



FASTEST FDR EVER!
from
INNOVATION Data Processing

Thomas J Meehan
SHARE Orlando
27 February 2008
Session Number 3069

Trademarks and statements: IBM, System Z, System z9, and z/OS are trademarks or registered trademarks of International Business Machines Corporation. DMX, SYMMETRIX and TimeFinder* are registered trademarks of EMC Corporation. FDR, FDRABR, FDRPAS, FDRMOVE, FDRERASE, FDRINSTANT and FDRCRYPT are service marks, trademarks or registered trademarks of Innovation Data Processing Corporation. All other service marks, trademarks or registered trademarks are the property of their respective owners.Inc.

Overview of FDRInstant Support for...



- IBM December 07 release of DS8KR3 mCode
 - Space Efficient FC (FC/SE)
 - Incremental FC
 - Dynamic Volume Expansion
 - FlashCopy performance improvements
- EMC January 08 release of Enginuity 5773
 - Compatible Native FlashCopy (CN/FC)
 - EMC TF/Snap to virtual volumes
 - TF Clone and Dataset Snap performance improvements
- HDS January 08 replication enhancements
 - FlashCopy V2 to TrueCopy (PPRC) primary...

Overview of FDRInstant support for...

December 07 release of IBM 2107 (DS8KR3) mCode

- Space Efficient FC (FC/SE)
[“under provisioned” virtual volumes backed by a pool real volumes only record tracks that change on the primary volume]
- Incremental FC
[hardware function to only copy changed tracks requires the volume to be left in a persistent relationship]
 - Incremental FlashCopy Establish is not defaulted.
User must code “type=incr”.
- Dynamic Volume Expansion
[a hardware console function to enlarge a volume while it remains online]
- FlashCopy performance improvement
 - Fast Recovery Restore is not defaulted, users must code “type=frr”

EMC recent replication enhancements

- Compatible Native FlashCopy (CN/FC)
- TF Clone and Dataset Snap performance improvements
- Auto-Swap...Employing FDRINSTANT to copy datasets to a target in an IBM or EMC system will result in the use of FlashCopy or TF SNAP hardware replication. Generally FDRINSTANT will always successfully complete with the datasets being moved to the target volume that is in a Hyper-Swap / Auto-Swap protected relationship.

HDS recent replication enhancements

- FCV2 to PPRC primary...FDRINSTANT had reverted to conventional IO if the target volume is in a Hitachi storage system. A new release of mCode available now for the USPv that has "Native FCTOPPRC Primary" facility that will provide the same support as IBM FlashCopy. Those of you with Hitachi equipment should check with the local HDS account team to see what hardware and mCode is in available in your account.

What Customers Can Expect...



Customers with FDR V5.4 L60+ can expect to see:

- 50% or more reduction in FDRInstant D2D backup and data set copy times depending on their environment.
 - FlashCopy V2 [EMC, HDS and IBM]
 - ShadowImage w/Quick Split [HDS]
 - TimeFinder/CLONE Full Volume and Data Set Snap Facility with Pre-Copy and parallel data set snap operations [EMC]

“Customer can expect...Anytime, Anywhere, Empowerment”

Reap the benefits of consolidating storage, servers and bringing up new applications with peace of mind; knowing your systems have end-to-end Business Continuity Protection that is fast, reliable, simple, automated, and non-disruptive.

Customers, with the latest releases of FDR, can expect to see immediate, performance improvements of up to an incremental 50% or more reduction in FDRInstant D2D backup and data set copy times depending on their specific FlashCopy, ShadowImage and TimeFinder CLONE/SNAP environments.

- (i.e. FDR V5.4 L60 or higher plus, FlashCopy V2, ShadowImage w/Quick Split, TimeFinder/CLONE Mainframe Snap Facility w/Pre-Copy and support for parallel data set snap operations).

Complement FDRInstant non-disruptive data protection with; FDRMOVE and FDRPAS...for the fastest non-disruptive volume migration and minimally disruptive volume consolidation solutions available.

FDRERASE...complements FDRMOVE and FDRPAS as the only z/OS data erasure solution available with US Government CCEVS certified secure erasure to quickly remove privileged data from mainframe disk when you are leaving a DR site and whenever mainframe disk storage is leaving your control.

