



IBM TotalStorage™

Enterprise Storage Server

Charlie Burger
San Jose Advanced Technical Support
cwburger@us.ibm.com

© 2004 IBM Corporation

IBM TotalStorage®



Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries or both.

Enterprise Storage Server, FlashCopy, GDPS, IBM TotalStorage, z/OS, zSeries

Other company, product, or service names may be trademarks or service marks of others.

Agenda

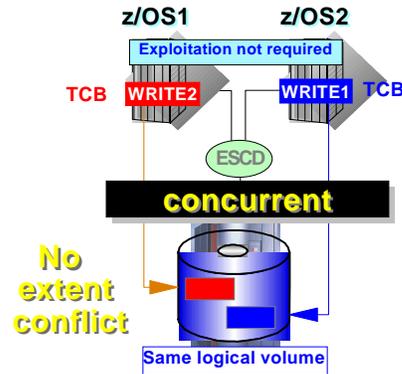
- Enterprise Storage Server Overview
- PPRC
 - Global Mirror for ESS
- FlashCopy V1
- FlashCopy V2
 - Multiple Relationship FlashCopy
 - Incremental FlashCopy
 - FlashCopy Consistency Group support
 - Inband support for FlashCopy
 - FlashCopy Across LSS/Across Cluster
 - FlashCopy establish time reduction
 - Data Set FlashCopy
- FlashCopy V2 Performance
- ESS 2.4.0 FlashCopy Enhancements

Enterprise Storage Server

- Multiple Allegiance
- Parallel Access Volumes (PAV)
- I/O Priority Queuing
- Advanced Copy Services

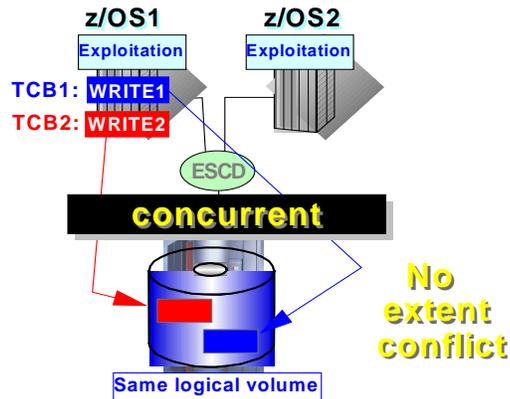
Multiple Allegiance

- Simultaneous access to same volume from multiple systems
- Multiple reads executed in parallel
- Multiple writes executed in parallel, serialized by extent specification in Define Extent command
- No specific software dependencies
- Sharply reduces performance impact of sharing volumes (PEND) between systems



Parallel Access Volumes (PAV)

- Multiple device addresses for the same volume
 - One base address, any number of alias addresses
 - New Unit Control Block (UCB types)
 - ✓ PAV base and PAV alias
- Defined in ESS configuration and zSeries HCD / IODF
- Multiple reads executed in parallel
- Multiple writes executed in parallel, serialized by extent specification in Define Extent command
- Sharply reduces volume contention (IOSQ) within a system
- Parallel Access Volumes (PAVs) managed dynamically by Workload Manager operating in Goal Mode
 - WLM dynamically reassigns alias addresses within LSS



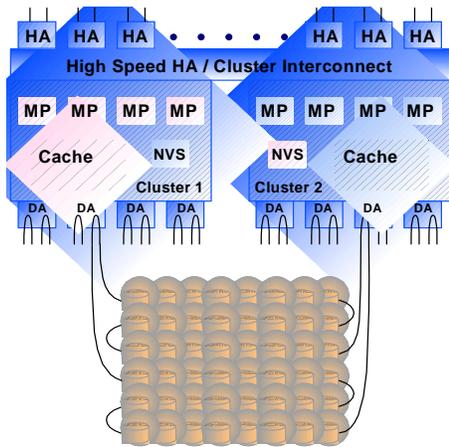
Static and Dynamic PAVs

- **Static PAVs**
 - Association between PAV-Base and its PAV-Alias is predefined and fixed
- **Dynamic PAVs**
 - Association between a PAV-Base and its PAV-Aliases in the dynamic pool
 - The WLM in Goal Mode manages the assignment of alias addresses
 - WLM instructs IOS when to reassign an Alias
- **Do not mix PAV types**

I/O Priority Queuing

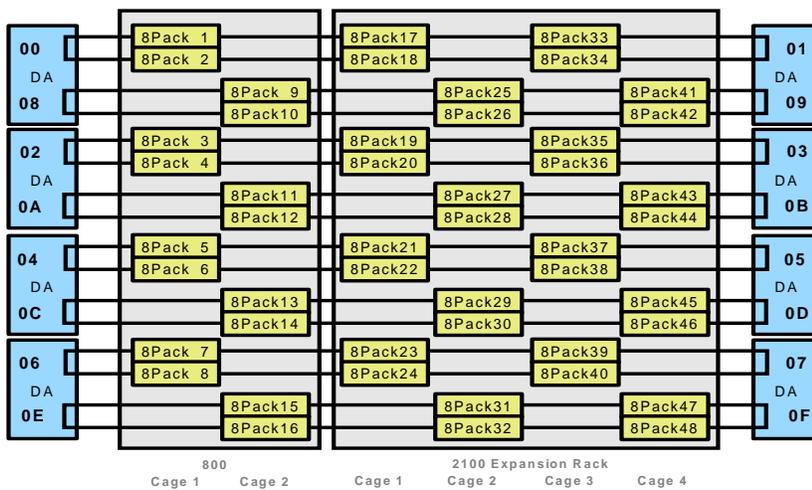
- **If I/Os can not run in parallel, they are queued**
 - ESS can internally queue I/Os
 - Avoids overhead associated with posting "device busy" status and redriving the channel program
 - Eliminates race conditions when one system is faster than the other
- **Priority queuing**
 - I/Os can be queued in a priority order
 - A priority byte can be set in the CCW to prioritize I/Os
 - OS/390's Workload Manager sets the I/O Priority Byte when running in Goal Mode
 - I/O priority for systems in a Sysplex
 - Fair share for each system

ESS Components



- 2 - 16 Host Adapters
 - ESCON or Ultra SCSI
 - 2 ports each (Total 32)
- Fibre Channel
 - 1 port (Total 16)
 - FCP or FICON
- 2 Clusters, each with
 - 4-way symmetric multiprocessors
 - Max 64 GB
 - 2 GB NVS
 - 4 Device Adapter pairs
 - Support up to 8 loops
 - Attach up to 384 DDMs
- 2 - 48 Disk Arrays
 - 8 DDMs /disk array
 - Configured as RAID-5 (or JBOD)
 - Enterprise Storage Server '8-Pack'
 - Attach up to 6 x '8-Pack' per loop

ESS 800 – Expansion Rack Loop Configuration



Advanced Copy Services

- PPRC
 - Metro Mirror (Synchronous PPRC) for ESS
 - Global Mirror (Asynchronous PPRC) for ESS
 - Metro/Global Copy (Asynchronous Cascading PPRC) for ESS
 - Global Copy (PPRC Extended Distance) for ESS
- XRC
 - Global Mirror (XRC) in zSeries environments
 - Metro/Global Mirror (three-site solution using Synchronous PPRC and XRC) in zSeries environments
- FlashCopy
 - FlashCopy full volume
 - Persistent FlashCopy
 - Incremental FlashCopy
 - Consistency Group FlashCopy
 - Data Set FlashCopy
 - FlashCopy Across LSS/Across Cluster
 - **Reverse restore FlashCopy**®

PPRC

Background

Dependent writes

- The start of one write operation is dependent upon the completion of a previous write to a disk in either the same subsystem frame or a different subsystem frame
- Basis for providing consistent data for copy operations

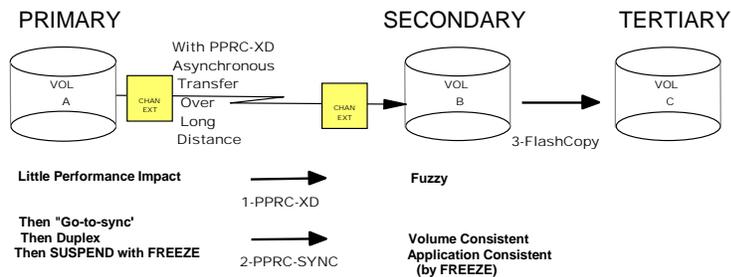
Consistency

- Preserves the order of dependent writes
 - For databases, consistent data provides the capability to perform a data base restart rather than a data base recovery
 - ✓ Restart can be measured in minutes while recovery could be hours or even days

Asynchronous processing

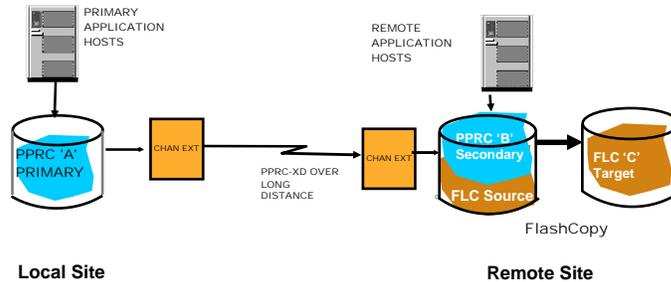
- The separation of data transmission from the signaling of I/O complete
 - Distance between primary and secondary has little impact upon the response time of the primary volume
 - Helps minimize impact to application performance
 - Examples are PPRC Extended Distance or XRC

Background – PPRC Extended Distance



- PPRC-XD (OPEN and z/OS® data)
 - Asynchronous transfer of updates over long distance through channel extenders with little performance impact on applications
 - Creates a fuzzy secondary copy
 - Consistency of data at remote not maintained
- Needs to "Go-to-Sync" (PPRC-SYNC) for a while to build consistency on secondary volume
 - After achieving sync, FREEZE to get application level consistency
 - FlashCopy onto tertiary to save consistent checkpoint

PPRC Global Mirror Overview

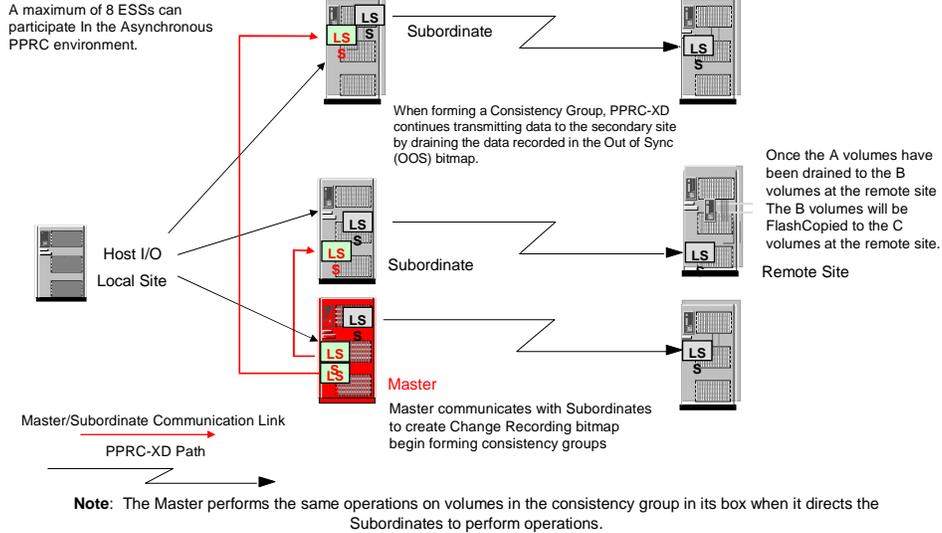


- PPRC Extended Distance from local to remote site
- Automatic periodic consistent copies created at remote site
 - Volume A uses change recording bitmap while consistency group is created
 - All data sent to remote and inband FlashCopy at remote creates consistent data on volume C
- Designed to adjust to peak workloads without impacting primary application

