

Enterprise Data Protection Considerations... Consolidation, De-duplication and Migration of Tape Data Files

Thomas J. Meehan
INNOVATION Data Processing

August 10, 2011
Session Number 9339

Executive Solution

Everyone is moving to the Cloud...

Mainframe customers...

- Are moving to the cloud with Linux on Z and z114/z196 zEnterprise Hybrid Computing

Now is the time to slip past conventional constraints...

- Technology is available that can provide non-disruptive z/OS and Enterprise Distributed Data Protection, that scales capacity and performance for “Big Data”.
 - Direct z/OS FICON access to Distributed Data, High-Capacity Disk and Tape Storage Devices, Storage Virtualization, Hardware Replication and VTL appliances with De-duplication

What is the Cloud...

What does the Cloud bring with it?

- **Cloud Computing is...**
an effort to make distributed computing more like a mainframe or put **Public/Private Clouds** on the mainframe.
- **What challenges does it bring with it?**
 - Economic Challenges
 - Multi-Platform Consolidation
 - Physical Storage Consolidation
 - **...too much data & it's difficult to adequately protect**
- **What does it need...a new way to provide...**
Multi-Platform Disaster Recovery & Distributed Data Protection.

Multi-Responsibilities

- Economic Challenge
 - Optimize Declining IT Budgets
 - Support Evolving Business Process
 - Safeguard the Data
- Multi-Platform Consolidation
 - Consolidate Platforms with Virtualization
- Physical Storage Consolidation
 - Larger Disk and Tape Volumes
 - Storage Virtualization
 - Hardware Replication
 - Virtual Tape Solutions
 - De-duplication

Multi-Platform Consolidation to System z/zEnterprise using Hybrid computing and mainframe infrastructure z114 and z196 Ensemble with zBX



**Linux on Z...UNIX on Z
z/OS Hypersockets...**

**AIX on Power...Linux on IBM X Series
Intra-ensemble Data Network (IEDN)...**

All INNOVATION products support zEnterprise Systems

**IBM Claims...14 New Z accounts: in 1Q 2011
50+ zBX units shipped; 300+ blades shipped as of 1Q**

Mainframe Resources for z/OS & Distributed Data Protection

Using System z Backup Storage

- New High-Capacity Tape
IBM TS1140 (4TB/295MBS) and ORACLE/STK T10K (5TB/240MBS).
- Physical Tape Library IBM TS3500 or ORACLE/STK SL8500.
- Virtual Tape Library IBM TS7700, ORACLE/STK VSM or EMC DLm.
- De-duplication Technology IBM TS7680 ProtecTIER or EMC DD890.
- EMC VMAX, HDS VSP, and IBM DS8800 disk.

Under the control of z/OS

- Auto-ops scheduling (CA-Scheduler, IBM TWS/OPC...)
- Tape management (CA1, RMM, TLMS...)
- Data access security (RACF, ACF2, Top Secret...)

INNOVATION is working closely with...



Enterprise Data Protection



CONSIDER...there are problems

- Consolidating to Linux on Z puts more data on the mainframe
 - ...so you need a good solution to protect it
- Enterprise backup congests communication networks
 - ...so take it off the communications network
- Backup takes too long
 - ...so you need to reduce the disruption of backup
- There is too much data and it's difficult to migrate to a VTS
 - ...so you need to reduce amount of backup data storage
 - ...so you need an easy to use tape to VTS migration tool

But there are solutions

Centralized Enterprise Data Protection using Mainframe Infrastructure Strengths

Enterprise Data Protection Solution Benefits

- z/OS controls backup data from OS server to z/OS tape
- Virtualization solves backup window problems
- Efficient tape drive and tape media utilization
- Centralizes scheduling, tracking and auditing
- Meet Disaster Recovery Restore Time Objectives

**FDR/UPSTREAM,
FDRSOS and FDR/UPSTREAM/SOS
are examples of Centralized Enterprise
Data Protection...**



Volume Level Business Recovery Protection File Level Enterprise Data Protection



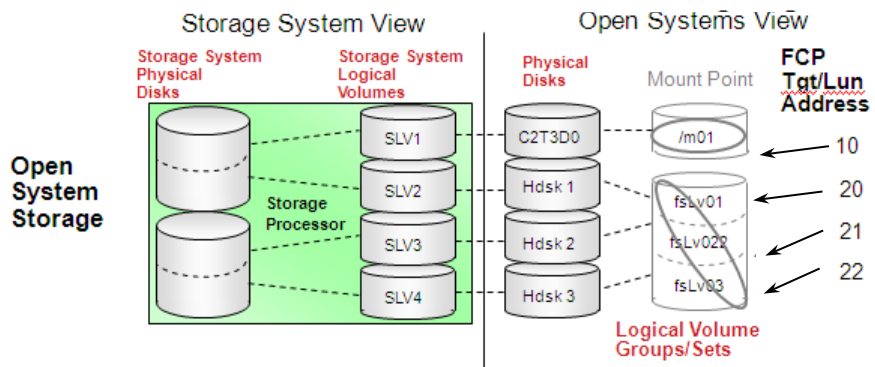
FDRSOS and FDR/UPSTREAM – are examples

- Business Recovery and Enterprise File Level Data Protection for Linux on System z, AIX, Linux, UNIX and Windows storage
- Employing z/OS existing Tape Management, Security and Scheduling
- Command data, meta data and backup data can all travel on a TCP/IP connection or a FICON Channel connection
- On-line data base support for DB2/UDB, ORACLE, LOTUS Notes, SQL Server, MS Exchange, etc.
 - Virtualizes full volume backup with “Deferred Full Merge”
 - Provides “Outboard Data De-duplication”...
Duplicate File Reduction, Duplicate Segment Reduction



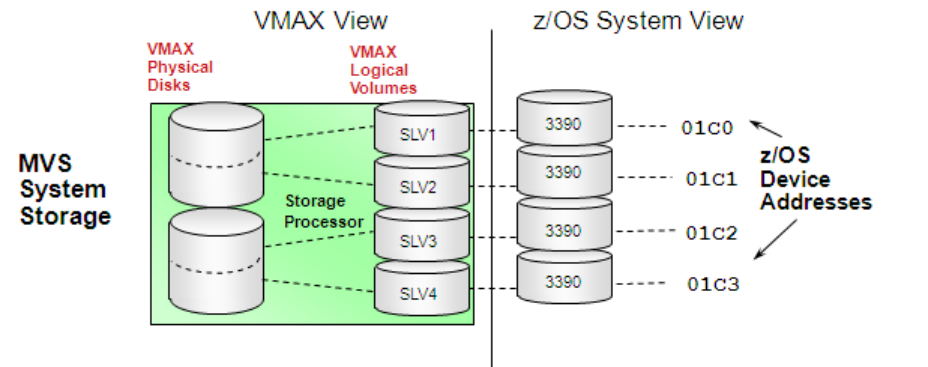
Conventional Disk Storage Architectures

Distributed System Disk Storage



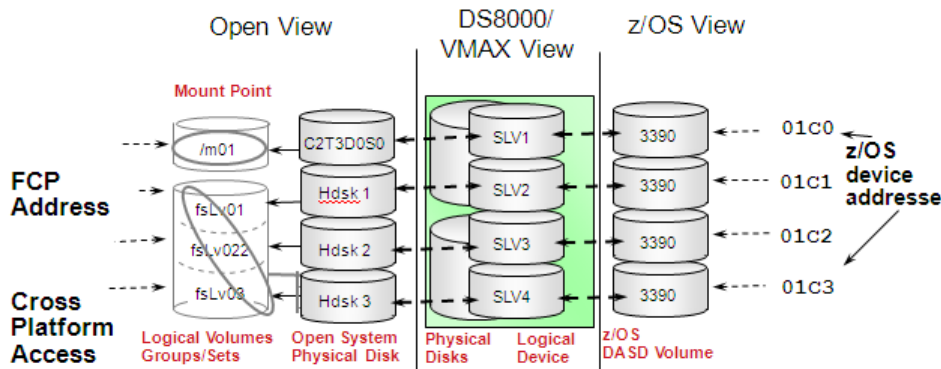
- Storage System Logical Device (i.e. Open System Physical Disk)
- Open System Disk
- Volume Groups

z/OS Disk Storage



- DS8000 / VMAX Logical Device (i.e. System z Physical Disk)
- CKD DASD volume addressed by an z/OS UCB

