

Soluciones de Cintas y Deduplicación de datos

IBM Briefing 2010 System Storage

...lo último en almacenamiento



Paulo Mauricio Pereira Jr
LA Deduplication Specialist



RIF: J-00019078-0



Agenda

- What's new
- LTO 5
- TS3500 High Density Frame
- Virtual Tape Library
- Data Deduplication
- ProtecTIER®
- Customer Case Studies
- Deduplication at a Glance



IBM Tape Density Achievement

35TB Tape

- Jan. 2010: IBM Research - Zurich recorded data onto an advanced prototype tape developed by Fuji Photo Film Co. of Japan
- Density of 29.5 billion bits per square inch -- about 39 times the data density of today's popular standard tape
- Could produce future tape cartridge holding up to 35TB
- New critical IBM technologies were developed:
 - Dramatically improved precision of controlling the position of the read-write heads
 - More than 25-fold increase in the number of tracks
 - New detection methods to improve the accuracy of reading magnetic bits
 - New low-friction read-write head
- Represents a step towards developing technologies to achieve tape areal recording densities of 100 billion bits per square inch and beyond



Tape's Future is Strong

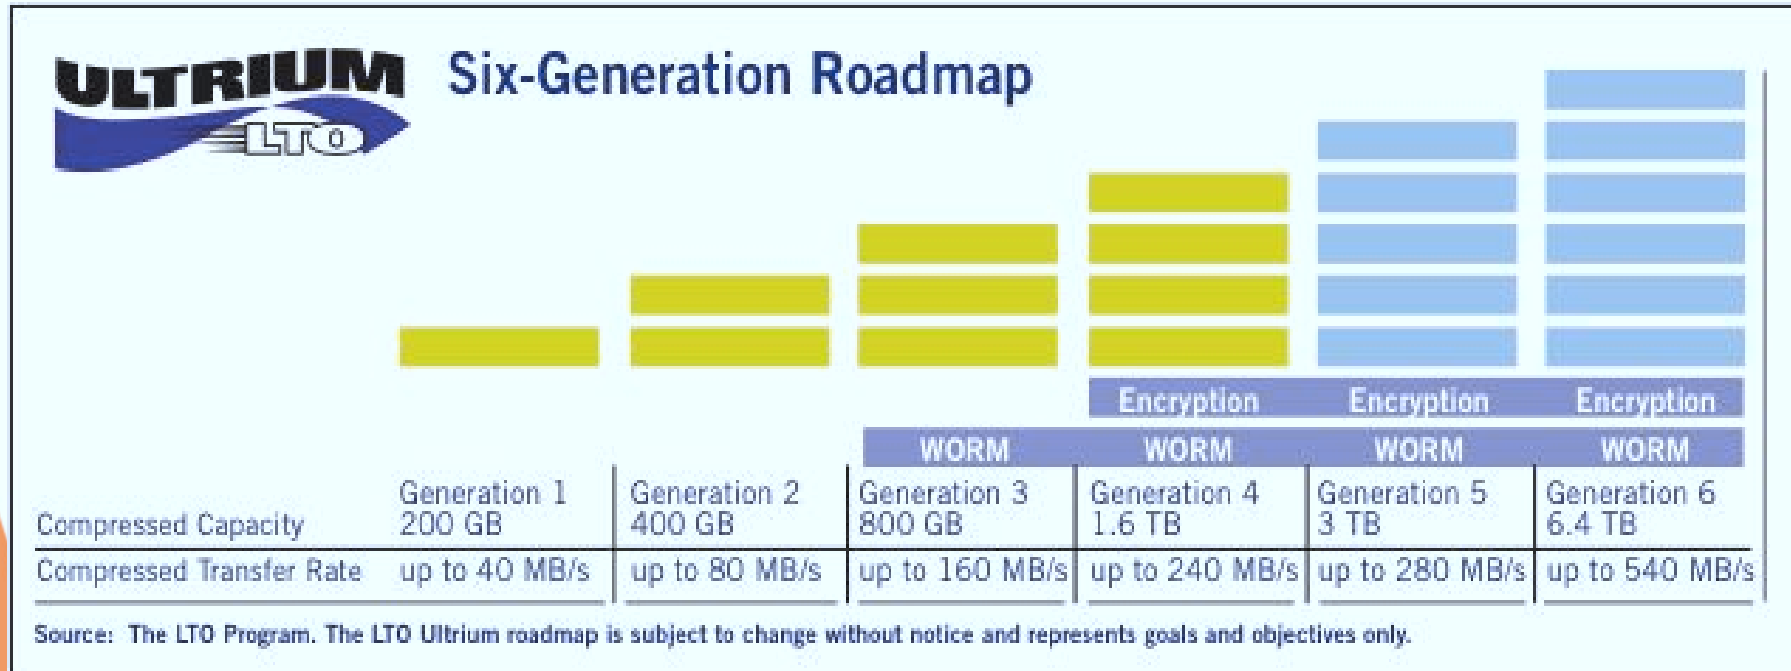


RIF: J-00019078-0




*Press release planned for Jan. 2011

LTO Tape generation



- As of December 2009,**
- **Over 140 million LTO tape cartridges sold**
 - **Over 3.2 million LTO tape drives sold**
 - **Four generations have been shipped**

New IBM LTO Generation 5



88% More capacity compared to LTO4 Cartridge.

1.5TB uncompressed Data.

14% Faster than the LTO4 Drive.

57% Faster than the LTO3 Drive.

140 MB/s Uncompressed.

48% Energy reduction than previous model on Idle mode.

Double the connection on the previous models.

Backward Compatibility LTO3 and LTO4

Native Encryption

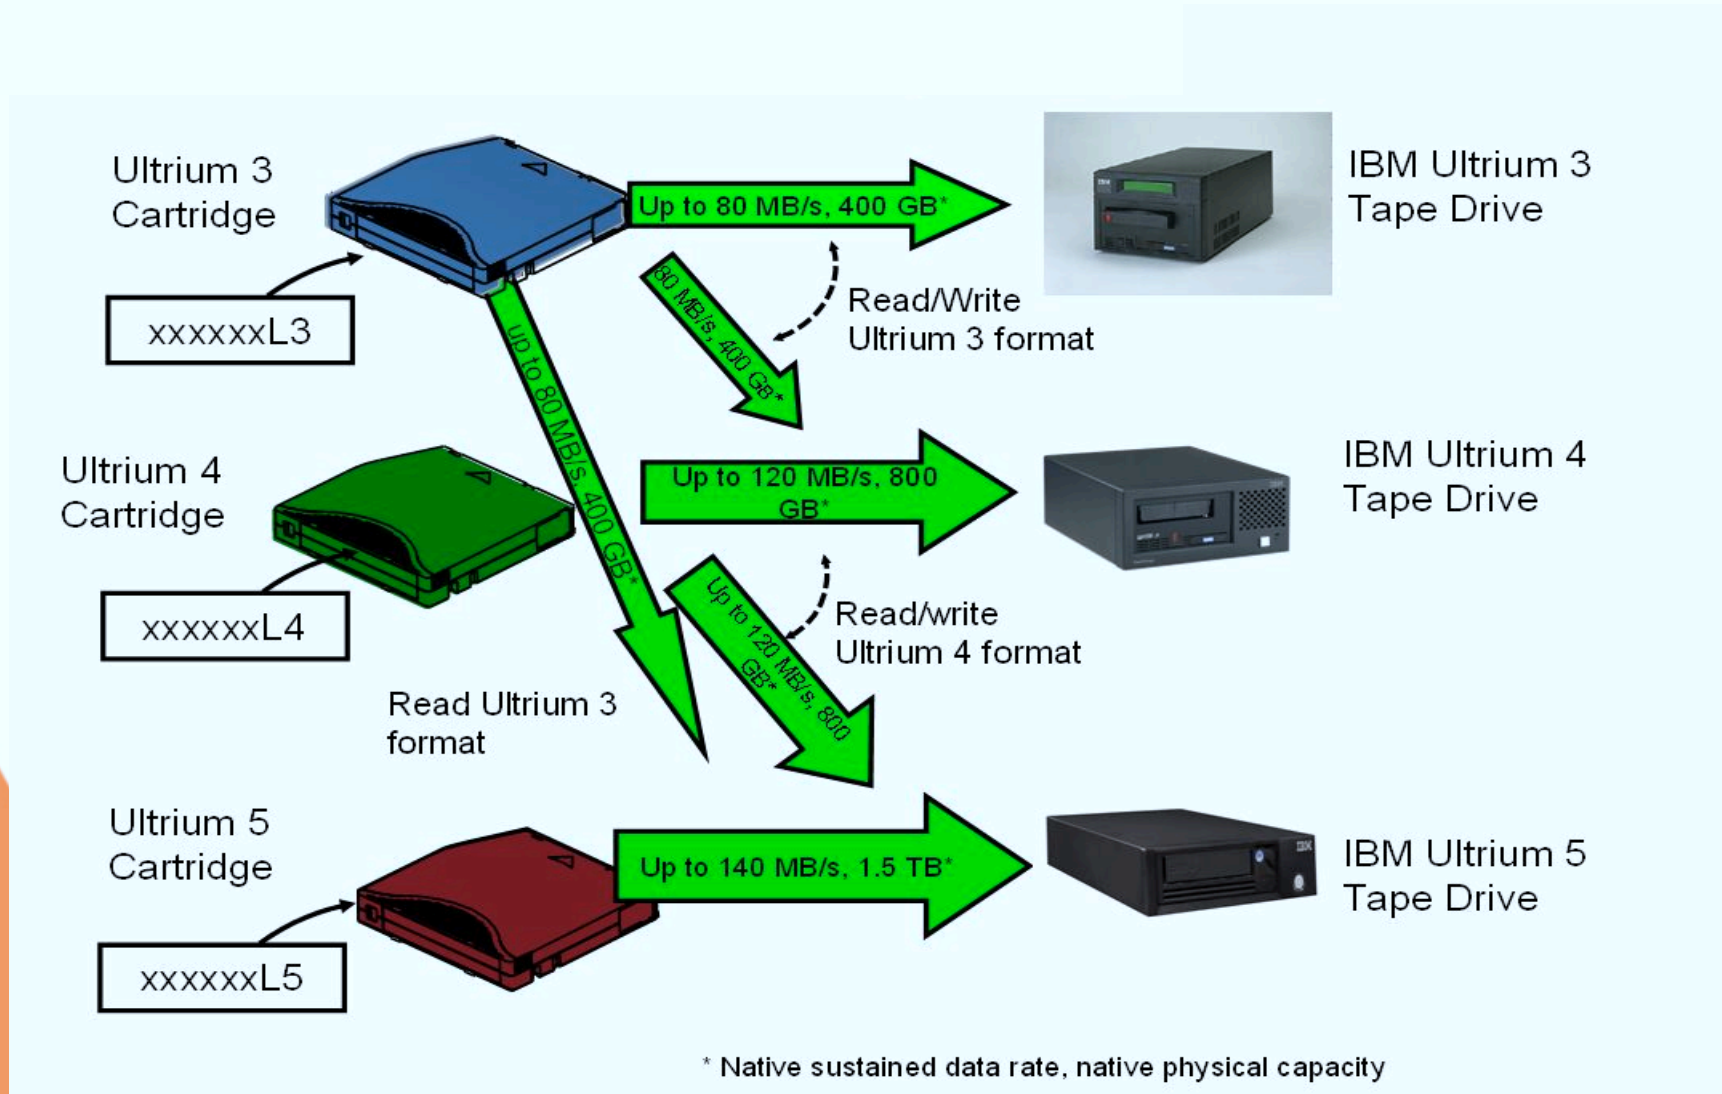
1GB copied on less than 8 seconds !