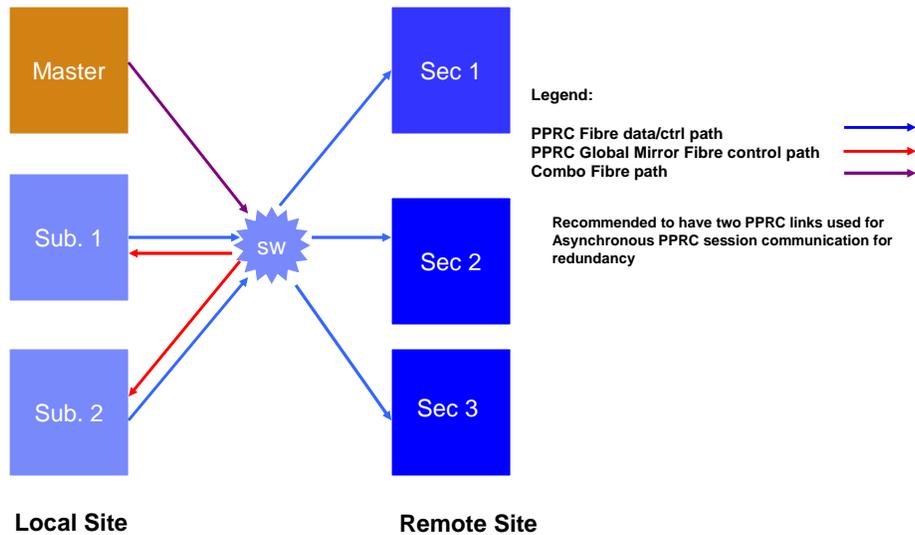
PPRC Global Mirror Capabilities

- Two site long distance disaster recovery and backup remote copy solution for open systems and zSeries environments
 - High performance, long distance remote copy solution using asynchronous technology
 - Updates made on the primary ESS (local) are asynchronously mirrored to a secondary ESS (remote)
 - Designed to automatically maintain a consistent copy of the data at the remote site, with currency of the remote site for many environments lagging an average of 3 to 5 seconds
 - Actual lag time depends upon a number of factors such as workload and bandwidth between the local and remote sites
 - Consistency can be managed across volumes residing within multiple ESSs
 - Support for virtually unlimited distance between the local and remote sites with little impact to host applications at the local site
 - Relieves 103 KM for ESCON PPRC links and 300KM for Fibre PPRC links in Synchronous PPRC environments
 - Limited only by the capabilities of the network and channel extension technologies
- A consistent and restartable copy of the data at the remote site
- Can exploit PPRC Failover/Failback
 - Helps reduce the time required to synchronize PPRC volumes after switching between sites during planned or unplanned outages by having Failover make the secondary a PPRC suspended primary (bit map) and with Failback only sending updates that have been made at the recovery site back to the production site when you are ready to return production to the local site
- Eliminates the requirement to do a manual and periodic SUSPEND at the local site and a manual FlashCopy at the remote site to preserve a consistent copy of the data at the remote site as is required with PPRC Extended Distance or 2 site Asynchronous Cascading PPRC

Typical PPRC Global Mirror Configuration



PPRC Global Mirror Network in a SAN Environment

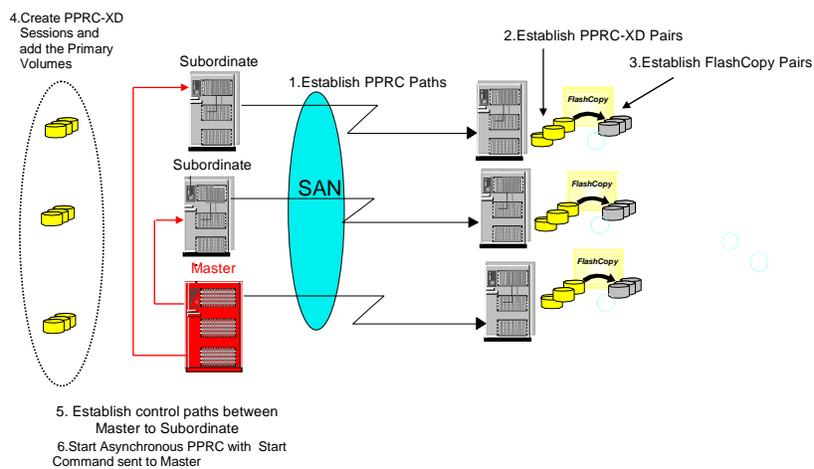


PPRC Global Mirror Network Prerequisites

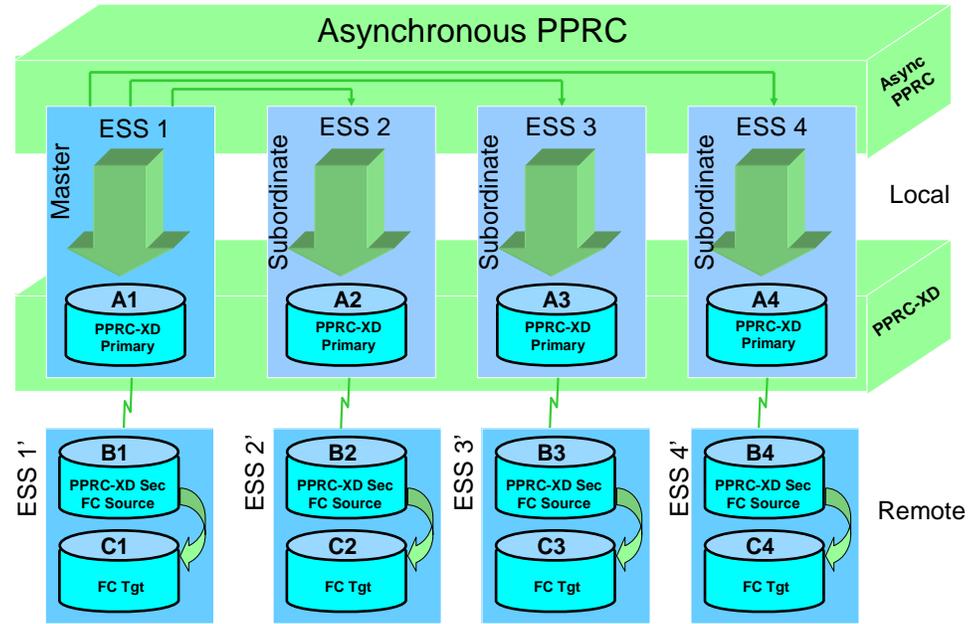
- An PPRC Global Mirror control network must be configured and specified by the Administrator using ESS Web Copy Services interface
- Master/Subordinate ESS communications occur over PPRC FCP links across LSS to LSS associations
 - PPRC paths must exist between Master and Subordinate LSSs before forming consistency groups
 - Must have one Master LSS PPRC path to each Subordinate LSS participating in PPRC Global Mirror environment

All Master management commands must be issued to the same LSS on the Master

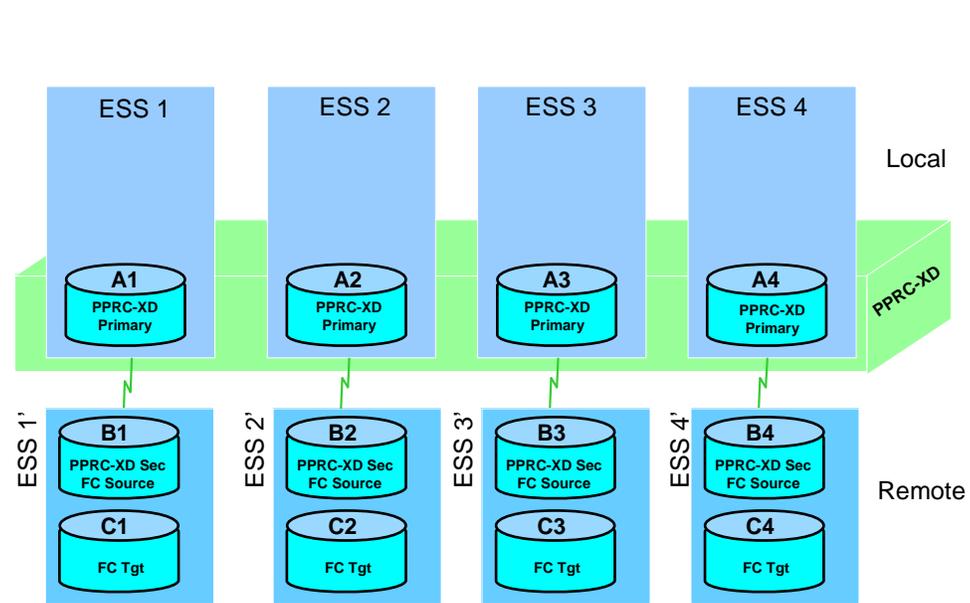
PPRC Global Mirror Initialization Process