IBM zDDB and EMC zSOS Multiplatform Disk Storage Architecture

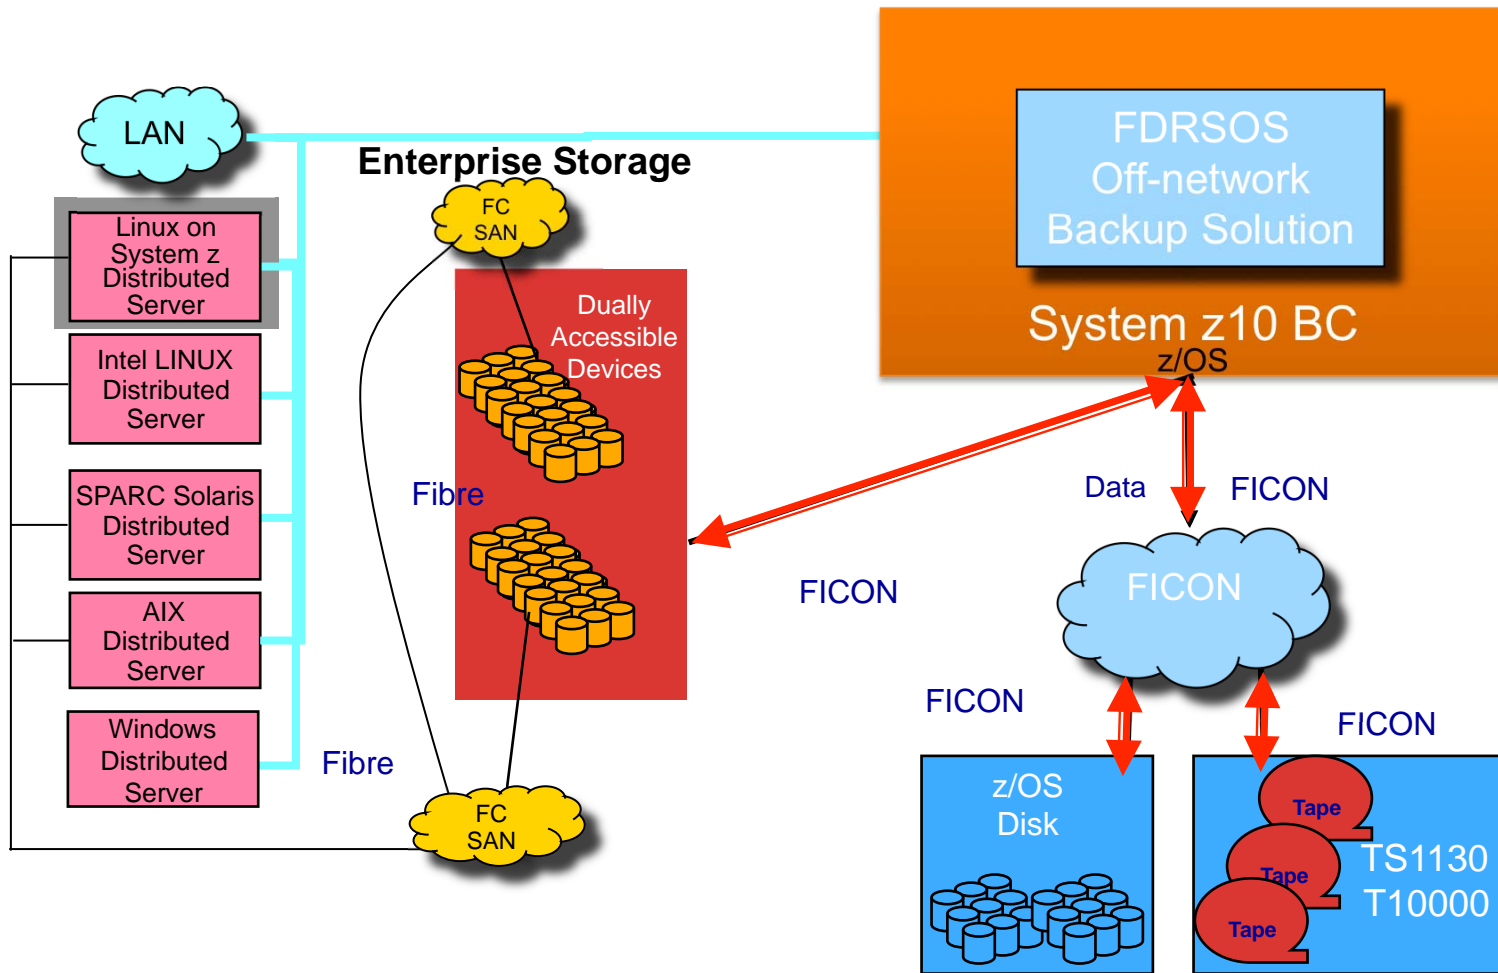


- Same DS8000 / VMAX Logical Device has two addresses
- addressed by a FCP address
- addressed by an z/OS UCB

IBM z/DDB and EMC zSOS Multi-platform Access...Benefits

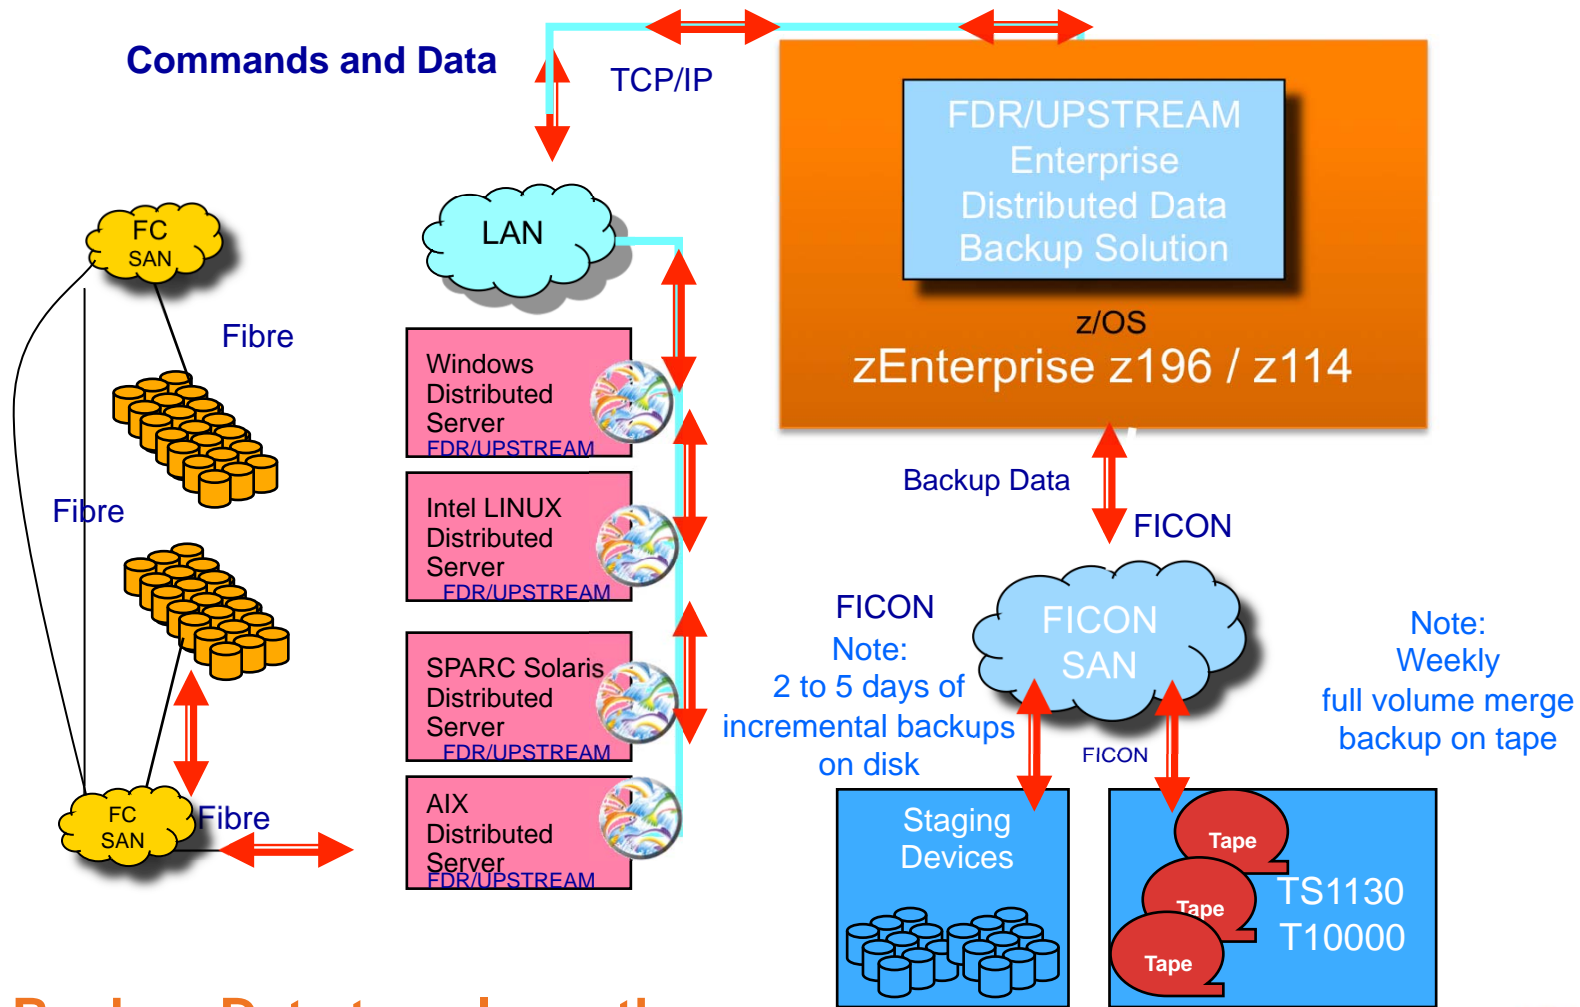
- Backup and recover large amounts of open systems data in less time with no negative impact on corporate networks.
- Production is not constrained by a shrinking backup window.
- Maximize uptime by avoiding costly customized integration.
- Protect your current mainframe hardware and software investments by using existing mainframe resources to protect Linux on Z and distributed enterprise storage.
- Extend mainframe security and automated operations to non-disruptive distributed data protection.

