Data Migration and Disaster Recovery: At Odds No More

Brett Quinn
Don Pease
EMC Corporation

Session 8036
August 5, 2010

Mainframe Migrations

Challenges…

- Disruptive
  - To applications
  - To Disaster Recovery
- Limited migration windows
  - Usually during the night and weekends
  - Planning required before migrations
- Costly
  - Require specialized skills
  - Off hours is cost prohibitive
- Complex
  - Infrastructure interdependencies
- Inflexible
  - Different tools required for volume and dataset migrations

Current Host-Based Migration Types

Simple – Volume-to-Volume

- Same size or smaller to larger
  - 3390-3 to 3390 Model 3
  - 3390-3 to 3390-9, or -27, or -54
  - 3390-9 to 3390-27, or -54
  - 3390-27 to 3390-54
- Customer preference of tool
- Currently licensed to use tool
- Familiarity with tools from specific vendor
- Limited scalability of migration process

Complex – Dataset Migration

- Volume consolidation
  - Several 3390-3s to a single 3390-9
  - Several 3390-9s to a single 3390-27 or single 3390-54
  - Several 3390-27s to a single 3390-54
- Dataset placement
  - Free up active packs to allow UCB reclamation
  - Multivolume dataset identification requirements
- Catalog cleanup requirements

Host-Based Migration Types

Steps to a Successful Migration

- Know Your data
  - Identify
  - Quantify
  - Verify
- Always use Volume Migration (z/OS Migrator, FDRPAS or TDMF) **first** for
  - Like-to-like volume copies (3-3,9-9)
  - Increasing volume size (3-9,9-27)
  - Watch for VTOC size as dynamic VTOC expansion not always possible with ICKDSF
  - Volumes with APF, Linklst, HFS/zFS, Catalogs
- Some volumes require special handling and might be candidates for reallocation
  - Page Volumes. Larger Page Dataset support with z/OS 1.8 (OA20749)
  - Spool volumes. LARGEDS option allow > 64K track spool dataset
  - Sysplex Couple Dataset volumes. Evaluate current allocations

Host-Based Migration Types

Steps to a Successful Migration

- Use DFDSS or FDR to move anything not in use (TSO, Batch, System)
  - Extent consolidation
  - Catalog updates kept to a minimum
  - Easy setup
- Let HSM handle ML1 volumes (ADDVOL/FREEVOL) if going to larger ML1 volumes
- Datasets that are redefined at regular schedule will move by attrition by changing SMS volume status (GDG's, Database Reorg etc)
- Use dataset level migration for what remains
  - Requires dataset name masking (include/exclude) to select exactly what is desired
  - High catalog activity can occur if not planned properly
  - Planning and setup can be **80%** of the work for the dataset-level migration

Host-Based Migration Types

Steps to a Successful Migration

- Analyze and clean the user catalogs as needed
  - EXAMINE and DIAGNOSE
  - Use Catalog Solutions or similar products
- Obtain lists of datasets to be moved
  - Attempt to sort them into groups where a new HLQ can be used for the move
    - Allows creation of a new catalog to be used only for the migration and deleted when complete
- SMS considerations
  - DISNEW all volumes you are moving from
  - Check ACS to ensure proper allocation of interim datasets to fall on the target volumes if needed
  - Allow enough target volumes for source datasets
- Work with small groups
  - No more than 500 datasets or 15000 extents (probably less)
  - Attempt to limit the source volumes online to as few LPARs as possible

But what about Disaster Recovery protection during Migration????

