Blue Cross Blue Shield of Minnesota - Replication and DR for Linux on System z

Brian Peterson, Blue Cross Blue Shield of Minnesota
Gail Riley, EMC

March 3, 2011

Objectives

At the end of this session, you will be able to

- Discuss replication in a Linux on System z environment, including z/VM and z/OS
- Describe managing replication from z/OS for z/VM and Linux as a guest virtual machine
- Understand different backup scenarios for z/VM and Linux on System z

Linux on System z
Disk Attachment Options

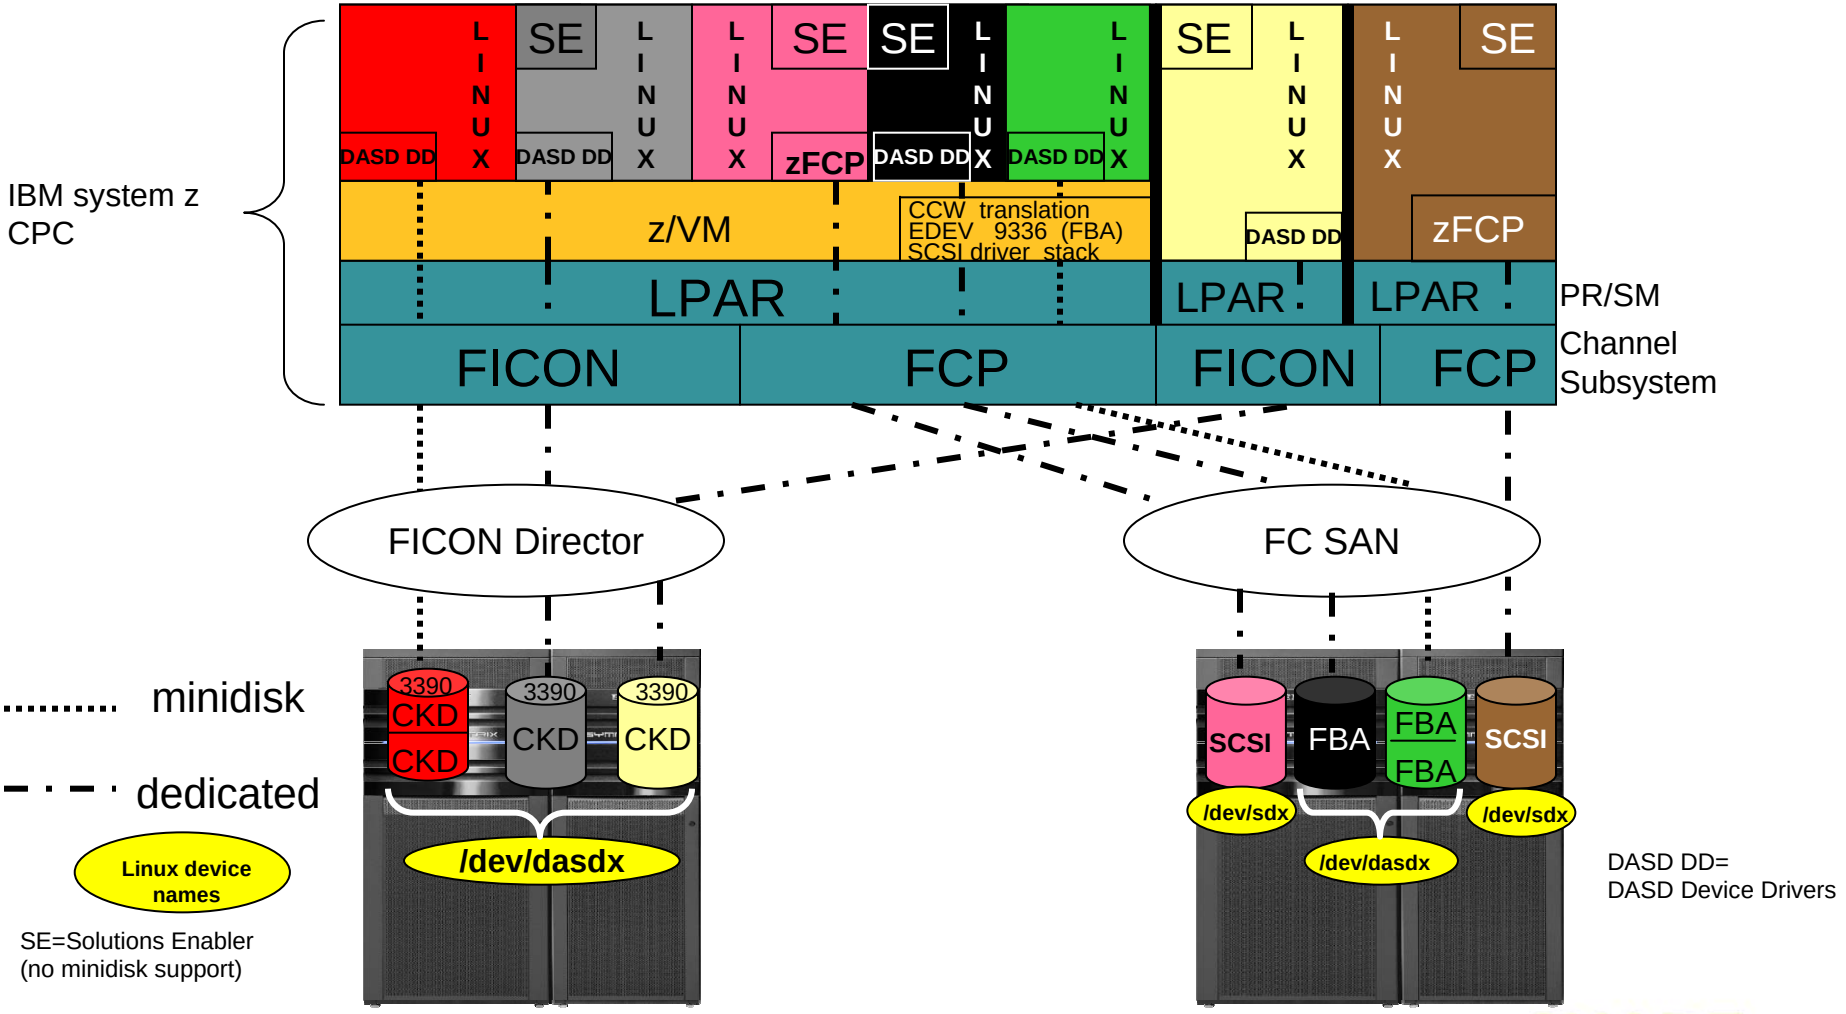

Linux on System z CKD Device Relationship Path

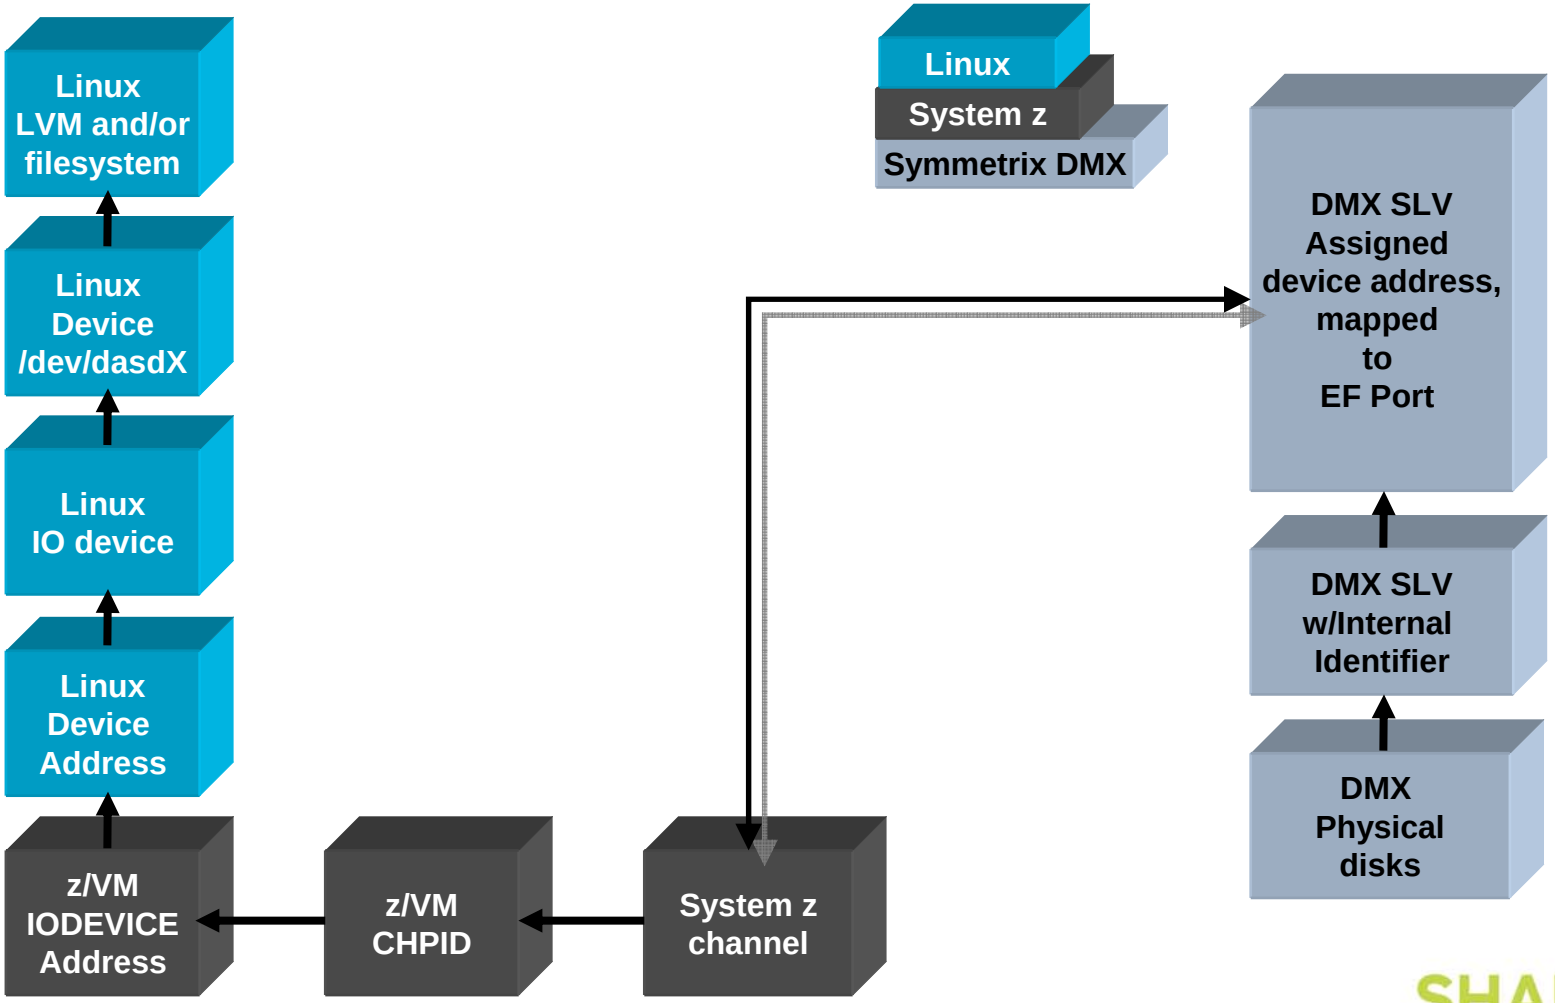

Replication – Business Recovery Tools

- TimeFinder – Local Replication
  - Clone – Full Volume copy, Source device size = Target device size
  - Snap - Pointer Based Replication,
    - Target Device is a virtual device housing a collection of pointers between the Source and a reserve area for a point-in-time view

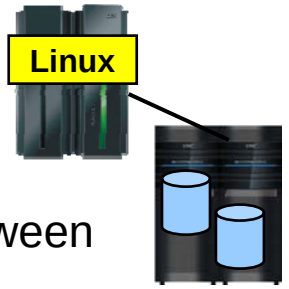

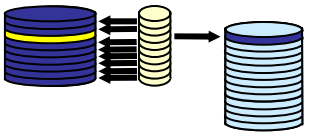

- SRDF – Remote Replication
  - Allows the movement of data between storage systems in the same room, to different buildings located across town, or thousands of miles apart