How FDRSOS® Works...Off-Network Business Recovery/Resiliency Solution



Data Travels Off Network

FDR/UPSTREAM®... Enterprise Distributed Data Backup

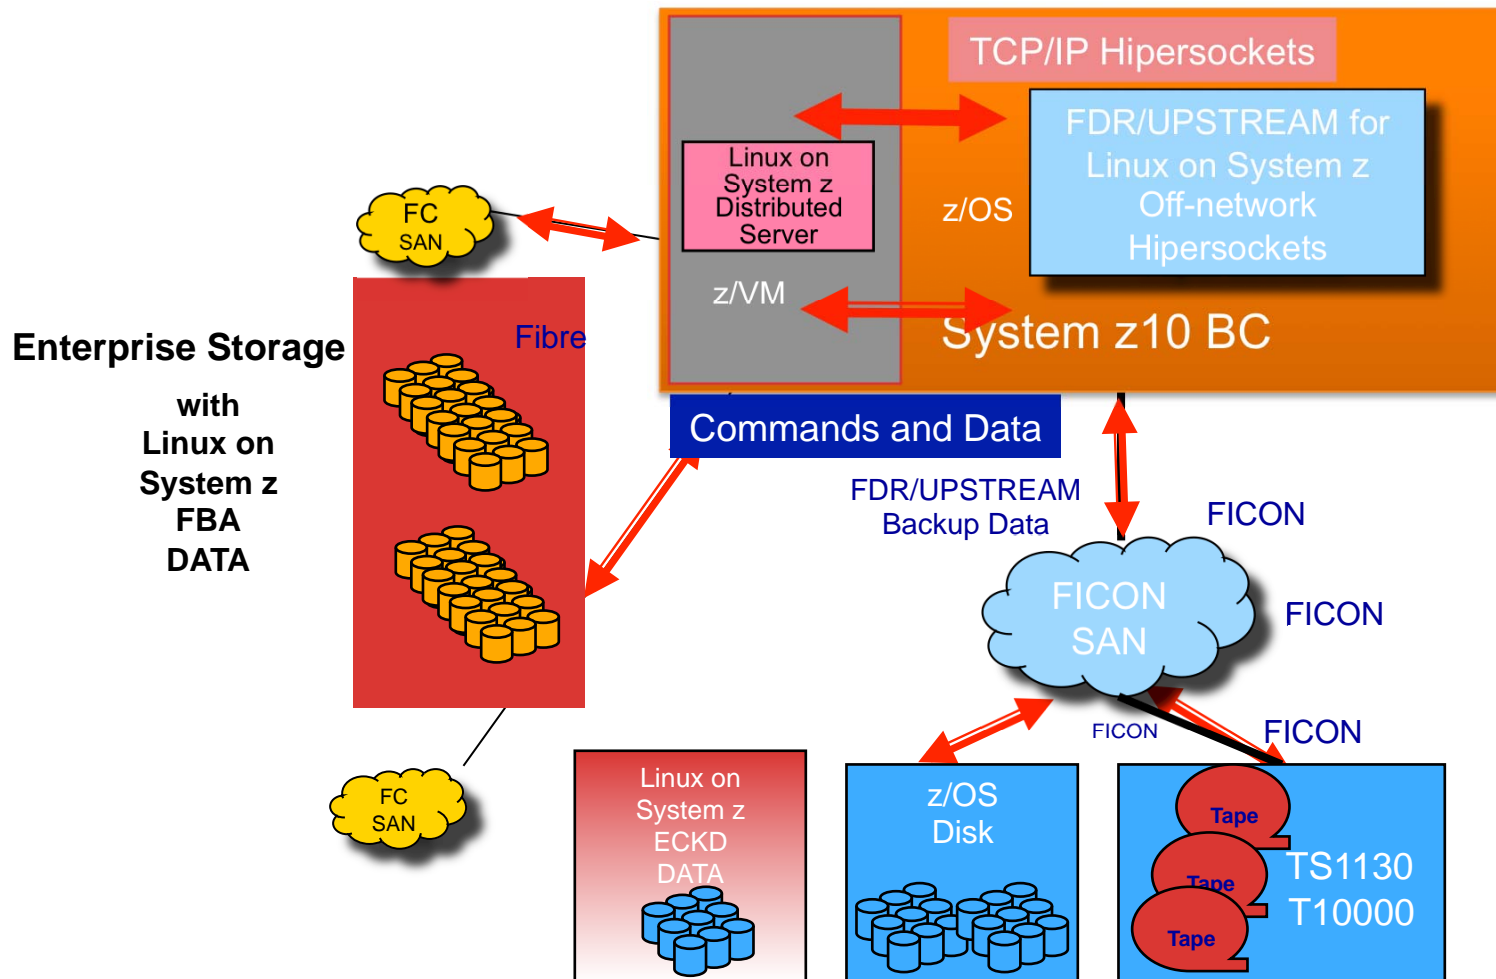


Backup Data travels on the Corporate Network



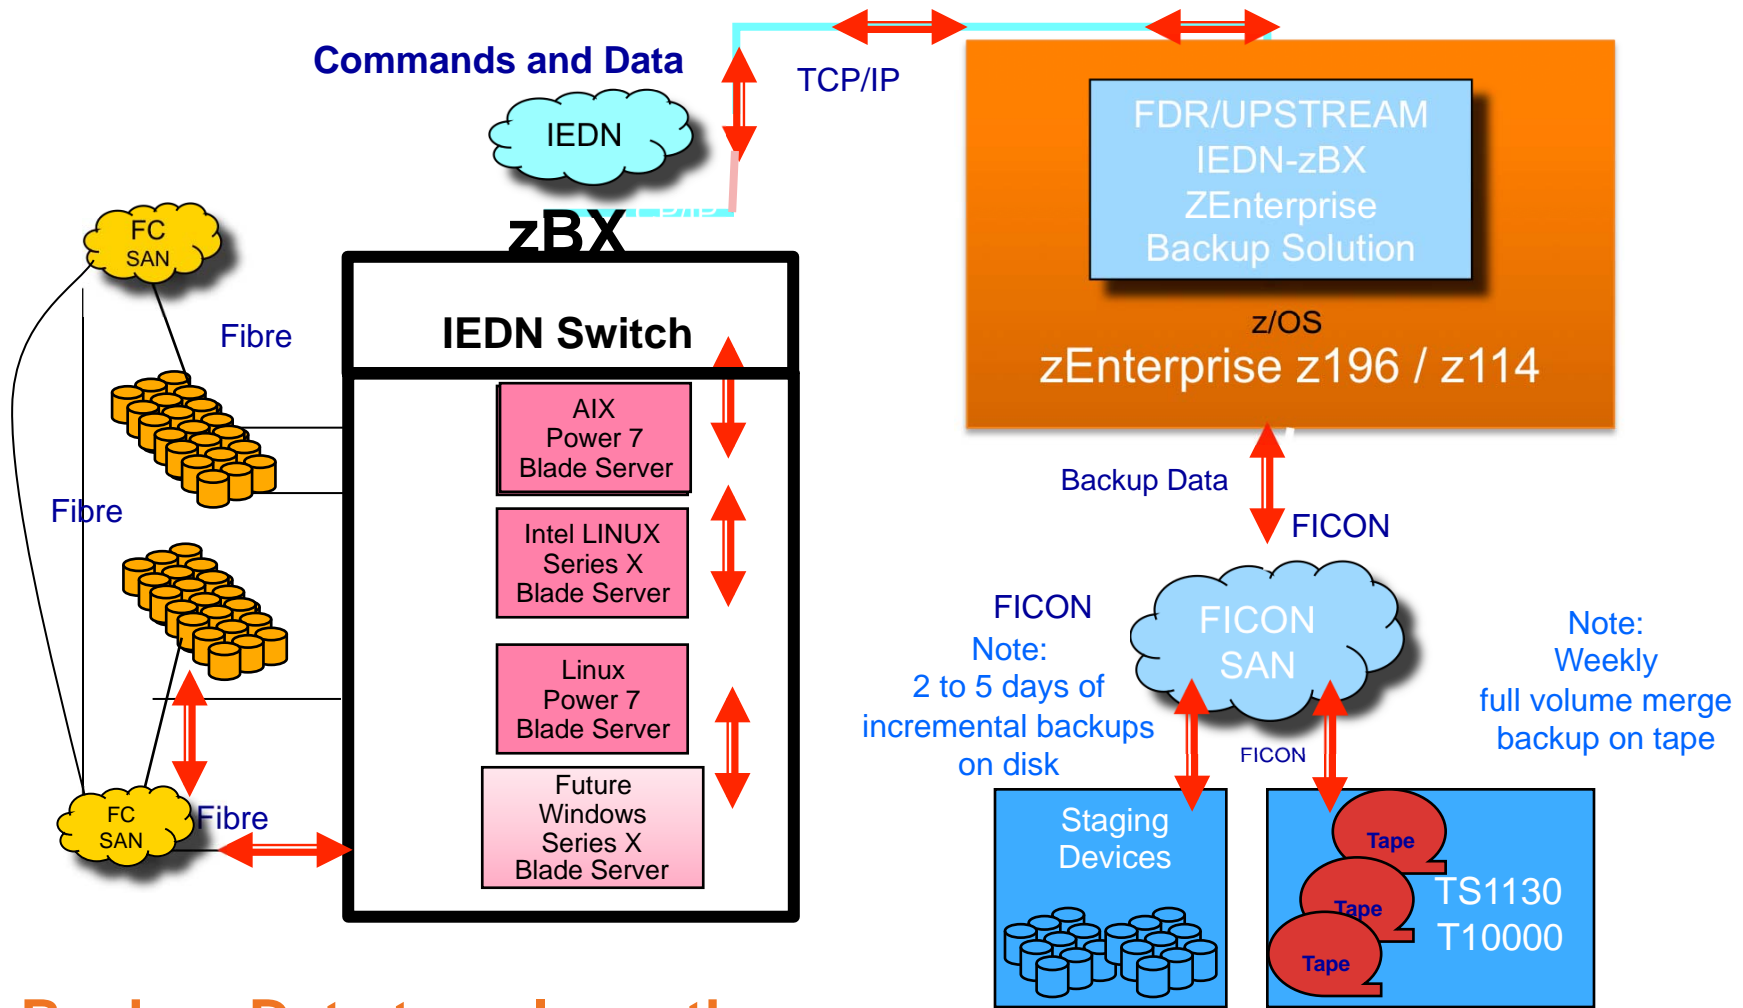
UPSTREAM for Linux on System z®

Off-network File Level Backup with Virtual Hipersockets



Data Travels Off Network

FDR/UPSTREAM® ... IEDN-zBX zEnterprise Backup




Backup Data travels on the
Intraensemble Data Network (IEDN)



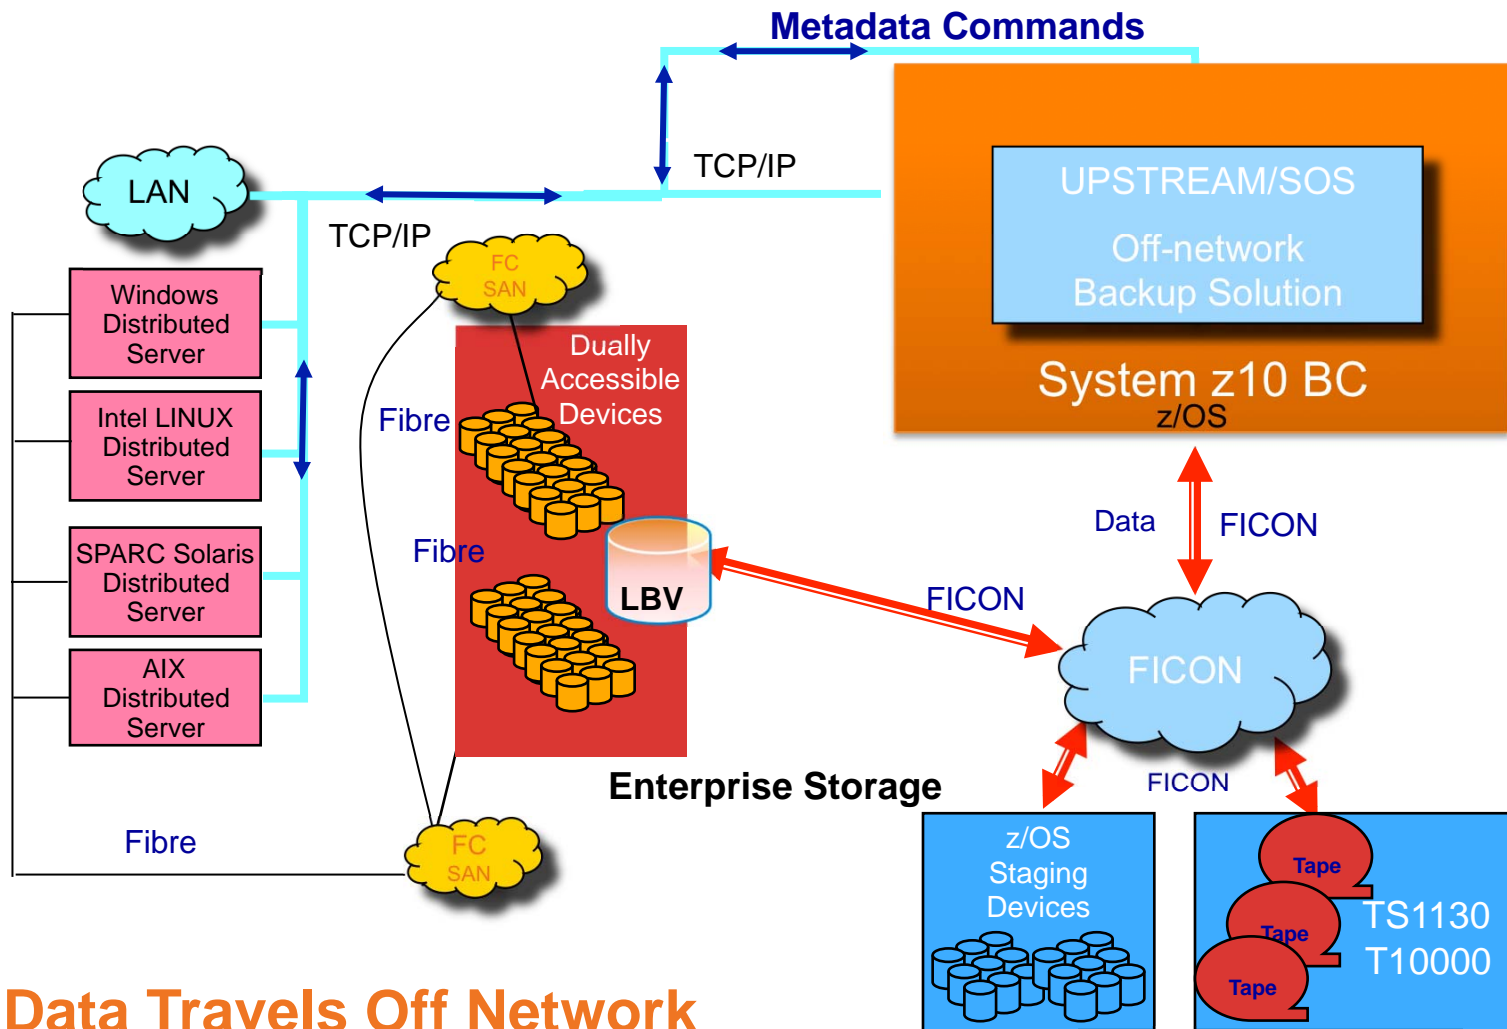
FDR/UPSTREAM/SOS[®] – Off-Network File Level Enterprise Data Protection