RIF: J-00019078-0



IBM LTO 5, 4 and 3 compatibility



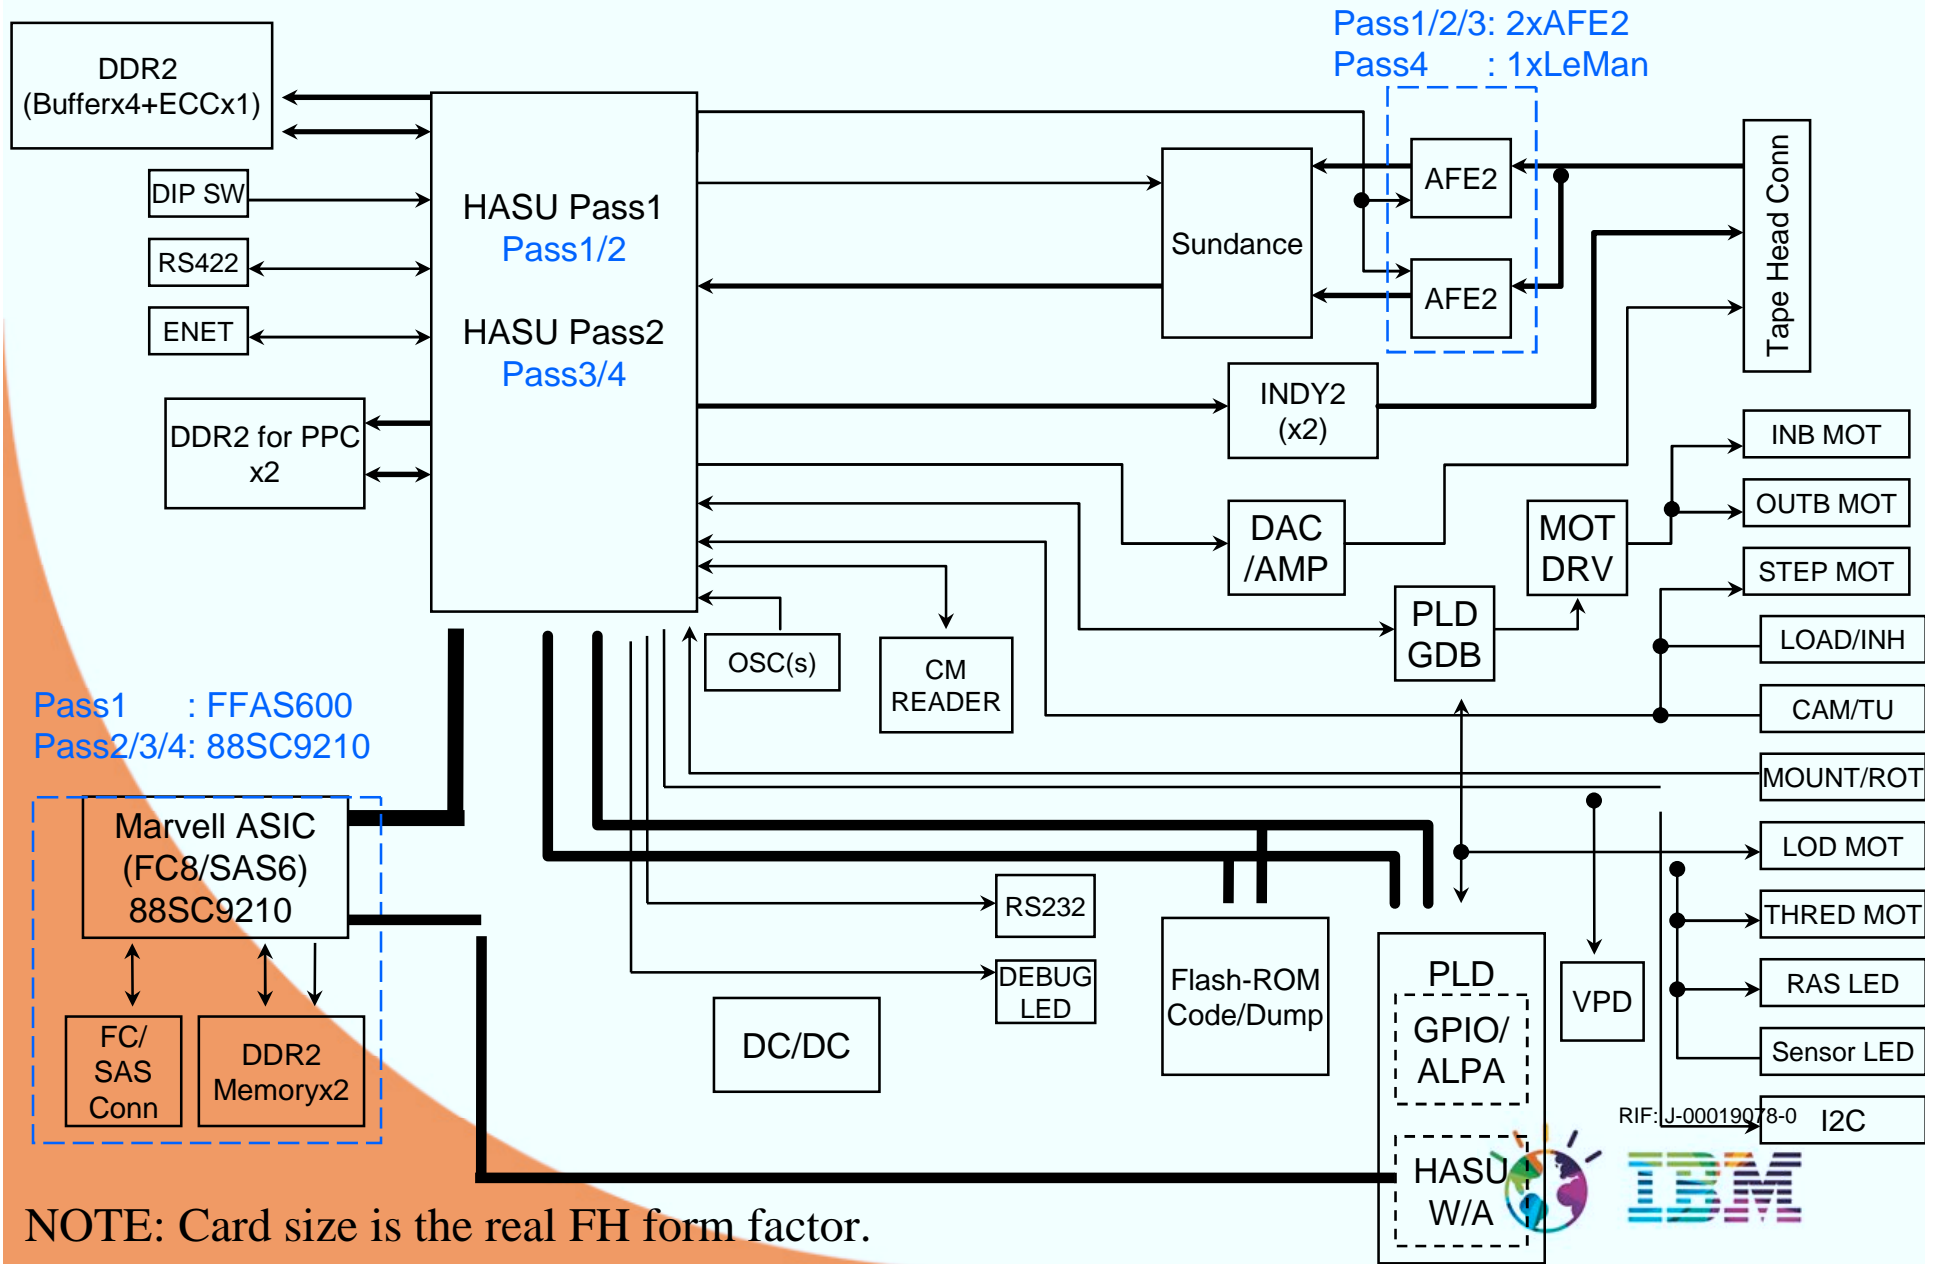
IBM LTO 5 Encryption indicator



- New Encryption LED
 - Only on LTO5 HH drive
- White key symbol
- Indicates encryption status
- Illuminated when a cartridge is loaded and all data on the cartridge is encrypted

A simple view of IBM LTO 5 tape drive

IBM LTO 5 Card block Diagram



NOTE: Card size is the real FH form factor.