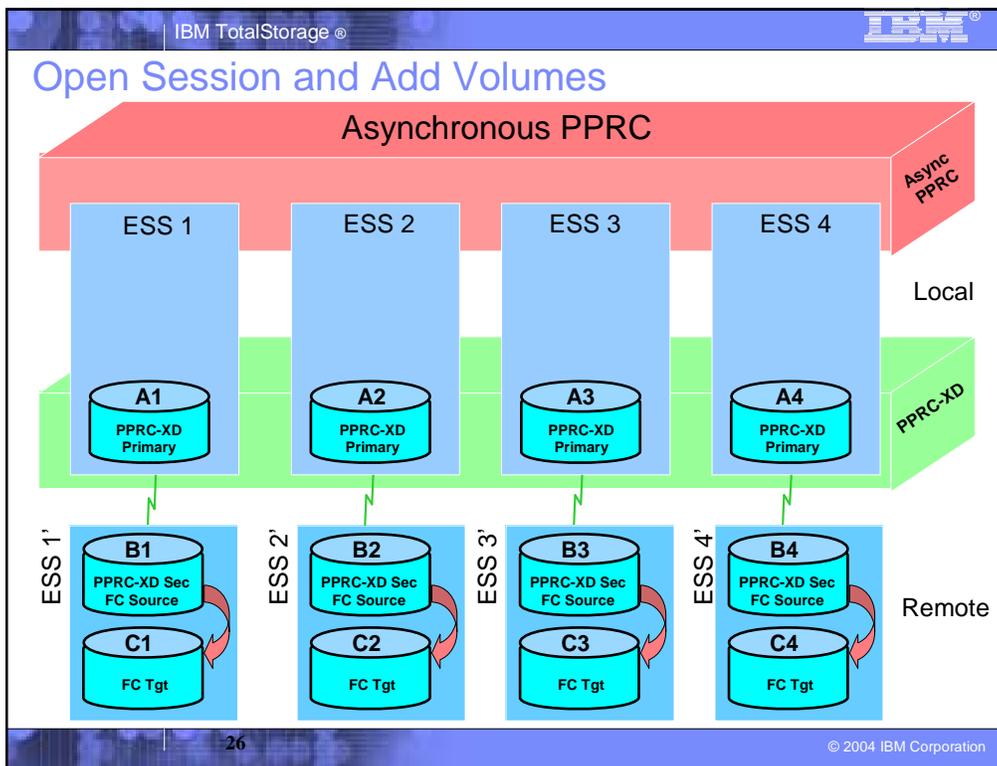
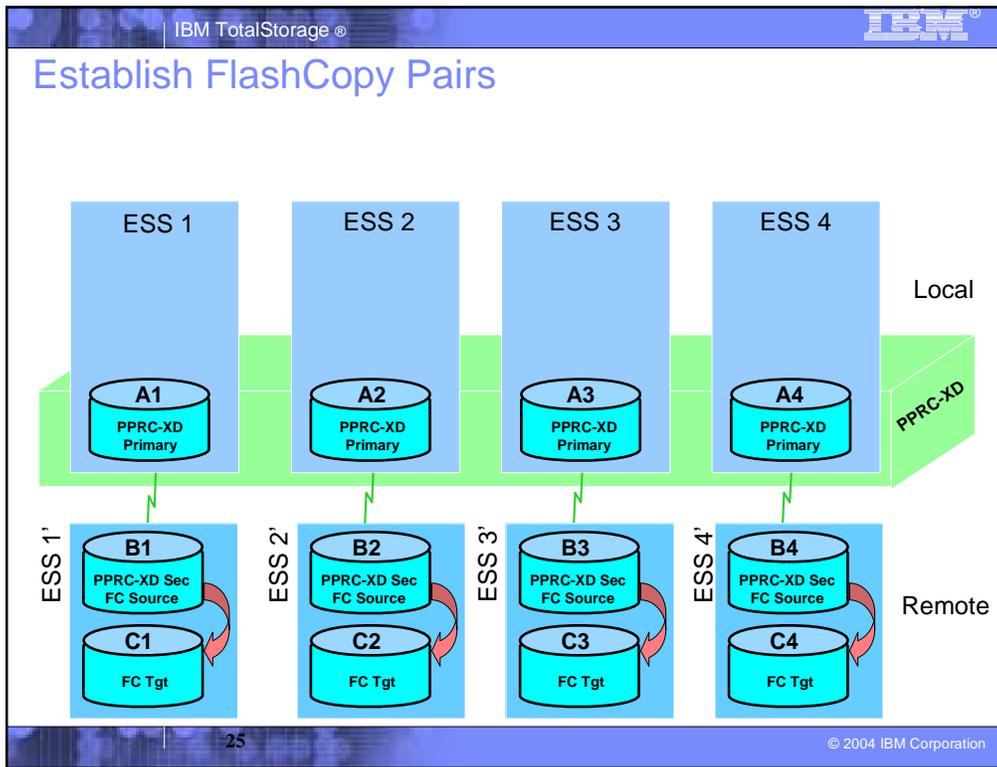


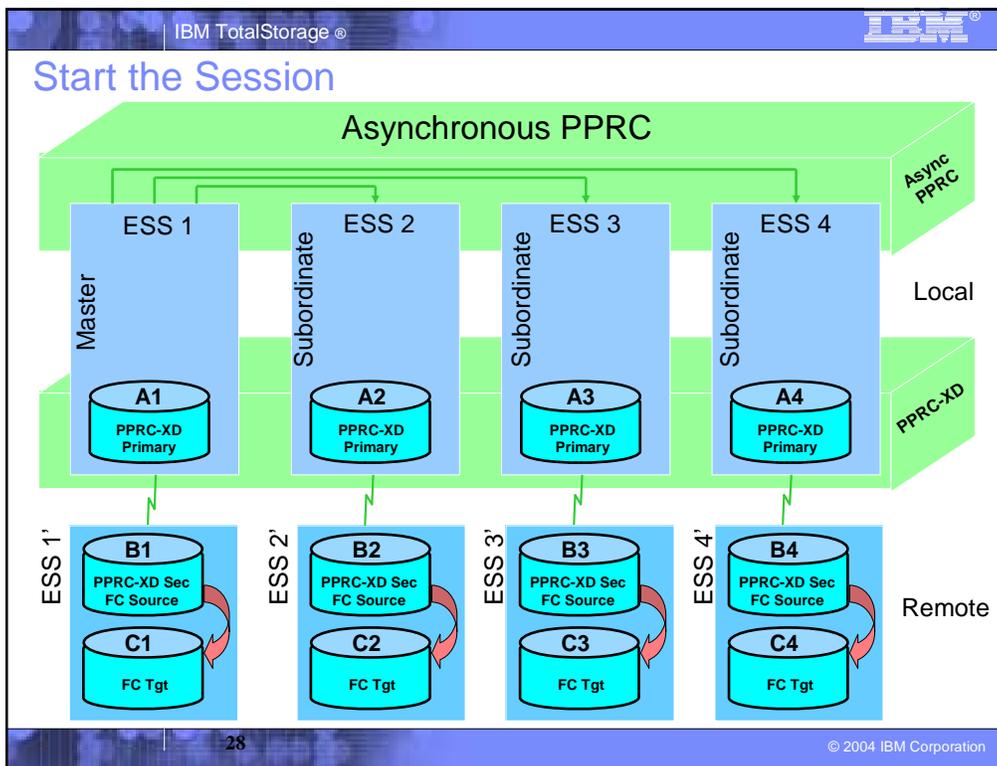
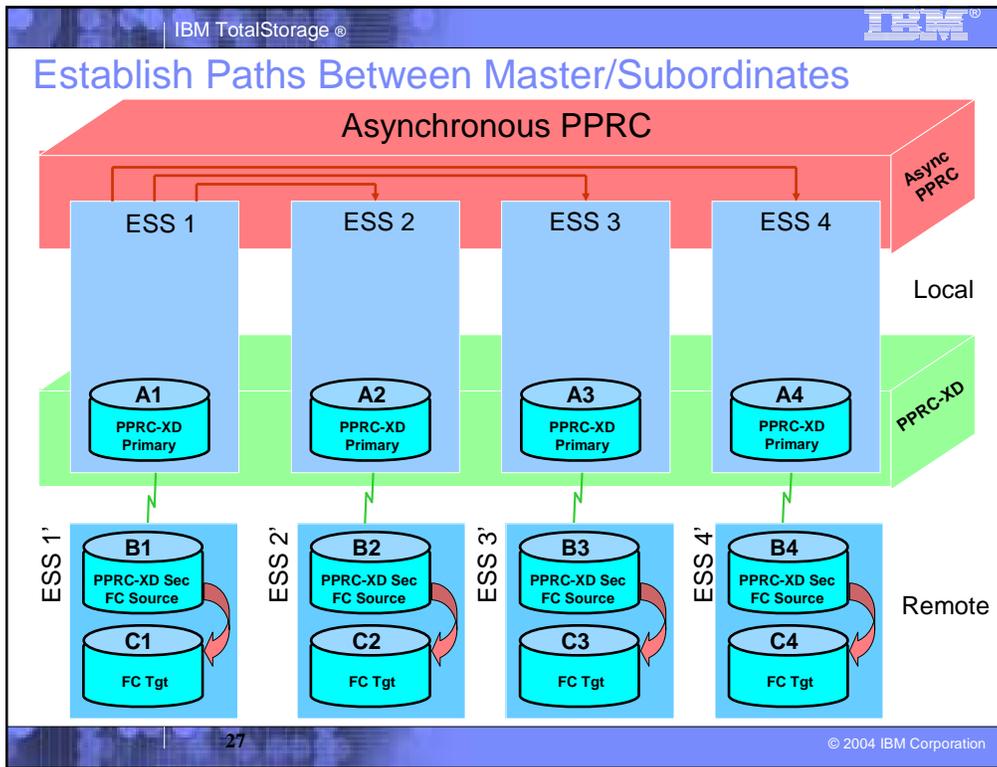
Global Mirror for ESS



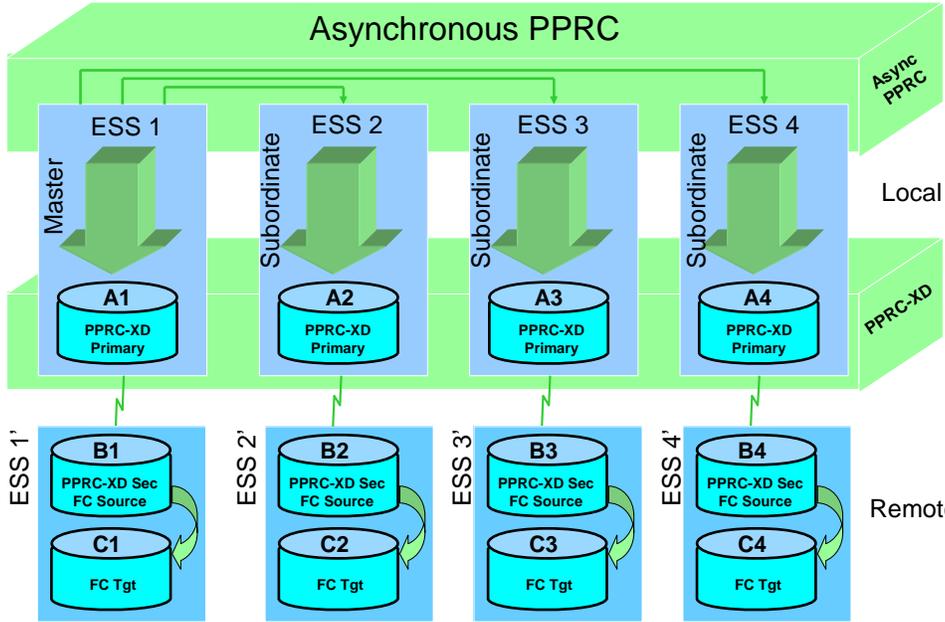
Establish PPRC Paths and Pairs with XD Option