Compelling... Hardware data replication technology



Full volume (FV) and extent level (DS) replication

- “Mirroring” – always fully provisioned (FV) persistent relationship, typically highest performance
 - e.g. EMC TimeFinder/Mirror, Hitachi ShadowImage
- “On Demand Copy” – fully provisioned (FV/DS) flexible relationship, w/high performance
 - e.g. EMC TF/Clone, EMC/HDS/IBM FlashCopy
- New “Space Efficient/Virtual – thinly provisioned (FV) flexible relationship, w/economical space-saving
 - e.g. EMC TF/Snap (FV), IBM Space Efficient FlashCopy

FDRInstant supports the use of all three types of hardware replication. Your storage subsystems may be licensed for any or all of these technologies, which have different names and sometimes significant differences in implementation depending on your hardware vendor of choice.


- “Mirroring” – persistent relationship w/high performance (FV)
 - e.g. *EMC TimeFinder/Mirror, Hitachi ShadowImage*
- “Standard” On Demand Copy – flexible relationship w/high performance
 - e.g. *EMC TF/Clone (FV/DS), EMC/HDS/IBM FlashCopy (FV/DS).*
- “Virtual Device/Space Efficient” On Demand Copy-economical space-saving
 - e.g. *IBM Space Efficient FlashCopy (FC/SE), EMC TF/Snap (FV).*

FlashCopy/SNAP/Clone Compelling... but Introduce Complexity



Hardware Point-in-Time Data Replication (PIT)

- Benefits...
 - Take backups and copy data sets without long interruptions
 - Eliminate unacceptable recovery time
- Complications...
 - Hardware replication produces exact duplicates
 - z/OS software is not tolerant of online duplicates
 - "Conditioning" or renaming require special procedures



Data Replication is a simple solution, to eliminate long interruptions for backup, unacceptable recovery times and production delays to copy data and hard downtime to install new disk.

Point-in-time replication allows you to instantly duplicate data sets and volumes.

However duplicate volume images have duplicate volume serial labels and duplicate files with identical names. Operating systems do not tolerate duplicate volumes online, so the duplicate is forced offline. Conventional software products cannot backup a duplicate directly because they only work against online volumes.

Some vendors provide software utilities for conditioning or re-labeling.

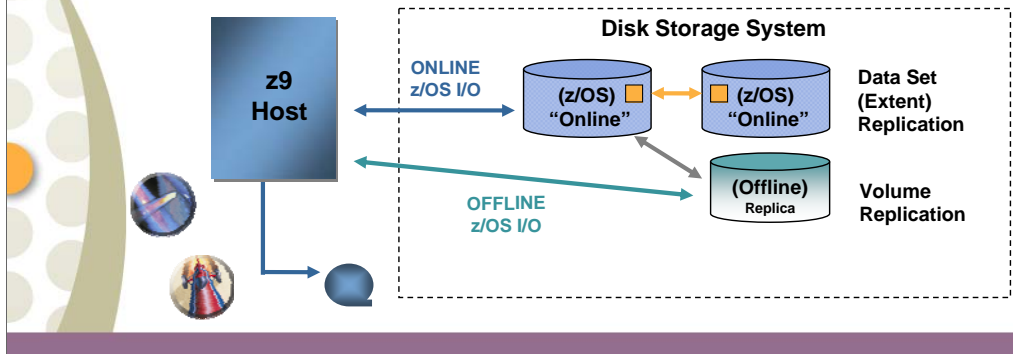
Re-labeling creates difficulties for system managed (SMS) volumes, VSAM and cataloged data sets. Data sets on re-labeled volumes appear as uncataloged, unless they are re-named and re-cataloged. A backup of a re-labeled volume appears to be from a new volume serial, not the original volume. This means special procedures during backup or restore in order to correct the volume serial (such as, restore volume B back to volume A and re-label it as A).

FDRInstant – Simplifies Hardware Data Replication



FDRInstant provides:

- **Simplicity...**Common interface to vendors different replication features.
- **Flexibility...**Volume and/or Data Set (Extent) granularity
- **Speed...**Hardware does quick copy to offline volumes or online data sets
- **No Disruption...**Backup reads offline target
- **Reliability...**Allocation, Catalog, Space Management under program control




Licensing the INNOVATION FDRInstantBACKUP product, sometimes called FDRInstant for short, enables FDR and FDRABR with “Instant” functionality; a simple common interface to the appropriate hardware data replication function no matter what mix of hardware vendor’s storage systems you may have in your site. ABRInstant, for example, automating the use of hardware data replication is the most advanced non-disruptive business data protection and business continuance solution available today for any z/OS mainframe storage system.

