Approach

- § **Avoid using TCP/IP on NSK due to CPU overhead**
- § **Use fastest I/O ports available on NSK system:**
  - 4 **FC on NS-Integrity series, SCSI on NS-S-series**
- § **Use tape I/O architecture**
- § **Use maximum block size**
- § **Maximize parallelism (number of backup streams)**
  - 4 **Use multiple target devices (within limits of I/O channels)**
- § **Maximize throughput of virtual tape subsystem**
  - 4 **Stage to fast disk**
  - 4 **Configure virtual tape platform for optimal I/O throughput**



## **BUT... Minimize Disk Size Req'd**

- § **Not cost-effective to replicate host's disk capacity**
- § **Must stage to physical tape in parallel w. NSK offload**
- § **Balance the # of physical tape drives against:**
  - 4 **Number of parallel streams from host(s)**
  - 4 **Sustained data rate of each stream to VT subsystem**
  - 4 **Maximum data rate sustainable by each tape drive**
- § **Optimize amount of staging disk required:**
  - 4 **Sequence virtual tape volumes onto physical tape**
  - 4 **Delete staged virtual tapes that have migrated to phys tape**



## The Role of TSM

§ Move staged virtual tape volumes to/from physical tape

§ Uses LAN-free approach:

### 4 TSM Server:

§ Manages pool of tape drives via FC fabric

§ Allocates tape drive to TSM client in BackBox

§ Controls tape library robotics

§ Manages tape media pool

§ Selects and mounts tape volume on drives & verifies label

§ Notifies TSM Client of ready drive/media

### 4 TSM Client:

§ Connects to tape drive via FC fabric when told by TSM Server

§ Writes Virtual Tape volumes to tape

§ Notifies TSM server at end of media or completion of backup



## The solution

- § **Uses 3 x 16-proc NS-Integrity systems (Itanium based)**
- § **Backup via 6 (5 in case of a failure) EZX Controllers**
  - 4 **Emulate 12 tape drives / FC connection / Controller**
  - 4 **At TMF sustained thruput/stream ave. 9MB/sec = **108 MB/sec****
  - 4 **Across 5 EZX controllers = 540 MB/sec**
  - 4 **Backup would take 3.7 hours**
- § **BUT ... this only considers transfer to BackBox disk**