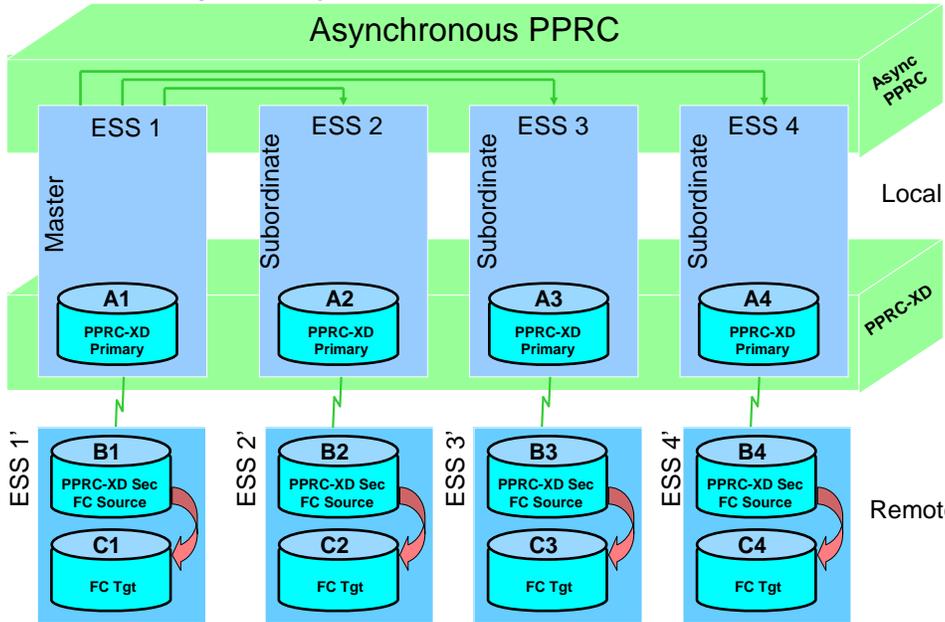




Form Consistency Groups



Consistency Group Formed



PPRC Global Mirror Setup Procedure

- Identify all volumes that will participate in the session
 - PPRC primary and secondary volumes
 - FlashCopy target volumes for the PPRC secondary volumes
 - ESS that will be Master
- Establish PPRC paths local to remote
- Establish PPRC pairs local to remote (A->B) with extended distance option
- Establish FlashCopy pairs (B->C) with the following options:
 - No background copy
 - Persistent FlashCopy
 - Start change recording
 - Inhibit target writes
- Open Asynchronous PPRC Session
 - See 'Opening or creating an Asynchronous PPRC session on an LSS'
- Add PPRC primary volumes (A volumes) to the Asynchronous PPRC session
 - See 'Adding or removing volumes on an Asynchronous PPRC session'
- Establish PPRC paths from Master to Subordinates
 - See 'Adding or removing volumes on an Asynchronous PPRC session'
- Start Asynchronous PPRC with the Start Command
 - See 'Starting or resuming an Asynchronous PPRC session on an LSS'

PPRC Global Mirror Session Management

A session is a collection of volumes in an ESS or in multiple ESSs that are managed together when creating consistent copies of data. When you start or resume a session, the PPRC Global Mirror configuration is specified or modified. This initiates PPRC Global Mirror to form a consistency group for each volume in the session ID.

- Master/Subordinate relationship topology defines a session
- An LSS can be assigned only to one session
 - Architecture allows for growth
- One Master/Subordinate relationship allowed within a single ESS
 - Architecture allows for growth
- Only one Asynchronous Master is allowed per ESS
- Maximum of 3 subordinates with LIC 2.4.0
 - An individual session ID consists of a Master and up to three primary ESS subordinates
 - The Master may have PPRC primaries but having primaries is not a prerequisite
 - Architecture allows for growth
- Maximum of 8 ESSs participating in Asynchronous PPRC environment
- PPRC Global Mirror mechanism will only act upon volumes which have been added to a session, other PPRC volumes will still be in PPRC-XD mode not included in the formation of consistency groups

PPRC Global Mirror Session Management Commands

- Manage Session
 - Open/Close
- Managing the formation of consistency groups with PPRC Global Mirror
 - Start/Resume
 - Pause/Terminate
- Manage Session Members
 - Add/Remove
- Query
 - Queries operation of the Asynchronous PPRC mechanism
 - Query the LSS in the ESS where the Start session command was issued
 - Query status of LSS
 - Query status of disk

PPRC Global Mirror Session Management – Open/Close

- Open or Close an Asynchronous PPRC session
- Select an LSS from the Logical Subsystems panel
 - Open
 - Opens an Asynchronous PPRC session for an LSS
 - ✓ A session is made available for the addition of volumes
 - Open a session for each LSS that you want to participate in Asynchronous PPRC
 - Close
 - Close a session when you no longer want the volumes in a LSS to participate in Asynchronous PPRC

PPRC Global Mirror Session Management – Add/Remove

- Add or Remove volumes from an Asynchronous PPRC session after PPRC-XD volume pairs have already been established
- Select a volume from the Volumes panel
 - Add
 - Add a volume to an existing session
 - Remove
 - Remove a volume from an existing session
- Select Manage Session Members
 - Select Add or Remove and session id

PPRC Global Mirror Session Management – Start/Resume

- Start or Resume a session
- Select an LSS from the Logical Subsystems panel
 - Start
 - Starts Asynchronous PPRC mechanism
 - ✓ Open a session for each LSS that will be participating in the Asynchronous PPRC environment prior to doing the Start
 - ✓ Add volumes prior to the Start
 - First Start command makes the selected LSS/ESS the Asynchronous PPRC Master
 - ✓ Can be performed on any device or LSS in the Master ESS
 - Resume
 - Resumes Asynchronous PPRC mechanism after a Pause command had been issued
 - Once the initial Asynchronous PPRC configuration is created and for as long as it exists, you can start or resume a session for the Master only on the same LSS where the initial support was performed
- Select Timer values
 - Modify or Do not modify
 - Consistency Group Timer (Seconds)
 - Maximum Coordination Timer (Milliseconds)
 - Maximum Consistency Group Drain Time (Minutes)

PPRC Global Mirror Session Management – Terminate/Pause

- Terminate or Pause an Asynchronous PPRC mechanism
 - Both commands stop the Asynchronous PPRC mechanism
 - Pause will not interrupt the formation of a consistency group before stopping the Asynchronous PPRC mechanism
- Select the LSS in the ESS where the Start command was issued from the Logical Subsystems panel
 - Terminate
 - Use when session topology is about to be changed (ex. Defining a new Master, adding or removing a subordinate)
 - Terminates Asynchronous PPRC mechanism
 - Removes Master/Subordinate relationship information
 - Once terminated, a new start command is required
 - Pause
 - Pauses Asynchronous PPRC mechanism

PPRC Global Mirror User Settings

Three simple parameters!

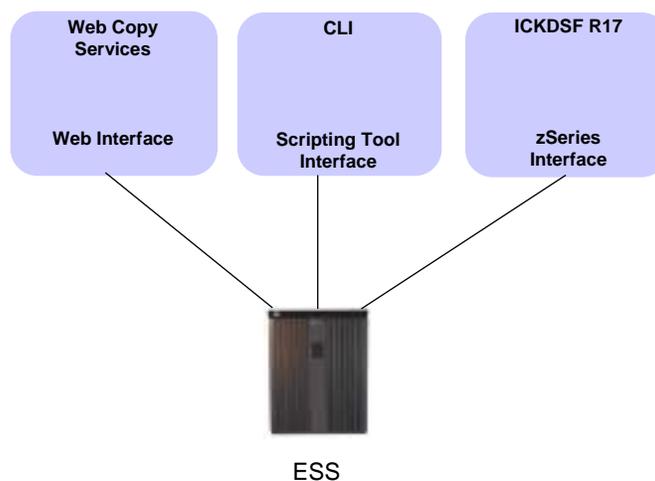
1. Specify Consistency Group Interval Timer (seconds)
2. Specify Maximum Coordination Timer (milliseconds)
3. Specify Maximum Consistency Group Drain Time (minutes)

**Defaults are available but
can be changed**

PPRC Global Mirror Monitoring

- Periodically issue SNMP alerts to host for failure to form Consistency Groups
 - Query available to monitor Consistency Group formation success
 - In MOST cases Asynchronous PPRC will continue to try and form Consistency Groups
 - ✓ Exception is if continued attempts will compromise the only existing consistent copy
 - If Consistency Groups cannot be formed within specified time (no error, but taking too long), they will be formed after some number of attempts
- Administrator may choose to add bandwidth, increase Consistency Group formation intervals, fix failing hardware, or remove failing components from the Asynchronous PPRC session to allow the successful formation of the Consistency Groups

Operational Management



ESS Remote Copy Positioning

- Designed for Recovery Point Objective of near 0
 - Data base restart will back out in-flight transactions
 - Synchronous PPRC
 - Two site, metro distance configuration for open systems and zSeries environments
 - Operational management via MDM Replication Manager and GDPS®, ICKDSF, Web Copy Services, CLI, TSO, or API
 - Asynchronous Cascading PPRC
 - Three site, long distance configuration for open systems and zSeries environments
 - Operational management via eRCMF, ICKDSF, Web Copy Services, or CLI
- Designed for Recovery Point Objective > 0
 - **PPRC Global Mirror**
 - **Two site configuration for open systems, zSeries, and mixed environments**
Operational management provided at this time via ICKDSF, Web Copy Services, CLI, and Multiple Device Manager Replication Manager targeted for 3Q/2004
 - XRC
 - Two site configuration for z/OS or zSeries Linux environments
 - Operational management via GDPS, TSO, or API
 - PPRC Extended Distance
 - Two site configuration suitable for data migration or periodic offsite point in time consistency
 - Operational management via ICKDSF, Web Copy Services, CLI, TSO, or API

PPRC Global Mirror Prerequisites

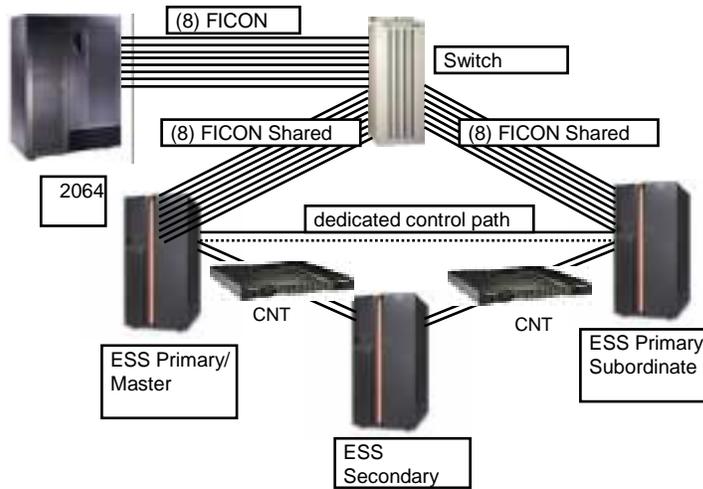
- ESS model 800 or 750 for the primary AND secondary subsystems
- PPRC V2 for primary and secondary ESSs
- Fibre PPRC links
- ESS LIC 2.4.0
- New ESS machines must have Feature code #9912 'PPRC Global Mirror Indicator'
 - Existing ESS machines do not require the new feature code
- FlashCopy V2 required at remote site
 - Any site with PPRC V2 that uses FlashCopy requires FlashCopy V2

PPRC Global Mirror Summary

- Enhanced business continuance solutions with PPRC GLOBAL Mirror
 - Two site remote copy at long distances with data consistency every 'Consistency Group Interval Time'
 - Helps minimize impact upon host application performance
 - Reduced complexity in managing remote copies