- IBM z/DDB and EMC zSOS Multi-Platform access extends FDR/UPSTREAM file level protection.
- Off-Network File Level Enterprise Data Protection for AIX, Linux, UNIX, Windows and Linux on System z.
- Control data travels on the TCP/IP connection.
- All backup data travels on FICON channels
Extends UPSTREAM File Transfer Facility with off-network capability.

FDR/UPSTREAM/SOS Data Flow, File Incremental Backups



FDR/UPSTREAM/SOS[®] Off-Network File Level Enterprise Data Protection



Data Travels Off Network

Non-Disruptive Business Continuance

Consider...

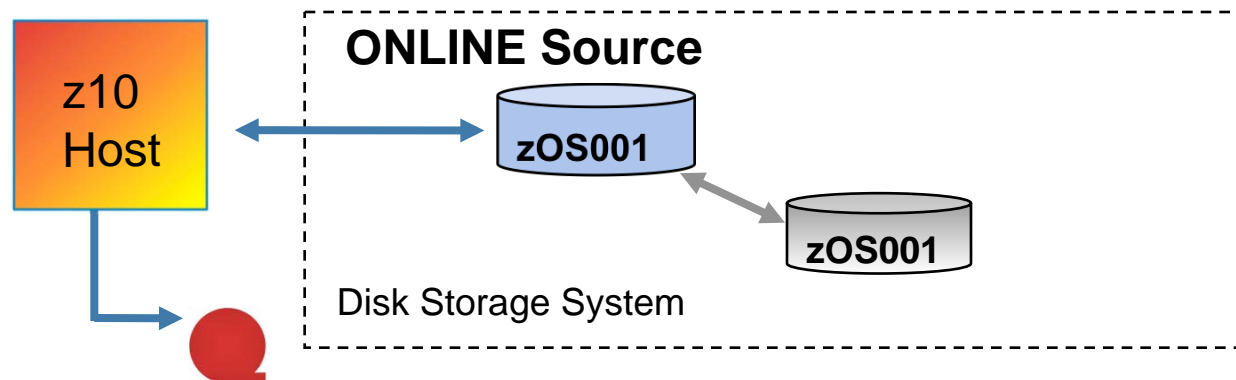
z/OS and Enterprise Business Continuance

- Backups are taking longer and longer.
- Larger disk volumes handle growth but require more time.
- Non-disruptive backup requires special DR procedures.
- Platform unique solutions make for costly integration.
- Business resiliency requires “rapid restore”.
- You need to reduce amount of backup data to store.

THESE ARE PAIN POINTS!

Consider...Virtualizing Backup with Data Replication for an Instant Point-in-Time copy

- No more problems meeting your backup window!
 - Take backups and copy data sets without long interruptions
 - Minimize CICS and WebSphere disruptions
 - Eliminate unacceptable recovery time
- Complications...create complexity
 - Hardware replication produces exact duplicates
 - z/OS is not tolerant of online duplicates
 - Recovery Procedures require special renaming or “Conditioning”
 - There are multiple replication technologies from different vendors

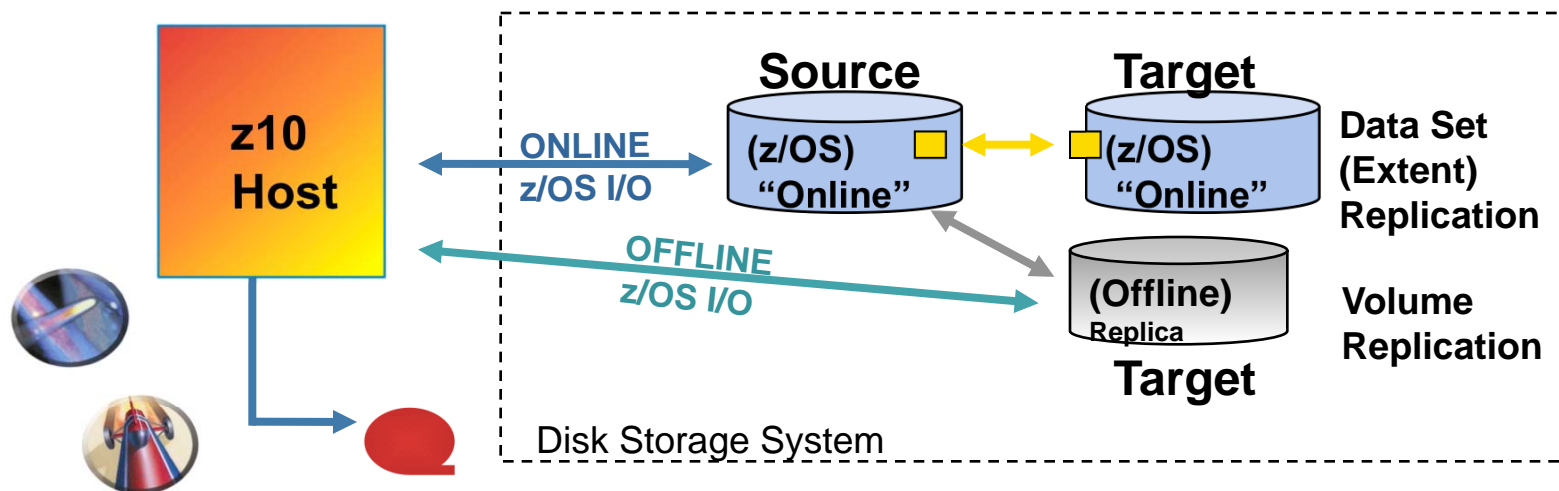


Ways to Reduce Complexity... of non-disruptive z/OS backup

Make something...

- Simple...provide a common interface
- Flexible...offer Volume and Data Set (Extent) granularity
- Fast...let hardware do a quick copy to offline volumes or online data sets
- Non-Disruptive...Backup the offline target
- Reliable...Allocation, Catalog, Space Management all under program control

FDRINSTANT is a software solution that does these things...



All FDR Products Support... All z/OS Storage Virtualization features

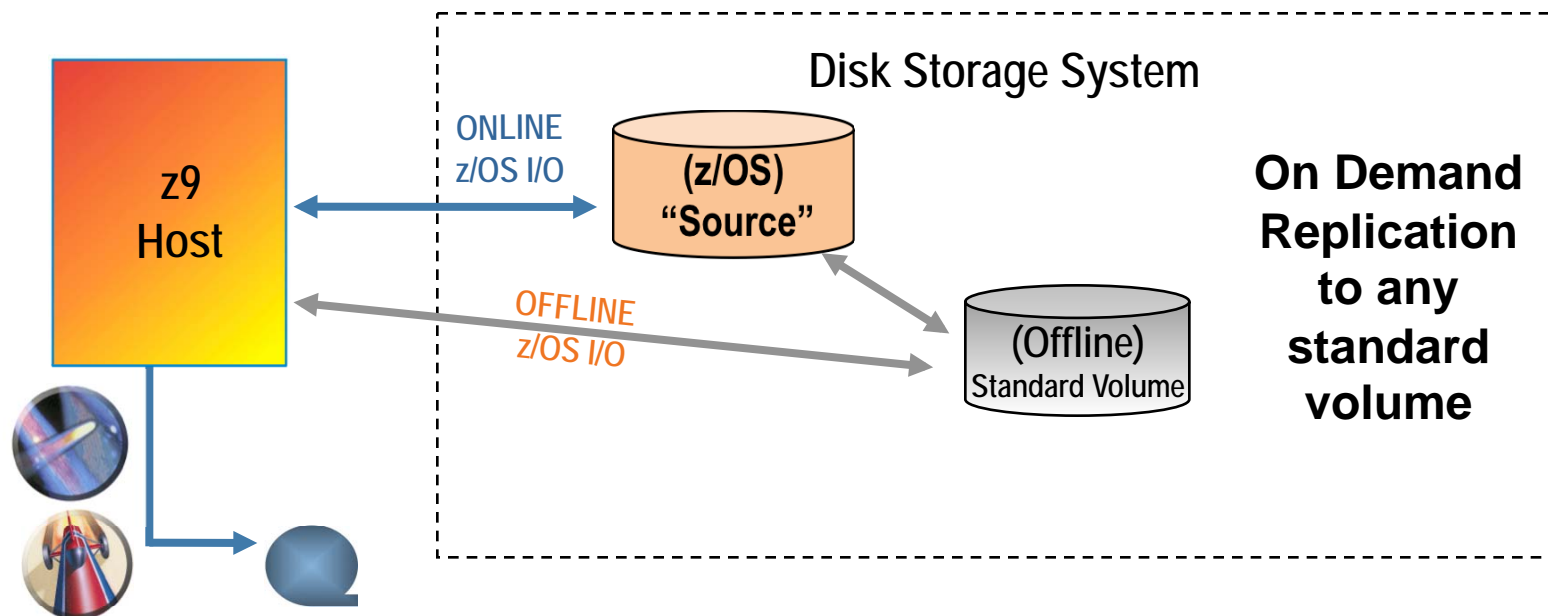
- IBM DS8700 / DS8800
 - FlashCopy, Space Efficient FlashCopy, Incremental FlashCopy, FRR (Fast Recovery Restore)
- EMC VMAX
 - TimeFinder CLONE/SNAP
 - Compatible Native FlashCopy (CN/FC)
 - TimeFinder SNAP to virtual volumes
- HDS USP-V
 - ShadowImage
 - FlashCopy



“Persistent / On Demand” FV replication (fully provisioned target volumes)

Split a persistent “mirror” relationship to create a (FV) PIT copy.