  - Offers various disaster recovery levels
  - Enables the following operations:
    - Disaster recovery, Disaster restart testing
    - Recovery from planned outages, Remote backup
    - Data center migration, Data replication and mobility

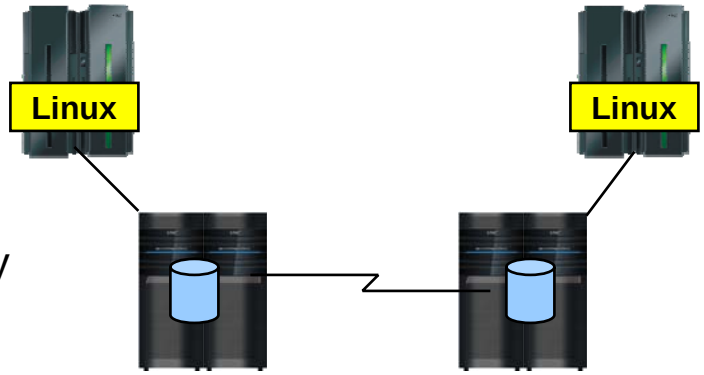

Symmetrix Remote Data Facility: Two Site Solutions

SRDF/Synchronous

- No data exposure
- Some performance impact
- Limited distance
- Source = Target

SRDF/Asynchronous

- Predictable RPO
- No performance impact
- Unlimited distance
- Only two copies of data required
- Source ≅ Target

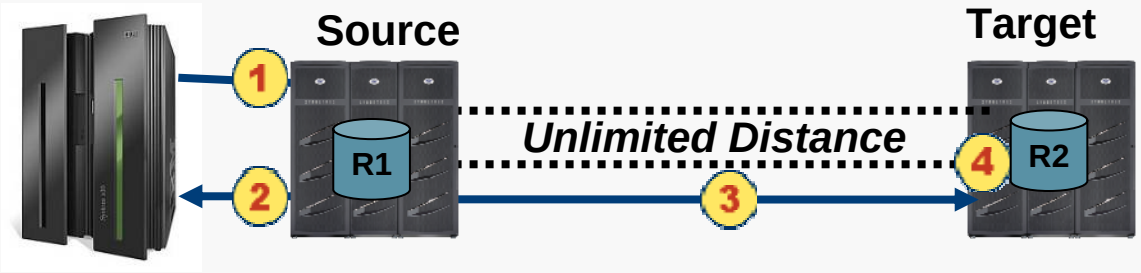

SRDF/AR

- Data Movement solution
- No performance impact
- Unlimited distance

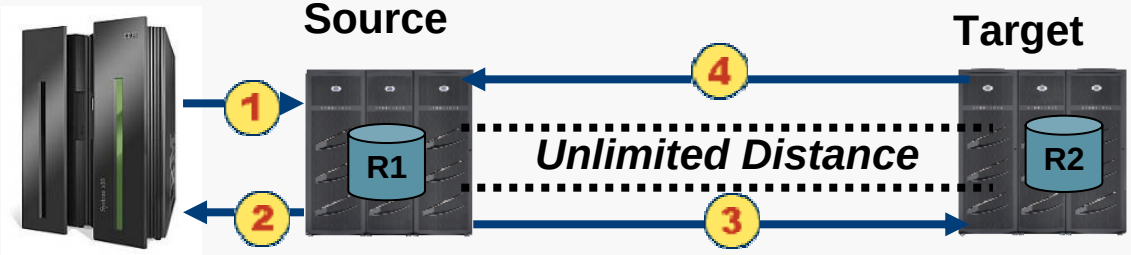

EMC Replication Management Options

- Solutions Enabler – 7.1 and later
  - Linux on System z – SUSE 10/11 and Red Hat 5.4+
  - Open Systems hosts – Windows, Linux, UNIX
- Mainframe Enablers
  - z/OS
- Symmetrix Management Console (SMC)
  - Windows, Linux (x86), UNIX
- EMC z/OS Storage Manager (EzSM)
  - ISPF-like panel menu interface
- EMC products for TPF
    - TimeFinder Controls for TPF
    - SRDF Controls for TPF
    - ResourcePak for TPF