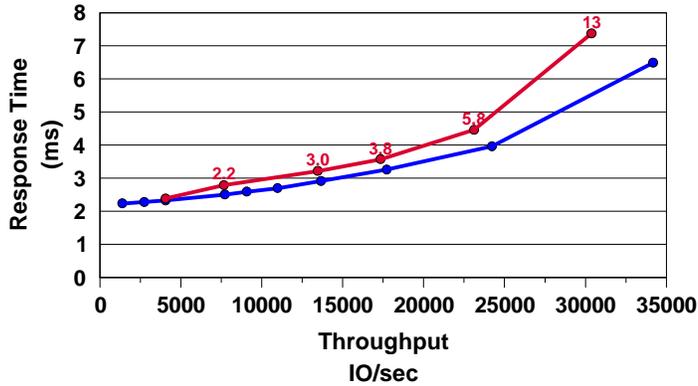
Global Mirror for ESS Performance

Asynchronous PPRC Test Configuration - z/OS



Asynchronous PPRC - z/OS (2) Primary, (1) Secondary ESS 800 and 2 FCP Paths

z/OS Cache Standard Workload



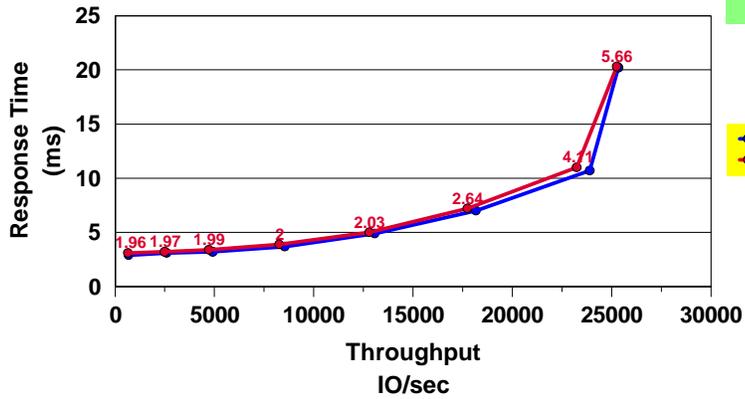
Workload has a 3:1 r/w ratio with 4K average transfer size

CGI Time/RPO in seconds is indicated on chart in RED

Asynchronous PPRC - Open (1) Primary, (1) Secondary ESS 800 and (2) FCP Paths

Open 70/30/50 OLTP Workload

Workload has a 2.33:1 r/w ratio with 4K average transfer size



CGI Time / RPO in seconds is indicated on chart in RED

Global Mirror for ESS Scenarios

Failover/Failback

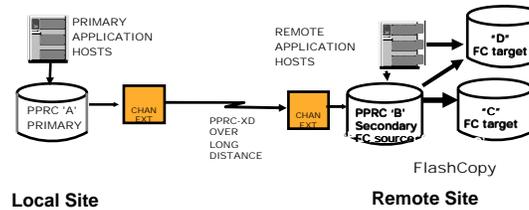
- Failover
 - The secondary volume to which the command was issued becomes a suspended PPRC primary
 - Volume has a bitmap
 - Establishes a new relationship between the volume the command was issued to and its PPRC primary volume
 - No communication occurs between the two volumes
- Failback
 - The volume to which this command is issued has its PPRC pair converted, if necessary, to a PPRC secondary
 - The volume to which the command was issued keeps track of changes and then begins to resync to its PPRC partner which becomes a PPRC secondary volume
 - Data is transferred from the volume to its secondary volume

PPRC Global Mirror Planned Outage -(Example 1)

- Pause session
- Freeze/run PPRC-XD between A>B
- Perform:
 - Hardware maintenance at remote site
- Reestablish paths A>B
- Reestablish PPRC-XD pairs A>B with resync
- Resume session

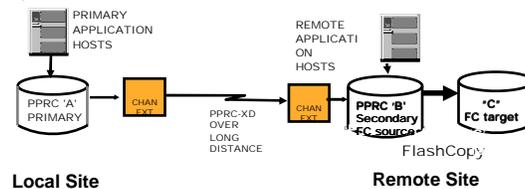
PPRC Global Mirror Planned Outage -(Example 2)

- Pause session
- Freeze/run PPRC-XD between A>B
- Run applications on A volumes
- FAILOVER issued to B>A
- Fast Reverse Restore (FRR) FlashCopy C>B
- Wait for background copy to complete
- FlashCopy B>C with the same options used during the initial Asynchronous PPRC setup
- FlashCopy B>D
- Perform on D volumes:
 - Disaster recovery testing
 - Application testing
 - Application development
- Reestablish paths A>B
- FAILBACK issued to A>B
- Resume session



PPRC Global Mirror Unplanned Outage - Recovery

- Terminate the session (if the Master is still available)
- Issue FAILOVER to PPRC B volumes to establish B>A
 - Volumes become SUSPENDED primaries
- Check FlashCopy status
 - Revert or commit current FlashCopy relationships depending upon status
- Fast Reverse Restore FlashCopy C>B
 - After command is issued the C volumes are not useable
- Wait for background copy to complete
- FlashCopy B>C with the same options used during the initial Asynchronous PPRC setup
- Recovery activity can now be resumed on the B volumes



PPRC Global Mirror Unplanned Outage - Returning

- Establish paths from remote site back to the local site
 - This step can be avoided if paths had originally been defined going both directions (FCP links are bidirectional)
- FAILBACK with XD to the B volumes
- Once XD is in a steady state the host activity should be quiesced
 - Data on the B volumes will drain to the A volumes at the local site
- If you wish, you could delete the paths going from B>A and reestablish the paths going from A>B
 - This step can be avoided if paths had originally been defined going both directions (FCP links are bidirectional)
- FAILOVER to A with XD to establish A>B
- FAILBACK to A volumes with XD to force B volumes to become secondaries
- Establish Asynchronous PPRC paths from the Master to the Subordinates
- Start the Asynchronous PPRC session
- Resume host activity on the A volumes

Revert or Commit? (Master/Subordinates Available)

- Determine the increment status
 - Command issued to local volumes

Indication	Specification on FlashCopy Withdraw
Revertible or Not Revertible and Equal	Commit
Non-Revertible and Equal	No Withdraw required
Revertible and Not Revertible	Revert
All Revertible	Revert
Consistency Group Corrupted	Manual recovery required

Revert or Commit? (Master/Subordinates Not Available)

- Determine the increment status
 - Query relations issued to each PPRC secondary

Revertible	Sequence #	Specification on FlashCopy Withdraw
All	Equal	Revert to all FlashCopy source volumes
At least one	Equal	Commit to all FlashCopy source volumes
None	Equal	No Withdraw required
Mixed	Revertible equal and non-revertible equal but different from revertible	Revert to all FlashCopy source volumes

FlashCopy

Agenda

- FlashCopy V1
 - FlashCopy Review
 - FlashCopy Enhancements November, 2002
 - FlashCopy Nocopy to Copy
 - Persistent FlashCopy
- FlashCopy V2
 - Multiple Relationship FlashCopy
 - Incremental FlashCopy
 - FlashCopy Consistency Group support
 - Inband support for FlashCopy
 - FlashCopy Across LSS/Across Cluster
 - FlashCopy establish time reduction
 - Data Set FlashCopy
- FlashCopy V2 Performance
- ESS 2.4.0 FlashCopy Enhancements

FlashCopy V1

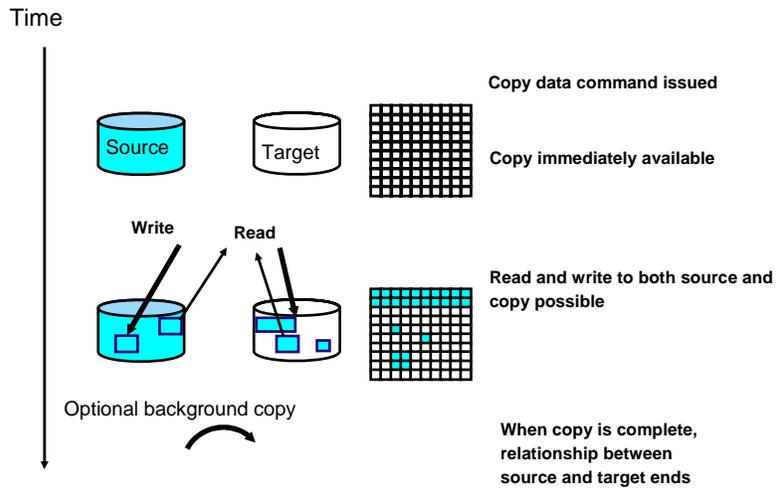
FlashCopy Overview

- "Instant" T0 (Time 0) copy
 - Source and target volumes immediately available for processing with full read/write access
- zSeries and Open
 - A hardware solution invoked by software
 - z/OS: DFSMSdss, TSO, or IBM ESS Web Copy Services
 - VSE/ESA: IBM ESS Web Copy Services or host commands for VSE/ESA
 - z/VM: IBM ESS Web Copy Services or host commands for z/VM
 - IBM ESS Web Copy Services for all other platforms
 - Copy Services CLI (Command Line Interface)
 - Check the Command-Line Interfaces User's Guide for servers that are supported to run CLI
 - Third party client application through CIM agent

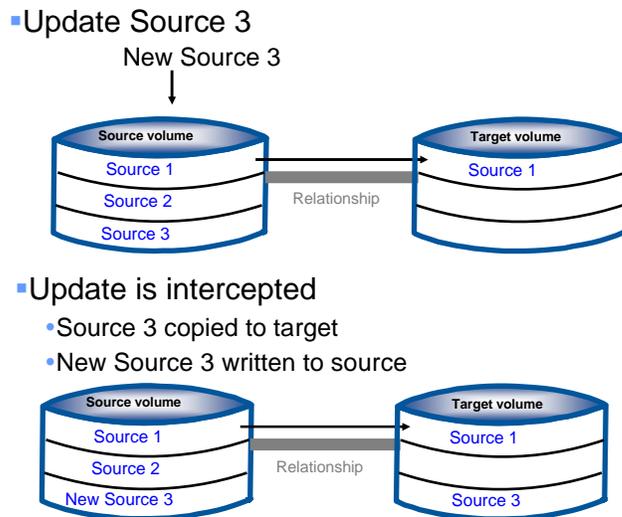
How does FlashCopy work?