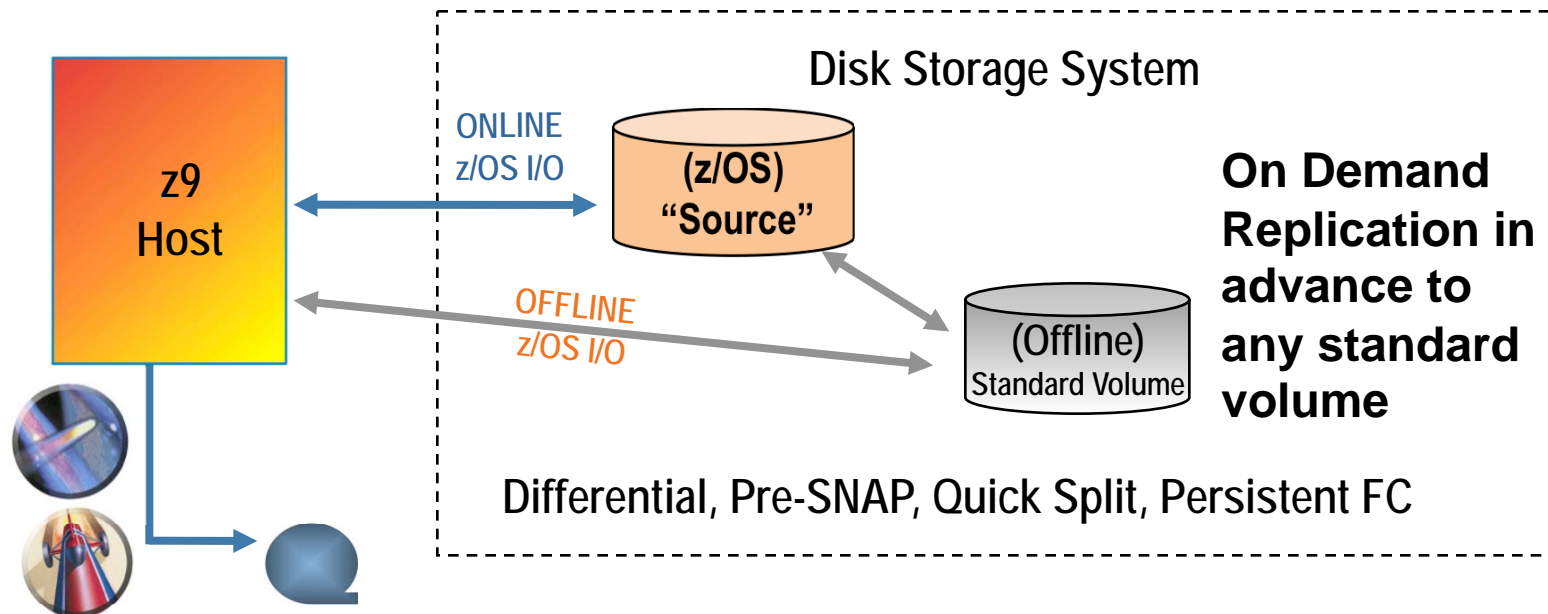
- EMC TF/Mirror – Hitachi ShadowImage – IBM Persistent FC
- Dynamically create a full volume PIT copy on demand.
- EMC TF/Clone – EMC/HDS/IBM FlashCopy – ShadowImage – Quick Split



Enhancements for Faster Full Volume Replication

FDR, ABR full volume and ABR incremental backup enhancements

- ...Creates an On-demand relationship in advance
 - TimeFinder/Clone (Full Volume operations) – Pre Snap, Activate, Post Snap
 - IBM Persistent FC (FV) – Supports FRR Restore

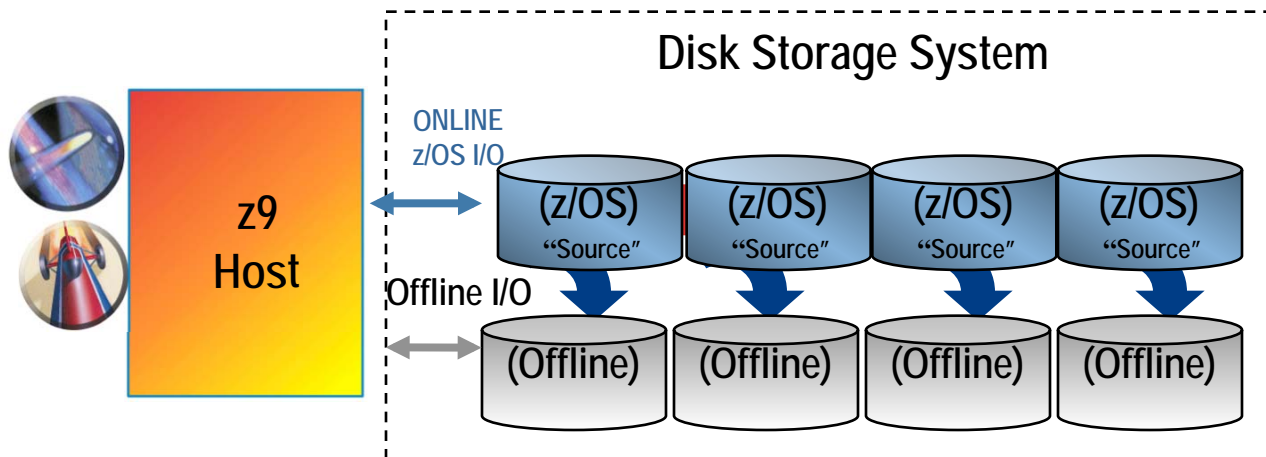


Enhancements for Faster Consistent Multi-volume Replication...

ABRInstant...

Creates multiple disk copies at a single point-in-time in parallel.

- Multiple volume point-in-time full volume replication – Consistent FlashCopy, Consistent ShadowImage, TF/Mirror CONSPLIT and TF/Clone CONSNAP
- Full-volume and incremental backup from offline “Consistent” (PIT) copies.



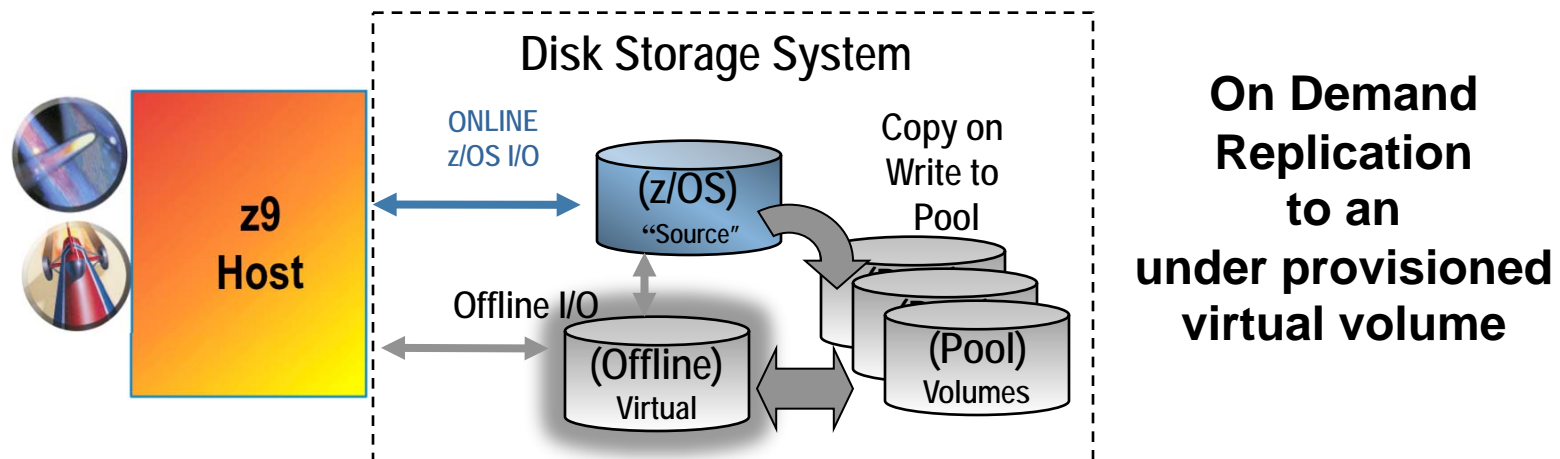
Thin Provisioning

Using Less Capacity Means

Fewer Drives, Less Power, More GREEN

Puts any standard volume into a “copy on first write” (COFW) relationship with a “space efficient” virtual device backed by a pool of physical volumes, i.e. only copies data tracks that are updated.

- IBM Space Efficient FlashCopy [FC/SE] (FV)
- EMC TF/SNAP (FV)



**On Demand
Replication
to an
under provisioned
virtual volume**

Considerations for Virtualized Storage (Full vs Thin Provisioning)

Full Provisioning (Traditional) Usage Suggestions

- When a full physical copy of data on disk is needed (for example to protect against hardware failures)
- Longer-term copies – FlashCopy/SNAP to be retained for more than 24 hours, as a backup or archive
- Copies of data with heavy update activity >20%
- Whenever performance is more important than economy