FDR and FDRABR backup and restore processing, as well as FDRCOPY data set copy and move at FDR V5.4 L60 and above now contain allocation, catalog and disk VTOC processing performance improvements. The introduction of more efficient CCW chains to access VTOC data, reduces I/O processing time. Multiple reads of the input volume’s VTOC for backup and output volume for restores are gone. Employing an IBM service to identify the end of active data in a VTOC, FDR now reads only as much of the VTOC as necessary. This significantly reduces VTOC read time especially for customers allocating larger VTOC sizes to accommodate today’s larger devices, e.g. 3390-27 and 3390-54.

Although these improvements apply to all FDR and FDRABR backups and restores, the improvement is most noticeable in FDRInstant data set FDRCOPY. When instant replication techniques, such as SNAP and FlashCopy, are used to copy the data, allocation, catalog and disk VTOC processing is a significant portion of the total elapsed time and that is exactly where these changes provide significant benefit.

FDRInstant – Supports Mirroring Replication



“Splits” a persistent mirror relationship to create a PIT in copy.

- EMC TimeFinder/Mirror (FV)
 - Consistent BCV Split & Concurrent BCV
- Hitachi ShadowImage (FV)
 - Consistent Split

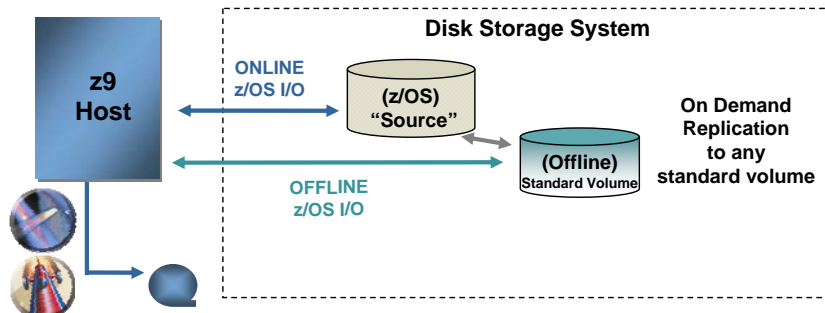
TimeFinder/Mirror and Hitachi ShadowImage – Similar to PPRC remote mirroring, TF/Mirror and ShadowImage employ a specially configured local volume (i.e. within the same storage system) as a persistent mirror of a standard (online) disk, that can be “split” from the standard volume to instantly create a point-in-time copy of that source volume. TimeFinder/Mirror is often referred to as BCV TimeFinder. Returning the volumes to their relationship after the backup allows them to remain in sync until the next “split”.

FDRInstant V5.4 L60 and above continues to support traditional TF/BCV and HDS ShadowImage. FDRInstant supports Concurrent BCV (two BCVs assigned to one standard volume) as well as Consistent Split (splitting multiple BCV or ShadowImage volumes at the same point-in-time).

The original and often referred to as an ultra-performance data replication technology, traditional TF/BCV and HDS ShadowImage are giving way to newer more flexible data replication technologies.

Although this presentation addresses ShadowImage, FDRInstant also supports Hitachi Synchronous Remote Copy (HRC, also called Truecopy), a similar PPRC like mirroring replication technology most often used between two subsystems; one local and one remote, that can also be used within a single disk subsystem. Truecopy does not support Consistent Split or Quick Split which we will talk about later...

FDRInstant – Supports “On Demand” FV Copy



Dynamically creates a full volume copy on demand.

- EMC TF/Clone (FV)
- EMC, HDS, IBM FlashCopy (FV)


TimeFinder/Clone and EMC, HDS and IBM FlashCopy (V1) (Full Volume operations) – employing standard disk volumes as targets, ABRInstant can “Clone/Flash” an image of an entire full volume onto an offline volume in an instant.

Available for access immediately as a point-in-time (PIT) copy of the original, tracks are copied to the replica in the background. This is a more flexible replication technology as it has no requirement for limited use specially configured volumes.

The ABRInstant default for backup FlashCopy NOCOPY improves FlashCopy system performance as it means only changed tracks are copied in the background to the target, while ABRInstant reads all unchanged tracks from the original online source volume.

Though the very first TimeFinder/Clone of any volume to a given target will always need to copy the entire volume, differential replication ABRInstant (i.e. SNAP=USE parameter), improves DMX performance in that it insures only updated tracks are copied to the TimeFinder replica for any subsequent CLONE operation.