## The solution contd.

### § Migration to tape

4 Uses LTO-3 Tape drives, capable of **~80 MB/sec** sustained

4 If use 1 x LTO-3 drive per EZX controller at a time:

§ End-to-end limit becomes 80 MB/sec rather than 108 MB/sec

### § Enough disk space to buffer the difference in transfer rate.

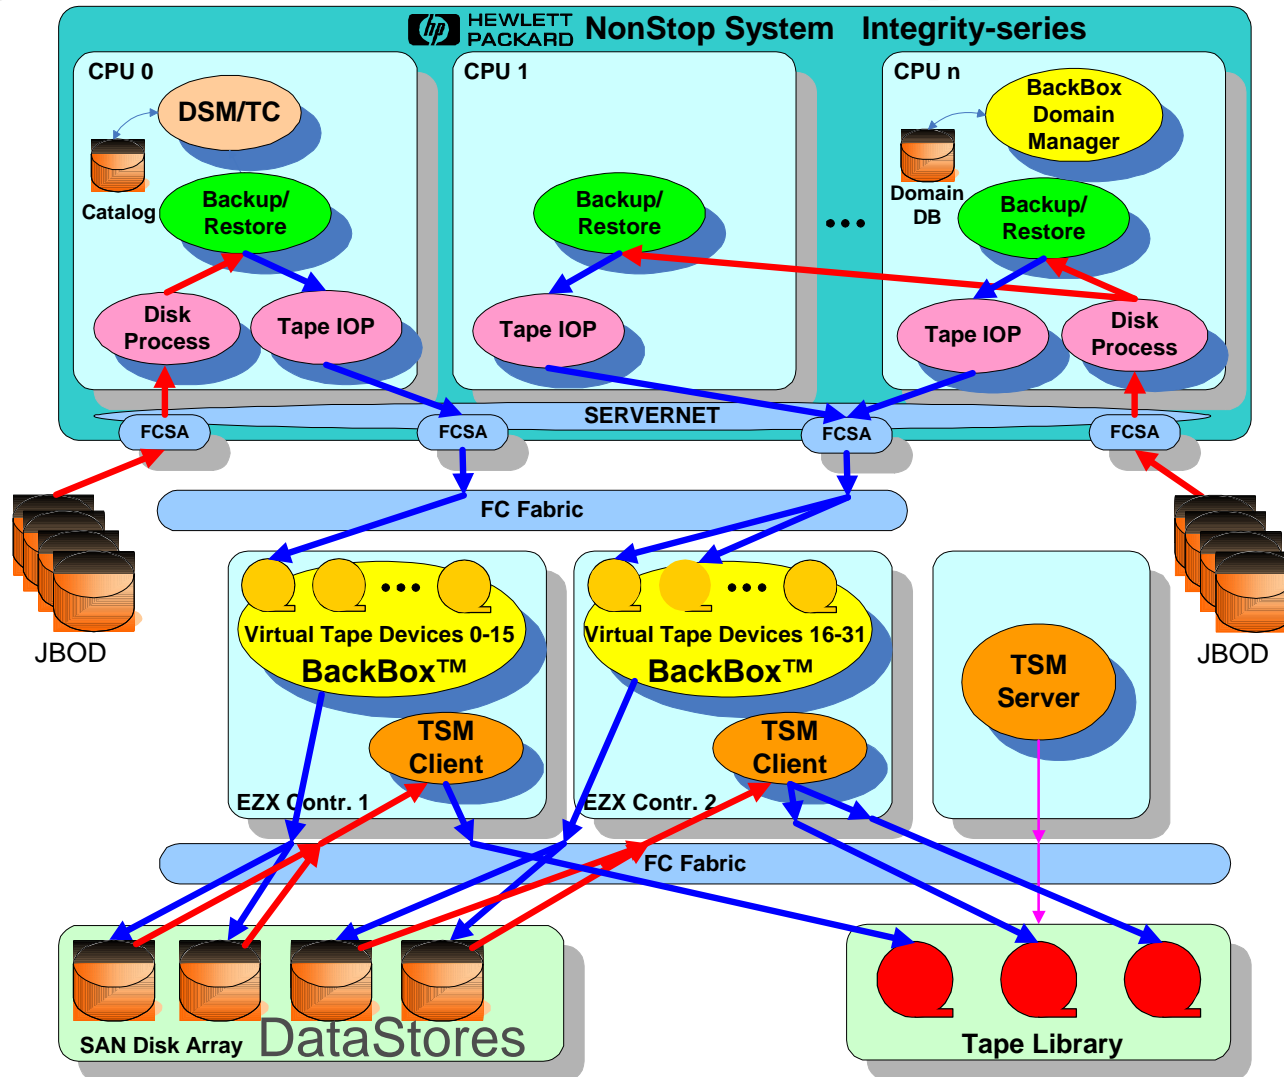
4 6 EZX controllers w. 1 tape drive each = **480 MB/sec**