IBM LTO 5 Connections

SAS/FC ports (2)




Ethernet port

IBM Long Term File System


The Smarter tape

- Efficient: can utilize new dual-partition planned for **LTO Gen-5** (next TS1100)
- Index partition and data partition: Mount a tape as if it was a hard drive
- Self describing tape enabling hierarchical directory structure, file names, file properties, metadata files, fast search indexes, domain-specific information




Potential Applications:

- *Video archive and individual scene access*
- *Design and drawing content for manufacturing, architecture, and more*
- *Medical / Health industry images*




RIF: J-00019078-0

Flangeless Tape Path




- Previous IBM LTO drives (and competitors' drives) use flanges on the tape guide rollers to limit Lateral Tape Motion (LTM).
- Constraining tape this way has the potential to damage tape edges and deposit debris on guide roller flanges.

RIF: J-00019078-0




Flangeless Tape Path



IBM LTO5

- **IBM LTO5 Full High Drive introduces completely flangeless tape path.**
- **Since tape is no longer constrained by roller flanges, media edge damage is minimized which reduces debris generation.**

RIF: J-00019078-0



IBM LTO G5 and G4 comparison

Function	IBM LTO 4 FH	IBM LTO 4 HH	IBM LTO 5 FH	IBM LTO 5 HH
Capacity (Native)	800 GB	800 GB	1.5 TB	1.5 TB
Data rate	120 MB/s	120MB/s	140 MB/s	140 MB/s
Max Sustained transfer rate	140MB/s(SCSI) 240MB/s (SAS) 280MB/s (FC)	240MB/s (SAS) 280MB/s (FC)	450MBS/s(SAS) 500MB/s (FC)	400MB/s (SAS) 450MB/s (FC)
Host Interfaces	SCSI 160&320 SAS 3Gb FC-4Gb	SAS 3Gb SCSI320 FC-4Gb	SAS 6Gb FC-8Gb	SAS 6Gb FC-8Gb
Speed Matching	30 to 120 MB/s	30 to 120 MB/s	40 to 140 MB/s	40 to 140MB/s
Encryption/WORM/Partition	Yes/Yes/No	Yes/Yes/No	Yes/Yes/Yes	Yes/Yes/Yes
Buffer	256 MB	256 MB	512 MB	256 MB
MTBF	250K Hours	250K Hours	250K Hours	250K Hours
Load/Unload MSBF	100K	100K	100K	100K
Media to Load / Unload	12 sec / 17 sec	12 sec / 17 sec	12 sec / 17 sec	12 sec / 17 sec
Load/Unload Life *	3x	1x	3x	1x
Rewind speed	8 m/s	6.4 m/s	10 m/s	6.4 m/s
Locate speed	8 m/s	6.4 m/s	10 m/s	6.4 m/s
Acceleration	10 m/s^2	5 m/s^2	10 m/s^2	5 m/s^2
Back hitch time	2.5 s	4.5 s	2.5 s	4.5 s
Average file access	57 s	72 s	50 s (TBD)	72 s

IBM LTO G5 HH Product family - 2010



TS2250 (3580) H5S External Tape Drive

- One **LTO Ultrium 5 HH** Drive
- SAS attach
- One Cartridge
- Desktop or Rack Mount
- LTO **Gen 5 Media**



TS2900 (3572) 1U Tape Autoloader

- One **LTO Ultrium 5** Drive
- HH SAS
- 9 Data Cartridge Slots
- 1-slot I/O
- Standalone or Rack
- LTO **Gen 5 Media**
- LTO **Gen 5/4 Encryption**



TS3100 (3573) 2U Tape Library

- 1-2 **LTO Ultrium 5 HH** Drive Feature
- HH SAS or Fibre
- 24 Data Cartridge Slots
- 1-slot I/O
- Standalone or Rack
- LTO **Gen 5 Media**
- LTO **Gen 5/4 Encryption**



TS3200 (3573) 4U Tape Library

- 1-4 **LTO Ultrium 5 HH** Drive Features
- SAS or Fibre
- 3-slot I/O
- 48 DataCartridge Slots
- Standalone or Rack
- Multi-Path Architecture
- LTO **Gen 5 Media**
- LTO **Gen 5/4 Encryption**



RIF: J-00019078-0



IBM LTO G5 FH Product family - 2010



TS3200 (3573) 4U Tape Library

- 1-2 LTO Ultrium 5 Drives
- Features
 - HH or FH SAS or Fibre
 - 3-slot I/O
 - 48 Data Cartridge Slots
 - Standalone or Rack
 - Multi-Path Architecture
 - Path Failover
- LTO Gen 5 Media
- LTO Gen 5/4 Encryption



TS2350 (3580) S53 External Tape Drive

- One LTO Ultrium 5 Drive
- FH SAS
- One Cartridge
- Desktop or Rack Mount
- LTO Gen 5 Media



TS3100 (3573) 2U Tape Library

- 1 LTO Ultrium 5 Drive
- Features
 - HH or FH SAS or Fibre
 - 24 Data Cartridge Slots
 - 1-slot I/O
 - Standalone or Rack
- LTO Gen 5 Media
- LTO Gen 5/4 Encryption



TS3310 (3576) Tape Library

- LTO Ultrium 5 Drive
- Features
 - Dual-Ported Fibre
 - Standalone or Rack
 - Multi-Path Architecture
 - Path Failover
- Model L5B – Base Control Module
 - 1-2 drives
 - 30 carts
 - 6 I/O
- Model E9U – Expansion Module (Max of 4 for 41U)
 - 1-4 drives
 - 80 carts
 - 12 I/O
- LTO Gen 5/4 Encryption




TS3500 (3584) Tape Library

- L53 Library & D53 Expansion LTO Frames
- L23 Library & D23 Expansion 3592 Frames
- HA1 High Availability & Dxx Dual Accessor Frames
- Reduced Footprint
- Capacity on Demand entry configuration
- TS1050 LTO Ultrium 5 and TS1130 3592 Drive Integration
- Mixed LTO & 3592 Drive & Media in Same Library
- Multi-Path Architecture
- Path Failover
- ALMS
- Advanced Electronics
- 3592 E06 & LTO Gen 5/4 Encryption
- S54 & S52 HD Frames


RIF: J-00019078-0



High Density frames - Up to 3x Capacity



3PB in less than 1 square Meter



RIF: J-00019078-0



Virtual Tape Library

Data deduplication

Dedup



RIF: J-00019078-0

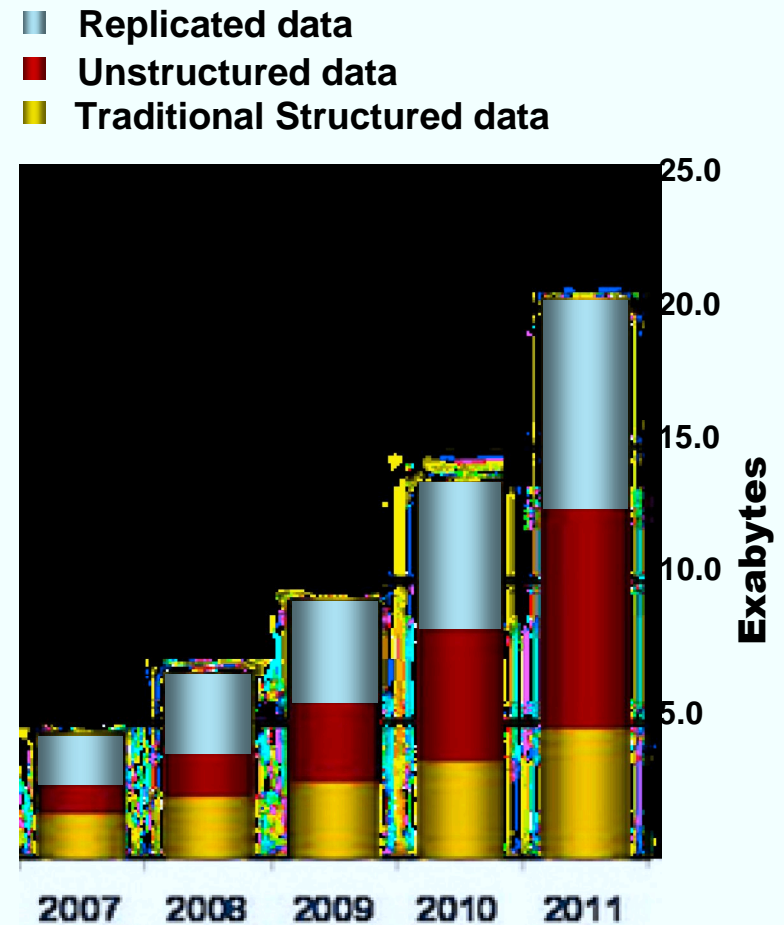


Information Growth


Information Infrastructure is doubling every 18 months

- **Structured data growing at 20%**
(databases for transactional workloads)
- **Unstructured data growing at 80%**
(such as user files, medical images, web and rich media content)
- **Replicated data growing at 50%**
(including backup, archive, business analysis, discovery and business continuance)

Sources: IDC worldwide enterprise disk in Exabytes from "Changing Enterprise Data Profile", December 2007



RIF: J-00019078-0



Symptoms

- Applications are slow
- Operational costs are increasing
- Legal costs are soaring
- Backups seem to take forever
- Batch jobs run into working hours
- Upgrading has become cost-prohibitive
- SLAs are being missed

63% of IT executives rate compliance with regulations a top challenge

Compliance

The average legal discovery request can cost an organization from **\$150K to \$250K**

Performance

Cost

Downtime costs can eat up to **16%** of revenue in some industries

Availability

Risk



RIF: J-00019078-0



IBM

Using the right balance of high density tape and high performance disk is the solution . . .