FDRInstant – Supports “On Demand” Consistent Multi-Volume Replication



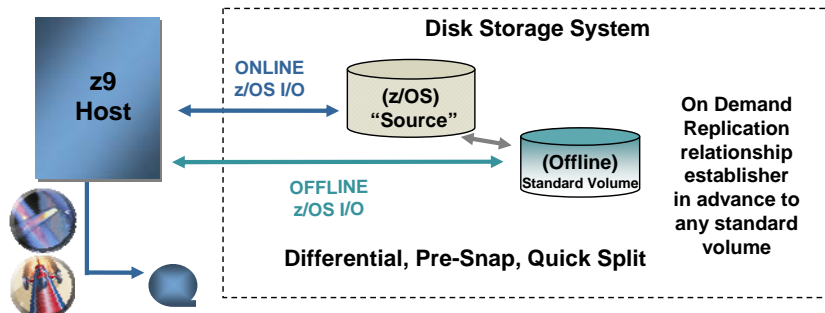
Create multiple, on demand copies, at a single point-in-time.

- ABRInstant concurrently creates multiple “Consistent” PIT backups in parallel and provides full-volume and incremental backup from the offline copies
- Multi-volume point-in-time full volume replication - Consistent FlashCopy, Consistent ShadowImage, TF/Mirror CONSPLIT and TF/Clone CONSNAP

Consistent PIT operations (SPLIT/CLONE/FLASH) allow the replication of multiple disks at a single consistent point-in-time. Ensuring their IO integrity and data consistency, by preventing applications from issuing dependent writes during the split/snap process. This allows the creation of point-in-time backups without quiescing database applications which typically employ journaling or logging to support recovery from this kind of crash consistent backup.

ABRInstant enhancements for consistent operations also significantly reduces the overall ABRInstant run time by allow the creation of the PIT backup replicas of multiple volumes to be a true parallel process, instead of a serial volume by volume backup process.

FDRInstant – Supports “On Demand” Fast Split/Clone FV Copy



FDR, ABR full volume and incremental support...

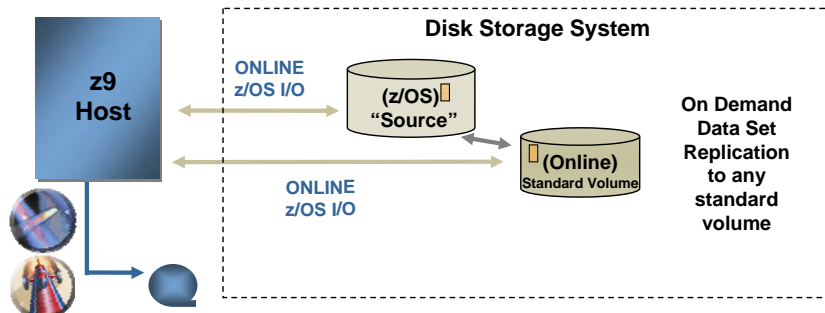
- Incremental FlashCopy and FRR [Fast Recovery Restore] (FV)
 - Incremental FC creates a persistent relationship in advance
- TimeFinder/Clone (FV)
 - Pre Snap, Activate, Post Snap
- ShadowImage (FV)
 - Quick Split

ABRInstant full-volume and incremental TF CONSNAP CLONE and ShadowImage backups see performance enhancements from the TF Pre-SNAP and SI Quick Split options.

A more flexible replication technology TF/CLONE has no specially configured volumes requirement, however it will benefit from maintaining a pre-existing relationship. Pre-SNAP processing, ABRInstant (i.e. SNAP=(USE,RET), improves FDRInstant performance as volumes are left in a “pre-copy” differential relationship. FDRInstant manages pre-snap, activate and post-snap processing. The result is that as tracks are updated on the source disk they are immediately replicated to the target disk, maintaining the target as a current copy of the source disk. There is little to no DMX background copy required with the CLONE target a current copy of the source disk,. This is conceptually similar to splitting a BCV mirror so on execution of a full volume operation the target disk is simply frozen and the ABR tape backup of the offline CLONE is able to complete in less time.

ABRInstant supports Hitachi ShadowImage Quick Split which eliminates the necessity to pre-establish ShadowImage pairs. Quick Split establishes a pair and splits it, with full volume integrity, in just a few seconds very similar to a FlashCopy operation. ShadowImage guarantees that the ShadowImage volume can be read as if were a traditional split point-in-time copy of the online volume.

FDRInstant – Supports “On Demand” Dataset Copy



FDRDSF and FDRCOPY support to...