4 But 5 (in case of 1 failure) EZX Controllers = **400 MB/sec**

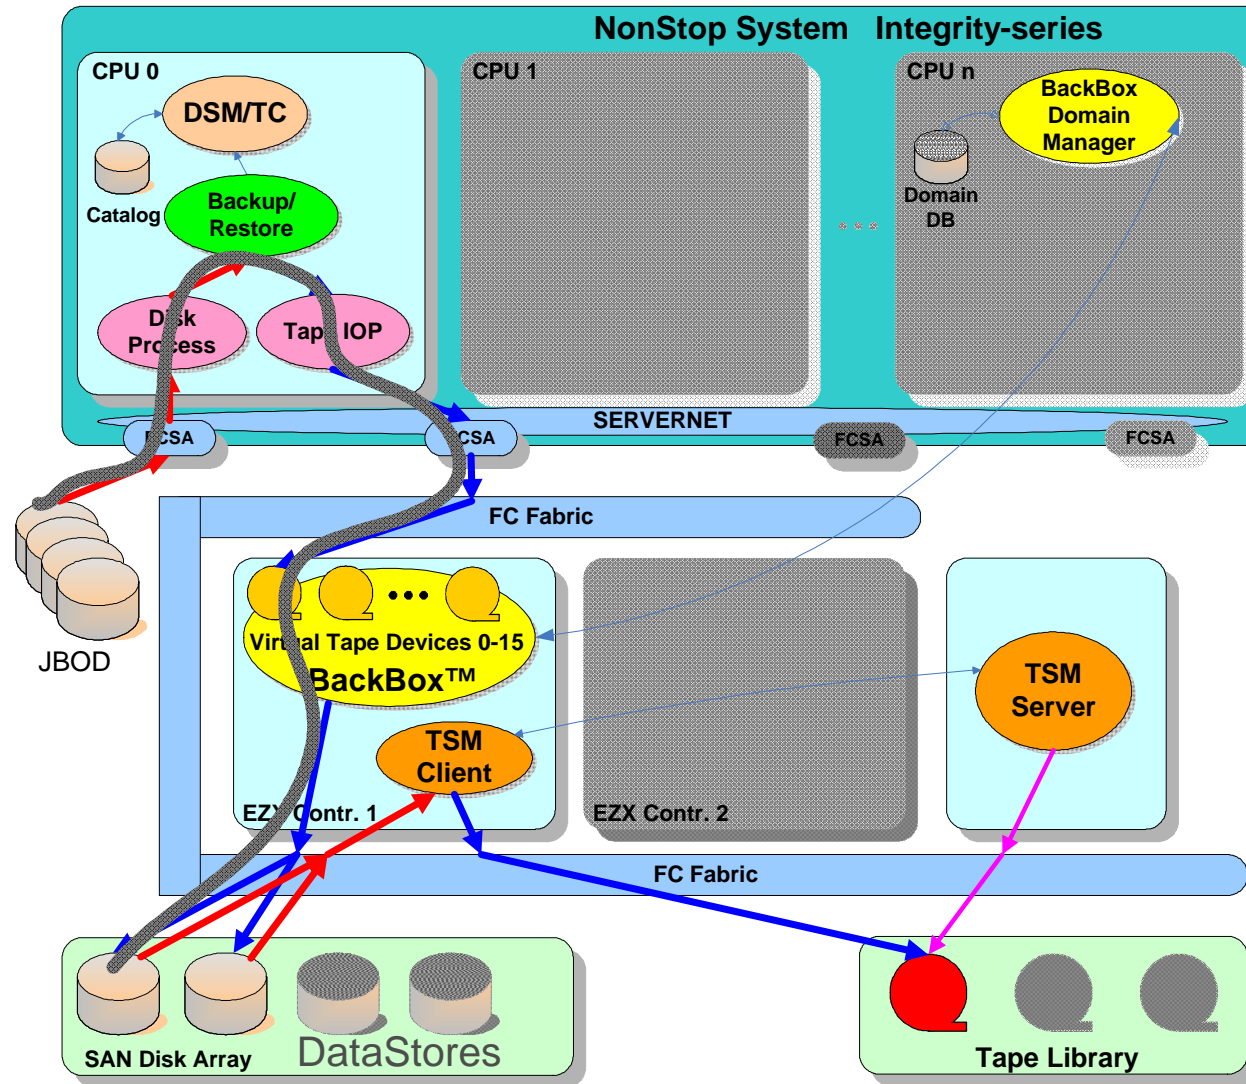
4 Backup of **6.7 TB** will take **~5.0 hours**