- Request copy from source to target
- ESS creates FlashCopy relationship between volumes
- ESS starts to copy from Source to Target
 - May be disabled if not required ("NOCOPY" option)
 - Typically when target copied to backup device
- Attempts to read/write data already copied proceed as normal
- Attempts to process data not yet copied intercepted
- "Copy on Demand" or read from source as needed
- Creates effect of instant copy

FlashCopy Implementation

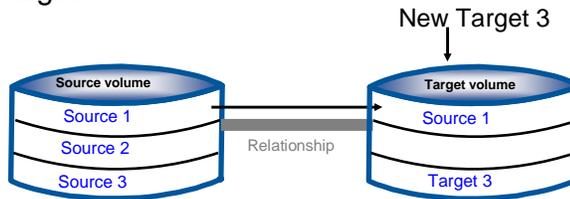


Update Source Example

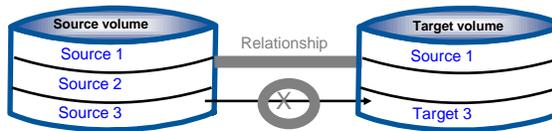


Update Target Example

- Update Target 3

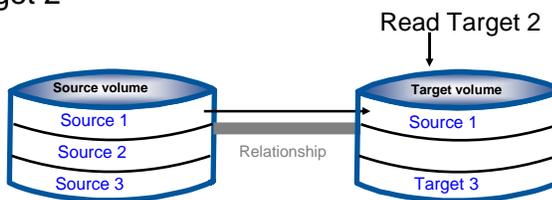


- Updated target is recognized
 - Source 3 is not copied to target

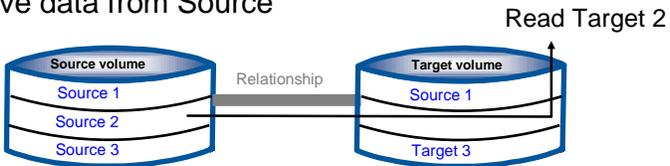


Read Target Example

- Read Target 2

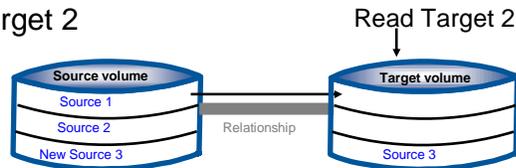


- Since Source 2 has not been copied to Target yet, retrieve data from Source

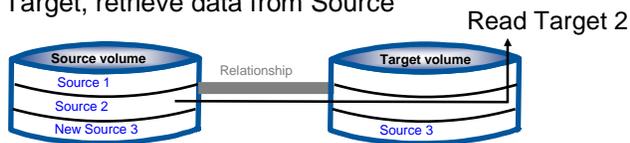


BACKGROUND NOCOPY

Read Target 2



- Since Source 2 has not been copied to Target because of the NOCOPY option, and corresponding track/sector has not been updated on Target, retrieve data from Source



- Above illustration also shows that Source 3 has been updated on the Source volume/disk with New Source 3 so Source 3 was copied to the Target volume/disk

FlashCopy V1 Requirements

- Source and target disks must be in the same logical subsystem (LSS)
- FlashCopy disks can only be in a single FlashCopy relationship at any time
- Source and Target disks require real disk space
- Source and target disks must be same format, target must be same size or larger than source
- FlashCopy operates at the disk level
- FlashCopy target can not be PPRC or XRC source

FlashCopy Considerations (1)

- FlashCopy only copies data that resides in the ESS
 - If there is modified data in processor memory, it will not be copied
 - Can result in data base logical errors when the FlashCopy target is referenced
- FlashCopy creates a track-by-track copy for CKD
 - How FlashCopy target is recognized by the same server or different server is operating system dependent
 - Duplicate volsers
 - VVDS volser mismatch
 - Catalog entries
 - etc

FlashCopy Considerations (2)

- If you have multivolume or striped data sets you must make sure that all volumes are copied
- Ensure that all updates are committed to disk and that all related data is at the same level of consistency
 - Quiesce applications, sync from operating system, stop database or shut down server
- Serialization is controlled through the application:
 - DFSMSdss will perform serialization on the volume or data set resource but will not flush buffers, etc.

DFSMSDss

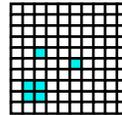
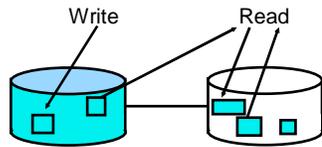
- DFSMSDss performs serialization
- DFSMSDss checks if source and target are eligible for FlashCopy
 - If not, DFSMSDss will do a normal copy
- Once the FlashCopy completes, the DFSMSDss job completes. It does not wait until the copy is physically complete
- DFSMSDss performs BACKGROUND(COPY) by default
 - APAR OW50446 provides NOCOPY and WITHDRAW support
 - BACKGROUND(NOCOPY) invoked with FCNOCOPY on COPY command
 - FCWITHDRAW specified on DUMP command will withdraw a relationship
- Once the BACKGROUND(COPY) is physically complete, the relationship between source and target is ended and you can set up a new FlashCopy for the source or the target
- You can specify COPYVOLID or leave it out
 - DFSMSDss dump conditioning APAR OW45674 allows this for SMS or non-SMS volumes

FlashCopy Enhancements - November, 2002

- FlashCopy NOCOPY to COPY
 - Can change background operation from NOCOPY to COPY
 - Establish FlashCopy relationship with NOCOPY
 - ✓ Possibly to minimize background operation impact to application I/O
 - Later change relationship to COPY
 - ✓ Possibly when application workload has gone down
 - This NOCOPY to COPY change will force background copy of remaining unchanged tracks
 - Invoked via:
 - ✓ Web Copy Services
 - ✓ CLI invoking a saved task
 - ✓ Third party client application through CIM agent
- FlashCopy Persistent Relationship
 - Help protect against inadvertent updates of recently created target volumes
 - Source volume is regularly copied to alternating target volumes
 - Persistent relationship will identify the target volume for the most recently completed FlashCopy
 - Relationship lasts until explicitly withdrawn
 - Invoked via:
 - ✓ Web Copy Services
 - ✓ CLI invoking a saved task
 - ✓ Third party client application through CIM agent

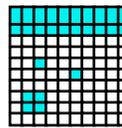
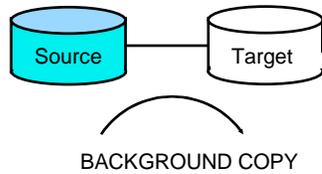
FlashCopy NOCOPY to COPY

Time



FlashCopy command issued with NOCOPY

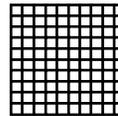
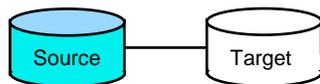
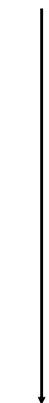
Change BACKGROUND from NOCOPY to COPY



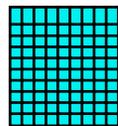
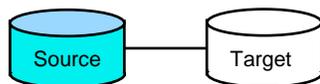
When copy is complete, relationship between source and target ends

Persistent FlashCopy

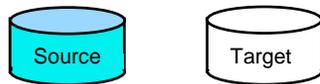
Time



FlashCopy command issued with COPY and PERSISTENT



When copy is complete, relationship between source and target persists



Explicitly withdraw the relationship

FlashCopy Enhancement Summary

- More flexibility with the means to create a permanent copy of the data even if the relationship was initially established with the NOCOPY parameter
- Increased protection from inadvertent updates to FlashCopy target volumes by maintaining the relationship which identifies the for the most recent FlashCopy

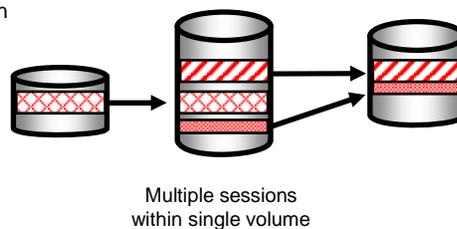
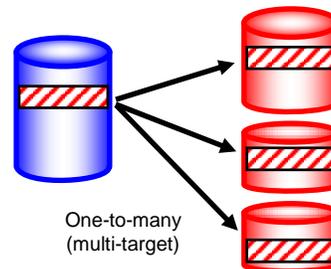
FlashCopy V2

FlashCopy V2

- Same functions as FlashCopy Version 1 plus:
 - Multiple Relationship FlashCopy
 - Incremental FlashCopy
 - Consistency Group FlashCopy
 - Inband FlashCopy
 - Target can be in any LSS within an ESS (Across LSS / Across cluster)
 - FlashCopy establish time reduction
 - Data Set FlashCopy
- Supported on Models 800 and Fxx
- Requires ESS LIC level 2.2.0, or later
- New feature codes:
 - 2105 Model 800 features #86xx plus 2240 Model FLC features #86xx
 - 2105 Models F20 and F10 features #186x
- FlashCopy Version 1 to Version 2 upgrades will be available
 - Feature exchange to same or higher capacity license with "delta pricing"
 - PPRC V2 requires FlashCopy V2

Multiple Relationship FlashCopy

- A volume can participate in multiple FlashCopy relationships
 - One source to many targets
 - A source can have up to 12 targets
 - A target can only have one source
 - A target cannot be a source at the same time
- A volume, LUN, or data set can be only a source or target at any given time
 - However, data set source and target can reside within a single volume
- Increased flexibility
- Improved capacity management and utilization



Number of Relationships

- The number of relationships per volume is a function of the volume's size
 - For a volume with n tracks, the maximum number of relationships allowed for the volume is slightly less than $n/64$ up to a maximum 64K
- The number of relationships for an ESS is also limited
 - For ESS with n tracks, the total number of relationships will be slightly less than $n/128$
- A background copy may be started for a relationship even if it was not requested by the user
 - For example, if a Data Set FlashCopy was done and the source data set is deleted, DADSM will start background copy to preserve the target data set

Incremental FlashCopy

- 'Change recording' tracks changes made to source and target volumes after establish of FlashCopy relationship

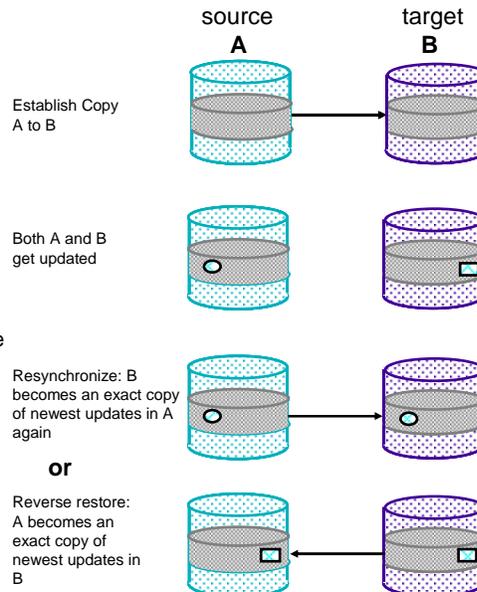
- Use in conjunction with background copy and persistent relationship options
- Supported only at full volume/ LUN level
- There can only be one incremental relationship per volume but can coexist with other non-incremental relationships

- During refresh:

- Only changed data is copied from the source to the target
- Source and target volumes can be reversed
 - Source volume (previously the target) will refresh the target volume (previously the source) with only changed data
 - Background copy must complete first

- Increased flexibility

- Faster completion time for physical copy



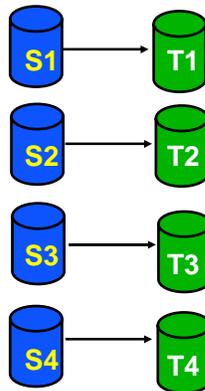
Incremental FlashCopy Support

- Initial Establish
 - Start Change Recording
 - Persistent FC (required)
- Increment FC
 - Establish FC task type
 - Increment FlashCopy
 - Start Change Recording
 - Persistent FC (required)