- **Long Term Retention**
 - Cost effective capacity
 - Removable & transportable
- **Compliance**
 - Meet financial & regulatory requirements
 - Data encryption, WORM

- **Short Term Retention**
 - Use disk for daily backup & restore operations
- **Performance**
 - Fast backups
 - Even faster restores
 - Meet “backup windows”

Why virtualizes Open System process

Improve the backup process

- Address network limitations
- Manage immense data growth

Improve the recovery time objective

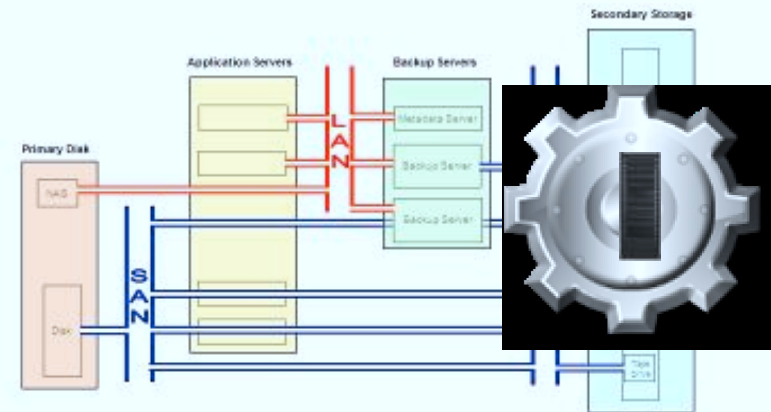
- Eliminate physical tape movement
- Reduce contention for resources

Improve the recovery point objective

- By creating incremental backups more frequently
- By writing incremental backups to cache

Augment the existing business continuance infrastructure

- Reduce bottlenecks and utilize tape assets more efficiently
- Optimize IT resources



RIF: J-00019078-0



Advantage of Virtual Tape Library

- **No dependency on streaming modes**
 - No tape cartridge operation
 - No start-stop with slow data streams
- **No limitations to number of libraries/tape drives/cartridges**
 - Can be used to parallelize backups or restores
 - ProtecTier®: 16 Libraries / 256 Drives / 500,000 Cartridges
- **Optimized utilization of cartridges**
 - If physical cartridges cannot be filled
 - Virtual cartridge sized can be defined according to environment
 - Advantage increases with parallel saves from system i
- **Require less FC connections for the tape drives / library**

Addressing The Data Center Challenge

- **Enterprise-class deduplication software technology**
- **Outstanding inline performance**
- **High capacity and scalability**
- **Significant User Acceptance**



A breakthrough in data storage management




RIF: J-00019078-0



ProtecTIER® & HyperFactor® Overview

Protect More. Store Less.™



ProtecTIER® reduces the required backup disk capacity by **up to 25 times or more!**

Data deduplication

How does it works



RIF: J-00019078-0

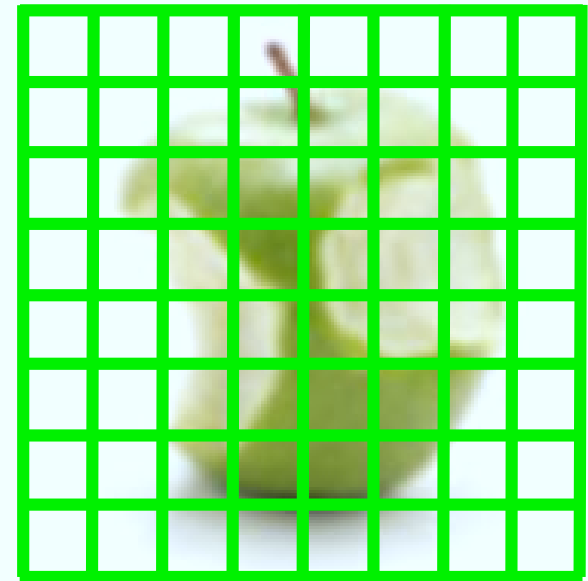


IBM

Data Deduplication – How does it work



RIF: J-00019078-0


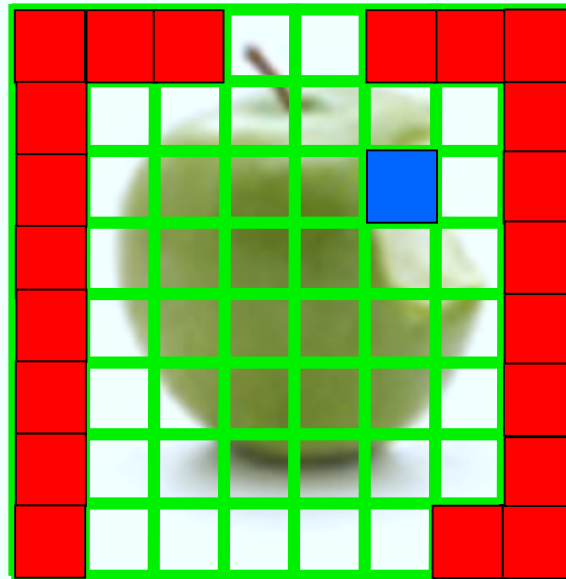

Data Deduplication – How does it work




RIF: J-00019078-0



Data Deduplication – How does it work





RIF: J-00019078-0



Data Deduplication – How does it work

New incremental backup example



RIF: J-00019078-0



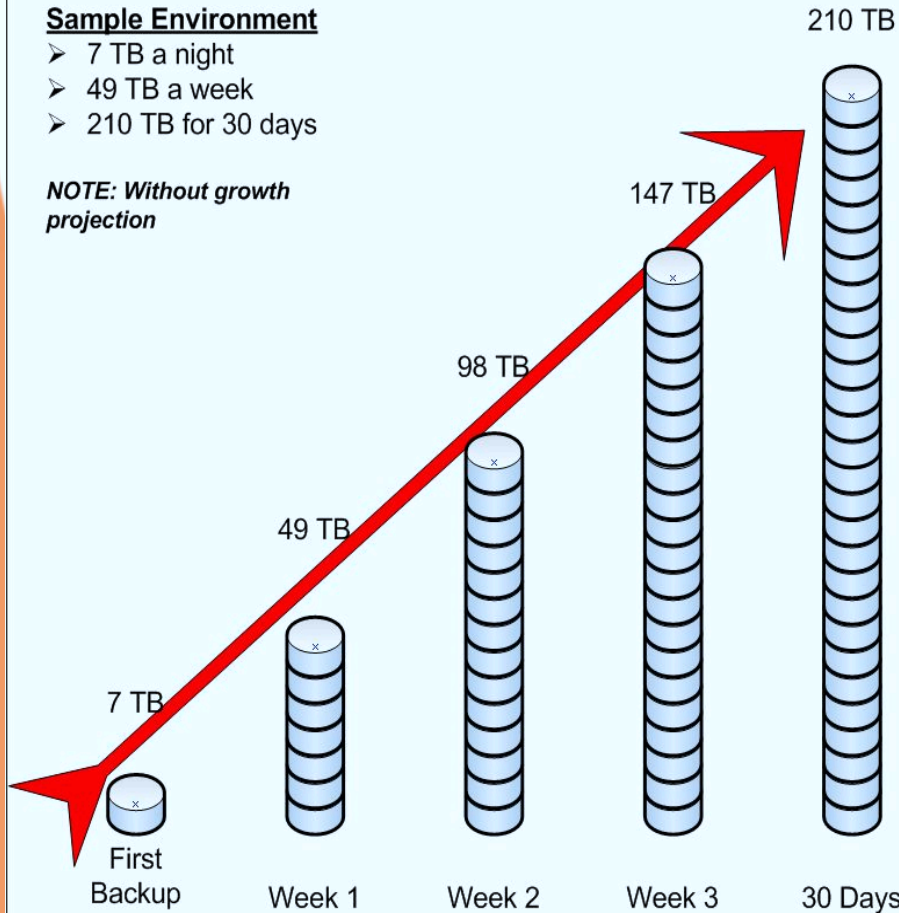
Data Deduplication is the key to using more disk more cost effectively!