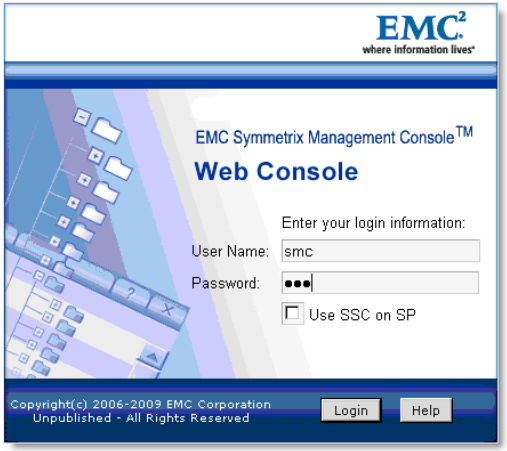

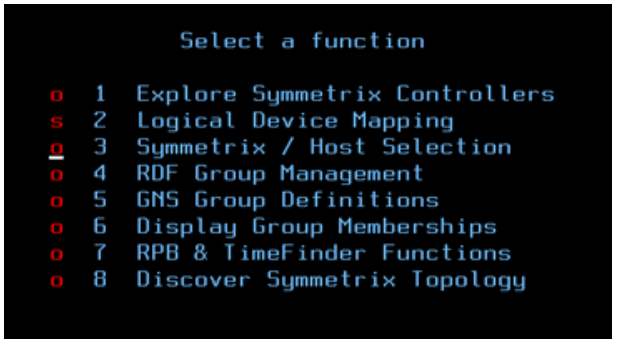

Blue Cross Blue Shield of Minnesota

Blue Cross Blue Shield Hardware - Today

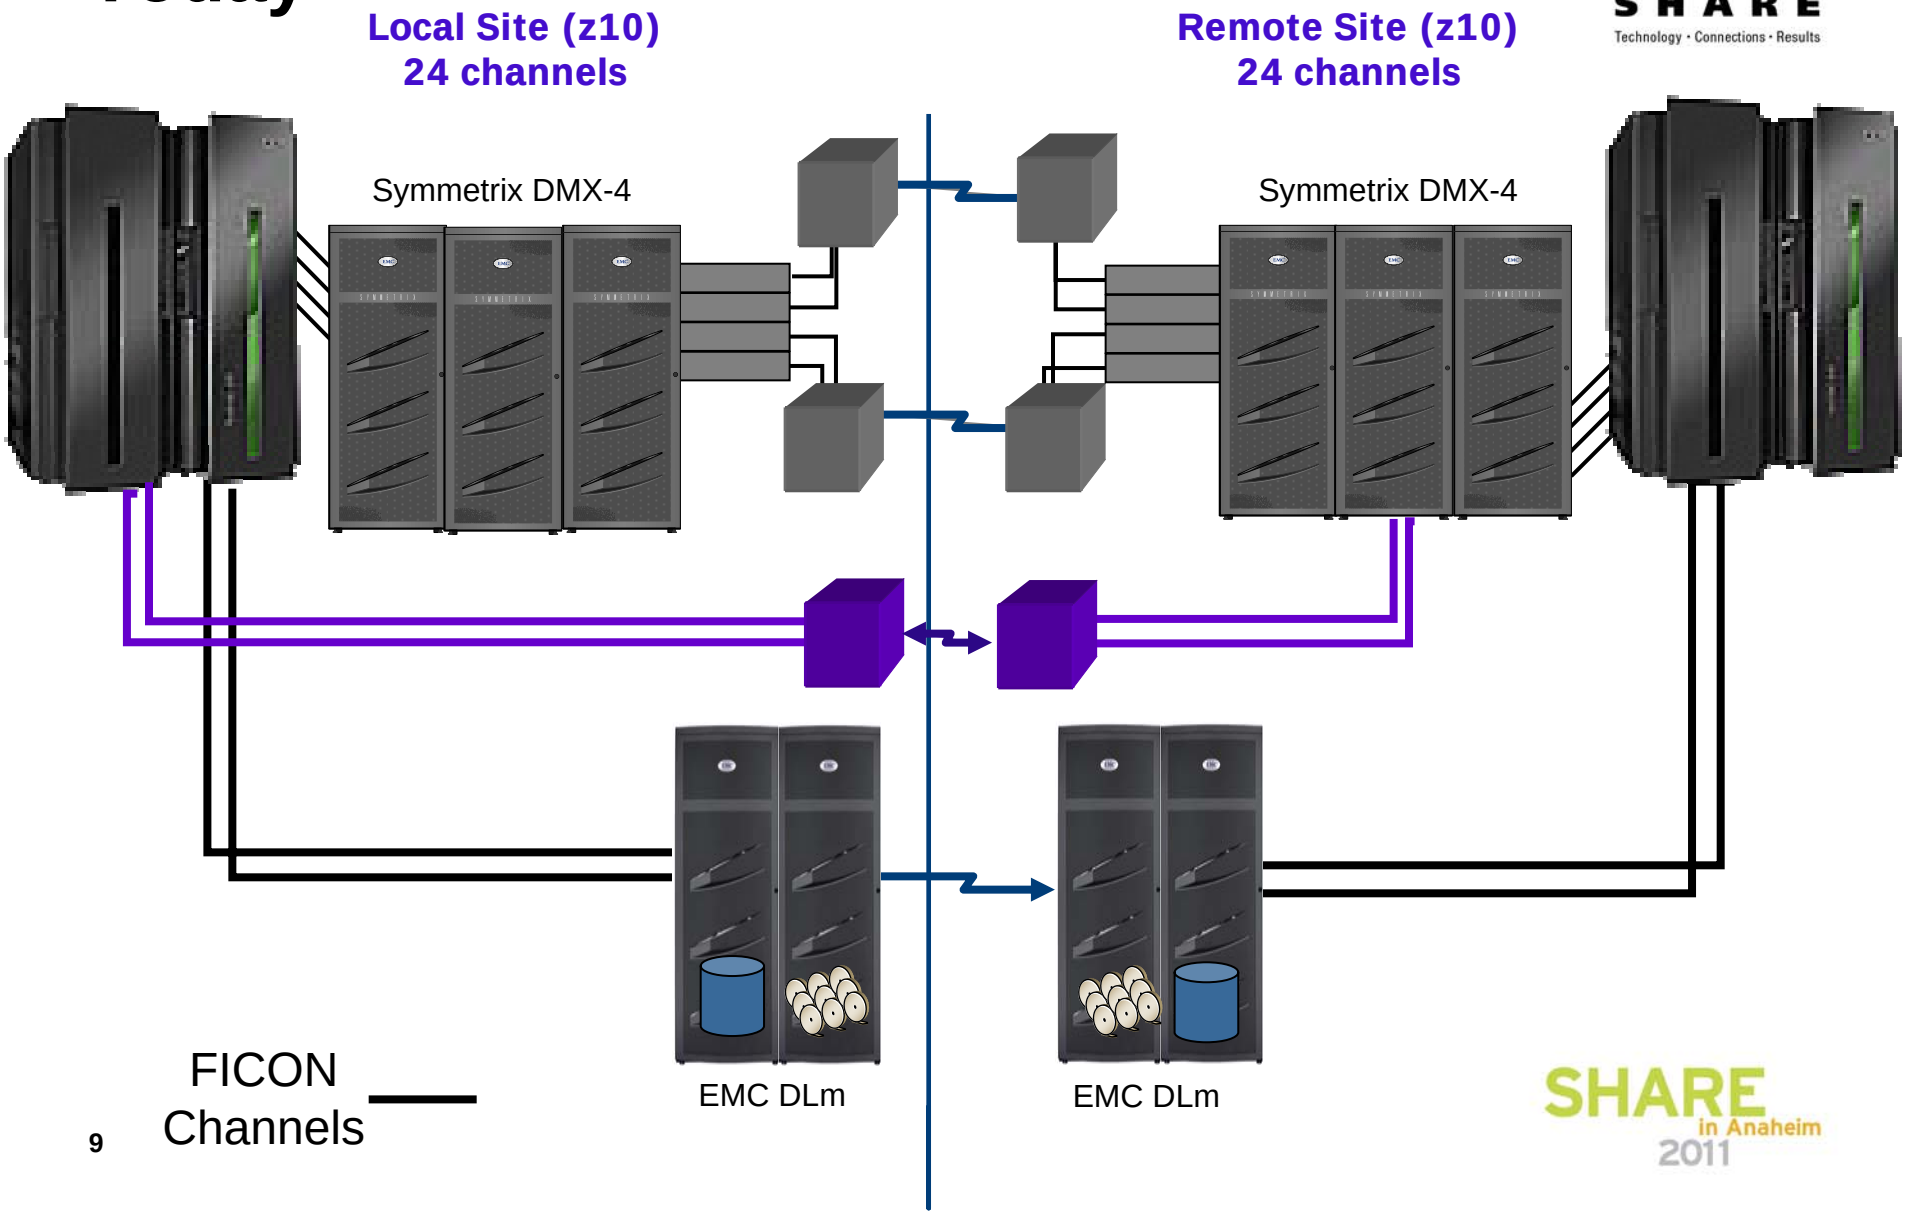

z/OS Environment

- z/OS LPARs – 7 Total
  - Production - 3
  - Test – 2
  - Backup - 1
  - Sandbox - 1
- z/OS 1.11
- CICS, IMS, DB2, SAP
- Symmetrix
  - CKD – mod-3, 9, and 27
- Remote Site
  - 14 z/OS LPARs normally deactivated