FlashCopy Consistency Group Support

- Hold off initiation / completion of write I/O to the source volumes until FlashCopy establish is completed
 - Select source and target volumes with 'freeze' option
- Create Consistency Group Created tasks to allow resumption of I/O
 - One per LSS
- Enables creation of a consistent point-in-time copy across multiple volumes with minimum host impact
- Target of each source volume are within one ESS but source volumes within a consistency group can span ESSs
- Source volumes can span multiple CS-Web domains if using CLI scripts



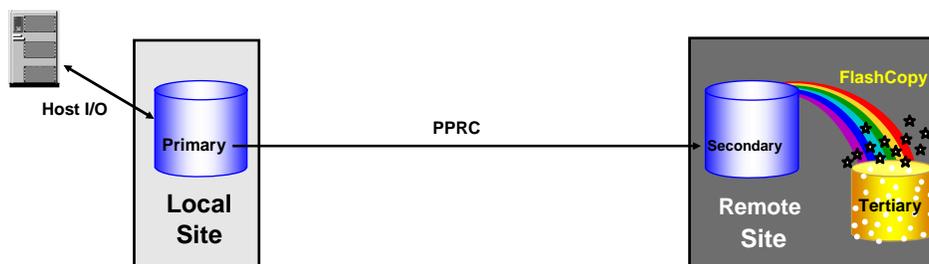
- FlashCopy S1 to T1
 - Writes cannot proceed on S1
 - Any writes occurring on S2-S4 are not dependent writes
- FlashCopy S2 to T2
 - Writes cannot proceed on S1 or S2
 - Any writes occurring on S3-S4 are not dependent writes
- FlashCopy S3 to T3 and S4 to T4
 - T1-T4 contain a consistent copy
- Issue Consistency Group Created
 - Writes may proceed to S1-S4.

FlashCopy Consistency Group Considerations

- If the Consistency Group source volumes are used with a journaled file system and the source LUNs are not unmounted before a FlashCopy, it is likely that fsck will have to be run on the target volume
 - Metadata describing changes to file system structures written synchronously to disk preserves file system integrity
- If the Consistency Group source volumes back a Data Base Management Subsystem (DBSM), inflight transactions will be backed out during restart using the logs
 - Log writes synchronously written to disk
 - Transaction not complete until log writes are complete
 - Database writes occur after log writes are complete

Inband FlashCopy

- FlashCopy commands are issued to the primary volume of a PPRC pair at the local site.
- PPRC pair acts as a conduit to the remote site for execution of the command at the remote site.
- Volume at remote site that executes the command must be a PPRC secondary volume



Other FlashCopy Enhancements

- Target volume can be in any LSS within the same ESS
 - Elimination of Logical Subsystem (LSS) constraint allows a FlashCopy source and target to reside in different LSSs of the same type (Open or CKD) within a single ESS
 - zSeries and Open
 - Improved flexibility
 - Simplifies capacity management and administration

- FlashCopy establish time reduction
 - Destage and Discard Scan are now done asynchronously
 - Improved performance reducing impact to application
 - Enables new FlashCopy applications (environments requiring rapid or numerous, repetitive establishes)

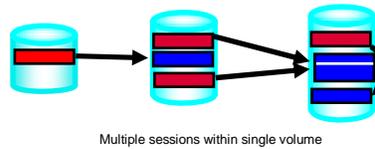
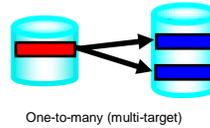
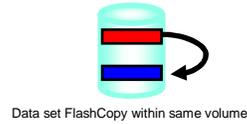
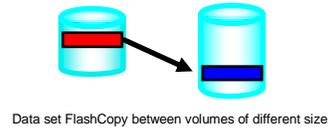
z/OS Enhancements

- DFSMSdss Enhancements
 - Data Set FlashCopy
 - Physical volume and tracks copy
 - Across LSS / Across cluster
 - FCWITHDRAW keyword supported for DUMP operations
 - DEFRAG use of FlashCopy if possible
- New keywords on COPY and DEFRAG commands
 - FASTREPLICATION(REQ | PREF | NO)
 - DEBUG(FRMSG(MIN | SUM | DTL))
- Improved informational messages and reason codes for cases when FC cannot be used

- TSO and DFSMSdss
 - Multiple Relationship FlashCopy

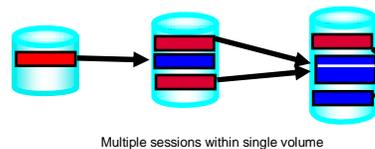
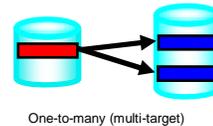
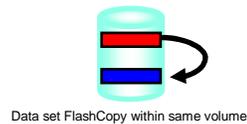
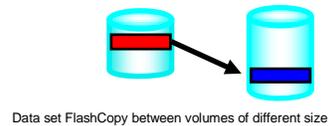
Data Set FlashCopy

- FlashCopy of individual data sets
 - OS/390 V2R10 and higher
- Source and target volumes do not need to be the same size
- Location within target volume does not need to match location within source volume
- Data set can be copied onto source volume at different location within the volume
- Improved granularity provides better capacity management and utilization



DFSMSdss Logical Data Set Copy

- Default is to use FlashCopy if possible
- FlashCopy eligible volumes preferred for new target allocation
 - SMS prefers FlashCopy eligible volumes in selected storage group
 - DFSMSdss prefers FlashCopy eligible volumes in output volume list
 - Same track geometry to be FlashCopy eligible
- FCNOCOPY keyword supported
- Cross LSS / Cross cluster
- All data set types supported when DFSMSdss called SnapShot supported by DFSMSdss calling FlashCopy!!!
 - DFSMSdss has never used fast replication for non-SMS multivolume indexed VSAM data sets copied to multivolume target
 - Fast replication not invoked when DFSMSdss calls a utility to perform the copy
 - SIBBATCH with DATAMOVER(DSS) will use FlashCopy
- Target data set can not be allocated on a PPRC primary volume on LIC 2.4.0 and lower



Invoking FlashCopy V2 Functions

- Multiple Relations (Transparent)
 - DFSMSdss
 - ANTRQST API
 - TSO Commands
 - Web Copy Services / CLI
- FlashCopy across LSSs and across clusters (Transparent)
 - DFSMSdss
 - ANTRQST API
 - TSO Commands
 - Web Copy Services / CLI
 - Third party client application via CIM agent
- Incremental FlashCopy
 - Web Copy Services / CLI
 - Third party client application via CIM agent
- Inband FlashCopy
 - Web Copy Services / CLI
 - Third party client application via CIM agent
 - GDPS (as part of Asynchronous Cascading PPRC support)
- FlashCopy Consistency Groups
 - Web Copy Services / CLI
 - Third party client application via CIM agent
- DataSet FlashCopy
 - DFSMSdss
 - ANTRQST API
 - TSO commands
 - Extents can be specified but VTOC is not updated
 - Could be used when performing data recovery

Software Prerequisites

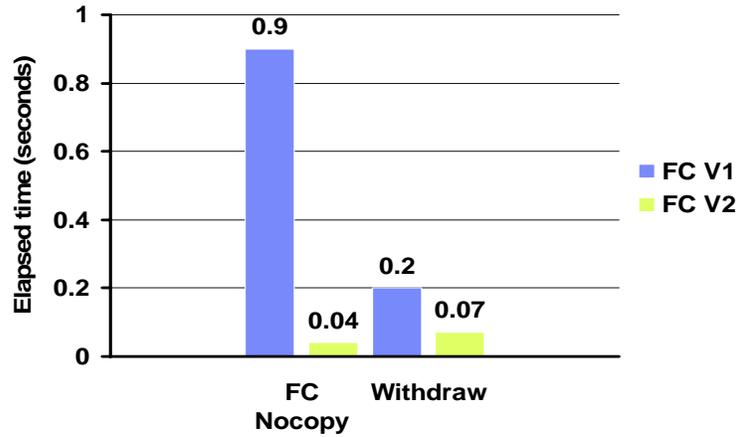
- **SPE will be available for**
 - OS/390 DFSMS V2R10 and all z/OS DFSMS releases
 - Systems without the V2 maintenance will still be able to invoke V1 functions on an ESS with V2 support but it is recommended all systems accessing the ESS with V2 have the support installed
 - ✓ If a volume has a V1 relationship initiated from a system without V2 support, no V2 function can be initiated until the V1 relationship is withdrawn
 - Applications using the ANTRQST SQRVDVCS interface can use Across LSS / Across Cluster functionality without change if the software support is installed and the ESS has the V2 LIC
 - IPL with CLPA (Clear Link Pack Area which is a cold start) will be required
- **DFSMS components enhanced**
 - DFSMSdss
 - SMS
 - SDM
 - AOM
 - Device Support
 - DADSM
 - DFSMShsm (coexistence only)
- **IXFP**
 - PTF L170770
- **RETAIN PSP Bucket**
 - Upgrade = 2105device
 - Select all of the MVS/ESA subsets

Microcode Code Load Issues for FlashCopy V2

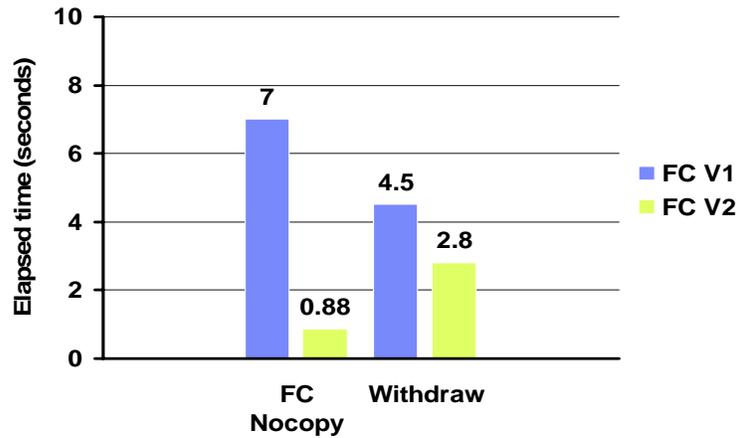
- Code load is concurrent BUT FlashCopy relationships will NOT be preserved during the code load procedure
- Recommendation is to either allow background copy operations complete and withdraw background nocopy operations or withdraw all relationships prior to code load