Homogenous migrations

Control unit replication based migration

Example:
Using SRDF FourSite for Data Migration

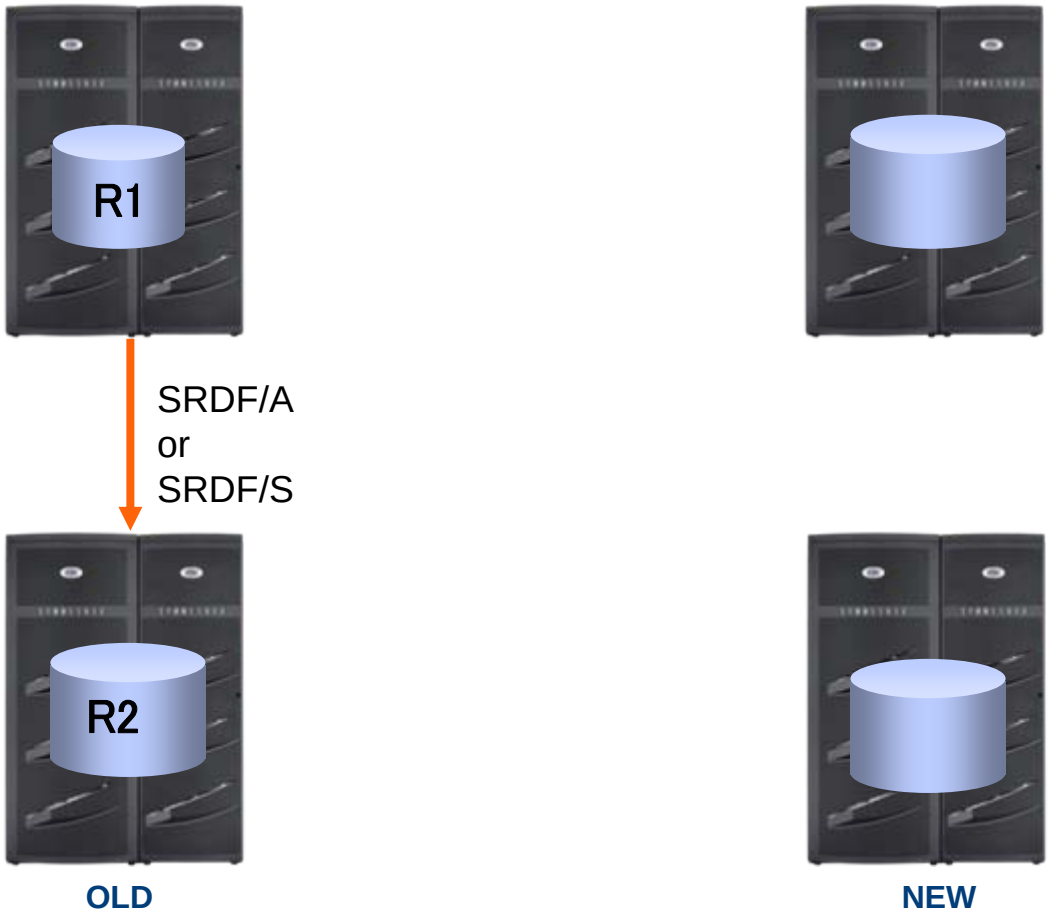

Example:
Using SRDF FourSite for Data Migration

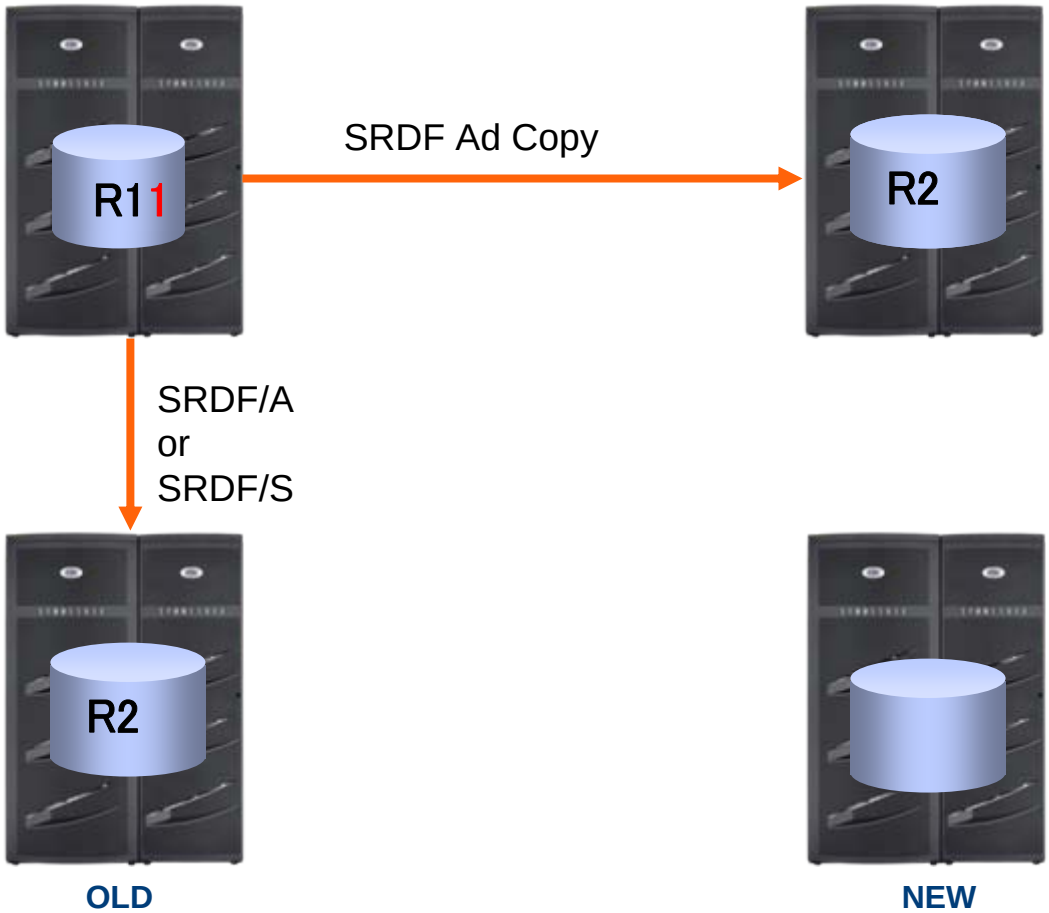

Example:
Using SRDF FourSite for Data Migration

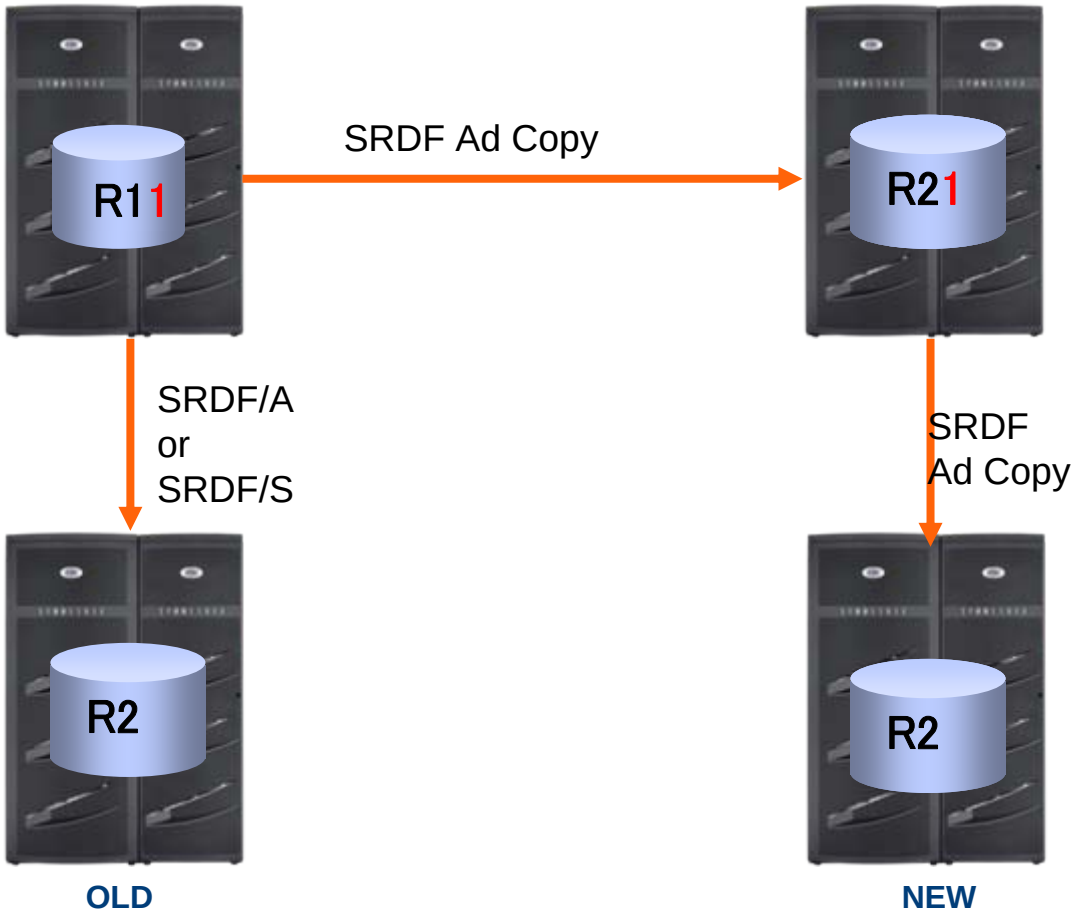

Example:
Using SRDF FourSite for Data Migration

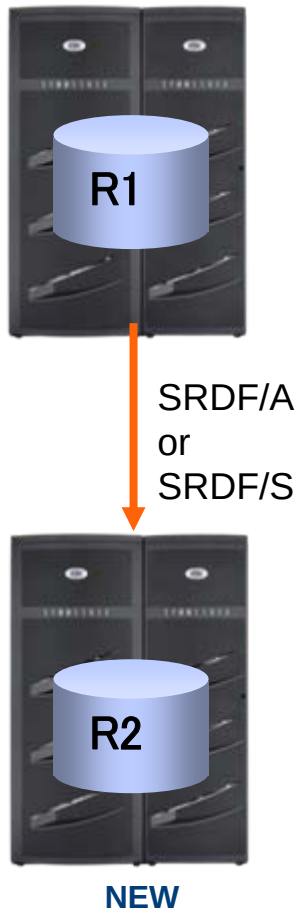

Homogenous migrations

Host based migration

Disaster Recovery and Host-based Migrations Maintaining Replication Solution

- Active Migration can impact ‘Continuous Availability’ protection (Hyperswap/AutoSwap)
  - AutoSwap / HyperSwap needs to be temporary disabled during migration’s UCB swap
  - EMC z/OS Migrator handles this automatically for AutoSwap only
  - IBM TDMF handles this for both AutoSwap and HyperSwap automatically
  - IDP FDRPAS provides jobs to automatically enable/disable Hyperswap/Autoswap
- DR consistency maintained when migrating within same replication type and consistency group
- Available DR bandwidth impacts speed of migration
- DR consistency can be lost when migration source and target use different replication type
  - TDMF can detect automatically and allow override
  - FDRPAS detects and issues warning then completes migration

Disaster Recovery and Host-based Migrations Maintaining Replication Solution

Consistency Maintained!

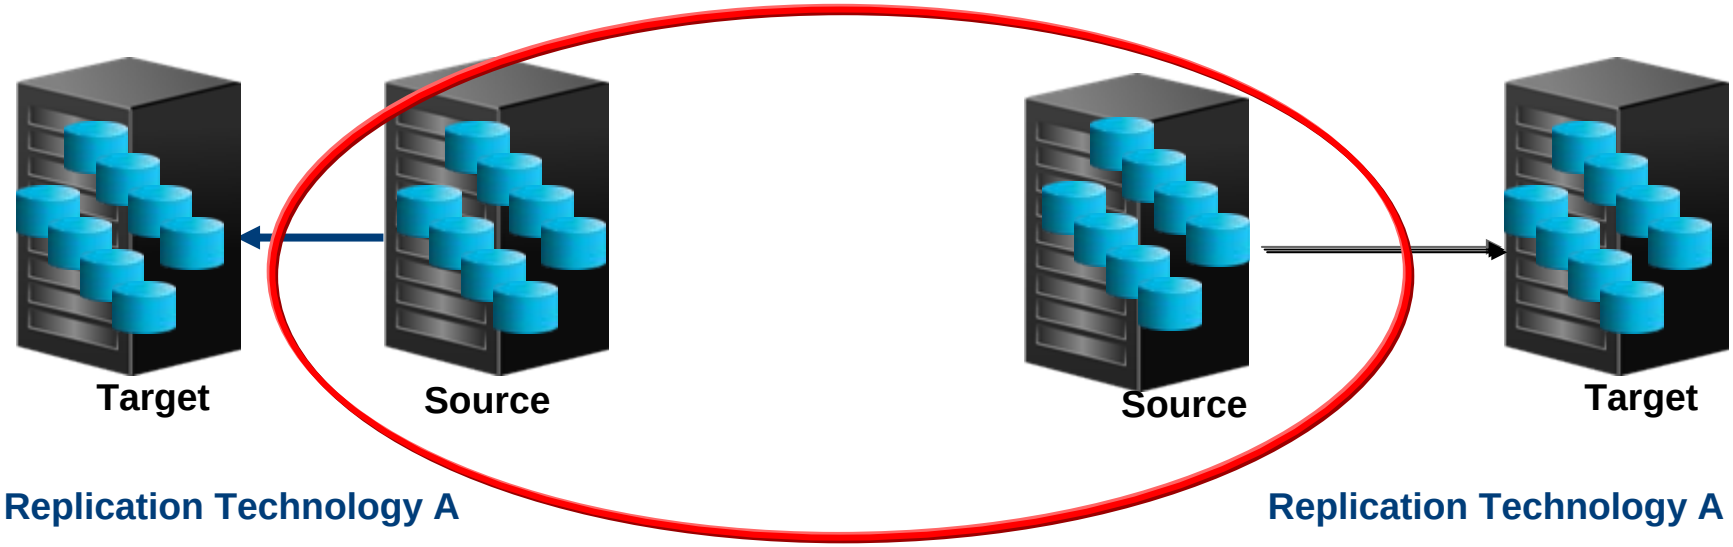

Disaster Recovery and Host-based Migrations Switching Replication Technology: The Problem

Consistency Lost!