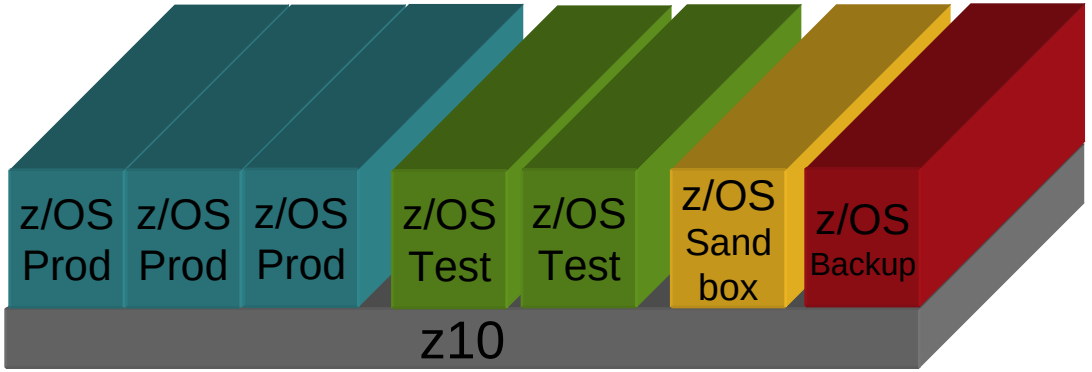

Local

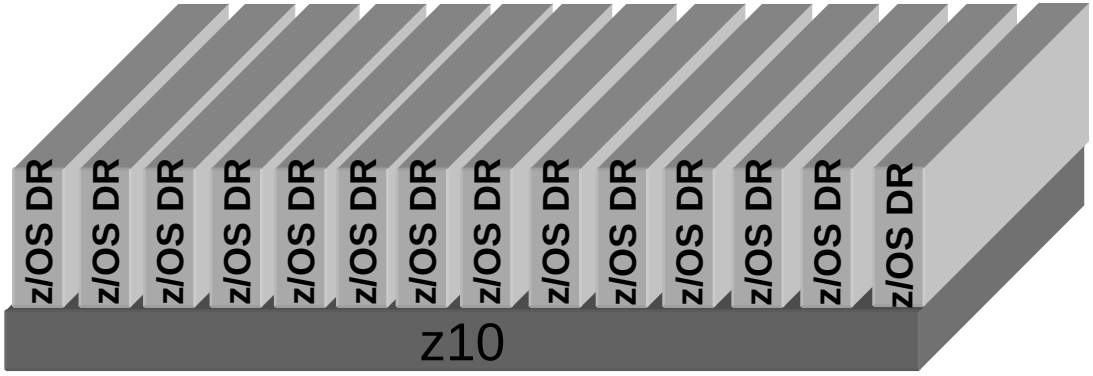

Remote

z/VM Environment

- z/VM LPARs – 5 Total
  - Production - 3
  - Test – 2
- z/VM 6.1
- 200 Linux virtual machines across 7 LPARs and 2 sites
- Symmetrix
  - CKD only
  - All minidisks, no dedicated devices to guest virtual machines
- No backup software lives on z/VM

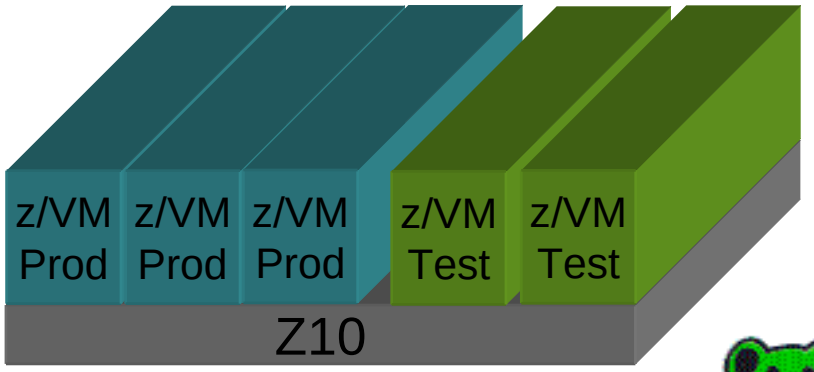

Local

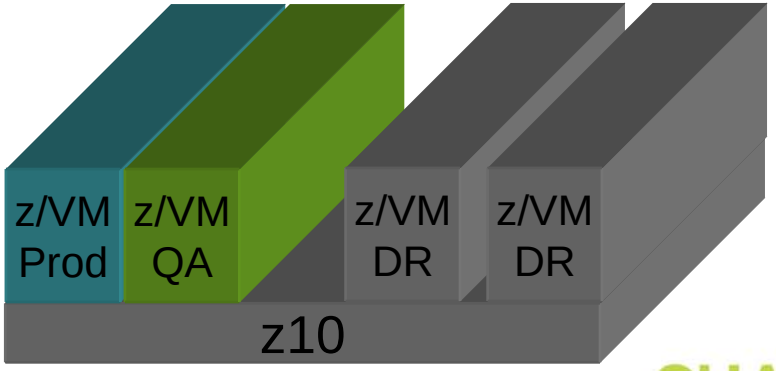

Remote

Linux Environment

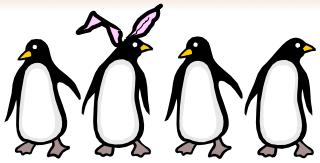

- All Linux under z/VM
- Linux is on CKD minidisks
  - Full pack
  - Partial pack
- SLES 10 SP3
- 200 Linux guests and growing
  - subset of guests are active/active
- TSM backup agents on Linux
- Database communication to z/OS
  - Hipersocket
  - OSA

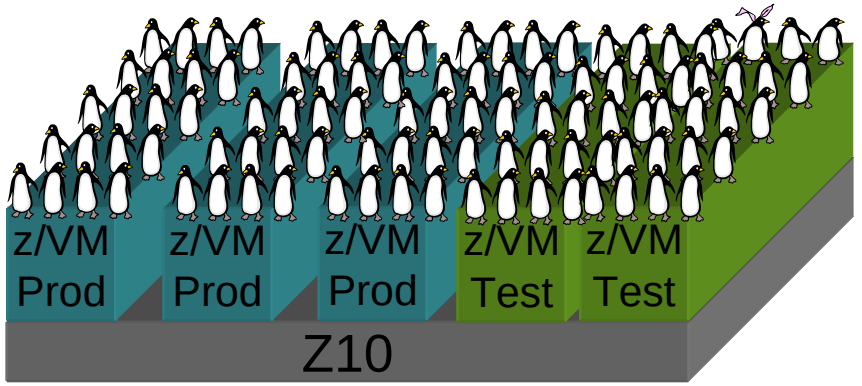

Local

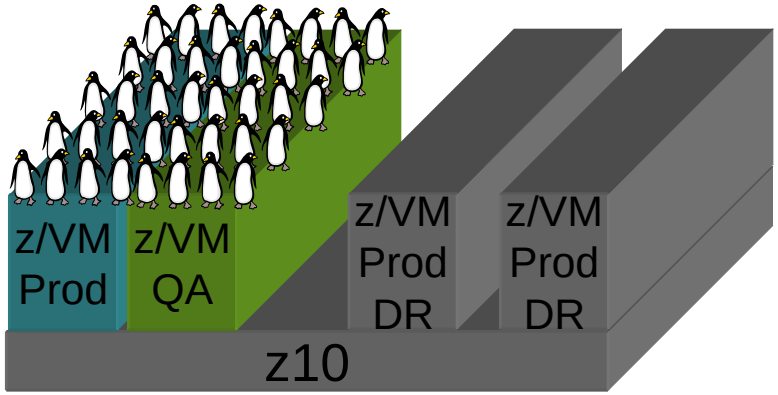

Remote

Why Linux on System z?