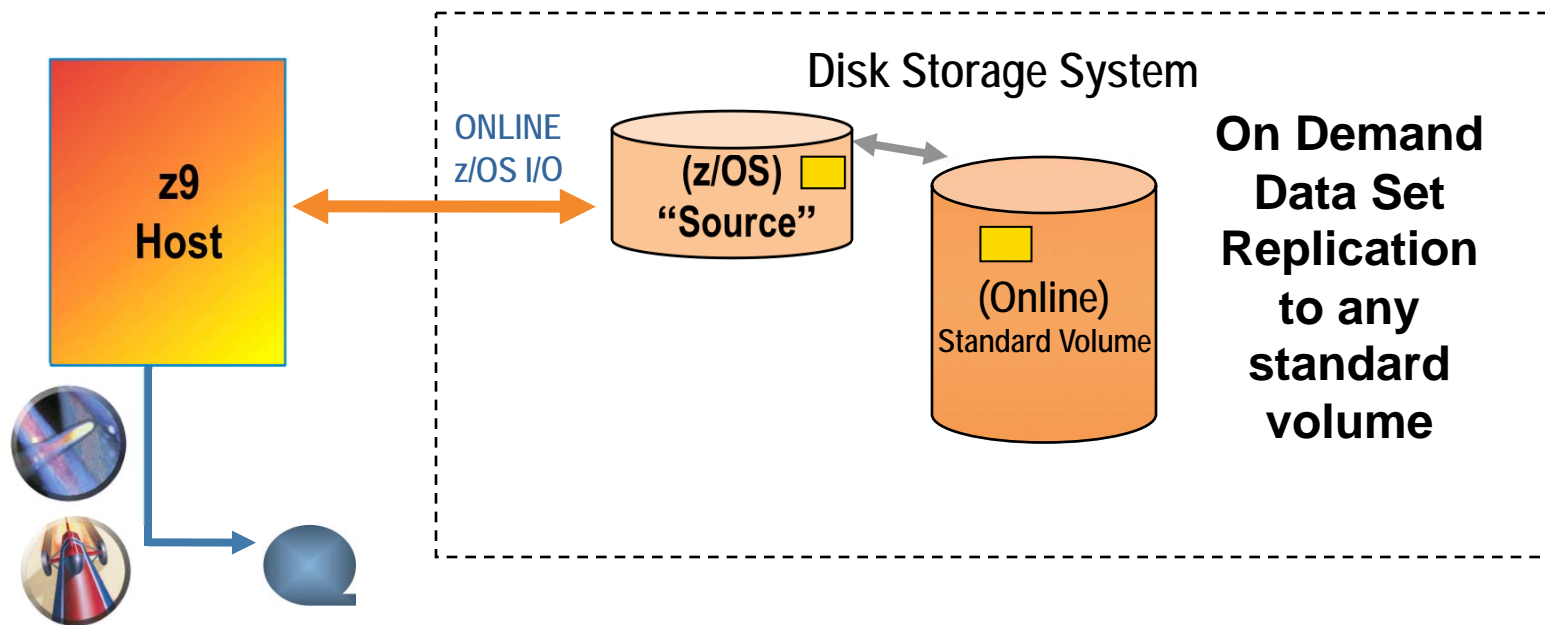
Thin Provisioning (Virtualized Storage) Usage Suggestions

- When a backup on tape or a full remote copy is available as the primary means of data protection
- Short-term copies – FlashCopy/SNAP to be kept for less than 24 hours for consistency testing or Data Mining
- Copies of data with light update activity <20%
- When a choice based on economy will not put the business at risk

“On Demand” Data Set Replication

Creates a data set (extent level) copy dynamically on demand.
 Copying or moving multiple data sets in parallel.

- EMC TF/Clone Data Set Snap Operations (DS)
- EMC/HDS/ IBM FlashCopy (DS)



Virtualize FDRINSTANT Data Set Operations for Performance

- **Objective**

- Improve the performance of four jobs that run concurrently to copy ~350 data sets, >111 million tracks.

- **Solution**


- Employ FDRINSTANT Copy using parallel sub-tasks to perform 20 concurrent copy operations. (MAXTASKS=5).

- **Results**

- Copy complete in <8 min (the longest of four concurrent jobs.)


FDRCOPY with FDRINSTANT
Copying 300 Data Sets (111 million tracks)

	FDRInstant parallel data set copy MAXTASKS=5
FDRINSTANT four concurrent jobs 5 data set operations per 350 Data Sets 111 million tracks	7.8 min 3.5 min 5.5 min 5.3 min



FDRMOVE using hardware replication consolidates... 25,000 DB2 Files to New Large Volumes in <8 minutes

- Objective
 - Consolidate 25,000 DB2 files to new larger volumes
- Task
 - Allocate ~ Catalog ~ Delete ~ 25,602 DB2 files
 - Copy (~ 1.4 million tracks)
- Challenge
 - Relocate the DB2 files with minimal impact
- Solution
 - Employ FDRMOVE with FDRINSTANT technology
- Results
 - Less than 8 minutes to consolidate 25,602 DB2 files to new larger volumes. FDRMOVE relocating data sets at the rate of 3,200/minute (i.e. moving data at about 150 MB/sec)



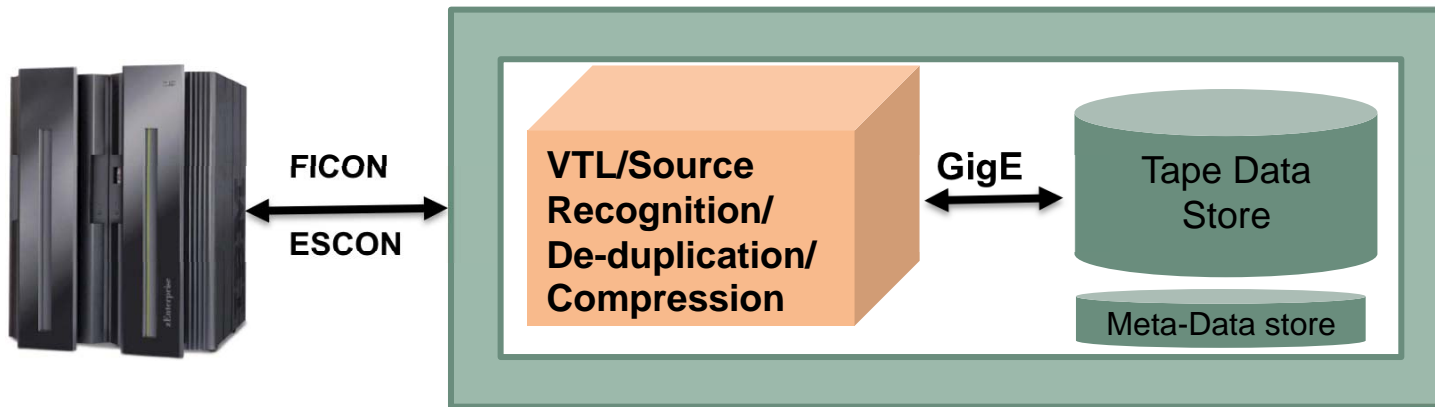
Benefits...with FDRINSTANT

Virtualizing Backup, Copy and Move...

- Maximum uptime avoiding costly customized integration projects
 - Proven technology that requires no changes for disaster recovery
- Freedom from the constraints of a constantly shrinking backup window
 - Non-disruptive backup solution for very large amounts of data
 - No negative impact on production schedules
 - Supports continuous business operation for Systems z mainframe data
- Protection for the investment in existing Systems z mainframe resources
 - Tape library management software, tape silos, VTL, tape drives and staff
 - Extend z/OS security to non-disruptive mainframe data protection
 - Extend z/OS auto-ops to non-disruptive mainframe data protection

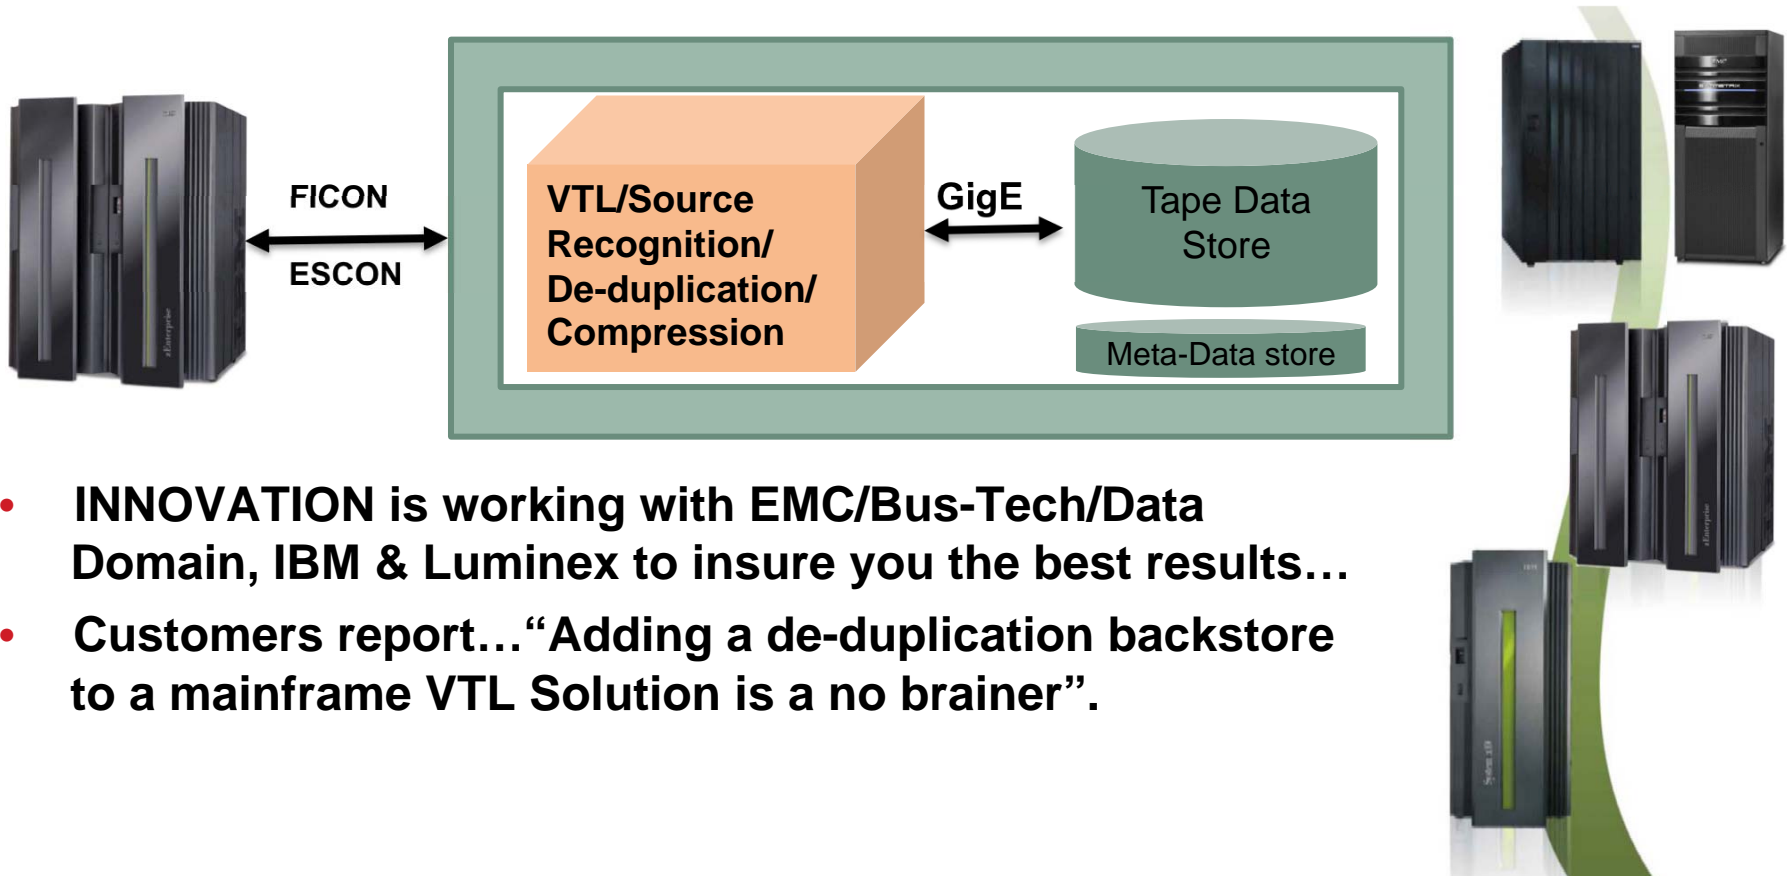
Storage Resource Optimization

De-duplication is a natural complement to z/OS mainframe and Multi-Platform z/OS Distributed Data protection...