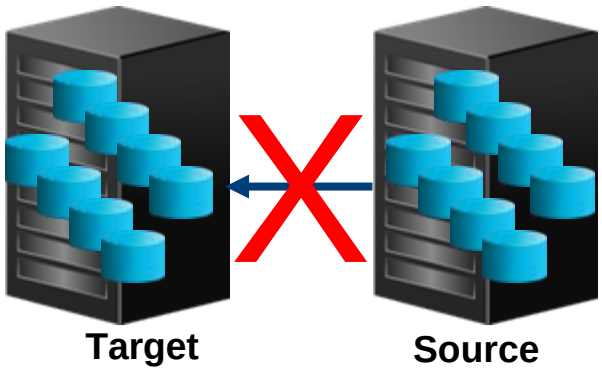

Replication Technology A

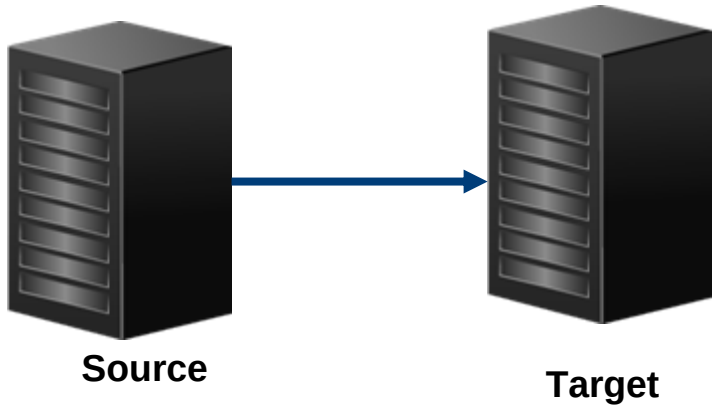

Replication Technology B

Disaster Recovery and Host-based Migrations
Switching Replication Technology: Current Solution

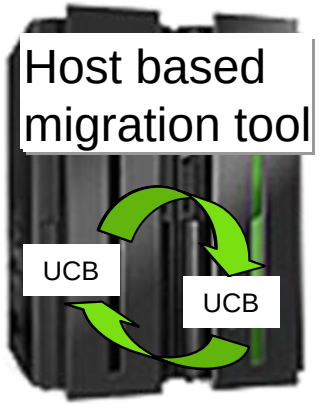

1. Suspend replication (aging RPO)

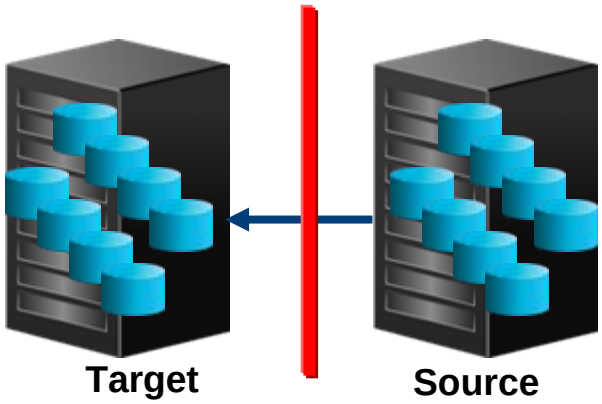

Replication Technology A

2. Migrate volumes in subsets

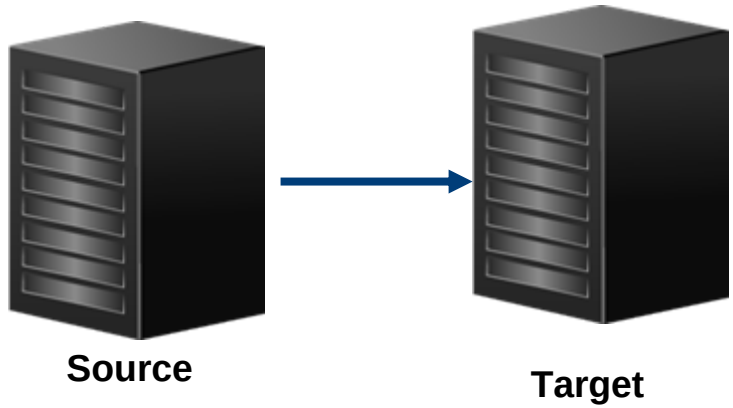

Replication Technology B

Disaster Recovery and Host-based Migrations Switching Replication Technology: New Solution

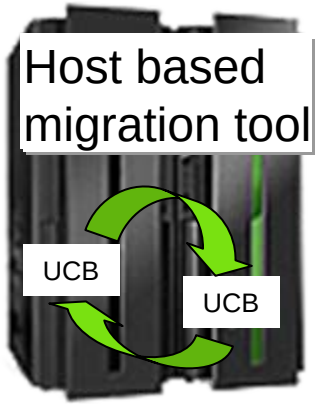

1. Copy volumes then mirror

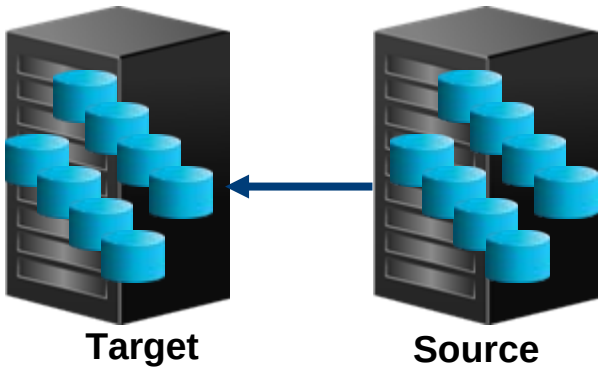

Replication Technology A

2. Swap all volumes at once

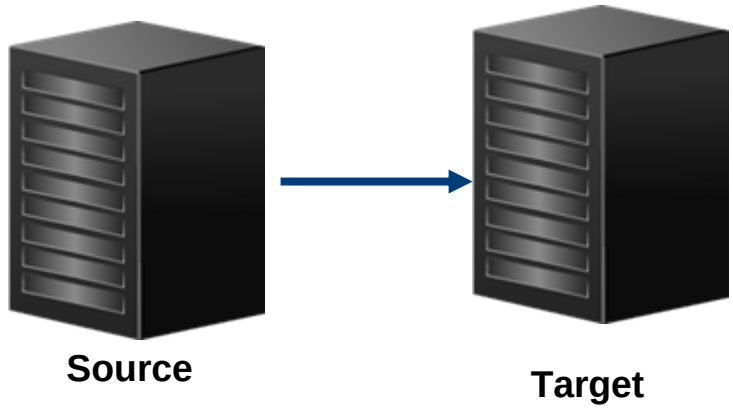

Replication Technology B

Disaster Recovery and Host-based Migrations Switching Replication Technology: New Solution

2 Non-disruptive to applications

3 Preserves remote replication at source—full disaster recovery operations

1 Host-based migration

5 Non-disruptive to DR when migrating between business continuity solutions

Example:

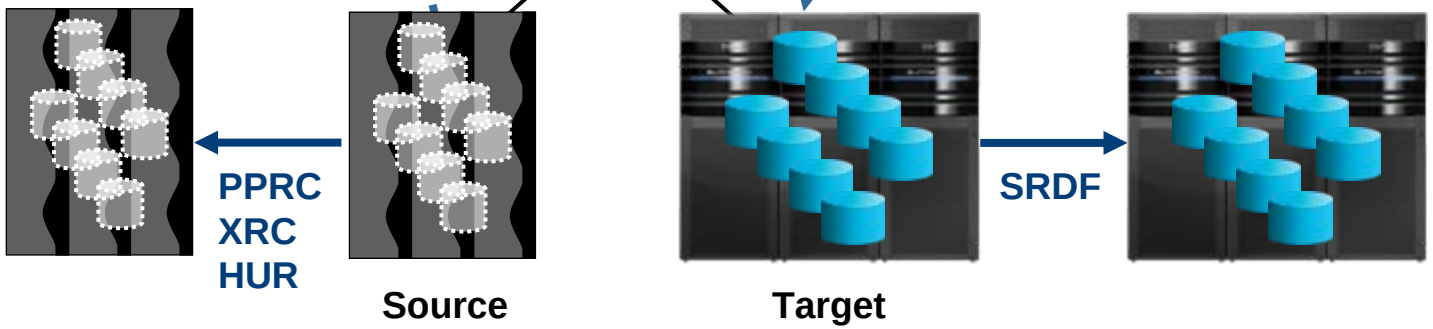

4 Supports heterogeneous storage

Host-Based Data Migration Products & Offerings

z/OS Migrator Functions and Capabilities

- Full volume data migration from source to target volumes
- Supported from any Count Key Data and ECKD control units to Symmetrix
- Swaps active volumes with few exceptions
- Operations are transparent to users and applications are unaffected
- Supports multiple system data-sharing environments

Host-Based Data Migration Products & Offerings

z/OS Migrator Functions and Capabilities

- Combines smaller volumes to larger volumes
  - MOD 3s to MOD 9s, or any size
  - Relieves UCB constraints with volume consolidation
- Transparent to application