- Create a data set (extent level) copy dynamically
- Copy multiple data sets in parallel
 - EMC TF/Clone Data Set Snap Operations (DS)
 - EMC, HDS, IBM FlashCopy V2 (DS)

TimeFinder/Clone (Data Set “Snap” Operations) and FlashCopy V2 (Data Set FlashCopy) – employing standard disk volumes, can instantly copy an image of multiple tracks from one volume to another. FDRInstant manages allocation, cataloging and uses TF or FC hardware extent level replication to copy tracks to the replicated image in the background after the “data set snap/flash” operation completes.

Prior to FDR V5.4 L60 FDRCOPY processed input disks serially. Allocating data sets concurrently to multiple output volumes, FDRCOPY would then process input disk volumes serially, i.e. one input disk at a time. A new ABRInstant operand, MAXTASKS=n (1-32), now allows FDRCOPY to concurrently process multiple input volumes in parallel when input data sets are selected through the catalog.

Parallel processing (MAXTASKS) can result in significant reductions in overall elapsed time when copying or moving data sets from multiple volumes when doing physical I/O to copy data, but provides the most performance improvement when FDRCOPY is used in conjunction with FDRInstant TimeFinder/CLONE DS SNAP and FlashCopy Data Set operations.

IBM views...FlashCopy Space Efficient (FC/SE) as Innovation that *Matters!*... for Storage, Clients and the Planet



Simplify and Increase Efficiency

- IBM FlashCopy Space Efficient capability (FC/SE) can lower costs by significantly reducing the disk capacity needed for copies.
 - Uses space to save just source data being updated
 - No need to allocate target copy capacity to match full source
 - Thinly provisioned target copies reduce capacity needed



Using Less Capacity Means

Fewer Drives

Less Power

More GREEN

Taken from an IBM marketing presentation

New ways to simplify and increase efficiency

IBM FlashCopy Space Efficient capability (FC/SE) can lower costs by significantly reducing the disk capacity needed for copies.

- Uses space to save just source data being updated
- No need to allocate target copy capacity to match full source
- Thinly provisioned target copies reduce capacity needed

Using less capacity means


Fewer drives

Less power

More GREEN

EMC TimeFinder/SNAP is a Similar
Space Efficient/Virtual Technology Solution

FDRInstant – Supports “Space Efficient/Virtual” FV Copy (Thinly Provisioned Target Volumes)



Relates a standard volume in a “copy on first write” (COFW) operation to a “space efficient” virtual device backed by a pool of physical volumes, i.e. only copies data tracks that are updated.

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

IBM Space efficient FlashCopy (FC/SE) and EMC TimeFinder/Snap –

employing a Vdev (Virtual Device) as a “space efficient” logical disk backed by a pool of “save disks”. FC/SE and TF/SNAP can instantly establish a “copy on write” relationship between the standard disk and the logical target volume, a Space Efficient FlashCopy to IBM and a Virtual Disk to EMC (sometime aka a skinny disk).

Data only moves i.e. tracks are copied to the logical target, when a write updates the content of the source volume, at which time FC/SE or TF/SNAP copy the “before” image of the track to a track in the pool backing the target. Space Efficient FlashCopy and TimeFinder/Snap are often referred to as an “economical space-saving option”.

FDR/ABRInstant supports full volume and incremental backup from offline FC/SE and TF/SNAP logical volumes.

Note... Though not a suggested use these logical volumes volumes have no restrictions on them. A user can simply vary a FC/SE or TF/SNAP logical volume on line at any time and begin using it just as they would any other volume. This of course would consume tracks in the pool backing the logical device.

Using Thinly Provisioned Target Volumes (IBM FlashCopy SE – EMC TF/SNAP)



To Employ FlashCopy Space Efficient or TF/SNAP Virtual Devices:

- Create a repository of volumes (an extent pool)
 - Repository space is to hold the pre-updated images of source volume tracks
- Define thinly provisioned target volumes in the extent pool
 - Thinly provisioned target volumes take very little space
 - Tracks are actually just pointers, until used
 - Space is freed in the repository when FlashCopy/SNAP relation is withdrawn, reestablished or when a SE/VDev volume is deleted
- Establish FC/SE or TF/SNAP from the source to the target volume
 - There is near-instant access to source and target after request
- Most Replication functions can be performed on a thin volume
 - **Regular, Consistency Groups, FRR can all be used**
 - **“Full volume” copy requires traditional FlashCopy or CLONE**
 - **Dataset FlashCopy coming soon**

Taken from an IBM marketing presentation

To Create a space efficient copy:

- Create a repository (multiple repositories can be created, max one per extent pool)
- Create a space efficient target in the extent pool of the repository where you want the source updates to be saved
- Very little space is actually used for SE target volumes – they are just pointers
- Establish the FlashCopy SE between the source and the FlashCopy SE target volume