Storage Growth – No Deduplication

Sample Environment

- 7 TB a night
- 49 TB a week
- 210 TB for 30 days

NOTE: Without growth projection

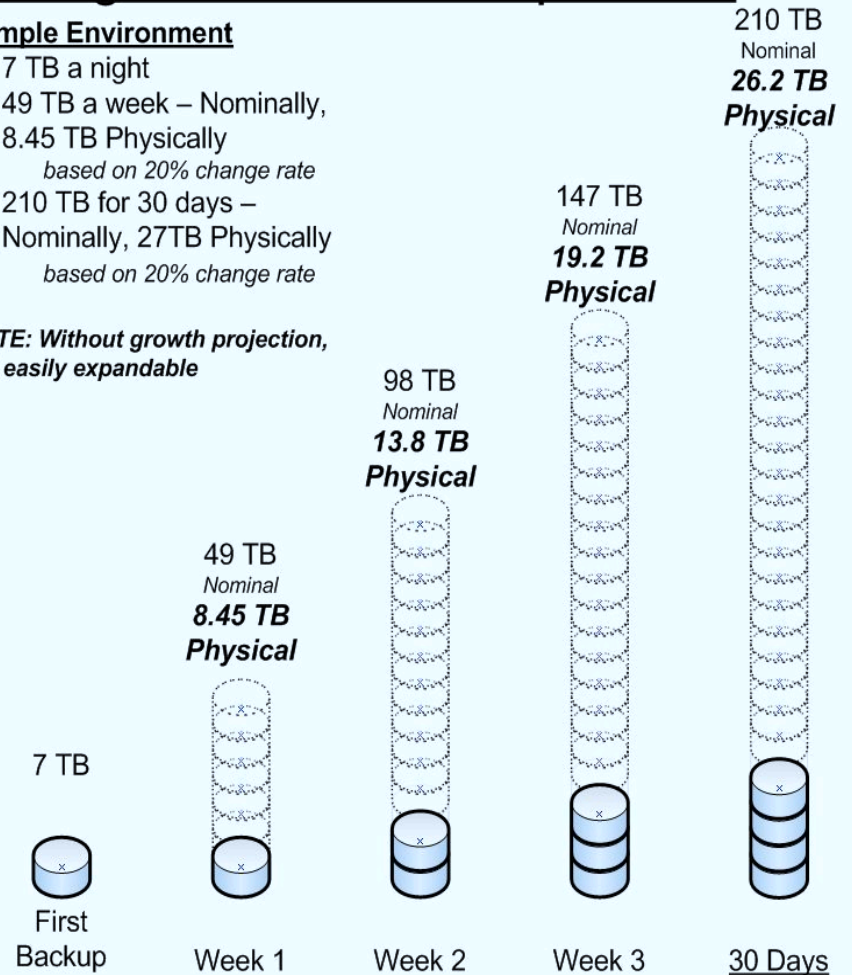


Storage Growth – Deduplication

Sample Environment

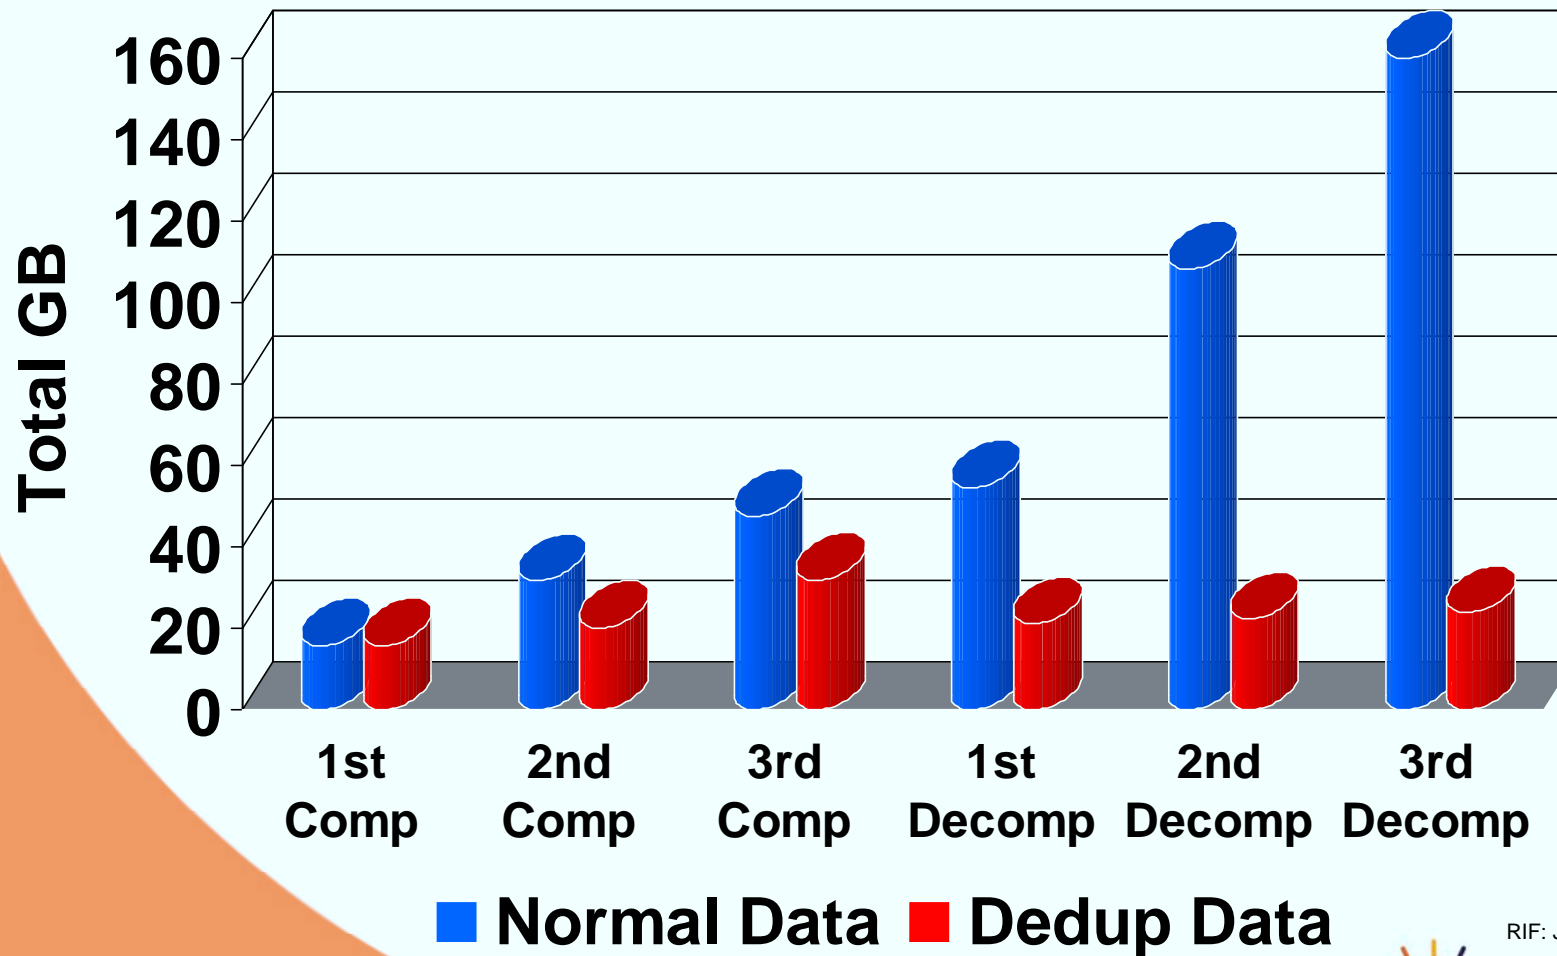
- 7 TB a night
- 49 TB a week – Nominally, 8.45 TB Physically
- based on 20% change rate
- 210 TB for 30 days – Nominally, 27TB Physically
- based on 20% change rate

NOTE: Without growth projection, but easily expandable



ProtecTIER® - Proof of Concept

Real example Customer Data




RIF: J-00019078-0




ProtecTIER® Vision and Design criteria

1. Data-agnostic factoring of up to 25 times or more
2. Unmatched performance up to 1000 MB/s or more running inline deduplication
3. Unequaled scalability: up to 1 PB physical data
4. Enterprise-class data-integrity: Not hash-based




*Most elements of new-data backed up today
already exist in previous backups*





ProtecTIER® Vision and Design criteria

1. Data-agnostic factoring of up to 25 times or more
2. Unmatched performance up to 1000 MB/s or more running inline deduplication
3. Unequaled scalability: up to 1 PB physical data
4. Enterprise-class data-integrity: Not hash-based
5. Simple, non-disruptive deployment
6. Supported in most hardware and software environments




*No other dedupe technology meets all these criteria!
Most elements of new data backed up today
already exist in previous backups*