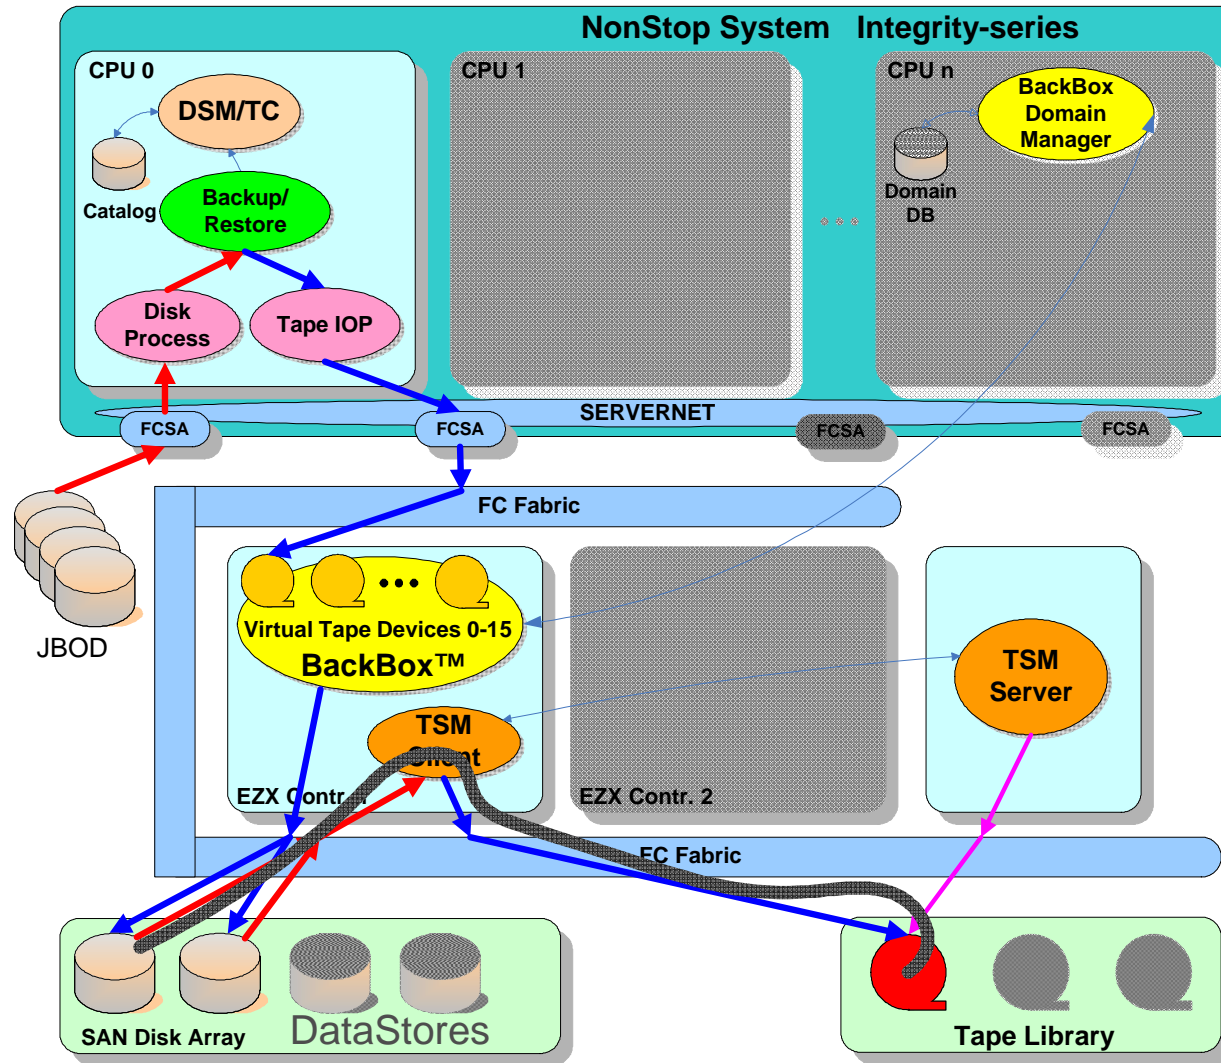
# Staging with TSM LAN-Free to Tape



# Data Flow: Staging to Virtual Tape



# Data Flow: TSM - Virtual to Physical Tape





## Key Points

- § **Approach is linearly expandable**
  - 4 **Add EZX controllers + physical tape drives**
  - 4 **Grow BackBox subsystem to match host DB growth**
- § **Architecture is fault-tolerant**
  - 4 **Failure of an EZX Controller or FC connection is reconfigured around by the BackBox Domain Manager**
- § **TSM Server platform:**
  - 4 **Lightly loaded since no data flows via it**



## Key Points (cont'd)

- § **Optimization of EZX controller platform important**
  - 4 **I/O throughput (instantaneous) during backup is:**
    - § **108 MB/sec Input via FC from NSK host(s) to SAN**
    - § **80 MB/sec Input from disk (TSM Client) to TSM controlled physical tape drives.**
- § **Similarly, disk subsystem configuration important**
  - 4 **Overall Write rate = 6 x 108 MB/sec = 648 MB/sec**
  - 4 **Simultaneously, overall Read rate = 6 x 80 MB/sec = 480 MB/sec**
  - 4 **Overall I/O rate is ~22,000 per sec**



## EZX/BackBox Benchmark Results

### § TMF (rates are per EZX Controller)

- 4 Online Dump: 12 x 9.7 MB/sec ave. = 108 MB/sec
- 4 Dump Restore: 12 x 6.2 MB/sec ave. = 70.0 MB/sec

### § Environment:

- 4 Itanium 16 processor system, 1 x FC connection
- 4 Tape IOP in each of 12 processors, 12 emulated tape devices
- 4 52 KB Block size, 10 GB Virtual tape size, ~230 GB dumped
- 4 Virtual media on StorageWorks XP12000 array

**Note:** Throughput includes overhead times for all mounts, etc.  
(actual transfer rates= 116 MB/sec Write, 75 MB/sec Read)



## EZX/BackBox Summary

- § **The only truly integrated virtual tape for NSK. *Interface with Mediacom, DSM/TC and TMF through standard Guardian API***
- § **Itanium & S-Series compatible**
- § **Ultra-fast and linearly expandable**
- § **Fault-tolerant: configured for no single point of failure**
- § **Can automatically reclaim storage space when Mediacom or TMF/DR set a virtual tape to SCRATCH.**
- § **Emulates a multi-tape drives silo with an infinite number of slots.**
- § **Supports direct connection to the TSM Server. Certified READY for IBM TIVOLI Software.**
- § **Supports ANSI, BACKUP, NL, TMF, IBM tape labels.**
- § **Easy configuration. Customer configurable. Number of devices, target storage can be changed by the customer. Very flexible.**
- § **Licensed per EZX BackBox and NonStop node. Any number of port or tape drive can be configured by the customer.**
- § **BackBox is context free**
- § **Disaster recovery based on proven standard NonStop operation procedure.**



## New Alliances – New Solutions

**EMC<sup>2</sup>**  
where information lives

**data domain**

BackHome/TSM XP for TMF



**EMC<sup>2</sup>**  
where information lives

**EZX** backbox



## Legato Datastore

The logo for EMC² | legato, featuring the text "EMC²" in a bold, sans-serif font, followed by a vertical line and the word "legato" in a lowercase, sans-serif font, all set against a red background.