- Prototyping Linux on System z for years
- Traction took hold 3-4 years ago when implementing SAP
- Moved it from x86 talking to z/OS DB2 to mostly Linux on z talking to z/OS DB2
- This was our first production scale Linux application implemented
- Very successful!
- Continuing to look at additional workload
  - WAS

Replication

Disaster Recovery Environment - 2005

One Site (z900s)

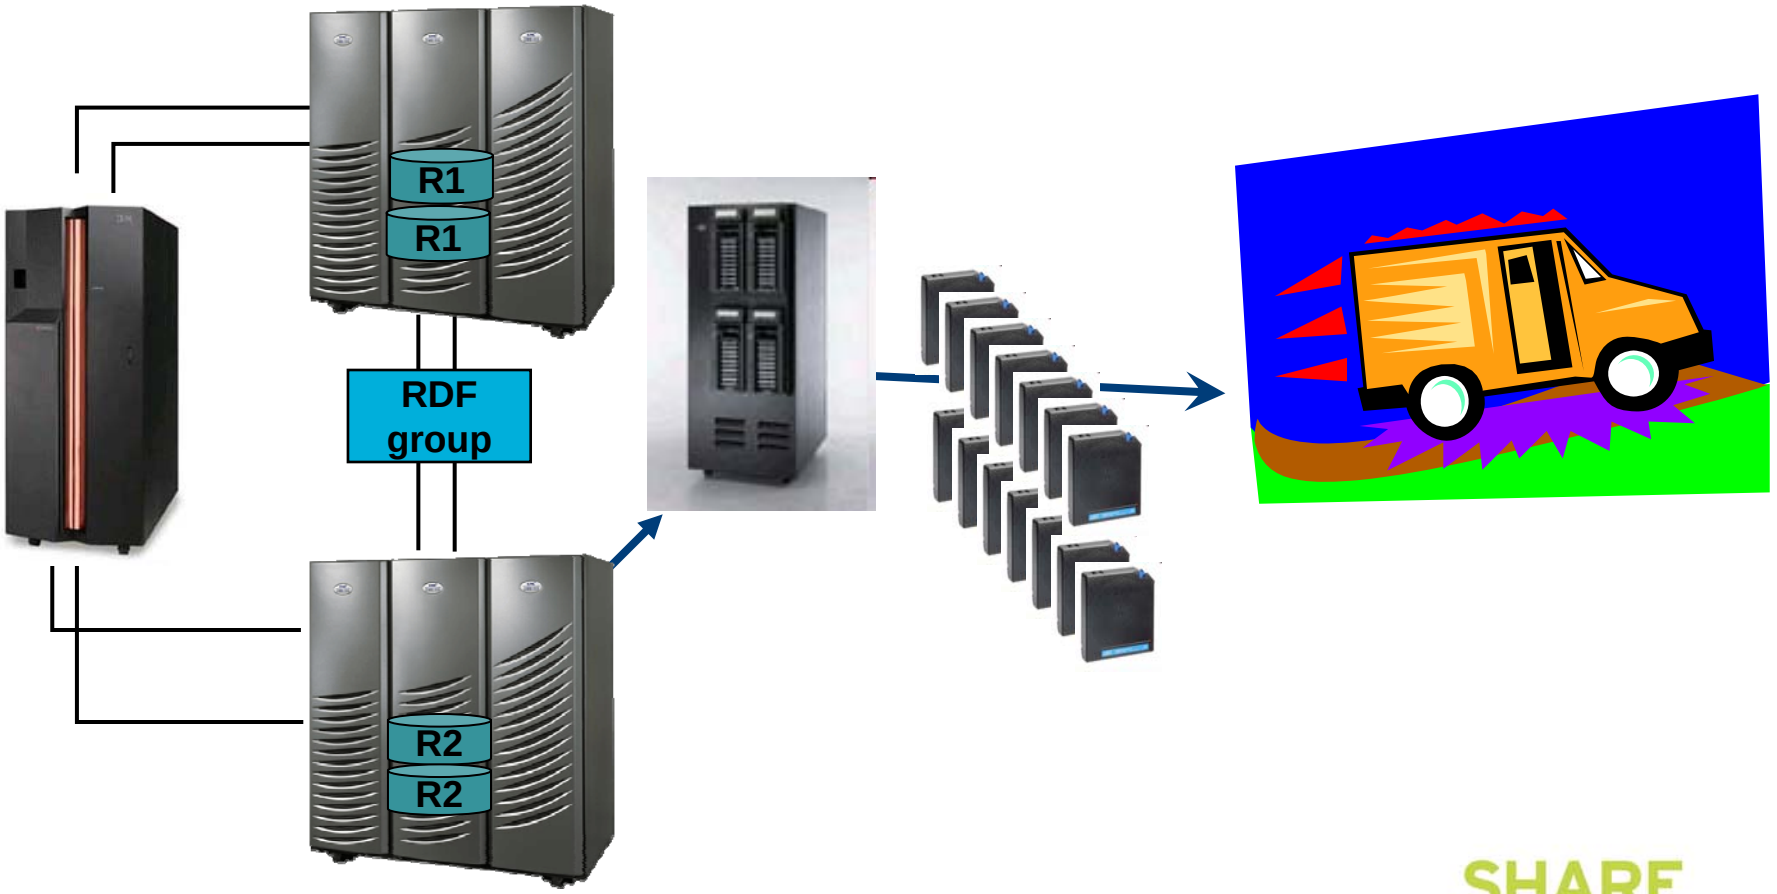

SRDF - 2005

- One RDF Group in 2005
  - Included everything - z/OS, z/VM and Linux on System z
  - Normal operation – SRDF AR from DMX3000 to DMX3000
  - Strictly for Disaster Recovery
- Backups occurred each 24 hour period via a script
  - Switch to SRDF Synchronous mode which enforces consistency across all devices maintaining application interrelationship consistency
  - Once invalid tracks reached zero, performed ConGroup trip, splitting off all R2s
  - RDF Group is suspended
  - Performed backup to tape
  - Switch back to Asynchronous

The Future Beckons…

What is the right path for our Disaster Recovery ?