FlashCopy V2 Performance

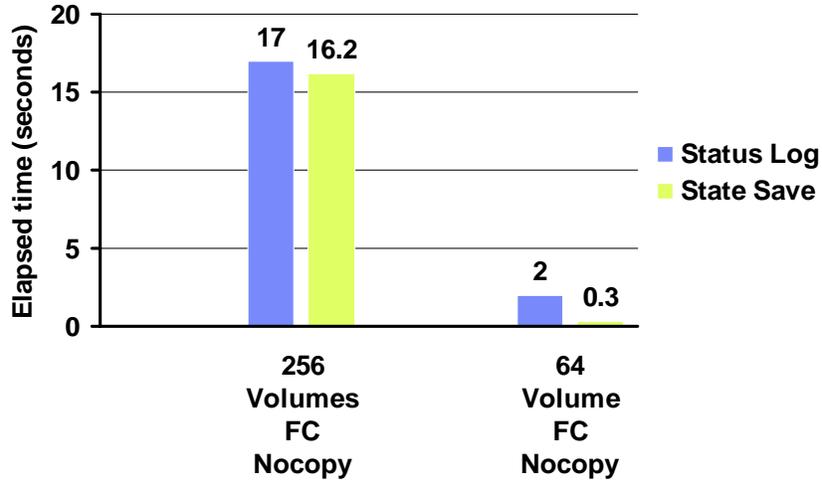
Improvement In FlashCopy Establish & Withdraw Time (z/OS using TSO commands – 1 x 3390-3 Volume)



Improvement In FlashCopy Establish & Withdraw Time (z/OS using TSO commands – 256 x 3390-3 Volume)

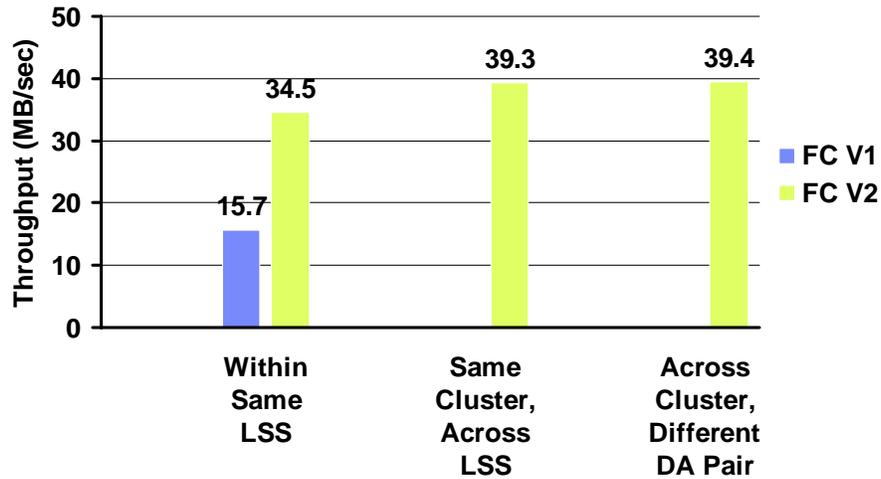


FlashCopy V2 Establish Time (Open Systems using Copy Services Server)



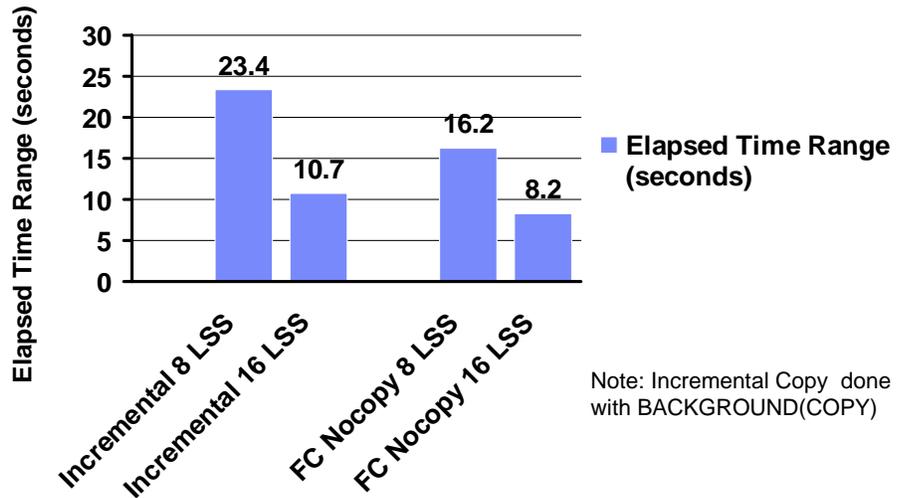
The status log has a one second granularity so State Save analysis can provide a more accurate timing measurement
 For this test, each volume is 6.5 GB
 For greater than 64 volumes, the Copy Services Server introduces a 2-3 second delay between subsequent groups of 64 volumes

Improvement in FlashCopy Background Copy Data Rates (Single 3390-3 Copied)



FlashCopy V2 w/Freeze (Minimum Timing to Complete using FC Manager Tool – z/OS)

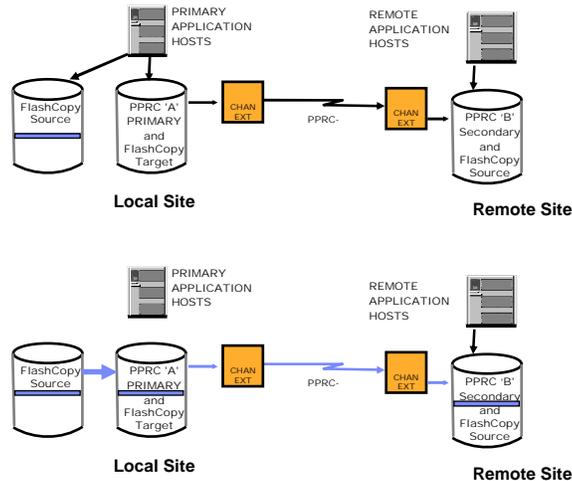
256 x 3390-9 (7600 cylinder) Volumes



ESS 2.4.0 FlashCopy Enhancements

New FlashCopy Function

Established PPRC primary in any PPRC environment (even suspended) can now be a FlashCopy target



97

© 2004 IBM Corporation

FlashCopy Target to a PPRC Primary

- Permits FlashCopy relationship if the FlashCopy target is also an established PPRC primary volume
 - Enables FlashCopy to create a full or incremental point in time copy and then have PPRC copy the data to secondary site
 - Supports incremental, consistency, persistent, and data set FlashCopy
 - PPRC primary can be synchronous, extended distance, or Global Mirror
 - Could enable the implementation of a multi-hop remote copy solution as an alternative to other vendors solutions
 - Data at the secondary site will be inconsistent until PPRC transmits all of the FlashCopy data to the secondary (primary in "pending duplex" state)
 - DFSMSdss APAR OA06196 provides new keyword FCTOPPRCPPRIMARY or FCTOPPRCP to allow the FlashCopy to occur otherwise traditional I/O will be used
 - New option in Web Copy Services when after selecting the FlashCopy source and target (which could be a PPRC source) allows the user to FlashCopy to the PPRC source
- Operational management
 - All environments
 - Multiple Device Manager Replication Manager, Web Copy Services, CLI, or API
 - z/OS
 - GDPS, TSO or API
- Prerequisites
 - ESS Model 800 or 750
 - ESS LIC level 2.4.0

98

© 2004 IBM Corporation

Fast Reverse Restore

- Reverses the FlashCopy relationship and copies modified tracks from the target volume to the source volume
 - The background copy process must complete before you can reverse the order of the FlashCopy relationship to its original source and target relationship
- After a reverse restore establish completes successfully, the volume that was previously a source becomes a target volume and the volume that was previously a target becomes a source volume
- You cannot reverse an incremental FlashCopy relationship on the volume pair if the source volume is involved in more than one FlashCopy relationship at the same time (Multiple Relationship FlashCopy)
- Used in Global Mirror for ESS environment

Revertible FlashCopy

- Intended use in Global Mirror for ESS environment
- Initial FlashCopy relationship at remote site for volumes 'B' and 'C' are:
 - Persistent
 - Start change recording
 - No background copy
 - Inhibit target write
- Revertible FlashCopy issued Inband by Global Mirror
- Withdraw command determines whether or not the FlashCopy relationship reverts or is committed
 - Query status of FlashCopy pairs to determine whether to revert or commit

Reverse Restore Steps

1. Establish FlashCopy pairs using volumes in a source LSS and target LSS
2. Select the Start Change Recording and Persistent FlashCopy options
 - Allow the background copy process of the FlashCopy pairs to complete
3. Using the same FlashCopy pairs established in Step 1, issue a new FlashCopy establish task and select the Reverse restore option

Summary

- Improved granularity
 - Data Set FlashCopy
- Increased flexibility
 - FlashCopy Multiple Relationships
 - FlashCopy Incremental Refresh
 - FlashCopy Consistency Group
 - FlashCopy Across LSS / Across Cluster
 - FlashCopy Inband
 - FlashCopy to PPRC source
- Better capacity management and utilization
 - FlashCopy Multiple Relationships
- Faster FlashCopy completion times
 - FlashCopy Incremental Refresh
 - FlashCopy scan elimination
- Easier management
 - Data Set FlashCopy
 - FlashCopy Multiple Relationships
 - FlashCopy Across LSS / Across Cluster
 - FlashCopy Inband
 - Web Copy Services



Addendum

© 2003 IBM Corporation

IBM TotalStorage © 

Copy Services Matrix

Device Is May Become	XRC Primary	XRC Secondary	PPRC Primary	PPRC Secondary	FlashCopy Source	FlashCopy Target	ConCopy Source
XRC Primary	No	Yes	Yes	No	Yes	No	Yes
XRC Secondary	Yes	No	Yes	No	Yes	No	Yes
PPRC Primary	Yes	Yes	No	Yes Note 1	Yes	Yes	Yes
PPRC Secondary	No	No	Yes Note1	No	Yes	No	No
FlashCopy Source	Yes	Yes	Yes	Yes	Yes Notes 3 & 4	Yes Note 4	Yes
FlashCopy Target	No	No	Yes Note 2	No	Yes Note 4	Yes Note 4	No
ConCopy Source	Yes	Yes	Yes	No	Yes	Yes	Yes

Notes:

1. Only in a Metro/Global Copy for ESS (Asynchronous Cascading PPRC) environment.
2. FlashCopy V2 at LIC 2.4.0 and higher.
3. FlashCopy V2 Multiple Relationship.
4. FlashCopy V2 Data Set FlashCopy.

104

© 2004 IBM Corporation

References

- **Publications**
 - SC35-0395 Advanced Copy Services
 - Go to publication updates and select OA01616 & OA03569 which will take you to:
<http://www.storage.ibm.com/software/sms/sdm/sdmpubs.html#OA01616>
 - Click on:
<http://www-1.ibm.com/servers/eserver/zseries/zos/bkserv/>
 - Click on 'z/OS elements and features' and then select DFSMS to find SC35-0395
 - SC35-0394 DFSMSdss Storage Administration Reference
 - SC35-0393 DFSMSdss Storage Administration Guide
 - SC26-7448 Web Interface User's Guide
 - GC35-0033 Device Support Facilities User's Guide and Reference
 - SC26-7494 Command Line Interfaces User's Guide
- **RedBooks**
 - SG24-5757 Implementing in an Open Environment
 - SG24-5680 Implementing ESS Copy Services with IBM eServer zSeries

References *(continued)*

- **URLs**
 - <http://www.storage.ibm.com/software/index.htm>
 - <http://www.storage.ibm.com/software/sms/sdm/index.html>
 - <http://www.storage.ibm.com/software/sms/sdm/sdmtech.html>
 - <http://www-1.ibm.com/servers/storage/support/disk/2105.html>