- § EZX BackBox direct connection to Legato Networker
- § No disk staging
- § Data flow directly from the NonStop through the EZX BackBox and on Legato Networker storage
- § To be released 4Q06




# data domain






# Industry Shifting from Tape to Disk

Audio



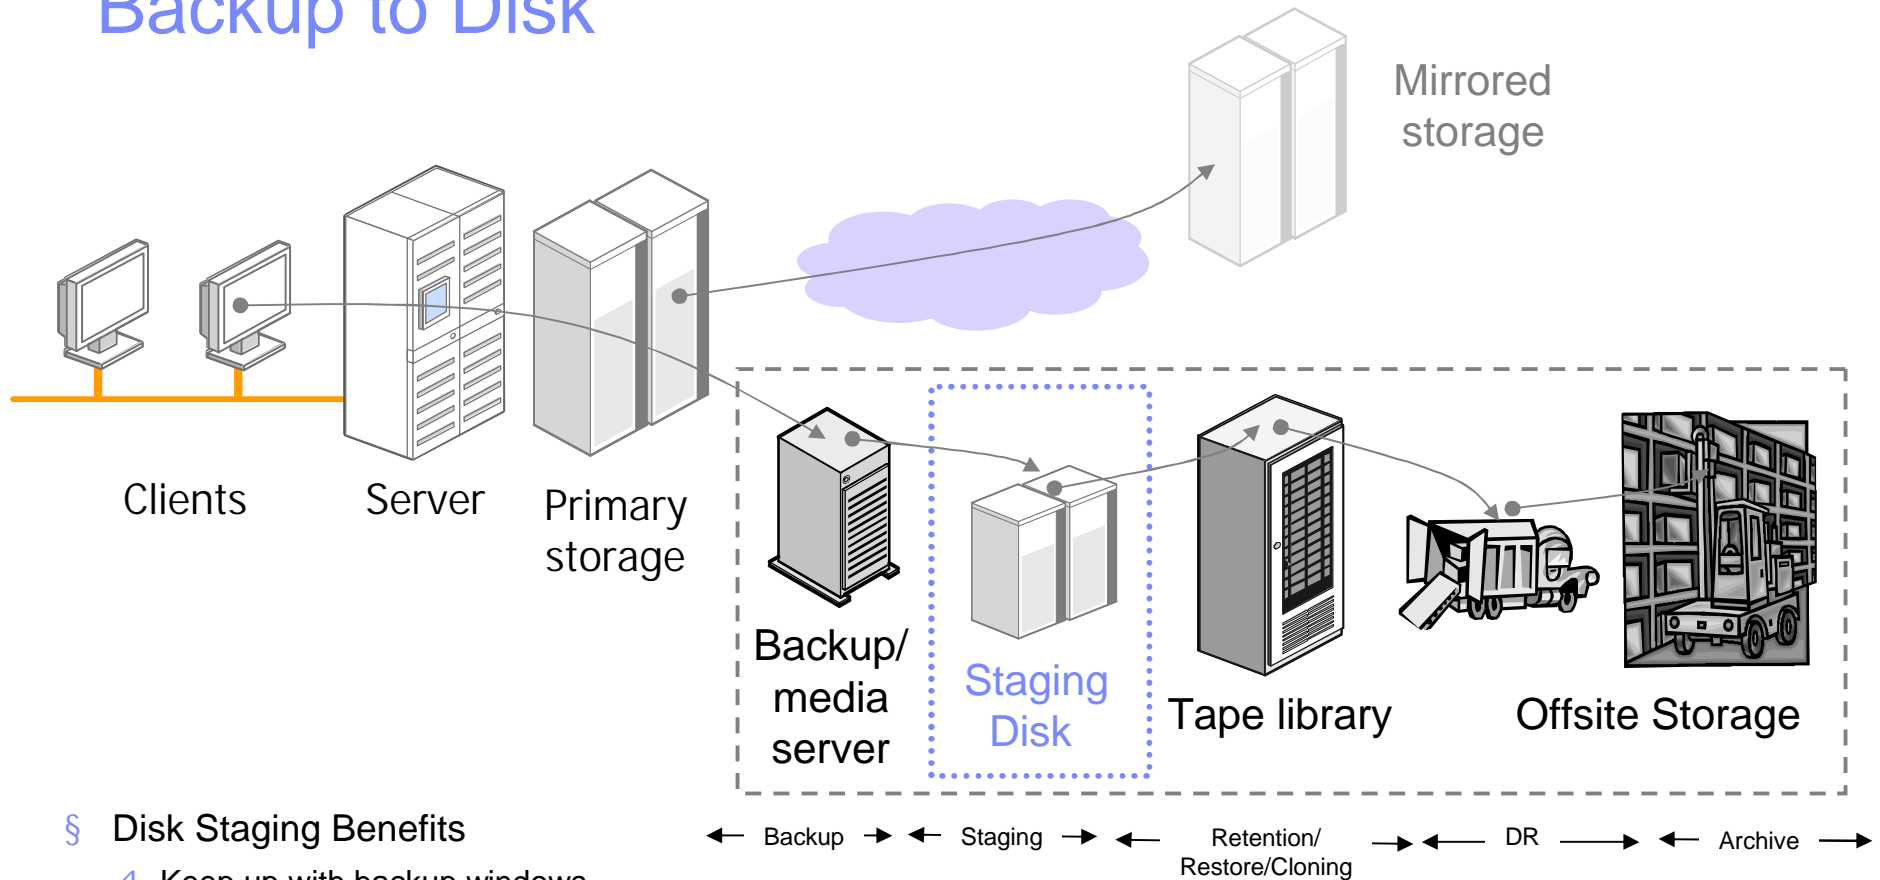
Video



Data Protection