RIF: J-00019078-0




ProtecTIER® Architecture Overview

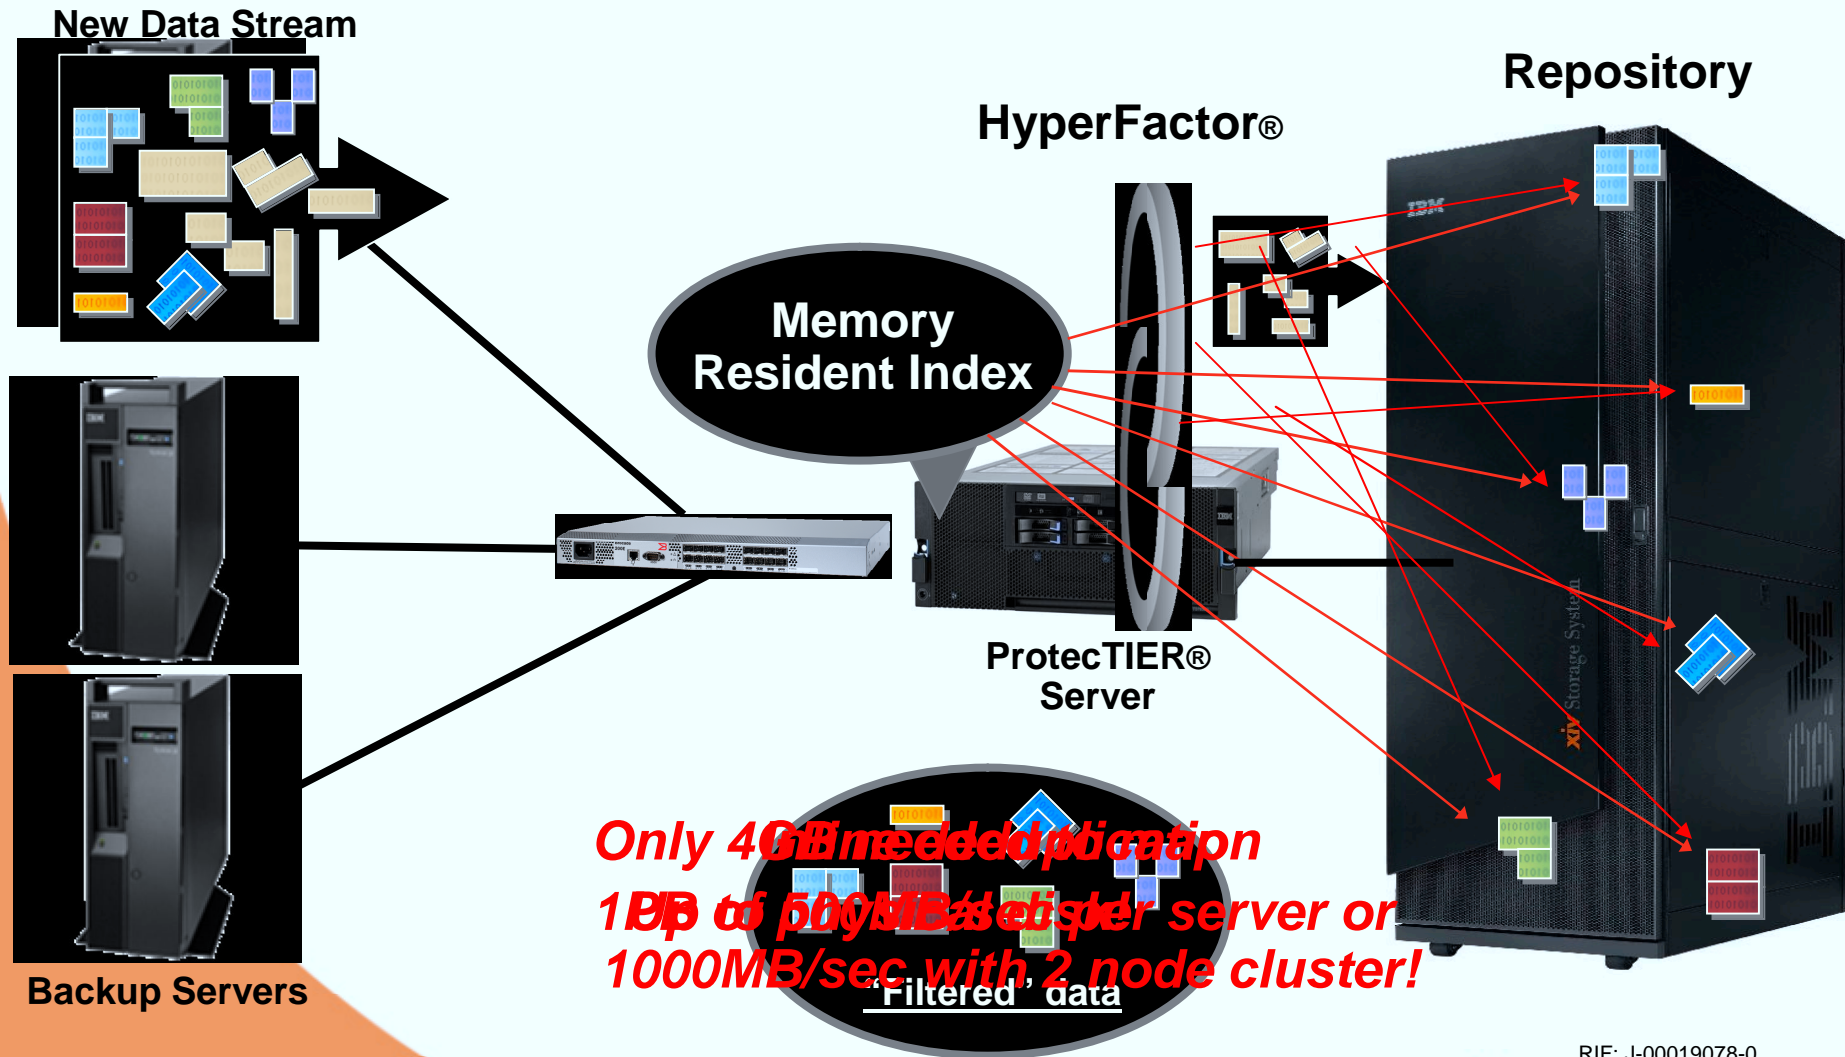


- **Linux server-based application**
- **Emulates a tape library unit, including drives, cartridges, and robotics**
- **Uses attached disk storage system as the backup medium**


RIF: J-00019078-0



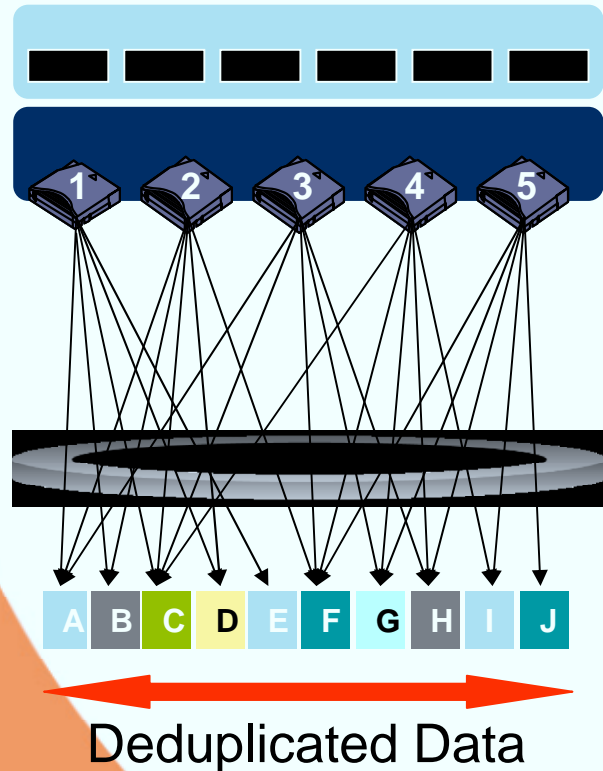
How ProtecTIER® works




RIF: J-00019078-0



Overview of ProtecTIER® operations




- Backup application writes data to ProtecTIER® as it would to tape
- Data goes through HyperFactor® deduplication engine
- Only unique data is stored
- Existing duplicate data is referenced
- When a data object expires or is overwritten, references are removed
- Free space is reclaimed and reused




RIF: J-00019078-0

Deduplication Methods

Inline Processing

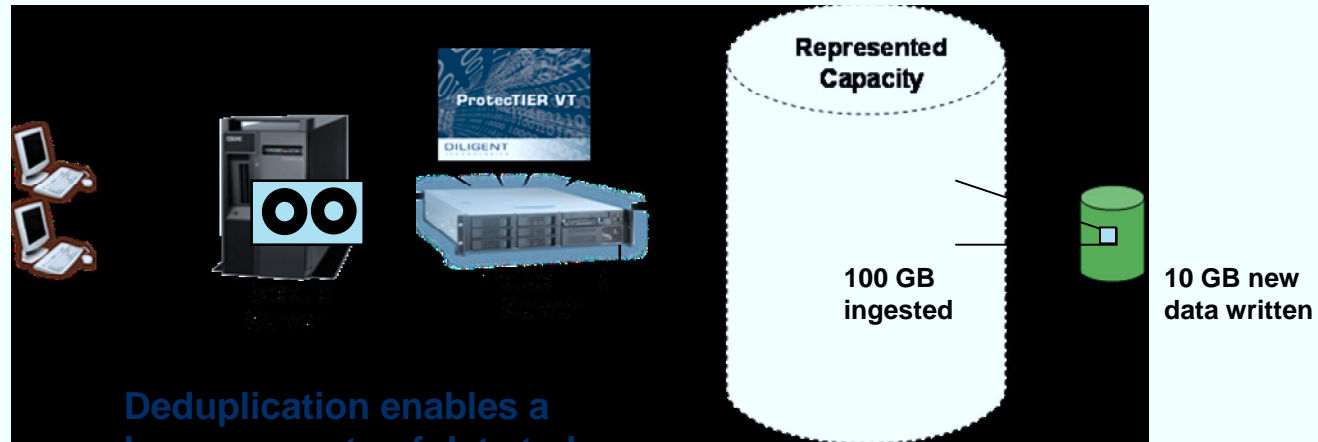


Post Processing



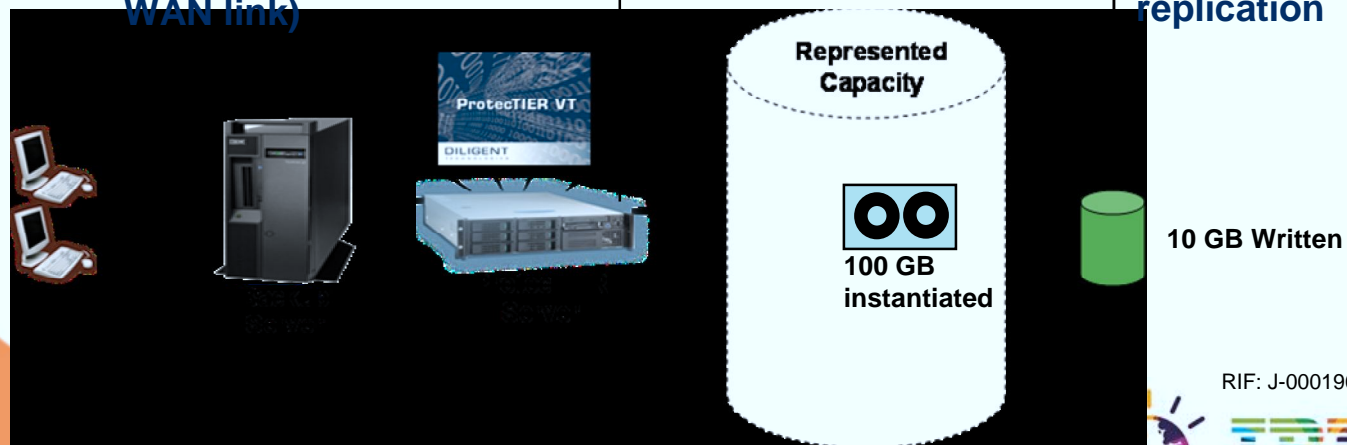
Significantly reduces replication bandwidth

Primary Site




Deduplication enables a large amounts of data to be replicated with significantly less bandwidth (IP based WAN link)

Secondary Site



Not a Disk based replication

RIF: J-00019078-0



ProtecTIER® GUI

IBM ProtecTIER Manager [-] [Maximize] [Close]

File Cluster Management System Node Repository VT View Tools Help

Refresh [Refresh Icon] Node [Add Node Icon] [Refresh Icon] [Refresh Icon] [Refresh Icon] [Refresh Icon] Repository [Refresh Icon] [Refresh Icon] [Refresh Icon] Library [Refresh Icon] Cartridges [Refresh Icon] [Refresh Icon] [Refresh Icon] [Refresh Icon]

Navigation Pane

Systems Nodes

- broadway (9.11.200.151)
- speedway (9.11.200.150)

speed_broad

Repositories

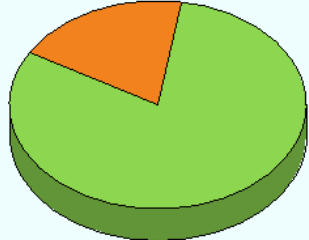
- Repository speedway

Services

- Library Red_LTO

Repository speedway
Status: Ok System: speed_broad

Capacity
Total capacity: 958.0 GB




Used:	184.1 GB
Allocable:	773.7 GB
Fragmented:	172.5 MB

[View resources](#)

Total utilization

Used space: 184.1 GB
Nominal data size: 1,218.1 GB

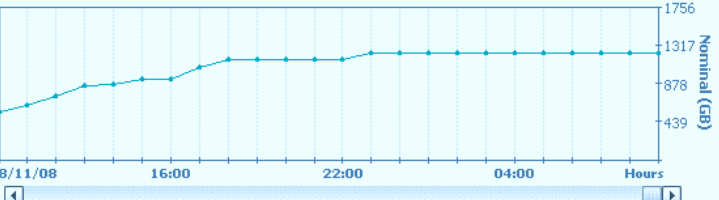


Repository configuration

Configured size (in TB)	1
Configured factoring ratio	14
Configured peak throughput	400
Configured MD raid	SATA-7.2K 2+2