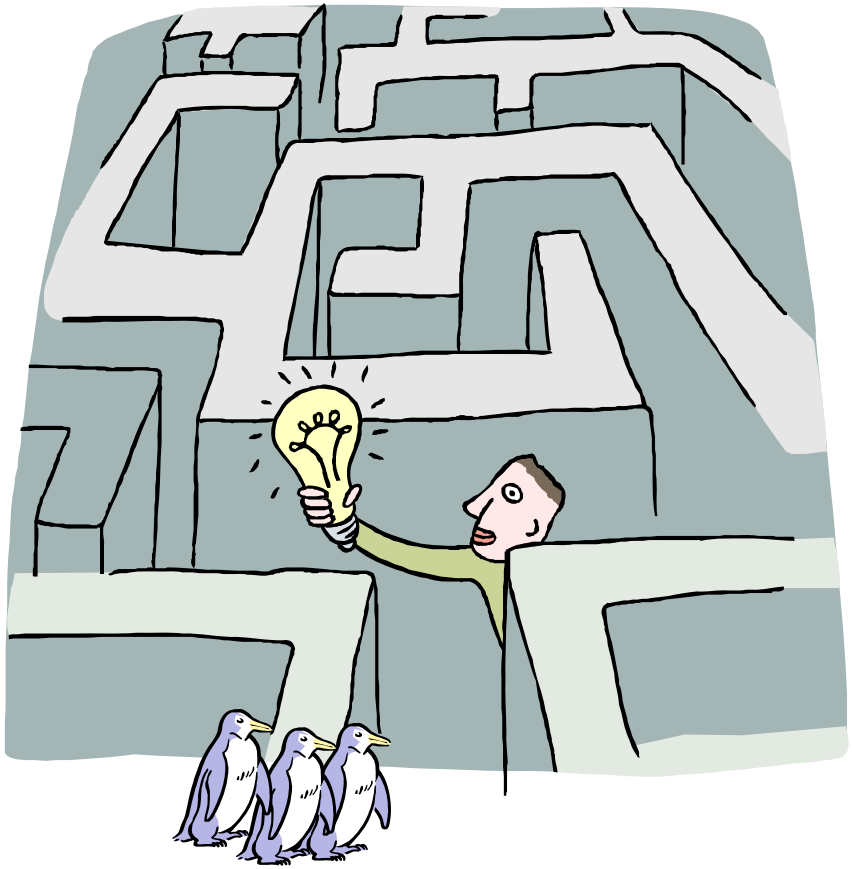

Remote Replication Environment - 2011

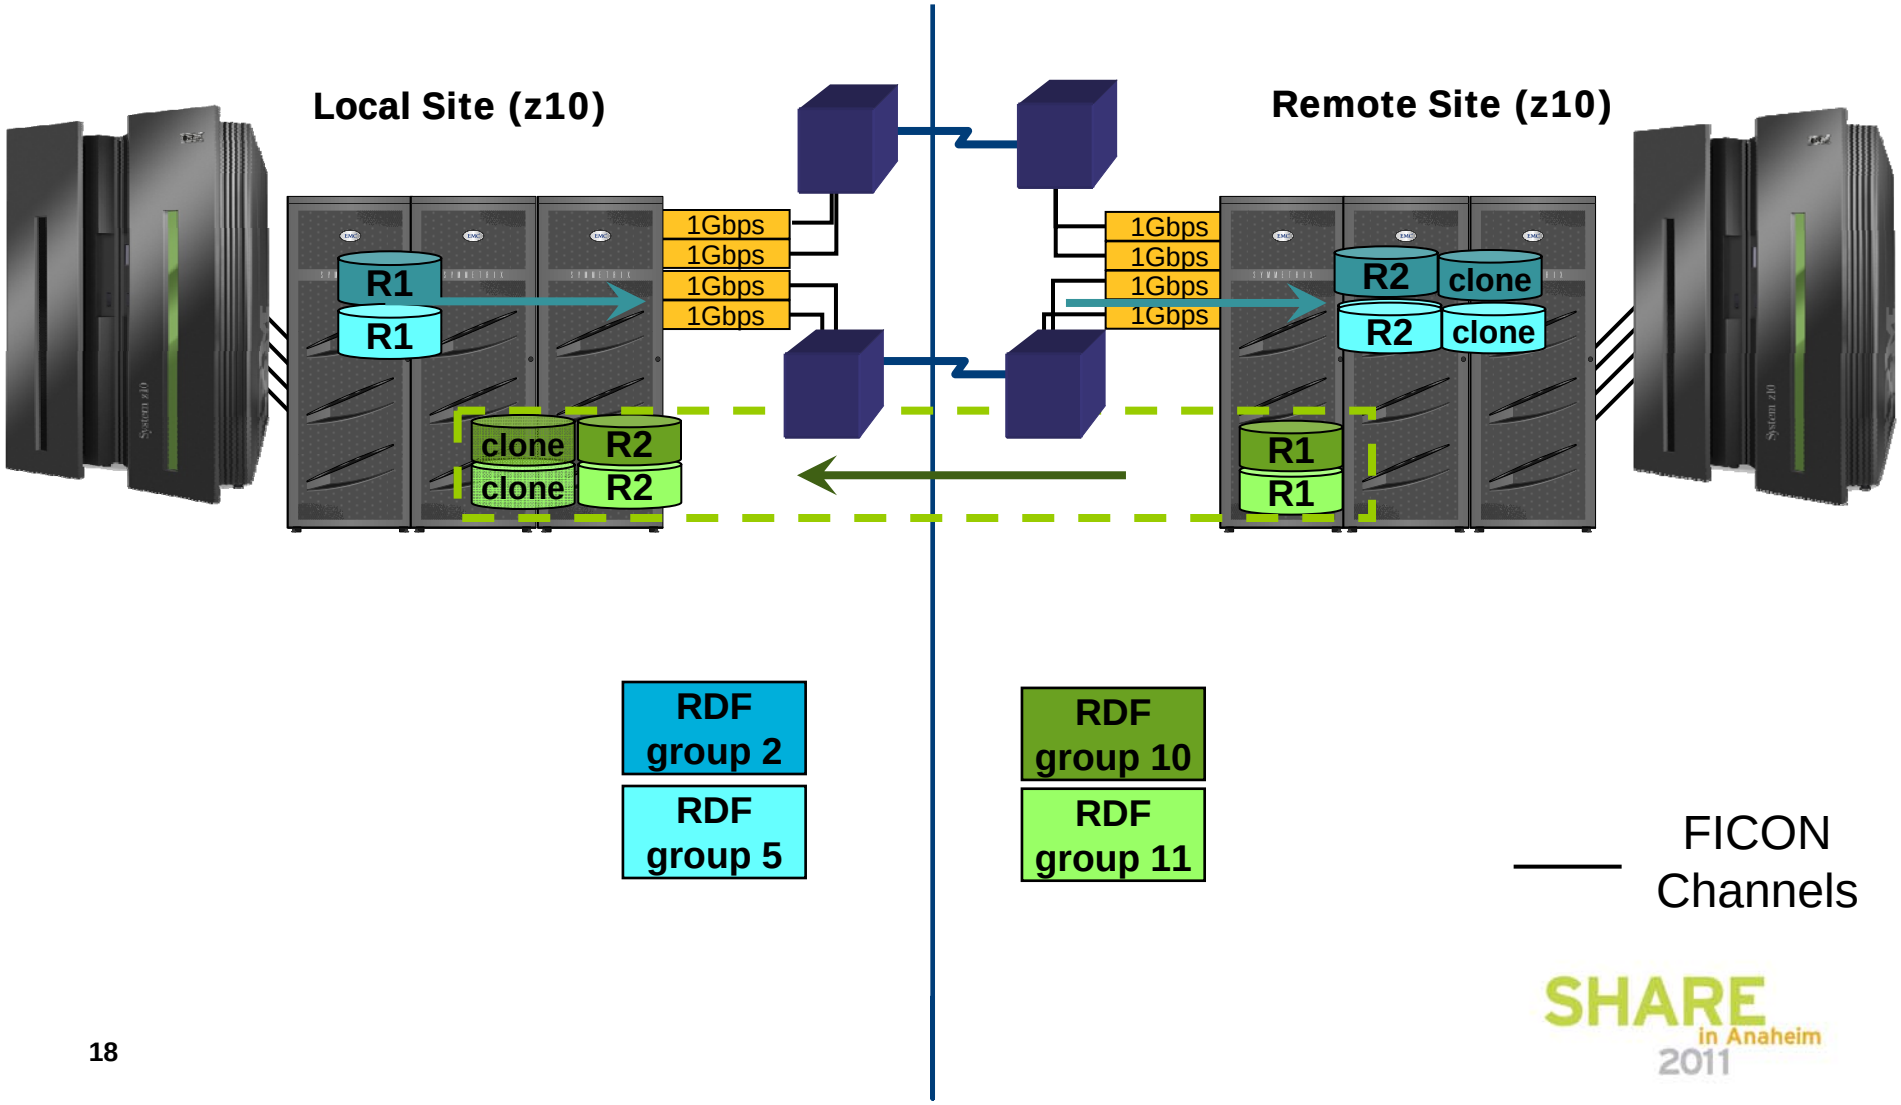

SRDF Groups – Local to Remote Site Setup

- RDF Group 2
  - SRDF/A (asynchronous) from Local to Remote Site
  - Includes z/OS, z/VM and Linux on System z
  - Consistency maintained across application environment
  - ~53TB

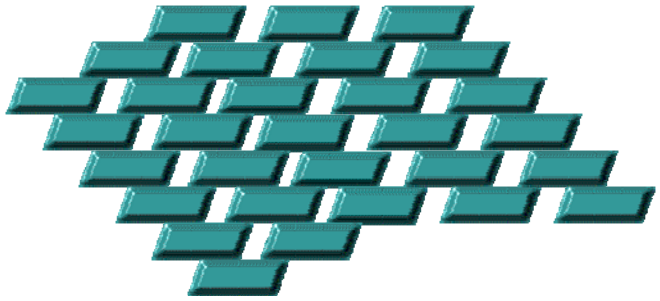

- RDF Group 5
  - Normally suspended
  - Includes z/OS, z/VM, Linux “Work” devices – page, swap, etc.
  - Data structures of volumes required, but not day-to-day data
  - Synchronized when changes are made
  - Data Distribution mode from Local to Remote Site
  - ~ 2TB

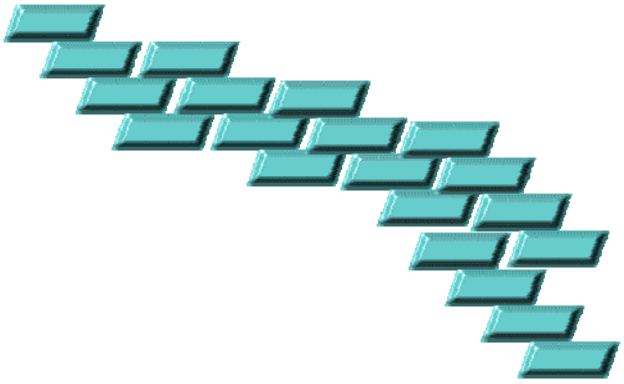

SRDF Groups – Remote to Local Site Setup

- RDF Group 10
  - SRDF/A (asynchronous) from Local to Remote Site
  - z/VM and Linux on System z
  - Consistency maintained across application environment
  - ~4TB
- RDF Group 11
  - Normally suspended
  - Includes z/VM, Linux “Work” devices – page, swap, etc.
  - Data structures of volumes required, but not day-to-day data
  - Synchronized when changes are made
  - Data Distribution mode from Local to Remote Site
  - ~ 0.5 TB