# Backup to Disk



## § Disk Staging Benefits

- 4 Keep up with backup windows
- 4 Optimize tape automation
- 4 More reliable backup
- 4 Some restores from disk



## Barrier to Disk Recovery Storage: *Capacity Needed*

**Primary Data**

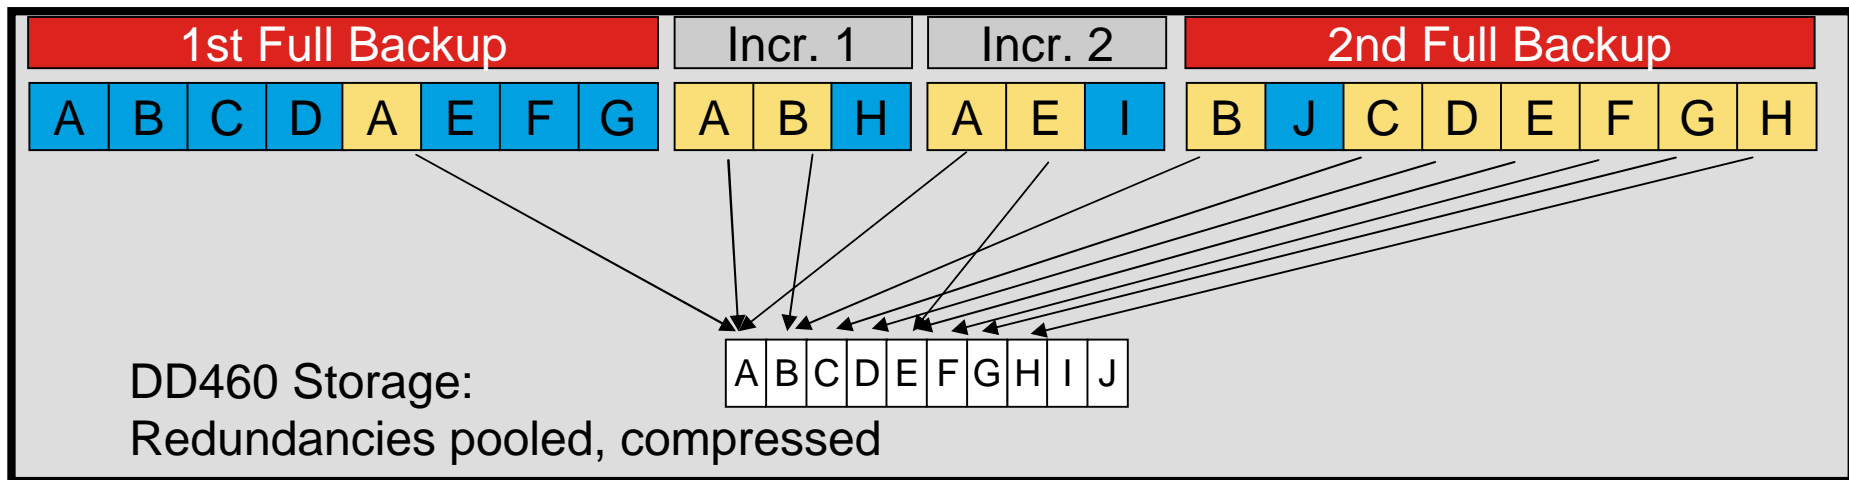
**Tape Library  
Capacity:  
5x - 10x Primary**