- Media generating headlines vs. analysts setting expectations
- z/OS backup already offers sophisticated data reduction...
 - Incremental Backup – Merge Backup – Last Tape Stacking
- De-duplication works...and dedupe appliances do more...

FDR, FDRABR, FDRSOS, UPSTREAM and RESERVOIR...all benefit.



- **INNOVATION** is working with EMC/Bus-Tech/Data Domain, IBM & Luminex to insure you the best results...
- Customers report...“Adding a de-duplication backstore to a mainframe VTL Solution is a no brainer”.

Efficiency Testing...

Shadow z/OS and Distributed Data Backup

50 z/OS volumes w/130GB (3390-1, 3, 9 and 27)

- z/OS IPL packs, System & ISV executables, JCL, Source, Listings, sysdump, FTP server, application work, DB2 systems & databases, BookManager and product softcopy manuals.
- ABR daily incr (7% changing) & weekly fulls.


50 distributed servers w/2 TB of data

- Windows Servers, Solaris on x86, SuSE, & RedHat Linux on x86 and on System z.
- UPSTREAM daily incr (4% changing) & a weekly full merge backup.

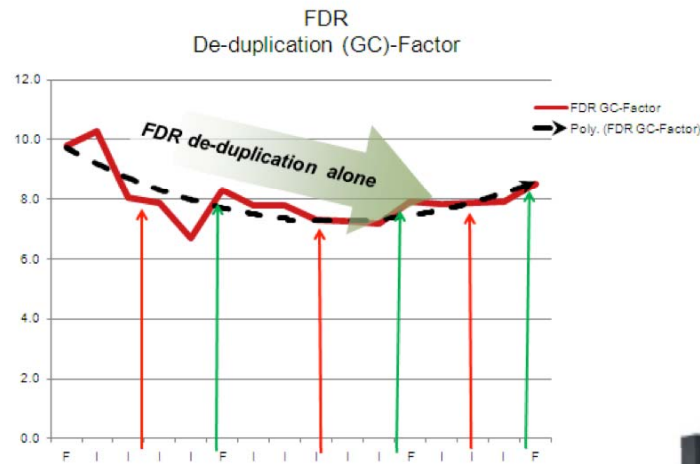


De-duplication by itself...

De-duplication by itself reduces
UPSTREAM backstore requirements...7 to 1

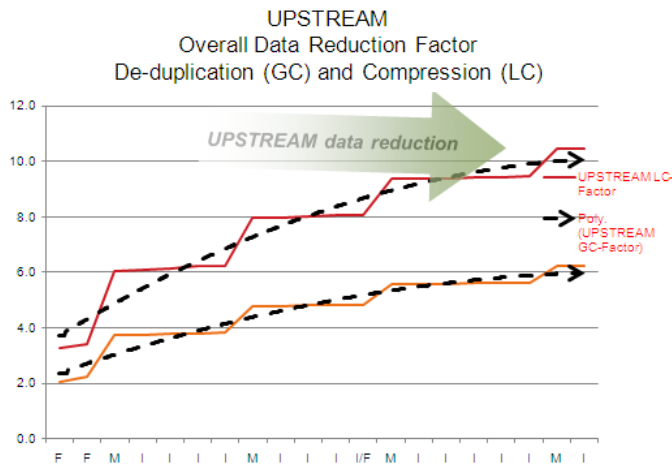


De-duplication alone reduces
FDR backstore requirements...8 to 1

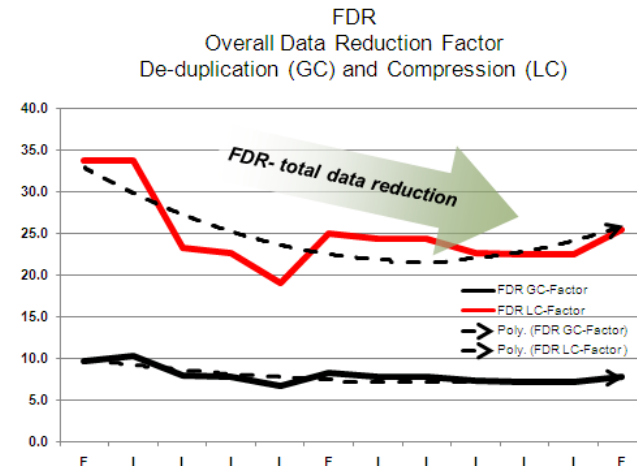


Leveraging compression...

further reduces
UPSTREAM backstore to...10.5 to 1



further reduces
FDR backstore to...25 to 1



Understanding Reduction Ratio and % Saved...



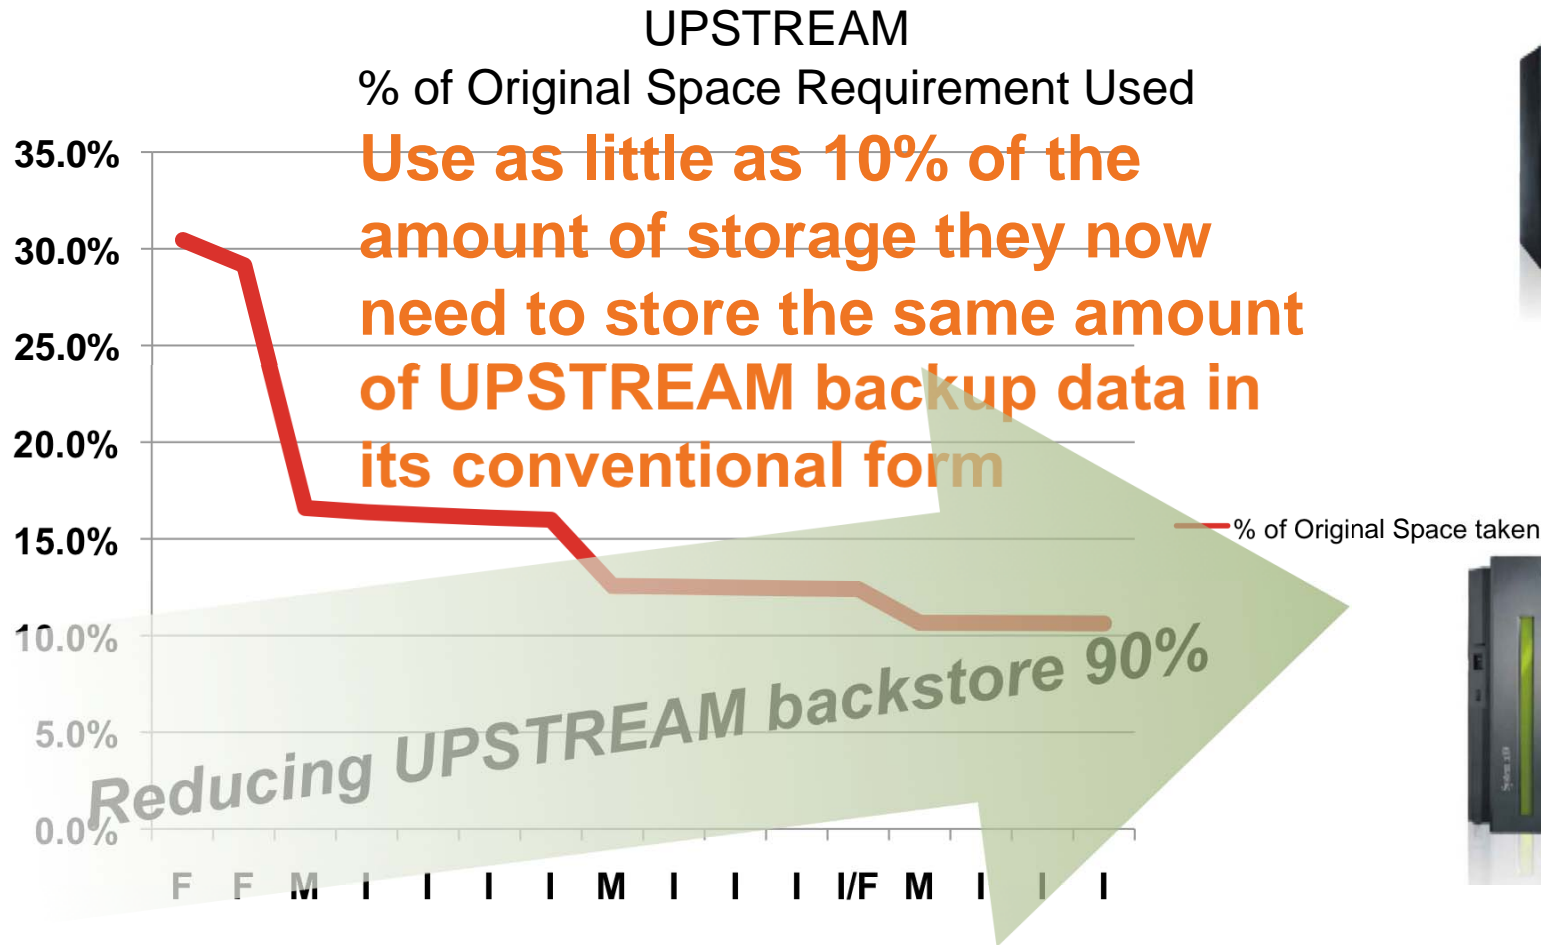
- 10 to 1 data reduction is a backstore savings of 90%
- Greater than 10 to 1 can be only be a bit more...
- How much is a 20 to 1 reduction?