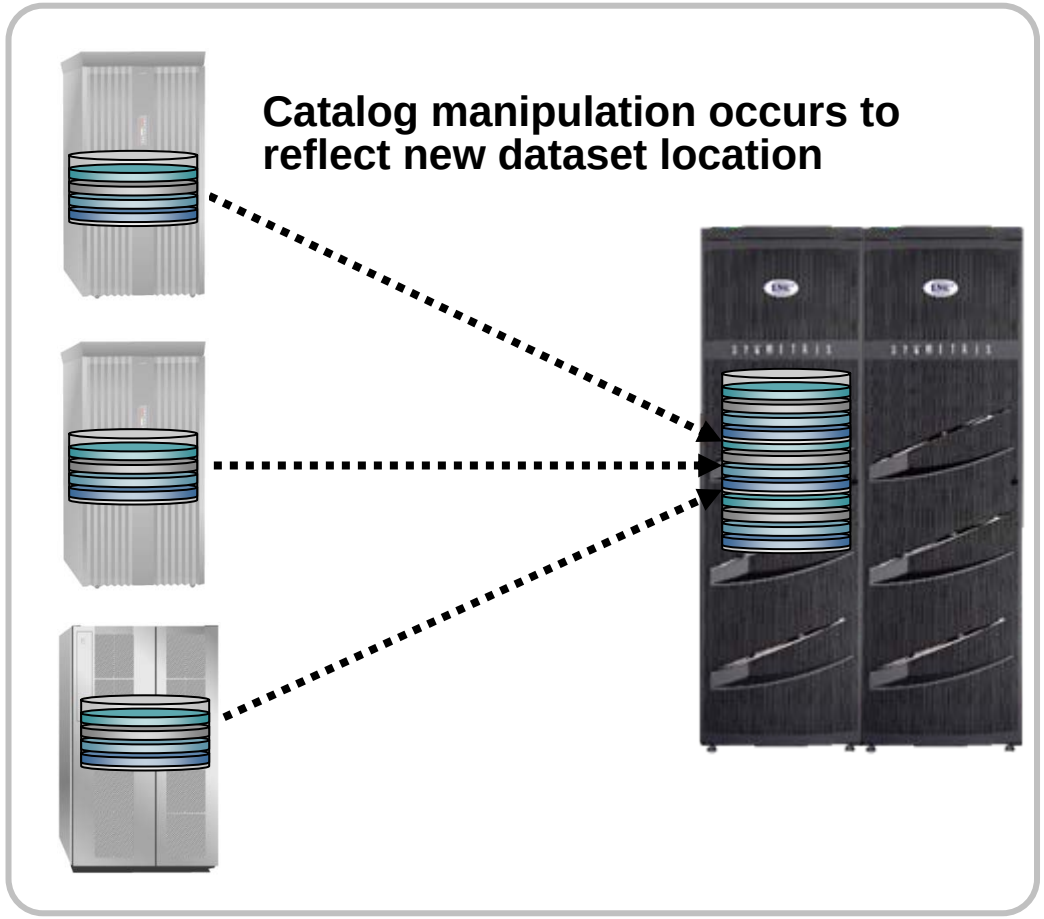

Host-Based Data Migration Products & Offerings

z/OS Migrator Dataset-Level Functions and Capabilities

- Logical data migration
  - Enables UCB reclamation without application outage
    - Also known as ‘active completion’
  - Provides the ability to migrate data at the logical dataset (extent) level from one set of volumes to another
    - All data movement is accomplished without application downtime
  - Updates the catalog automatically while applications continue running
  - Supports all CKD/ECKD control units
  - Migration sessions are parameter-driven
  - Allows for full system sharing throughout the datacenter
  - Uses both **copying** and **mirroring** techniques to achieve source/target synchronization

Host-Based Data Migration Products & Offerings

z/OS Migrator Dataset-Level Functions and Capabilities

- Logical data migration (continued)
  - Works with pairs of datasets
  - Transparent to applications
  - Vendor independent of source and target storage arrays
  - SMS considerations
    - Storage group and other ACS parameters are honored. Target datasets are allocated using source dataset names.
    - No option to bypass SMS
      - *Must ensure the ACS is properly coded to do what you want*
    - No effect on last reference date; it is maintained
    - As with all SMS datasets, migrating datasets **must** be cataloged

Host-Based Data Migration Products & Offerings

z/OS Migrator Dataset-Level Limitations and Restrictions

- Application bounce required for removal of z/OS Migrator from the system but **not** to free source UCBs for reuse
- Generally limited to datasets catalogued in shared user catalogs
- The target controller must be at an **equal or higher** technology level and be downward compatible
- Extensive planning required when working at the dataset level
  - Pre-implementation catalog ‘health check’ required

Host-Based Data Migration Products & Offerings

z/OS Migrator Technical Benefits

- Provides the ability to introduce new storage subsystem technologies with minimal disruption of service
- Allows users to easily reclaim z/OS UCBs by simplifying migration of datasets to larger volumes (combining volumes)
- Ensures all metadata always accurately reflects the location and status of datasets being migrated
- Supports datasets larger than 64K tracks

Host-Based Data Migration Products & Offerings

z/OS Migrator Business Benefits

- Eliminates limited migration windows
  - Migration windows typically take place on weekends
  - Time allotted is usually very small, possibly several hours
    - During this time, applications are taken down and data can be moved
  - Allows you to move the data at leisure at **any time of the day**
    - Transforms a ‘once-a-week’ 2-hour migration window to an unlimited ‘full week’ (168-hour) migration window
- Eliminates business outages
  - By eliminating the migration window, businesses no longer need to plan to have extended periods of application down-time
- Lowers cost
  - Allows most migration activity to be performed during normal business hours
- Reduces hardware expenditures
  - Frees up storage resources more quickly

z/OS Migrator: A look inside

Volume-Level Migration

Volume-Level Migration -- 1

- Multi-phased process
  - Initialization (create and promote)
  - Activation
  - Copy
  - Refresh

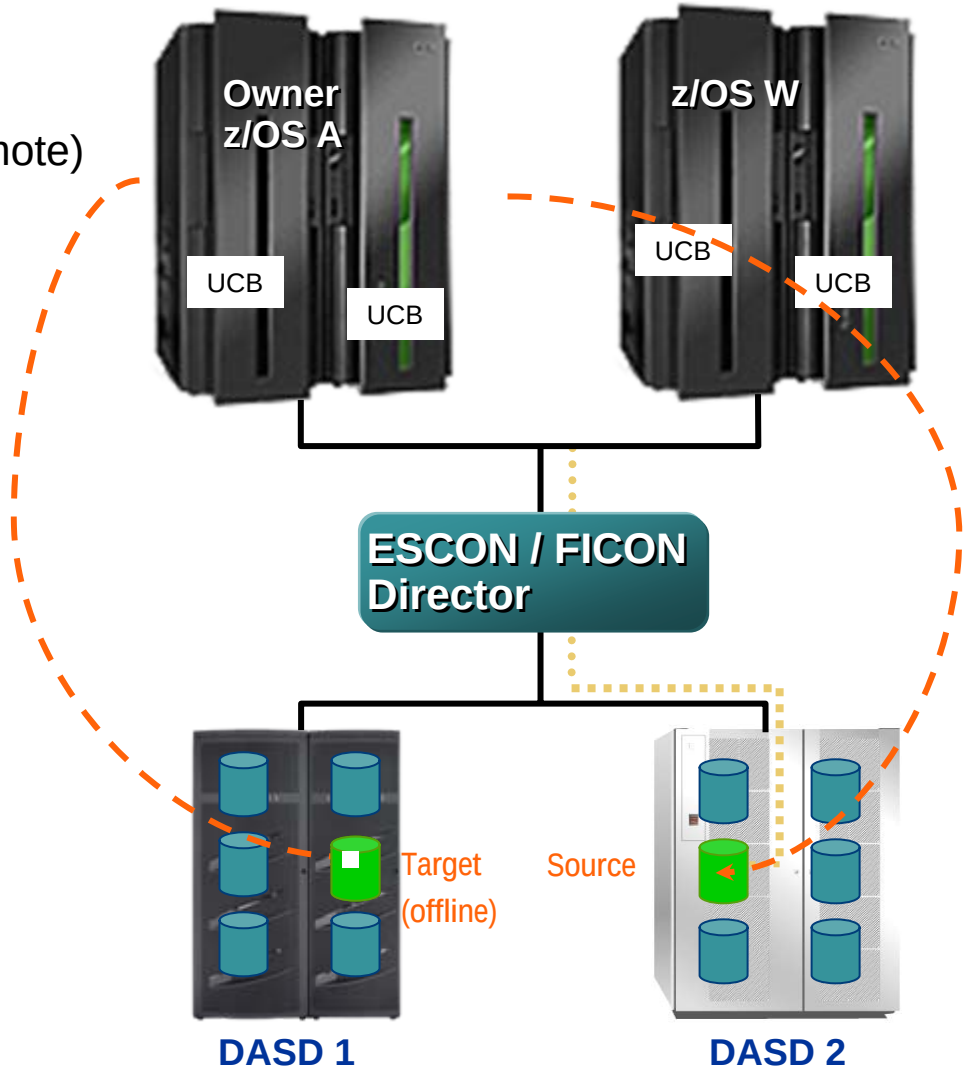

Volume-Level Migration -- 2

- Multi-phased process
  - Initialization
  - Activation
  - Copy
  - Refresh
  - Quiesce
  - Synchronization