Repository space is allocated only to hold the pre-updated images of source volume tracks

Space is freed in the repository when FlashCopy relation is withdrawn, reestablished or when a SE volume is deleted

There is near-instant access to source and target after copy request and after restore

Most FlashCopy functions can be performed on a FlashCopy SE volume

- Regular, Consistency Groups, Reverse Restore, FRR, Revertible can all be used
- Dataset FlashCopy coming soon
- Full volume copy requires traditional FlashCopy

Application Considerations for Fully vs Thinly Provisioned PIT Copy



Fully Provisioned (Traditional) Suggested Use

- Disaster Recovery...when a full physical copy of data is needed
 - Use local disk copy as protection against hardware failures
- Longer-term copies...relationships that last more than 24 hours
 - Production or long-term test copy
- When source data volume will be updated significantly (>20% updates)
- When performance is more important than economics

Thinly Provisioned (FC/SE – TF/SNAP) Suggested Use

- Disaster Recovery...when creating a sequential backup on tape
 - Use tape backup as protection against hardware failures
- Short-term copies – relationships that last less than 24 hours i.e. Data Mining
 - For qualification and consistency testing
- When source data volume will only have a small amount of updates <20%
- Economics should never put reliable recovery at risk
 - Thinly provisioned volumes only contain changes...the source volume contains unchanged data

Taken from an IBM marketing presentation

FlashCopy (traditional) Suggested Use

- Whenever a full physical copy of data is needed (for example to protect against hardware failures or to use as a production or long term test copy)
- Longer-term copies – FlashCopy relationships that last more than 24 hours
- Copies for source data that will be updated significantly >20% updates
- Whenever performance is more important than economics

FlashCopy SE Suggested Use

- Short-term copies – FlashCopy relationships that last less than 24 hours
- Copies of data with a small amount of update activity <20% during time they are active
- Copies for testing such as qualification and consistency testing
- Copies for Data Mining
- Backup on tape or other media (sequential read of target)

IBM views Dynamic Volume Expansion as Innovation that Matters! For Growth with Ease



Dynamic Volume Expansion

- Simplifies management by enabling easier, online, volume expansion to support application data growth.
- Yields more highly available, simplified volume expansion
 - No longer need to backup data, bring volume offline, delete it from the hardware configuration and then recreate volume to expand it
 - Can now be done online with a single command or GUI screen
 - Copy services relationships must be removed before expansion
 - *Note: Some operating systems-specific actions are needed on the z/OS server side before applications can see larger volumes...*
 - *INNOVATION has the tools to help here...*




Taken from an IBM marketing presentation

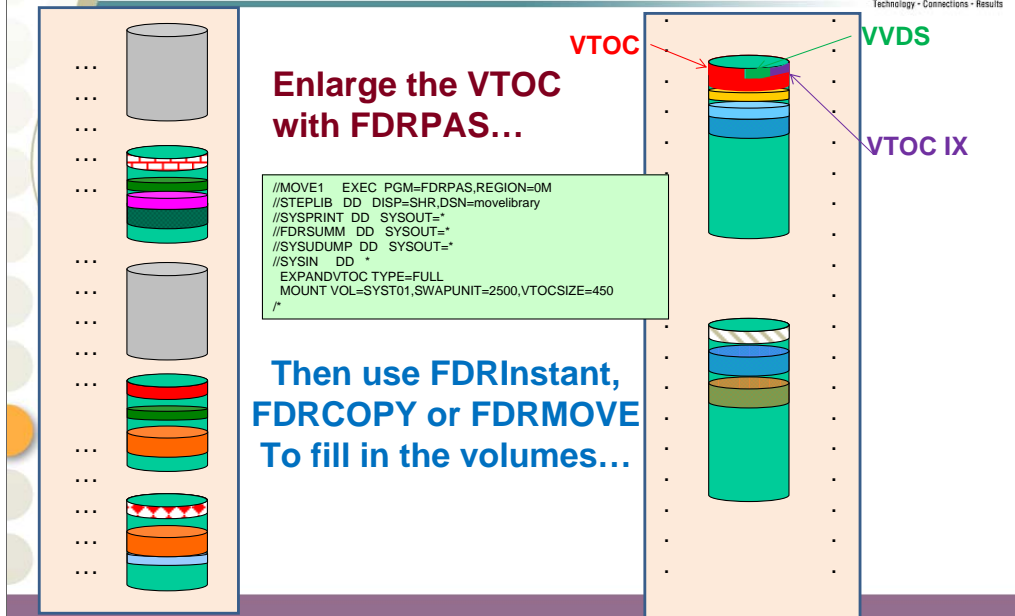
Dynamic Volume Expansion

- Simplifies management by enabling easier, online, volume expansion to support application data growth.
- Yields more highly available, simplified volume expansion
 - No longer need to backup data, bring volume offline, delete and recreate volume to expand it
 - Can be done on-line with a single command or GUI screen
 - Copy services relationships must be removed before volume expansion
 - Some operating systems-specific actions may be needed on server side before applications can see larger volumes
- ***INNOVATION Has Tools to Help Here...***