# Use Disk Longer

NFS or  
CIFS

View from Backup Software:  
Up to 230 TB

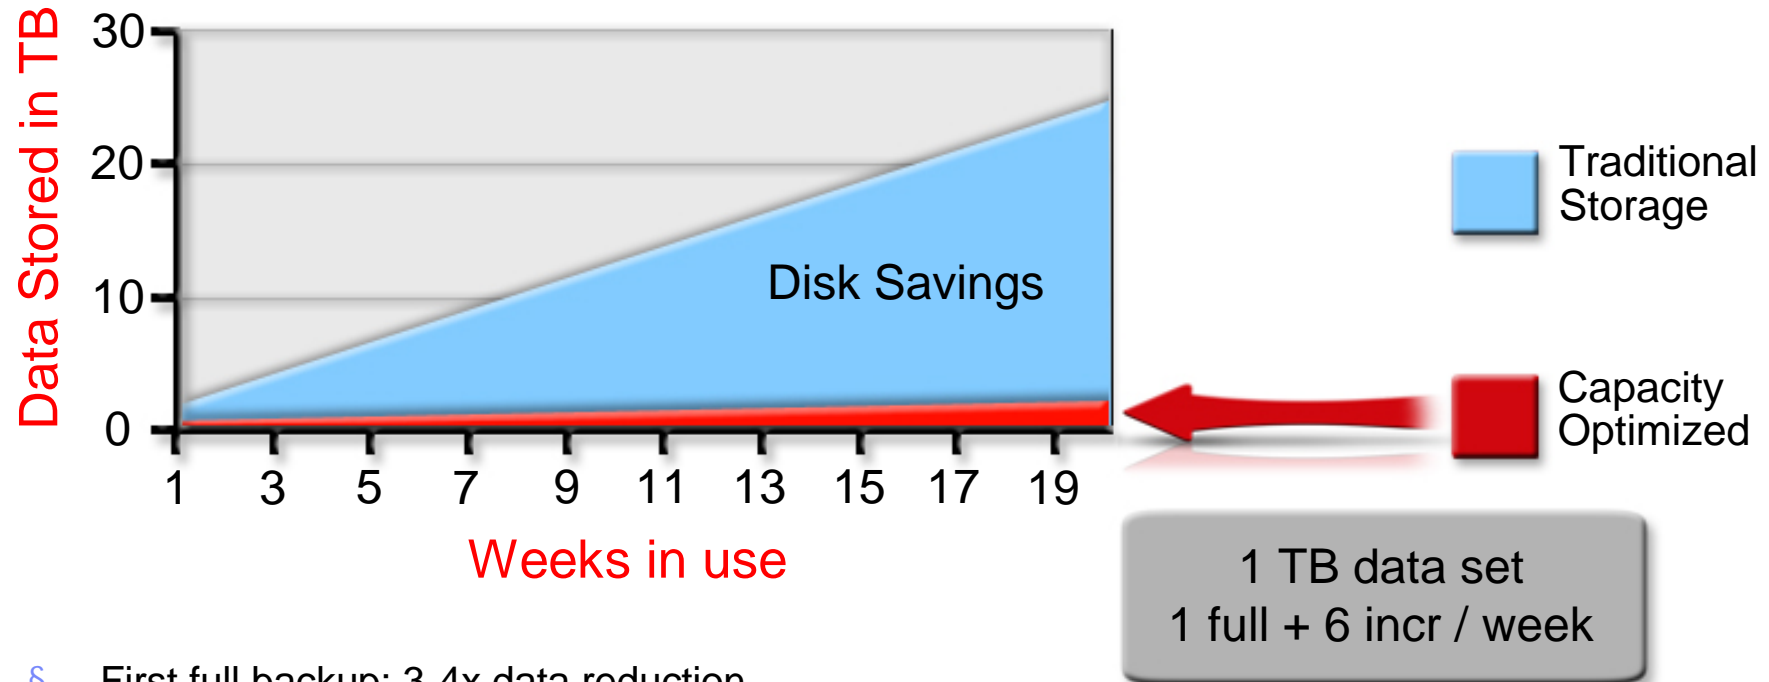


## DD460 Global Compression

- = Unique variable segments (4KB-12KB)
- = Redundant data segments
- = Compressed unique segments



# Data Domain Capacity Optimization Process



- § First full backup: 3-4x data reduction
- § File-level incrementals: 6-7x
- § Subsequent full backups: 50-60x
- § Aggregate with Weekly Fulls, Daily Incrementals: 20x



# data domain

Direct connection to DataDomain controller

First tests gave a 50x reduction. YMMV – Maybe higher

Next version of EZX BackBox and next version of DataDomain 4Q06





ETI-NET, Inc.