Nominal data size graph

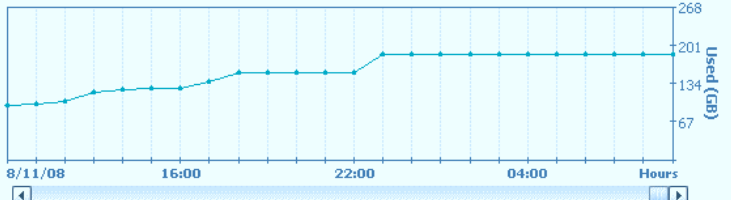
Recent days Recent hours



8/11/08 16:00 22:00 04:00 Hours

Utilization graph

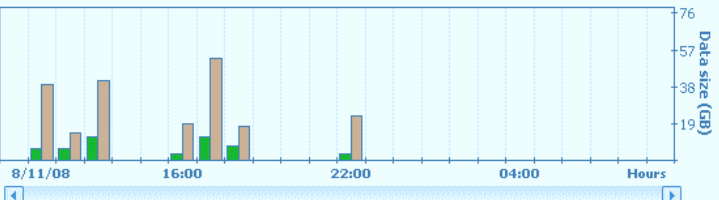
Recent days Recent hours



8/11/08 16:00 22:00 04:00 Hours

Marginal HyperFactor ratio graph

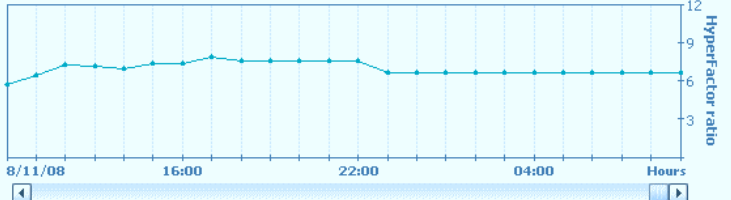
Physical Nominal



8/11/08 16:00 22:00 04:00 Hours

HyperFactor ratio over time graph

Recent days Recent hours



8/11/08 16:00 22:00 04:00 Hours

VT Online abbe logged in to speed_broad as administrator

0.00 (MB/Sec) 0.00 (MB/Sec) 9:23 AM Alerts Events log

ProtecTIER® differentiators

■ Performance

- Sustainable 500MB/s per node (1000MB/s two node cluster), performing INLINE deduplication.

■ Capacity

- Up to 1PB of physical capacity in single node or cluster configurations.
- IBM Disks or OEM Disks

■ Enterprise-Class Data integrity


- Binary diff process during deduplication designed for the highest data integrity.

■ Non-Disruption

- Inline deduplication eliminates need for significant secondary processing.
- Integrates well with existing backup environment and infrastructure

■ Replication

- Many to one replication from primary to DR site.
- Flexible policy driven replication



ProtecTIER® Deduplication family



**TS7610
ProtecTIER
Appliance
Express**

Good Performance
Entry Level
Easy to Install



Up to 80 MB/sec
4 TB and 5.4 TB
Useable Capacity

**TS7650
ProtecTIER
Appliances**

Better
Performance
Larger Capacity
Scalable



Up to 500 MB/sec
7 TB to 36 TB
Useable Capacity

Highest Performance
Largest Capacity
High Availability

**TS7650G
ProtecTIER
Gateways**



Up to 1000 MB/sec
Up to 1 PB
Useable Capacity

Scalable Capacity and Performance



RIF: J-00019078-0

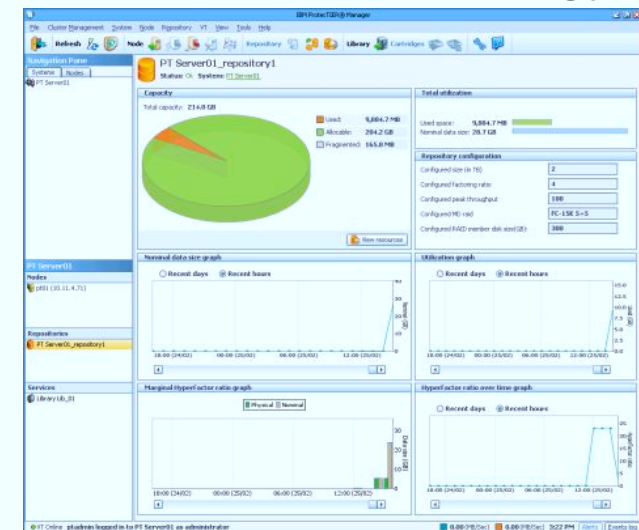
IBM TS7610 ProtecTIER® Deduplication Appliance Overview



- Fast 80MB/s backup performance
 - Even faster restores!
- Inline deduplication
 - No post processing burden!
- Designed for 100% data integrity
 - No chance of hash collision!
- Two configurations: 4.0 TB & 5.4 TB
 - Useable capacity, not RAW
 - Field upgradeable by customer
- IBM quality components - High RAS
 - No low cost commodity parts




**Same enterprise-proven
ProtecTIER® technology**




RIF: J-00019078-0



IBM TS7650 ProtecTIER® Deduplication Appliance Overview



- Pre-configured for rapid deployment into existing backup environments
- Extremely powerful solutions, featuring:
 - IBM ProtecTIER® software with patented HyperFactor® deduplication technology
 - IBM System x Server – multi-core server for enterprise-level performance
 - IBM Storage Controller with Fibre Channel drives - Proven reliability and performance
- Complete solution that includes rack, cables, switches, and everything that is needed



IBM TS7650 Appliance configuration

- **Ideal Customer for 7TB ProtecTIER® Appliance**

- 1 TB or less incremental backups per day
- 1-3 TBs full backups each week
- Experiencing average data growth
- Needs a cost effective solution

- **Ideal Customer for 18TB ProtecTIER® Appliance**

- 3 TBs or less incremental backups per day
- 3-6 TBs full backups each week
- Experiencing rapid data growth
- Needs good performance to meet backup window

- **Ideal Customer for 36TB ProtecTIER® Appliance**

- 5 TBs or less incremental backups per day
- 5-12 TBs full backups each week
- Additional growth expected
- Meeting the Backup window is an issue - higher performance needed



RIF: J-00019078-0




IBM TS7650 ProtecTIER® Deduplication Gateway Overview



**TS7650G ProtecTIER®
Deduplication Gateway**

- Powerful and Flexible solution, featuring:
 - IBM ProtecTIER® software with patented HyperFactor® deduplication technology
 - IBM System x Server – multi-core server for enterprise-level performance
- Supports both IBM & Non-IBM disk
 - IBM DS4000, DS5000, DS8000 and XIV
 - HDS, EMC and others
- And delivers:
 - Up to 500 MB/sec or more performance
 - Up to 25 times or more data reduction
 - Scalable to 1PB physical capacity
 - Enterprise-class data integrity




IBM TS7650 ProtecTIER® Deduplication Gateway Cluster Overview



**TS7650G ProtecTIER®
Deduplication Cluster**

- Our most Powerful and Flexible solution:
 - IBM ProtecTIER® software with patented HyperFactor® deduplication technology
 - 2 IBM System x Servers – multi-core servers for maximum performance & availability
- Supports both IBM & Non-IBM disk
 - IBM DS4000, DS5000, DS8000 and XIV
 - HDS, EMC and others
- And delivers:
 - Up to *1000 MB/sec or more* performance
 - Active-active cluster technology
 - Two nodes working together as one repository
 - Easily manageable yet highly scalable



IBM ProtecTIER® Capacity Planner

7650G Planner - [HyperFactor Sizing]

File Edit View Insert Format Records Tools Window Help Adobe PDF Type a question for help

Capacity Planner

Company: LA Class Example 2 Created: 18-Jun-09
 Workload: Updated: 18-Jun-09 By: JFB

Output Profile

Full Frequency: 7 Full TBs: 31.0
 Dailys per Full: 6 Daily TBs: 5.2
 Daily Retention: 12 TB or %: TB %
 Comp Rate: 2 Include GFS:

Change Rate New Data

Full: 12% 3.7 New TBs from Full
 Daily: 15% 0.8 New TBs from Daily
 Percent of: Daily Full
 Delta %: 0.0%

Results

Day	Cycles	GFS Days	User Data (TB)	Archive Data	Disk	Ratio
30	5	15	248	31	32	7.8
60	9	30	372	31	40	9.3
90	13	50	527	62	50	10.6
180	26	120	992	124	80	12.4

Roundup: Meta Data: Flat Rate

Performance

Backup Window

Full (Hrs): 13 Required: 695 MB/sec
 Daily (Hrs): 4 Source: Full Backup

Replication Needs

Physical Logical
 Payload %: 100%
 Time (Hrs): 1
 Bandwidth: 1,192 MB/sec Source: Full Backup
 Links: 2 OC-192 Links

Record: 3 of 3
Form View

RIF: J-00019078-0



IBM ProtecTIER® Performance Planner

7650G Planner - [Performance Planner]

File Edit View Insert Format Records Tools Window Help Adobe PDF Type a question for help

ProtecTIER Performance Planner

Company: CommVault?: Created:

Workload: Updated: By:

Target Performance Profile

Initial Perf: MB/sec
 Future Perf: MB/sec

Warning: The required performance suggests 2 nodes

Rep Size (TB):
 Factoring Ratio:

ProtecTIER Repository Configuration Options:

Meta Data:					User Data:				
Disk	RPM	Cap (GB)	RAID	Group	Disk	RPM	Cap (GB)	RAID	Group
FC	15,000	300	10	4+4	FC	15,000	300	5	4+1
Future Performance									
SATA	7,200	500	10	4+4	SATA	7,200	500	5	4+1
Future Performance									

Results:

Meta Data:

Meta Data:		User Data:		
HDD's	F/S's	HDD's	F/S's	Cap (TB)
<input type="text" value="24"/>	<input type="text" value="3"/>	<input type="text" value="175"/>	<input type="text" value="35"/>	<input type="text" value="41.0"/>
Total HDD's: <input type="text" value="199"/>		Raw Cap: <input type="text" value="58.3"/>		
Future Performance				
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Total HDD's: <input type="text"/>		Raw Cap: <input type="text"/>		
<input type="text" value="40"/>	<input type="text" value="5"/>	<input type="text" value="360"/>	<input type="text" value="72"/>	<input type="text" value="140.6"/>
Total HDD's: <input type="text" value="400"/>		Raw Cap: <input type="text" value="195.3"/>		
Future Performance				
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Total HDD's: <input type="text"/>		Raw Cap: <input type="text"/>		

Record: of 3

Form View

Customer Case Studies



Telecommunications

Business challenge

This legacy Diligent Technologies client is a leading wireless provider in the US serving millions of customers and currently ranks 10th among the 2008 F500 and 29th among the 2008 Global F500 companies. With the introduction of a new phone brand, this customer quickly experienced a data growth explosion growing from 3PBs in Feb 2007 to 6 PBs of data by July that same year. Their major downtown Dallas datacenter literally had no more floor space, no more power to consume and they couldn't expand. They needed to reduce their backup disk capacity in order to keep operational.