Using Dynamic Volume Expansion... For Consolidating to a Larger Volume and Finding the VTOC Too Small?



SHARE
Technology • Connections • Results



VTOC to small?

Enlarging a volume, ENSURE the VTOC on the source volumes enlarged by Dynamic Volume Expansion or moved with FDRPAS is large enough to accommodate the higher number of additional data sets that will eventually reside on the larger volume. Use FDRMOVE w/FDRPAS to expand the VTOC of a target volume as necessary.

ICKDSF is disruptive requiring shared DASD be put offline to all systems but one.

FDRMOVE w/FDRPAS includes EXPANDVTOC and SWAPBUILDIX useful utilities, employing the co-ordination techniques of FDRPAS™, that allow you to non-disruptively expand the size of the VTOC and rebuild a damaged or missing VTOC Index while the volume remains in use on all systems.

- **3390-9 use a VTOC of 150 trks (accommodates about 7,000 data sets)**
- **3390-27 use a VTOC of 450 trks (accommodates about 20,000 data sets)**
- **3390-54 use a VTOC of 900 trks (accommodates about 40,000 data sets)**

See the FDRMOVE Manual for notes on VTOC sizes and expanding VTOCs.

FDRInstant – Hardware Support



- Hardware Vendor Independent
 - Supporting all vendors – EMC, HDS IBM, SUN/STK
 - z/OS storage hardware replication feature Clone/Snap, FlashCopy, ShadowImage, SnapShot
- Storage System Independent
 - EMC Symmetrix/DMX
 - Hitachi USPv/Tagmstore/USP,9900V
 - IBM 2105 SHARK/2107
- System Feature Independent
 - Raid-S, Raid 10, Raid 5
 - PPRC, SRDF (ConGroup), TrueCopy, HUR
 - HyperSwap/AutoSwap
 - HiperPAV, PAV, non-PAV
 - Midaw
 - Cache Fast Write

EMC²
where information lives™

HITACHI
DATA SYSTEMS

IBM®


STORAGETEK™

FDRINSTANT Hardware Feature Support

Disk Vendor Hardware Independent...FDRINSTANT enables near instant copy of z/OS data and is the fastest and most reliable solution available for full volume and incremental non-disruptive backup of point in time copies of active volumes while applications continue updating them in real time.

FDRINSTANT recognizes, coexists and works in conjunction with PPRC, TrueCopy, Hitachi Universal Replicator (HUR), SRDF, Cache Fast Write, MIDAW and PAV (Parallel Access Volumes) on EMC Symmetrix/DMX, Hitachi USPv/Tagmstore/ USP, 9900V, IBM 2105 SHARK / 2107z/OS disk storage systems.

Employing FDRInstant (and FastMOVE) to replicate volumes and datasets to a target that is a PPRC/SRDF primary in an IBM or EMC system will result in the use of FlashCopy or TF SNAP hardware replication.


In a Hitachi storage system, FDRInstant (and FastMOVE) will revert to conventional IO if the target volume is a TrueCopy/Hitachi Universal Replicator primary without mCode that provides a "Native FCTOPPRC Primary" facility, e.g. the same support as IBM standard FlashCopy.

Performance Improvement for ABRInstant Full Volume and Incremental Operations




- Objective
 - Reduce the impact on the nightly production schedule of a job that must perform full volume and incremental backups of a mix of 654 (3390-3/9/27) volumes.
- Challenge
 - No job structure or job JCL changes.
- Solution
 - Employ FDR V5.4 L60+
- Results
 - Reduction in replication time from ~14 minutes to <2 minutes


	Prior	FDR V5.4 L60
FDRInstant Replicate 654 volumes	~ 14 min	<2 min



- **Objective** – Reduce the impact on the nightly production schedule of a job that must perform full volume and incremental backups of a mix of 654 (3390-3/9/27) volumes.
- **Challenge** – No changes to the job structure or job JCL.
- **Solution** – Employ FDR V5.4 L60 ABRInstant backup.
- **Results** – FDR V5.4 L60 using ABRInstant consistent, differential (and pre-copy replication) reduced the impact on the customer’s nightly production schedule of a job doing a mix of full volume and incremental backups of over 650 (3390-3/9/27) volumes, reducing D2D replication run time from ~14 minutes to <2 minutes.