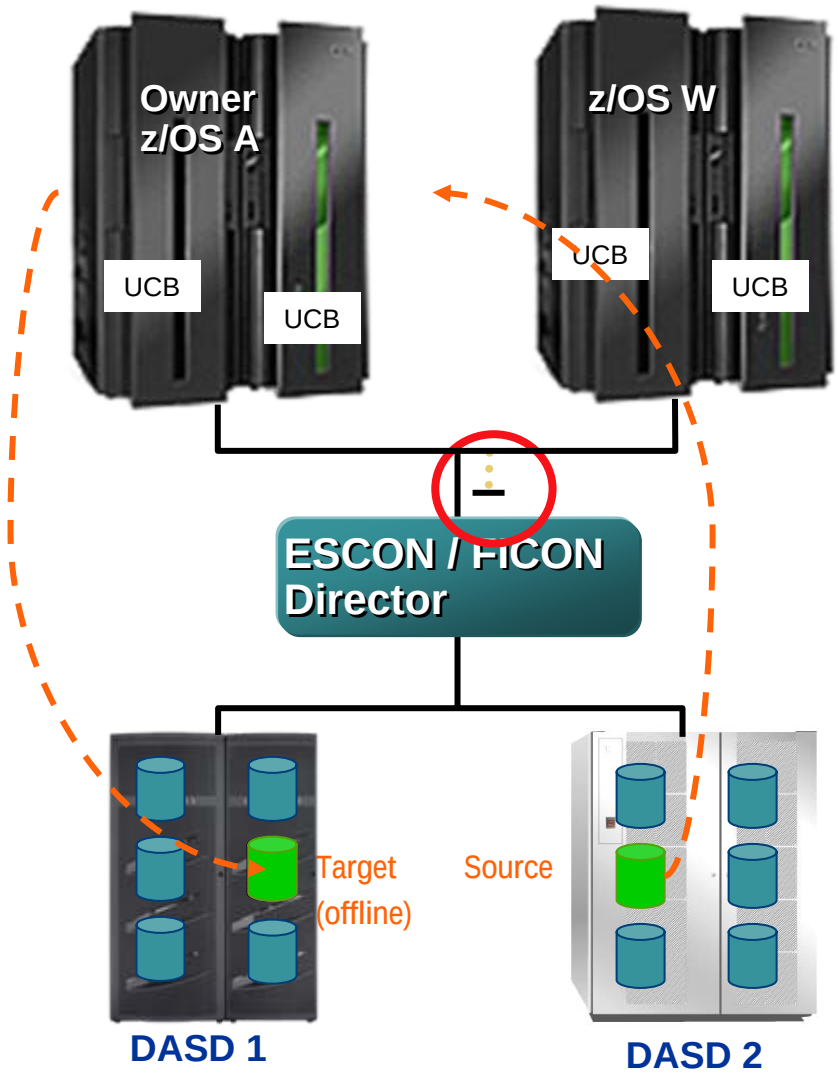

Volume-Level Migration -- 3

- Multi-phased process
  - Initialization
  - Activation
  - Copy
  - Refresh
  - Quiesce
  - Synchronization
  - Swap
  - Resume
  - Complete

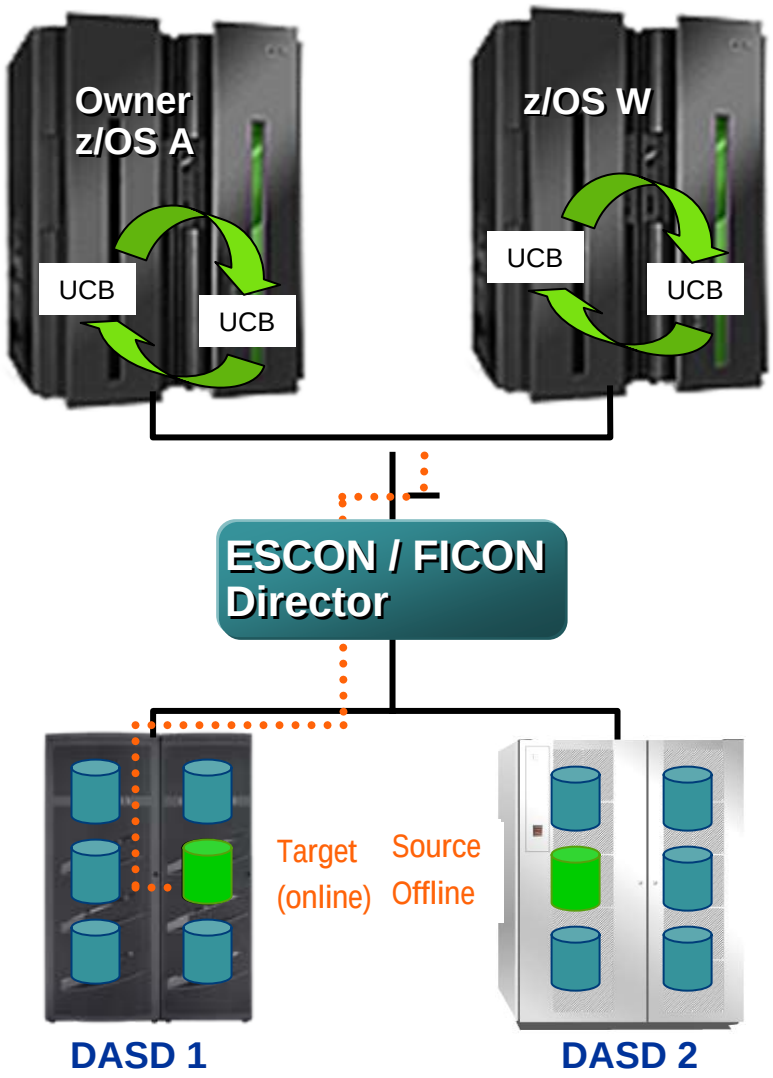

Volume-Level Consistent Migration

Volume-Level Consistent Migration -- 1

- Multi-phased process
  - Initialization (create and promote)
  - Activation
  - Copy
  - Refresh

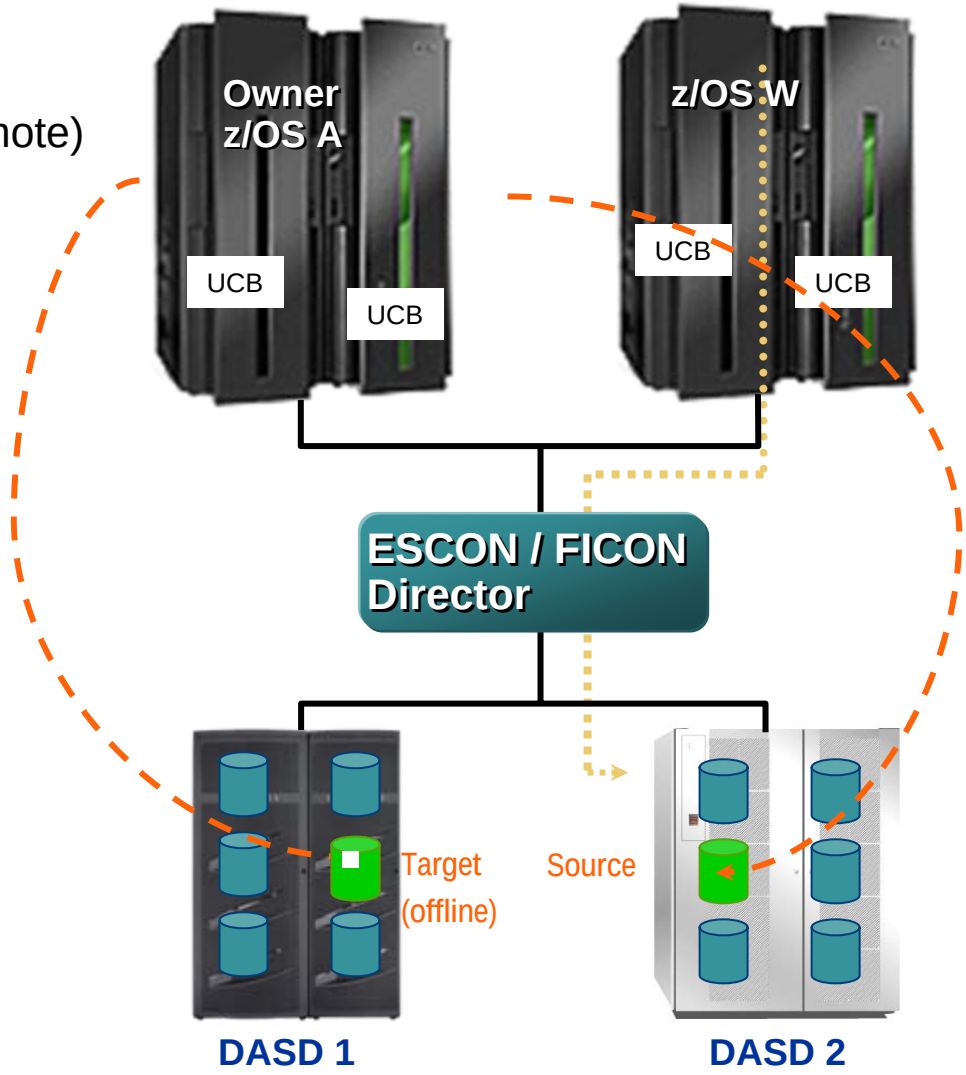

Volume-Level Consistent Migration -- 2

- Multi-phased process
  - Initialization
  - Activation
  - Copy
  - Refresh
  - Quiesce
  - Synchronization