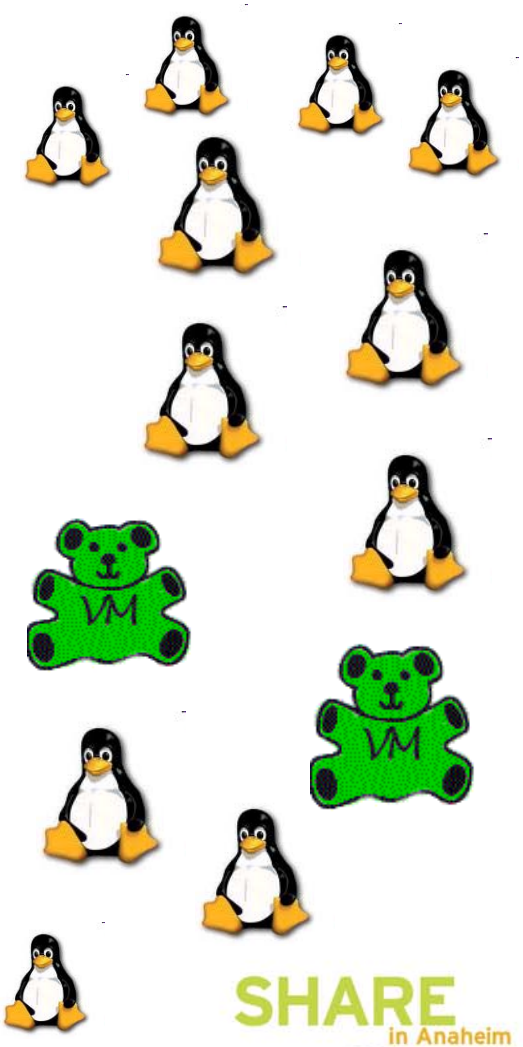

Recovery Objectives

- RPO and RTO in 2005
  - RPO was 3-4 days
  - RTO was ~ 72 hours
- RPO and RTO in 2008
  - RPO 30-60 seconds normally
    - 2-3 minutes, worst case
  - RTO 2-4 hours
- Includes z/OS, z/VM, Linux
- How often it's tested?
  - Reduced Test Set
    - Test at will –when changes occur with minimal personnel
  - Corporate Participation
    - Two times a year

Recovery Considerations

- z/OS and Linux considerations
  - SAP – distributed application across z/OS and Linux
  - Application environment can be spread across many Linux instances
  - Consistency maintained via SRDF group
  - z/VM environment is also maintained
- Currently manual failover and reintegration between site
- Exploring GDDR for future automation

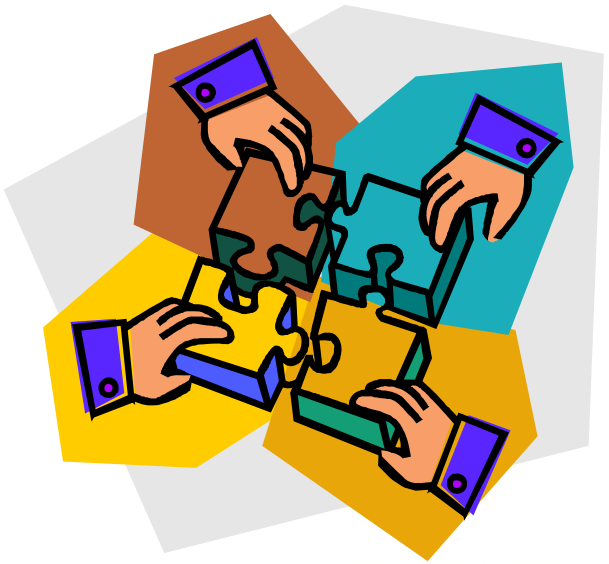

Backups for z/OS, z/VM and Linux on System z

Backups from z/OS

- Backups taken from local and remote TF clones
- Point in Time backups take 3 times per week from z/OS which includes
  - z/OS
  - z/VM
  - Linux on System z – contained in z/VM minidisks
- All z/VM R2 Clone devices are seen from the z/OS backup LPAR
  - Seen as 100% allocated, no free space
  - No datasets
- Backup products
  - DFDSS
    - use CPVOLUME option to backup z/VM DASD
  - DBS – DASD Backup Supervisor from OpenTech Systems
- 3592 tape drives