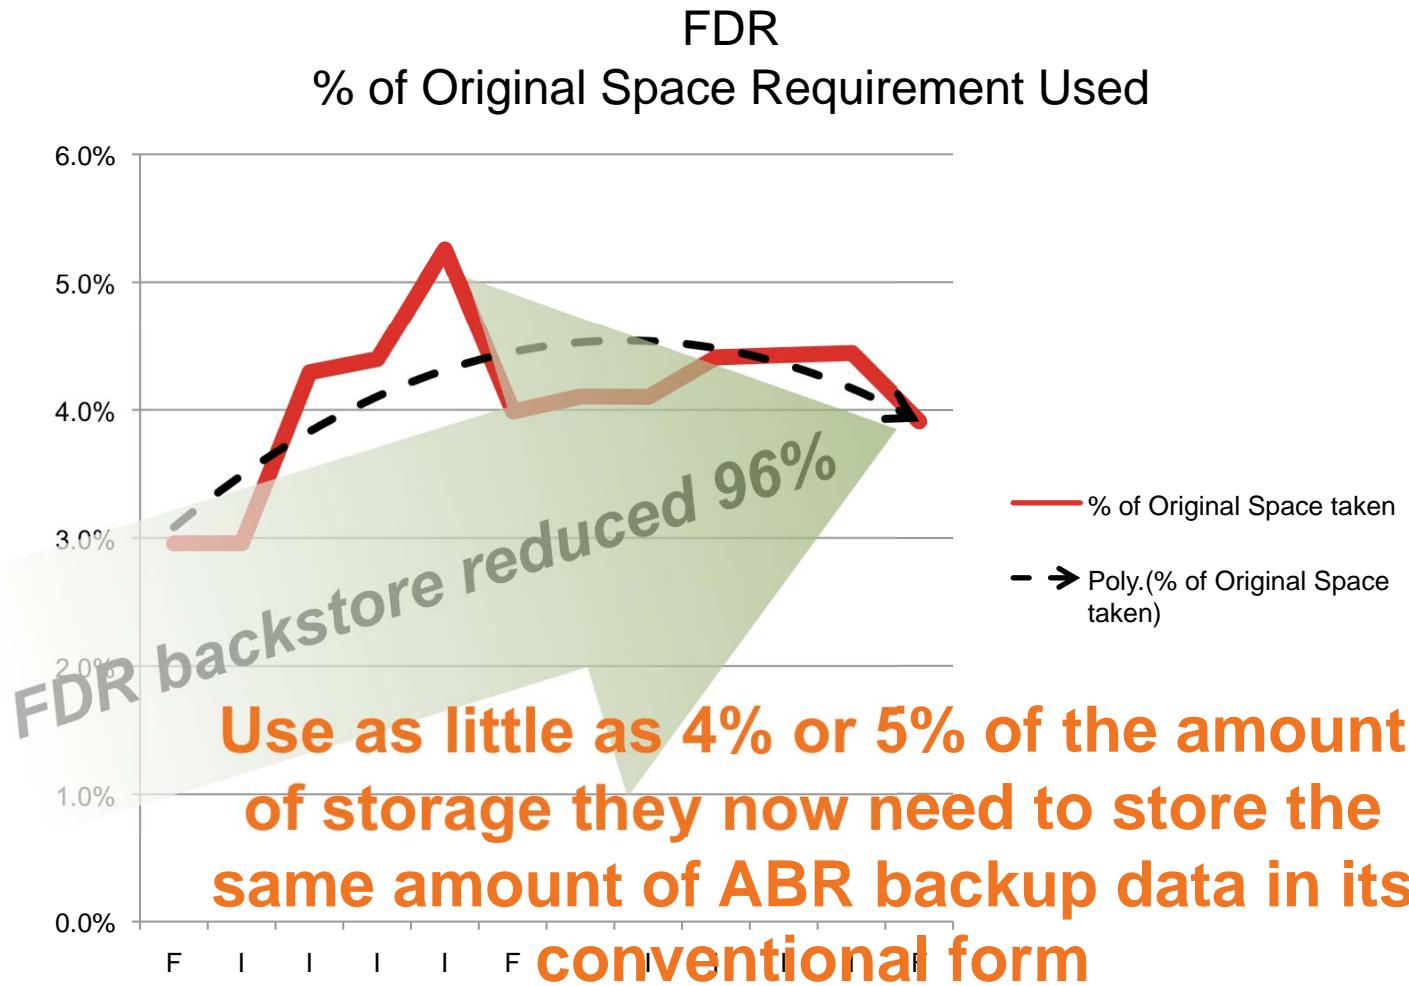
Ratio to 1	Saved	Stored
50	98%	2%
25	96%	4%
20	95%	5%
16	94%	6%
12	92%	8%
10	90%	10%

Customers Report Average

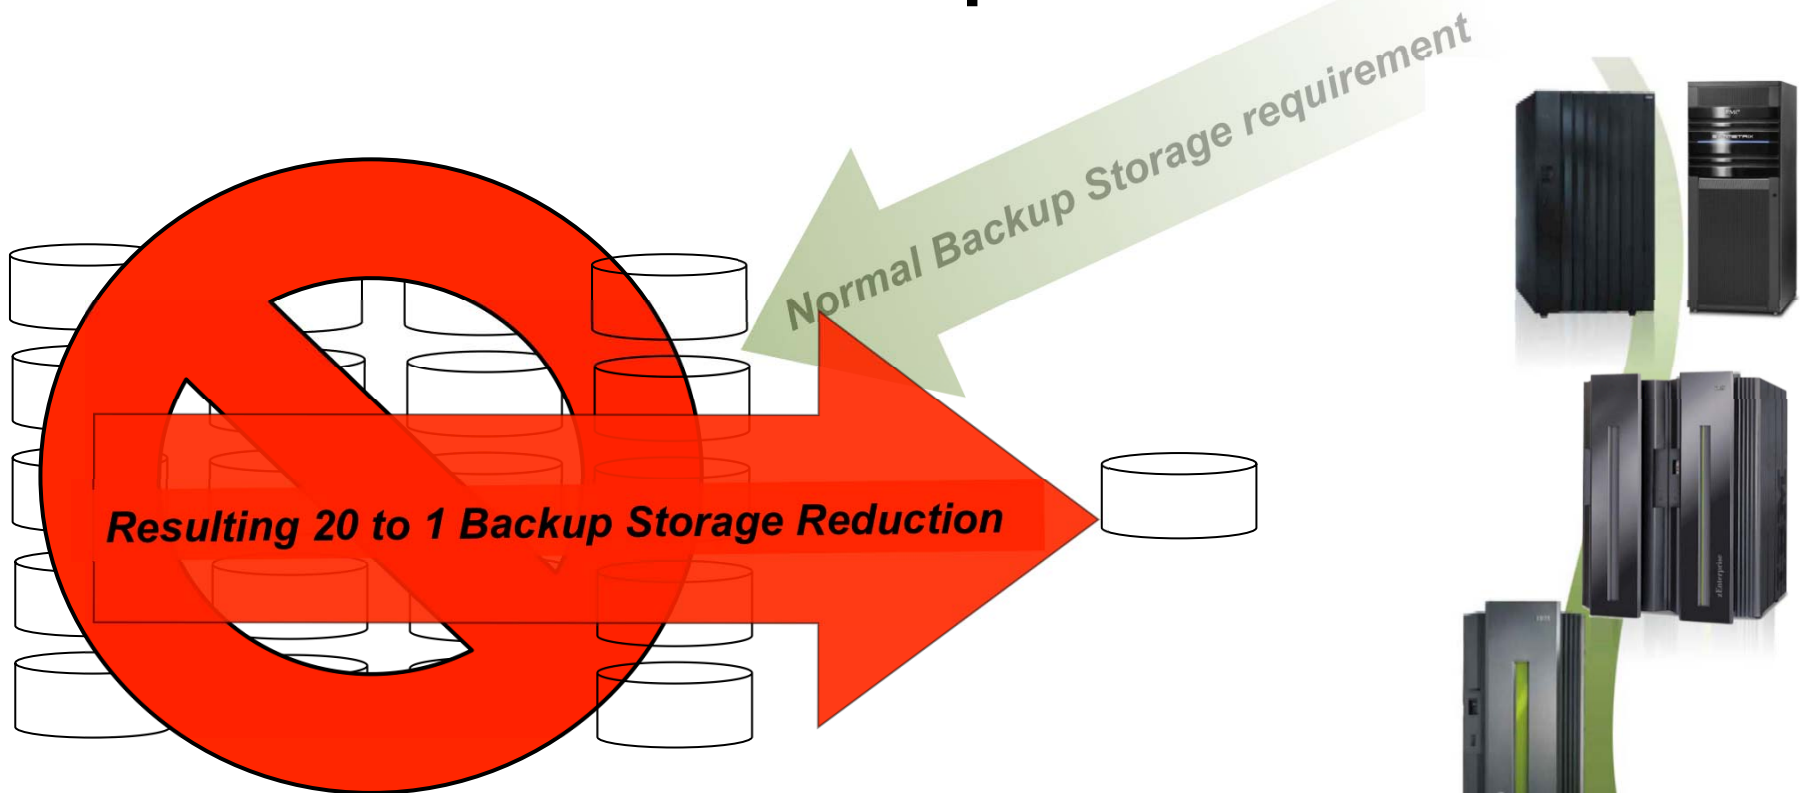
How much can de-duplication reduce the ...UPSTREAM backstore?



How much can de-duplication reduce the ...FDR backstore?



What does 20 to 1 mean in practical terms?



Results will vary but the addition of a de-duplication backstore to a mainframe Virtual Tape Solution is for certain a “no brainer” ...

How do you *automate consolidating physical and virtual tape system?*



Bring balance to tape conversion and stacking with...FATSCOPY™ FATS/ERASE...



TS1140 / T10000



EMC DLm



IBM TS7700



SUN VSM



EMC MDL

The ONLY Tape Migration and Stacking Tools you will ever need!

Summary

- **INNOVATION Data Processing leverages the latest technological innovations to virtualize, simplify and green all of its enterprise business resiliency solutions...**
- **INNOVATION Data Processing is a single source for enterprise business resiliency solutions addressing z/OS, Linux on System z, Linux on Power, Linux on x86, Novell NetWare/OES2, UNIX, Windows, Client/Server, SAN and direct attach storage providing:**
 - High-Performance Data Protection
 - Enterprise Data Protection
 - Non-disruptive Business Continuance
 - Storage Resource Optimization
 - Privacy Information Protection

Storage “Virtualization” makes for a... “Smarter...Planet...Business and...Data Center

Take away...

Mainframe Linux and z/OS users...Seize the opportunity

- **Shake off the conventional data protection constraints**
- **Meet Consolidation, Virtualization, and Cloud objectives**
- **Choose INNOVATION business resiliency solutions**
- **In conjunction with z platform advances, distributed enterprise storage cross-platform access, hardware replication and high capacity mainframe VTL deduplication.**



CORPORATE HEADQUARTERS: 275 Paterson Ave., Little Falls, NJ 07424 • (973) 890-7300 • Fax: (973) 890-7147
E-mail: support@fdrinnovation.com • sales@fdrinnovation.com • <http://www.innovationdp.fdr.com>

EUROPEAN OFFICES:	FRANCE 01-49-69-94-02	GERMANY 089-489-0210	NETHERLANDS 036-534-1660	UNITED KINGDOM 0208-905-1266	NORDIC COUNTRIES +31-36-534-1660
--------------------------	--------------------------	-------------------------	-----------------------------	---------------------------------	-------------------------------------