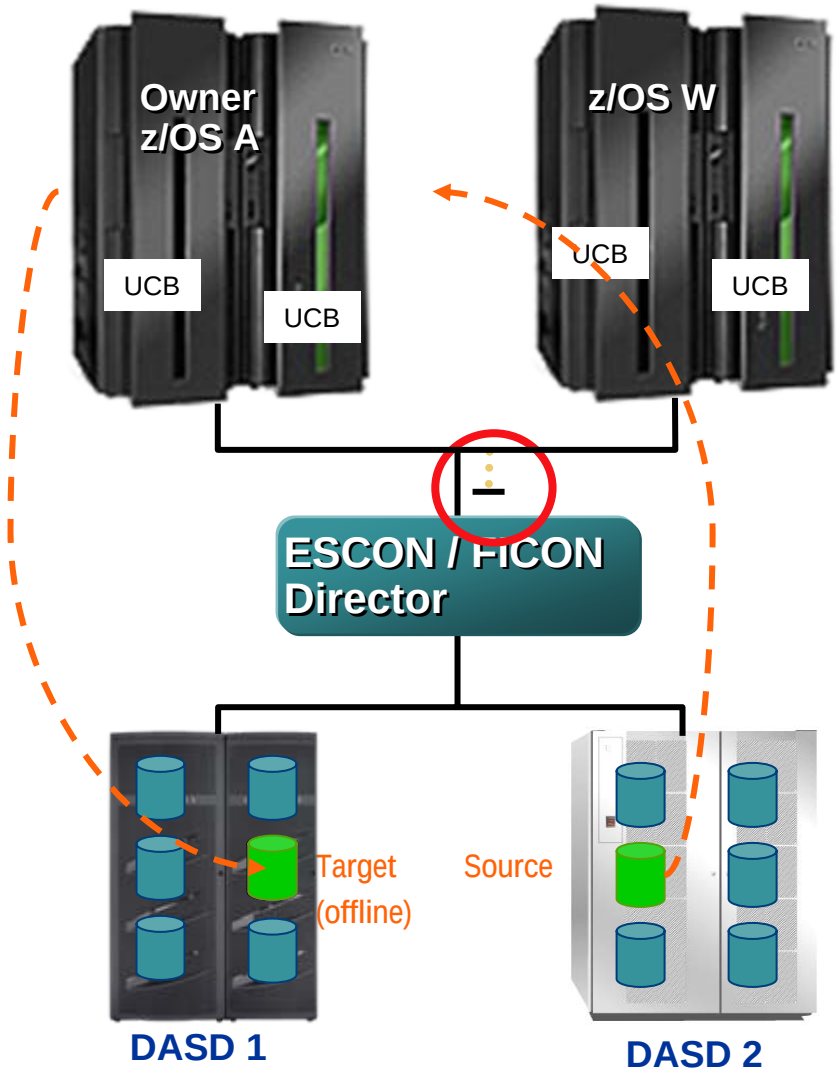

Volume-Level Consistent Migration -- 3

- Multi-phased process
  - Initialization
  - Activation
  - Copy
  - Refresh
  - Quiesce
  - Synchronization
  - Mirror

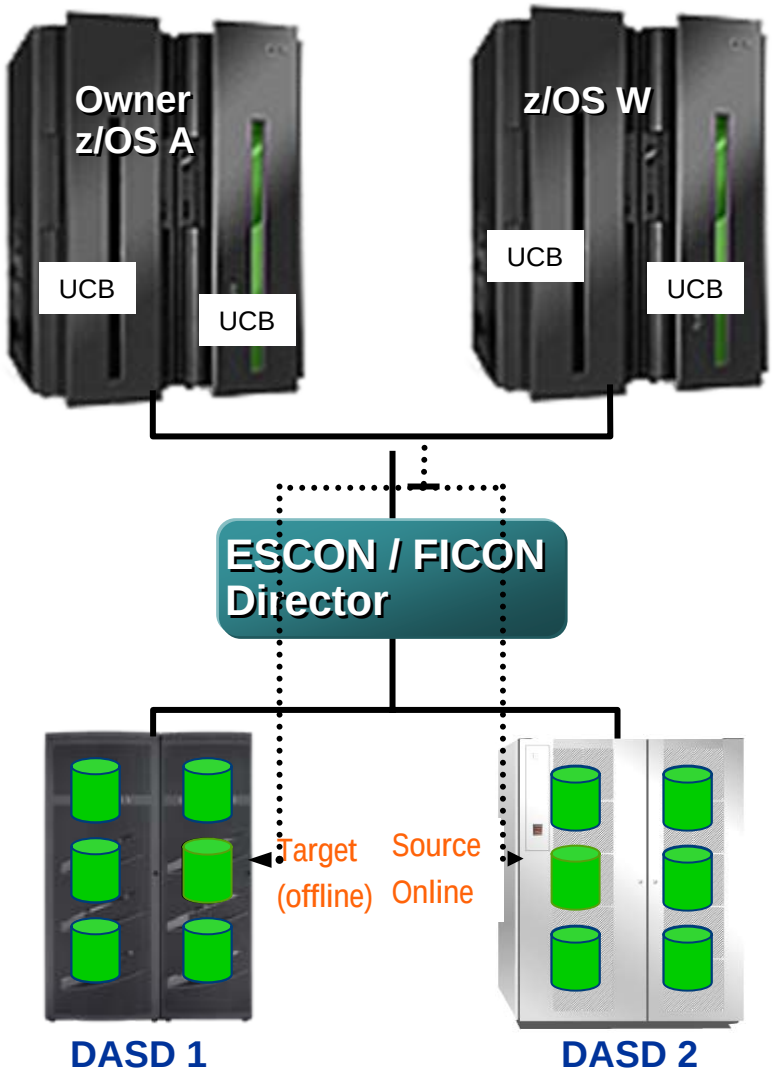

Volume-Level Consistent Migration -- 4

- Multi-phased process
  - Initialization
  - Activation
  - Copy
  - Refresh
  - Quiesce
  - Synchronization
  - Mirror
  - Consistent Swap
  - Resume
  - Complete

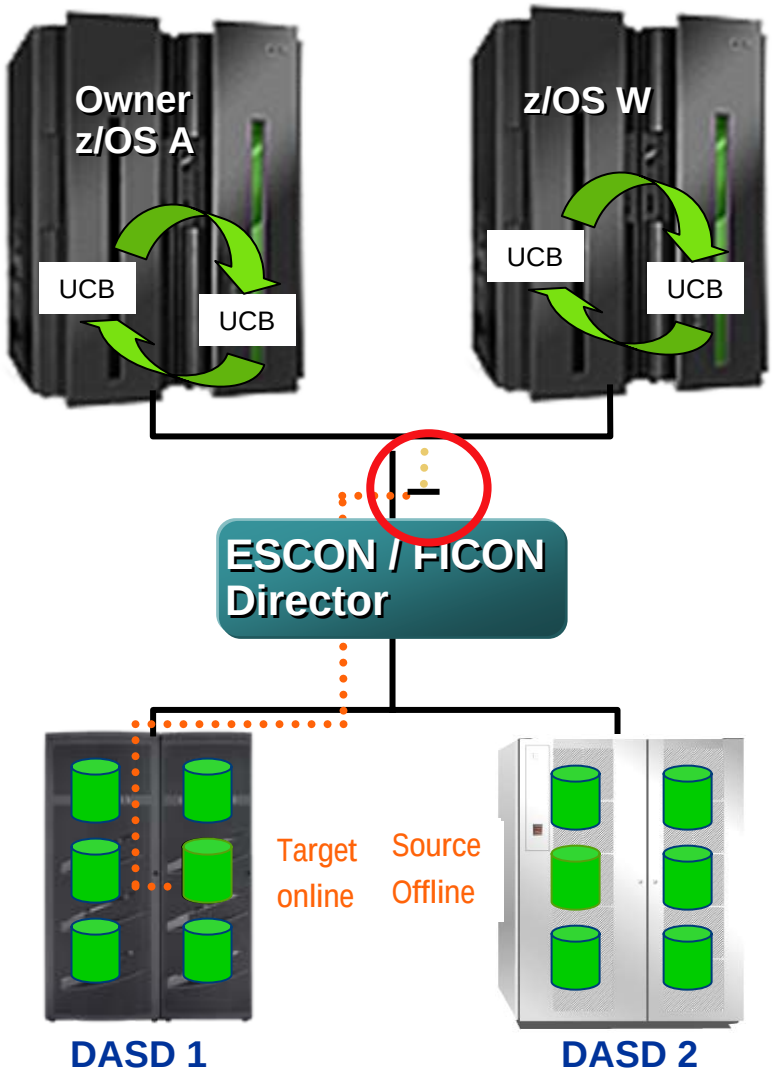

Volume Level Migration
Volume Support Considerations

- Volumes containing the following datasets cannot be migrated:
  - z/OS Migrator database dataset
  - Active Page datasets
  - Active CA-OPS/MVS datasets
  - Active Sysplex Couple datasets
- Volumes not supported as source or targets of a migration
  - Extended Address Volumes
  - FBA volumes
  - Virtual Devices (VDEV)
  - CSC (EMC Cross System Communication task) gatekeeper devices
  - zHPF attached volumes