Performance Improvement for FDRInstant Data Set Operations



SHARE
Technology • Connections • Results


- **Objective**
 - Improve the performance of four jobs that run concurrently to copy ~350 data sets, >111 million tracks.
- **Challenge**
 - No job structure or job JCL changes.
- **Solution**
 - Four FDR V5.4 L60+ jobs using parallel data set tasks to perform 20 concurrent copy operations. (MAXTASKS=5).
- **Results**
 - 70% time reduction across four jobs.

	Prior	FDR V5.4 L60 FDRInstant parallel data set copy MAXTASKS=5
FDRInstant four concurrent jobs	26.1 min	7.8 min
5 data set operations per 350 Data Sets	11.7 min	3.5 min
111 million tracks	18.4 min	5.5 min
	17.5 min	5.3 min



- **Objective** – Improve the performance of four jobs that run concurrently to copy ~300 data sets, totaling 111 million tracks.
- **Challenge** – No changes to the job structure or job JCL
- **Solution** – FDR V5.4 L60 employ FDRInstant data set operations (FDRCopy) using multiple parallel data set tasks to perform 20 concurrent copy operations. (MAXTASKS=5).
- **Results** – 70% time reduction in copy time across four jobs.


ABRInstant Copying 300 Data Set (111 million tracks)



Customer Experience: Move 25,000 DB2 Files at the rate of 3,200/minute



- **Objective**
 - Move 25,000 DB2 files to new volumes
- **Task**
 - Allocate ~ Catalog ~ Delete ~ 25,602 DB2 files
 - Copy (~ 1.4 million tracks)
- **Challenge**
 - Move the DB2 files with minimal impact on production systems
- **Solution**
 - FDRInstant technology completes the move in 8 minutes, at 3,200 datasets a minute.



- **Objective** – Move 25,602 DB2 files to new volumes.
- **Task** –
 - Allocate~25,602 DB2 data sets
 - Copy (~1.4 million tracks)
 - Catalog~25,000 DB2 data sets
 - Delete~25,000 B2 data sets
- **Challenge** – Move the DB2 files with minimal impact on production systems.
- **Solution** – FDRInstant technology completes the move of 25,602 DB2 files in 8 minutes, moving data sets at the rate of 3,200/minute (8.5 GB/min).

Executive Summary



- FDR V5.4 L60+ is the fastest z/OS business continuance solution available today!
- INNOVATION works closely with EMC, Hitachi, IBM, SUN/STK and other storage technology vendors to insure FDR data protection, business continuance and business resiliency solutions for z/OS resident data are simple, secure and fast.
- FDRInstant is z/OS software technology that leverages hardware replication technology for non-disruptive backup and rapid recovery.
- FDRInstant (V5.4 L60+) brings significant performance improvements to TimeFinder CLONE/SNAP, FlashCopy, ShadowImage and SnapShot.
- If you are not using INNOVATION FDRInstant, to improve your business resiliency, you should seriously consider joining the hundreds who are automating, simplifying and reducing their data protection and business continuance efforts.

What..., Why..., Who..., When..., Where...Executive Summary

Your staff is under pressure to do more, in less time, with fewer resources.

Companies today see continual growth pushing up the time it takes to complete their backups, while a new sense of customer service, legislative oversight and regulatory penalties combine to demand near continual 24 x 7 data availability.

You do not have the time nor do you want to try solving these problems piecemeal with individual products that you have no idea will work together. Understanding the squeeze this phenomenon puts on its z/OS customers, INNOVATION is working to deliver pre-tested new releases of its FDR Suite that will seamlessly employ technology advances to provide significant reductions in the time associated with data protection, business continuance and business resiliency, i.e. non-disruptive backup, data migration, storage consolidation, data set copy, move and rapid recovery...“Customers can expect to see hardware-based data replication and disk-to-disk backup time cut in half.”

**FDRInstant, ABRInstant & FDRMOVE with FDRPAS
Broader Support...Faster and Faster!**


Non-Disruptive FDRInstant Solutions for
Data Protection, Business Continuance
and Business Resiliency!

 **INNOVATION[®]**
DATA PROCESSING



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



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