## BackHome/TSM XP for TMF

***FAST and EFFICIENT***




# BackHome/TSM XP The differences

## BackHome/TSM



Compression happens here on 32KB blocks

## BackHome/TSM XP

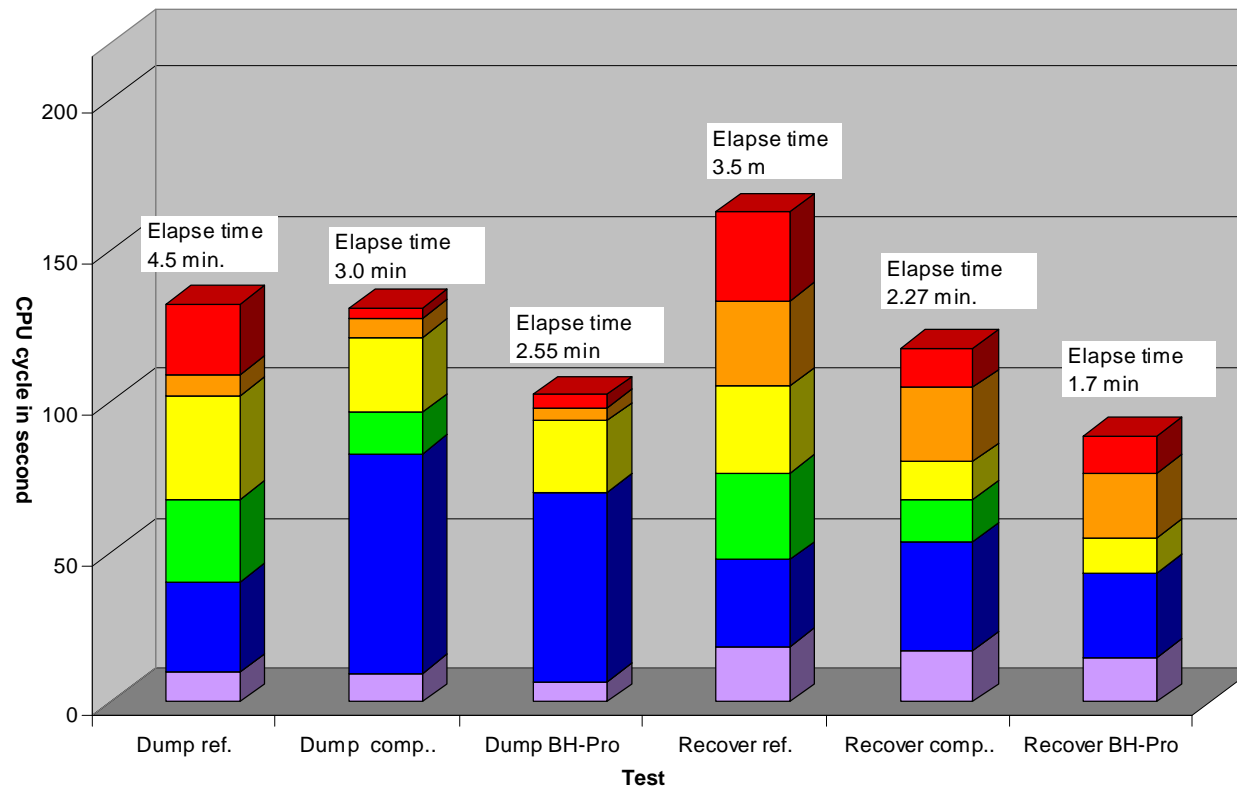


Compression happens here on 52KB blocks



# BackHome/TSM XP for TMF

CPU cycles used by processes



TMFDR
  ETMF
  BFSxxx
  TSMREQ
  DP2
  TCPMON

## BackHome/TSM XP for TMF – Some results

### § No compression

4 38% less CPU cycles

4 13% less time

### § Data compressible at 2.5

4 8% less CPU cycles

4 42% less time

### § Data compressible at 9.7

4 47% less CPU cycles

4 60% less time



# Summary

- § **EZX products from ETI-NET for:**
  - 4 **Cross-platform transfers (EZX/Gateway)**
  - 4 **Storage manager consolidation (EZX/BackBox)**
  - 4 **Virtual tape on local or SAN array storage (EZX/BackBox)**
- § **Ultra high-speed**
  - 4 **Fibre Channel + parallelism**
- § **Fault Tolerant**
  - 4 **Context-free controllers & NSK-based management database**
- § **Developed by experts in NSK, MVS & open systems**
  - 4 **Many years of cross-platform and backup consolidation experience**
- § **Lets you leverage your storage investment**



IT REALLY WORKS!!!



## Contact

**Mirko.Buzolitch@etinet.com**

**+1 514 484 3597**

**www.etinet.com**