z/OS Migrator: Dataset Migration

Dataset Migration -- 1

- Multi-phased process
  - Initialization
  - Activation
  - Copy

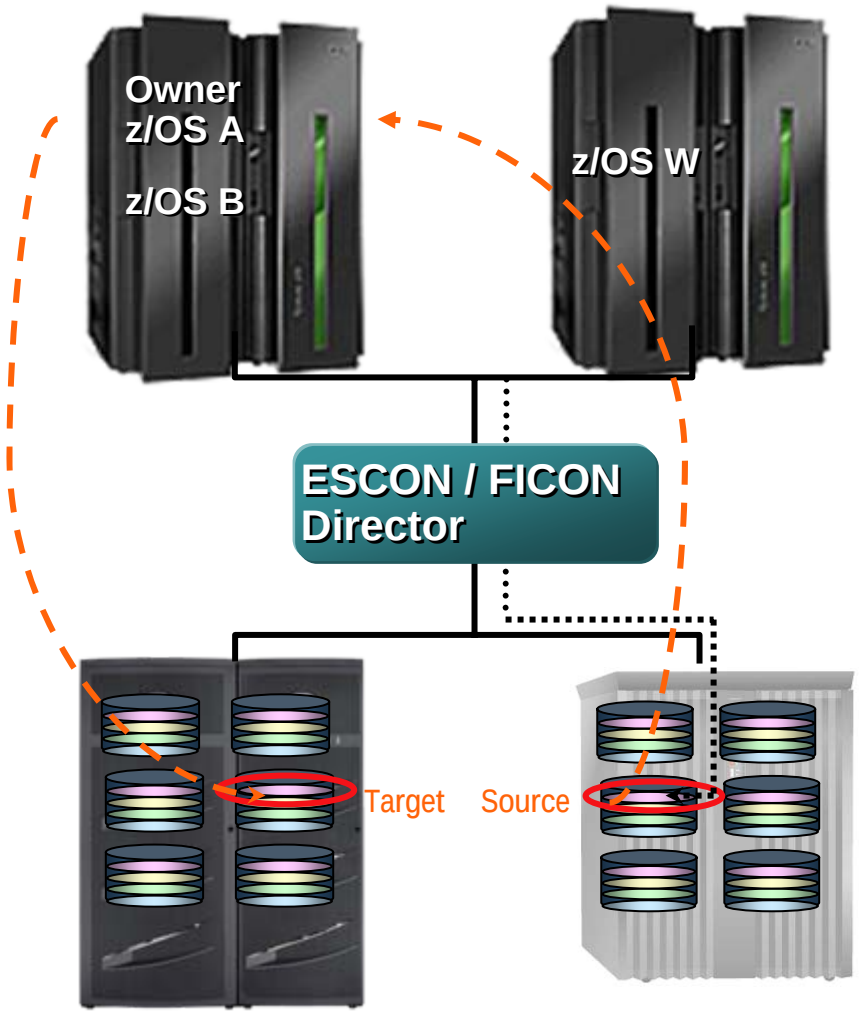

Dataset Migration -- 2

- Multi-phased process
  - Definition
  - Activation
  - Copy
  - Synchronization
  - Mirroring

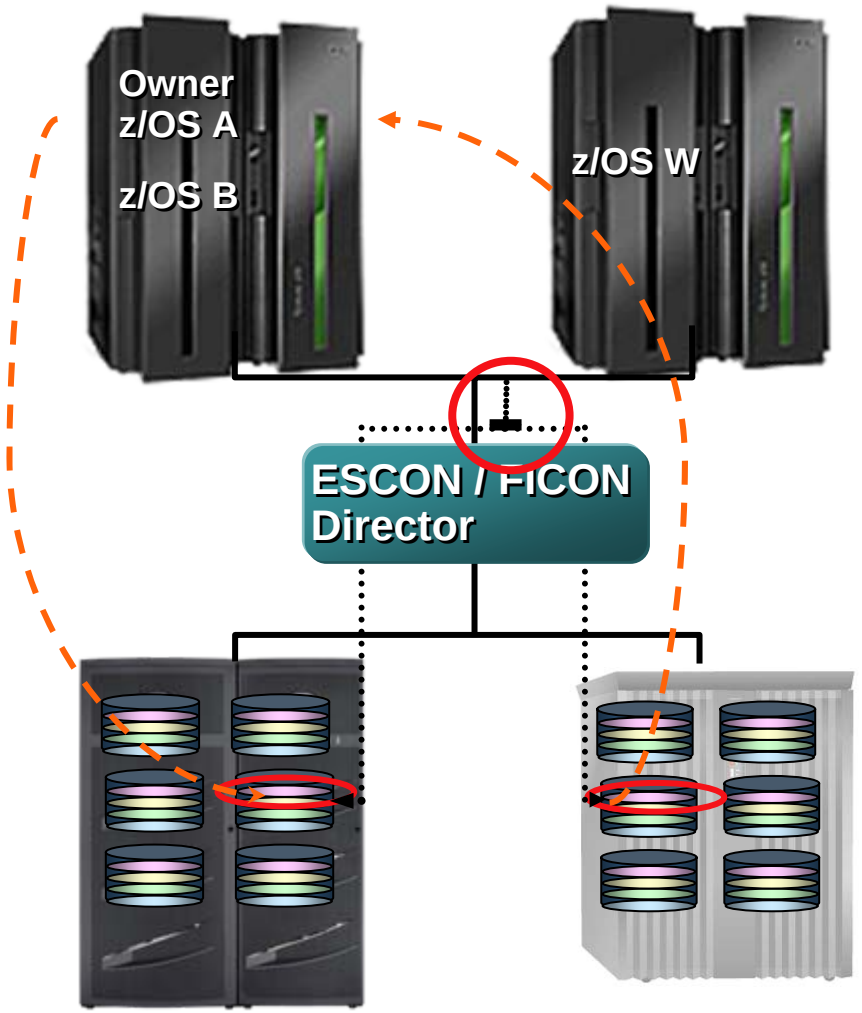

Dataset Migration -- 3

- Multi-phased process
  - Definition
  - Activation
  - Copy
  - Synchronization
  - Mirroring
  - Diversion

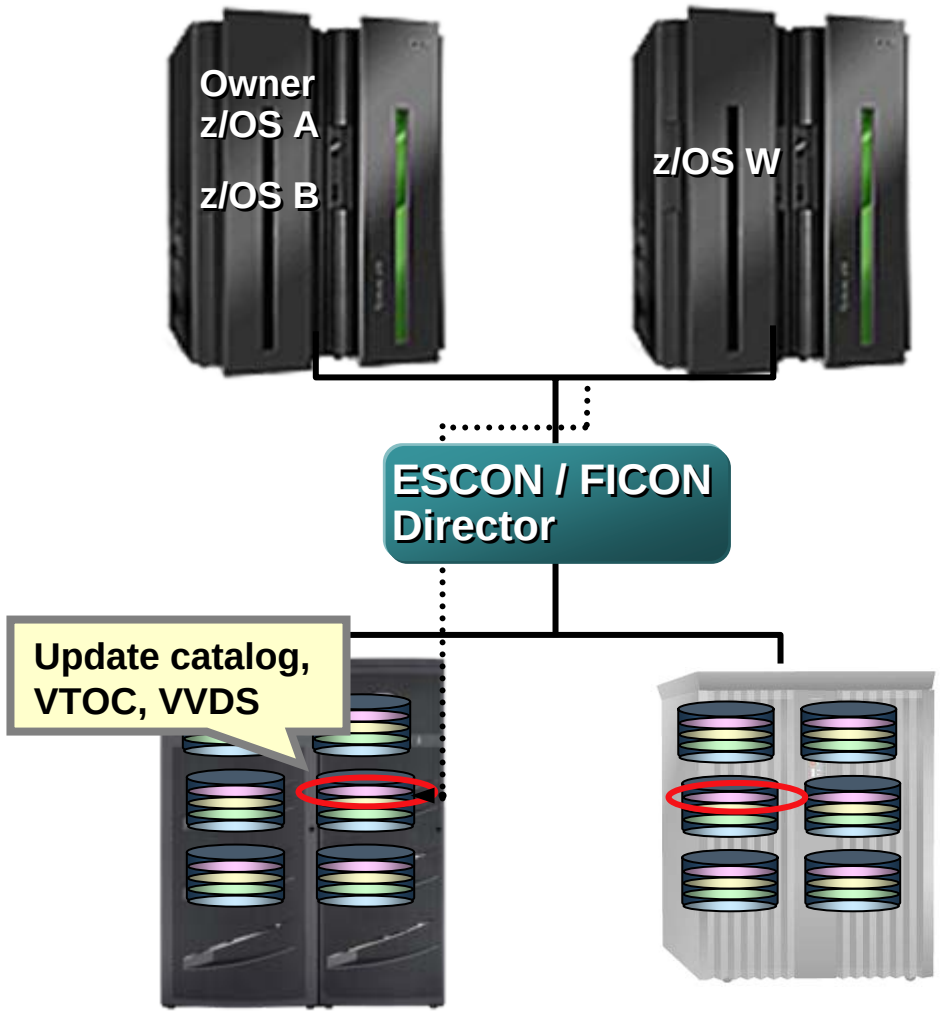

Dataset Migration -- 4

- Multi-phased process
  - Definition
  - Activation
  - Copy
  - Synchronization
  - Mirroring
  - Diversion
  - ‘Active’ Completion
  - Completion