Current Solution

- A win over Data Domain & NetApp
- IBM's ProtecTIER Deduplication Solution – now up to 40 servers with 32TB behind each server across multiple data centers
- Symantec NetBackup

Benefits

- ProtecTIER reduced the amount of backup data to free up much needed disk capacity to handle growing backups
- Seamlessly integrated with existing infrastructure (disk, server and backup application)



ProtecTIER is the only enterprise-class deduplication solution that enables customers to protect their existing investments while meeting requirements for performance, capacity and scalability



RIF: J-00019078-0

IBM

Healthcare

Business challenge

Founded in 1974, this client is one of the foremost cancer research facilities in the US. Handling 2-3TBs nightly of 300 NBU server clients, their backup windows were growing 12 to 14 hours, and recoveries were creeping up from hours to days. They were making extremely tough decisions on what servers to backup on a nightly basis. Wanting to move to a disk-based backup solution without changing their backup process, they chose IBM's ProtecTIER Deduplication solution.

Current Solution

- IBM's ProtecTIER software (3rd party disk)
- IBM x Series server
- Symantec NetBackup

Benefits

- Immediately reduced backup window from 14 hours to 6 hours
- Effectively managing their data growth and physical storage capacity through deduplication
- DR backup sets run during the day without impacting production servers
- Easily integrated with existing backup policies, practices and procedures



“The ProtecTIER Deduplication Solution enables us to complete our backups within the allotted timeframe now and we never have to make hard decisions on which servers to back up. We just do all of them very quickly now.”

Director, Systems & Operations



RIF: J-00019078-0



Financial Services - Banking

Business challenge

This client wants to be tapeless in 2-3 years. They had already moved to a disk-based backup and recovery solution using Data Domain DD690s but was hampered by their inability to provide performance, scalability, and capacity to meet their backup and recovery requirements. As new datasets were added and their environment continued to grow, performance and capacity suffered. With the current DD690 appliances, their only choice was to keep adding appliances to try and solve the problem. They decided not to invest any more time and money and opted for IBM's TS7650G deduplication solution in a clustered configuration to have a more robust, dependable solution that could guarantee performance and scalability.

Current Solution


- IBM's ProtecTIER TS7650G (6 in a clustered configuration)
- IBM DS4700 disk arrays (2)

Benefits

- Provides industry-leading performance, scalability and availability with true global deduplication technology
- IBM provided a solution of 6 TS7650 Gateways vs. DataDomain's 34 appliances to handle the same amount of data
- Enables management of their environment holistically and will enable them to meet their goal of going tapeless in the next 2-3 years



With industry-leading performance, scalability and capacity, ProtecTIER continues to exceed expectations on meeting customer requirements of all sizes



RIF: J-00019078-0

Financial Services

Credit Intelligence

Business challenge

This client is a global leader in providing consumer and business credit intelligence, portfolio management and fraud detection solutions. Although they are on the cutting edge of their market offerings, their tape-based backup and recovery infrastructure was lagging. With hundreds of terabytes of data to support and a growth projection exceeding a petabyte, the need to eliminate over a half a million (500,000) tapes across mainframe and open systems in rotation was the biggest challenge. Working with their partner Accenture, they chose IBM's ProtecTIER with XIV over EMC/Data Domain & HDS to dramatically improve the effectiveness of their NetBackup and Mainframe data protection infrastructure.

- TS7650G dual node cluster configurations
- (6) XIV arrays with 420TB; 210TB at each location
- TS7740 for mainframe data


Solution

Benefits

- Eliminate over 500,000 tapes in rotation saving time and money
- Improve backup and recovery performance
- Once deployed, enables cost effective replication of data for long-term DR strategy



ProtecTIER's data deduplication offers Equifax a cutting edge improvement over their existing tape-based backup and recovery infrastructure. The total IBM solution enables them to incrementally scale their environment when needed and eliminates extra cost and complexity



Insurance

Business challenge

This client is a leading international insurance organization operating in more than 130 countries. They backup about 60TBs nightly of SAP, Oracle and Exchange data and need to replicate it approximately 1700 miles from the east coast to the midwest. Their biggest challenges included reducing the dependency of tape and the need for a more reliable, cost-effective disaster recovery plan. Working with their partner Vicom, they chose IBM's TS7650G ProtecTIER with XIV over long-time incumbent EMC to deploy a robust and dependable data protection environment.

Solution – Winback over EMC!

- (8) TS7650G in two locations
- XIV arrays 95TB & 105TBs
- Legato Networker (existing user)

Benefits

- Provides scalability and performance to backup operations and improves application performance
- Easily increases capacity of disk storage through deduplication
- Provides cost-effective replication for improved DR requirements



ProtecTIER & XIV are a powerful combination offering customers the ability to improve backup/recovery operations and boost application performance in the smallest storage footprint possible



RIF: J-00019078-0



Consumer Products

Food & Beverage

Business challenge

This customer's main objective was to eliminate their use and dependency of physical tape for backup and recovery. Physical tape was too slow for the amount of data they backed up and the management and overhead of trying to recover from tape was complex and costly. They chose IBM's ProtecTIER TS7650G Data Deduplication solution over EMC/Data Domain to increase the scalability and availability of their data while reducing their overall storage footprint.

Solution 7650G's in 4 clusters on third-party disk. ProtecTIER is managing a repository of over 400TBs of data

- ProtecTIER's IP-based native replication for many-to-1 approach
- NetBackup (existing customer)

Benefits

- Deployed solution with 50% less systems than competition, saving the customer cost and complexity
- Provides scalable architecture increasing backup performance and data availability
- Provides highest level of redundancy and no single point of failure



*The ProtecTIER
TS7650G is the only
deduplication solution
that enables customers
the ability to easily
deploy disk-based
backup and recovery
with their choice of disk
giving them flexibility
and ease of installation
into their existing
processes*



RIF: J-00019078-0



Deduplication Market at a glance



DEDUPE TECHNOLOGY

ProtectTIER with HyperFactor

- ✓ Byte-level diff comparison
- ✓ Inline Deduplication
- ✓ Block Level Deduplication

PERFORMANCE

- ✓ Single node performance 500 MB/s
- ✓ Dual node Cluster performance 1000MB/s

RESOURCE UTILIZATION

- ✓ No disk staging area required
- ✓ Only 4GB RAM needed for a 100TB repository



DD880

RockSoft Hash-based

- ! Potential Hash collision
- ✓ Inline
- ✓ Block Level

- ✓ 1500MB/s using OpenStorage and 10GB Eth link
- ! Clustering not available

- ✓ No disk staging area required
- ! Over 300GBs of RAM!



DXi7500

RockSoft Hash-based

- ! Potential Hash collision
- ! Post process
- ✓ Block Level

- ! 130 MB/s
- ! Clustering not available

- ! Staging area > than the size of largest full backup
- ! Over 300GBs of RAM!



VTL 700

SIR Hash-based

- ! Potential Hash collision
- ! Post process
- ✓ Block Level

- ! 188 MB/s
- ! Clustering with Global Dedupe not available

- ! Staging area > than the size of largest full backup
- ! Over 300GBs of RAM!



S2100-ES2

DeltaStor

- ✓ Byte-level diff comparison
- ! Post process
- ⊘ File Level

- ! 160 MB/s
- ! Clustering with Global Dedupe not available

- ⊘ Staging area > twice the size of largest full backup
- ✓ 24GB of RAM
- ⊘ Not hash based



Deduplication Market at a glance



ProtectTIER with HyperFactor

data domain

**DD880
RockSoft
Hash-based**

Quantum.

**DXi7500
RockSoft
Hash-based**



**VTL 700
SIR
Hash-based**

SEPATON

**S2100-ES2
DeltaStor**

PRODUCT STABILITY

- ✓ IBM in business for nearly 100 years
- ✓ ProtecTIER in production since 2006
- ✓ Over 25PBs of in production

- ✓ Acquired by EMC
- ✓ In production since 2006
- ✓ Many small systems in production

- ! Over \$400 million in debt
- ! Post process
- ! Very few small customers

- ! Small struggling company
- ! GA October 2008
- ! Very few small customers

- ⊘ Acquisition or failure imminent
- ! GA May 2008
- ⊘ Almost no deduplication in production

CAPACITY-SCALABILITY

- ✓ Single system can scale to 1PB capacity
- ✓ Up to 16 virtual tape libraries
- ✓ Up to 512 virtual tape drives
- ✓ Up to 512,000 virtual tape cartridges

- ! 58TB Maximum useable capacity
- ! Limits not published
- ! Limits not published
- ! Limits not published

- ! Limited by rapid hash table growth
- ✓ Up to 64 virtual tape libraries
- ✓ Up to 160 virtual tape drives
- ✓ Up to 130,000 virtual cartridges

- ! Limited by rapid hash table growth
- ✓ Up to 128 virtual tape libraries
- ✓ Up to 1024 virtual drives
- ✓ Up to 64,000 virtual cartridges

- ! Limited by huge storage requirements
- ✓ Up to 192 virtual tape libraries
- ✓ Up to 192 virtual tape drives
- ✓ Up to 5.3 million virtual cartridges

MEETS ENTERPRISE REQUIREMENTS?

✓ YES

! NO

! NO


! NO

! NO



Respaldar puede ser fácil...

**Pero a veces el camino del Recovery
puede ser muy largo y agonizante.....**





For more information on IBM's ProtecTIER®

The main ProtecTIER® Web Page

www.ibm.com/systems/storage/tape/protectier

The ProtecTIER® Gateway Web Page

www.ibm.com/systems/storage/tape/ts7650g

The ProtecTIER® Appliance Web Page

www.ibm.com/systems/storage/tape/ts7650a



RIF: J-00019078-0



Questions ?!?



Thank You

*Paulo Mauricio Pereira Junior
Storage ProtecTIER® LA Specialist
paulojr@br.ibm.com*



RIF: J-00019078-0

IBM