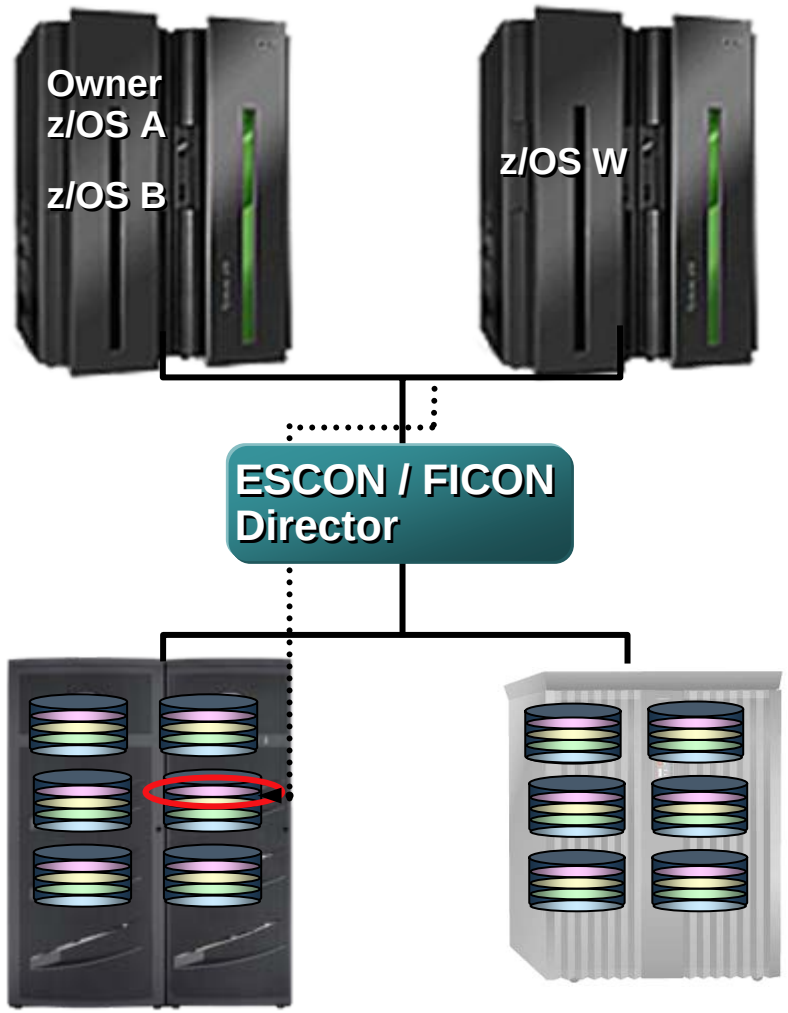

Logical Migration
Dataset Support Considerations

Supported Datasets

- Direct Access (DA) datasets
- Extended format sequential datasets
- Extended Partitioned datasets (PDSE)
- GDG base names and GDG datasets
- Partitioned (PO) datasets
- Physical Sequential (PS) datasets
- Striped Sequential datasets
- BDAM datasets
- Extended format VSAM KSDS
- VSAM datasets
  - Alternate Index (AIX)
  - ESDS
  - KSDS
  - Linear
  - RRDS
  - Spheres
    - KSDS | ESDS + Paths + AIX
  - VRRDS

Unsupported Datasets

- VSAM datasets with the following options specified
  - IMBED
  - KEYRANGE
  - REPLICATE
- Catalogs
- ISAM datasets
- Individual PDS members
- Page and swap datasets
- HFS / zFS datasets
- Unmovable datasets (DSORG=U|PSU)
- VTOC, VTOCIX and VVDS datasets
- Temporary datasets (&&dsn) and other uncataloged datasets

- **NOTE: EAV and zHPF volumes not supported at this time**

Summary:
z/OS Migrator Use Cases

Technology Refresh

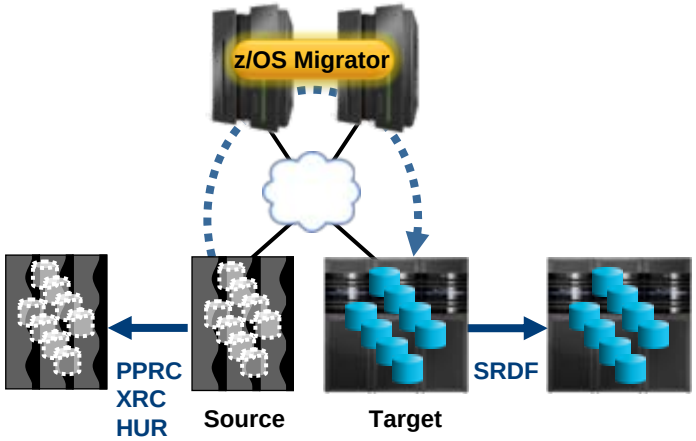

Migrate to arrays non-disruptively to both business continuity solutions and applications

Move from Smaller to Larger Volumes

Reclaim Unit Control Blocks (UCBs)

Storage Consolidation

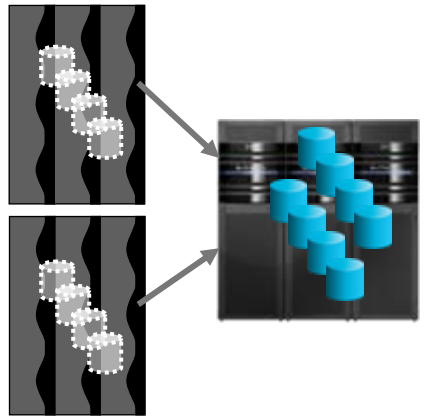

Multiple storage arrays into one high-capacity array

Symmetrix Migrator Software Package

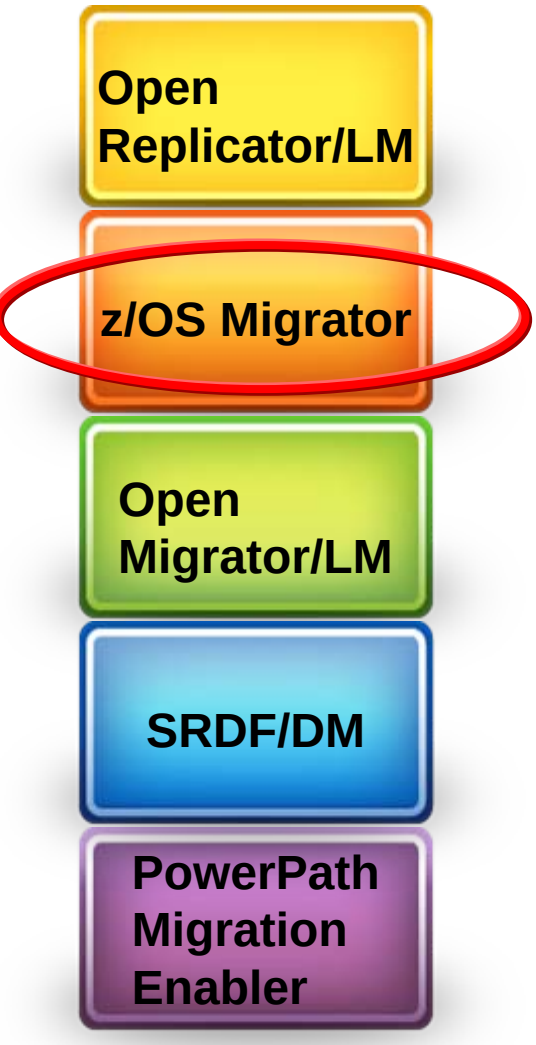

- Available at no-charge to Symmetrix owners
  - Symmetrix DMX and V-Max models
- Deploy the right technology to quickly and efficiently migrate data to where you need it most
  - Host based migrations for smaller projects
  - Network based migrations for environment-wide projects
  - Array based migrations for large scale system consolidation
- Comprehensive environment support
  - Open Systems
  - Mainframe
  - Mixed Open Systems/Mainframe
  - EMC and qualified third party arrays

Choice and flexibility to deliver against your business and IT migration objectives


EMC²
where information lives®