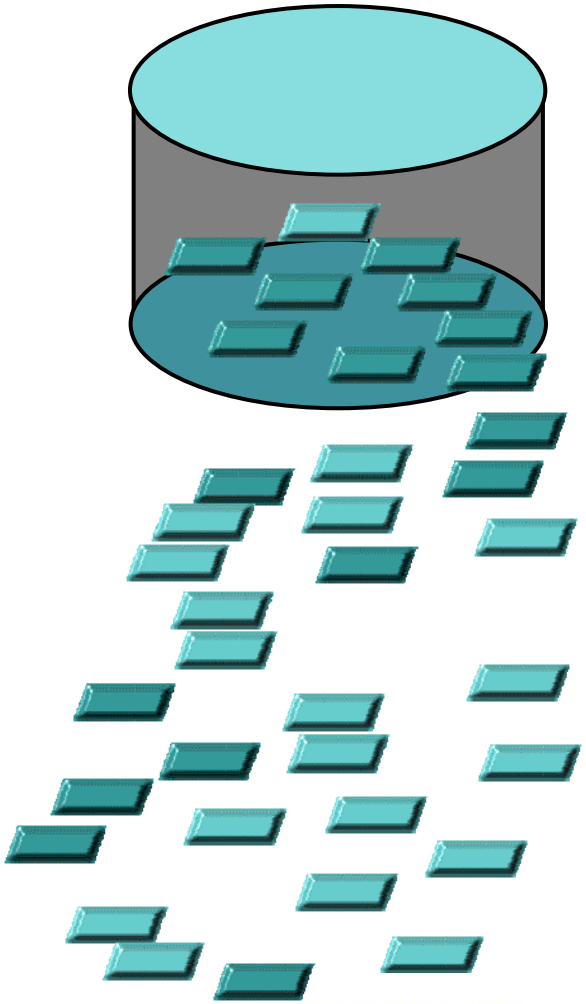

Backup Environment - 2011

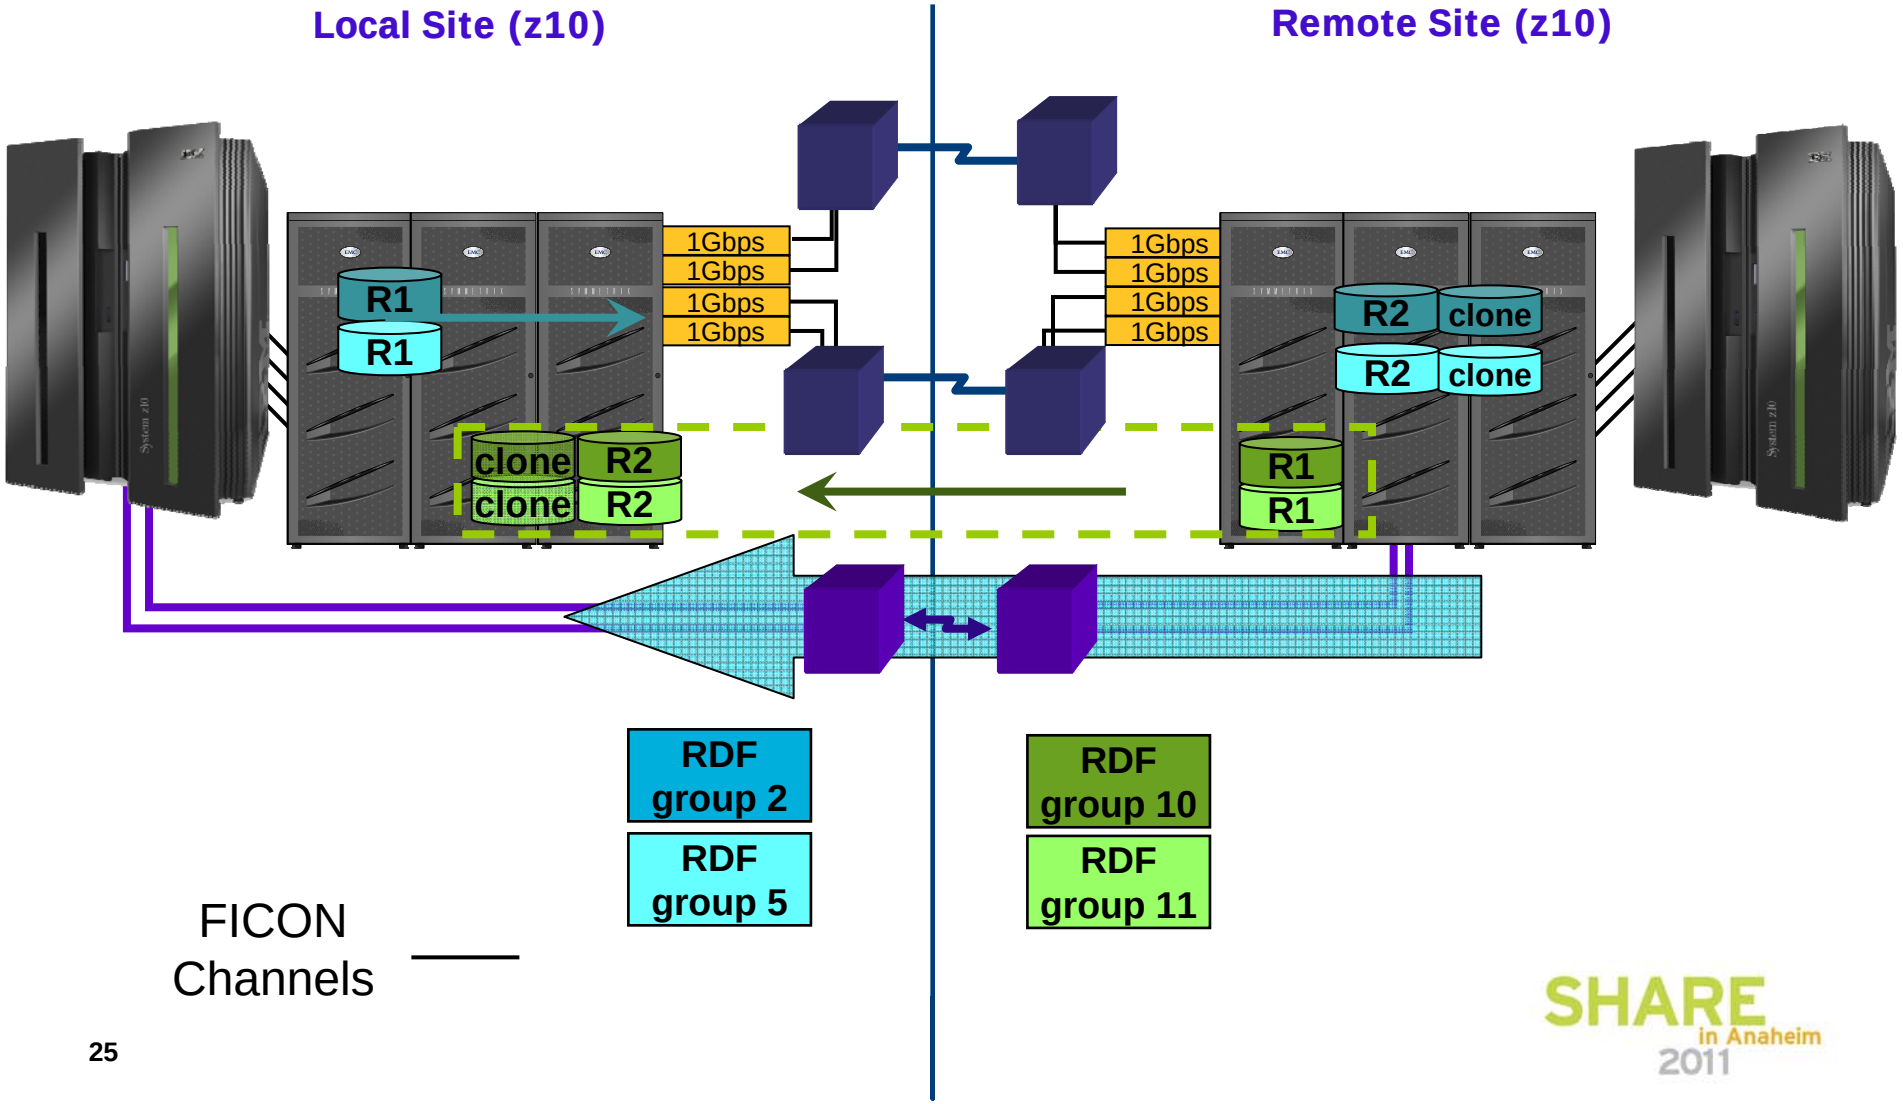

WAN Backup – Remote Clone to Local Tape Drives

- Split Clone volumes in remote DMX-4
- Bring remote Clone volumes online to local z10
- Run backup jobs
- Bring remote Clone volumes offline to local z10
- Re-Establish Clone volumes in remote DMX-4

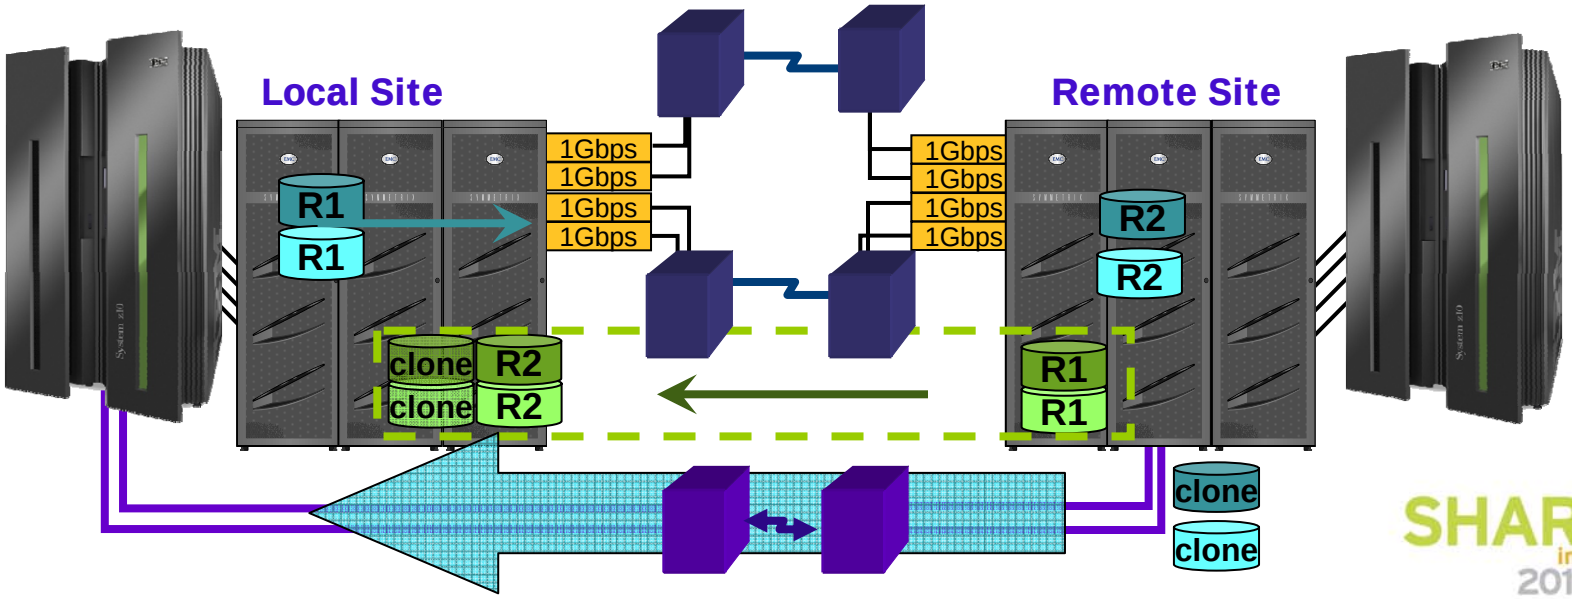

Restoration Options from z/OS

- z/OS
  - Full volume
  - Datasets
- z/VM and Linux on System z
  - Full volume
  - Minidisk - restoration via specific cylinder range
- Restores full volumes encompassing Linux environment
  - Ability to access minidisk, restart Linux and pull information if necessary

Advantages of "WAN Backup"

- Allows us to dual-purpose remote Clone volumes
- Original purpose was DR Tests
- Now used for DR Tests as well as three-times-per-week backups to tape
- Uses network "white space" as most replication traffic was Local --> Remote, while "WAN Backup" traffic is Remote --> Local

Backups from Linux on System z

- Tivoli Storage Manager agents live on Linux
- Incremental file level backups performed nightly
- Tivoli Storage Manager Server lives external to the mainframe environment
  - Current Tivoli Storage Manager 5.4.4
  - Converting to TSM 5.5.4.1
  - Investigating moving to Linux on System z
- Advantage of Linux backups
  - Ability of Linux administrators to restore individual files

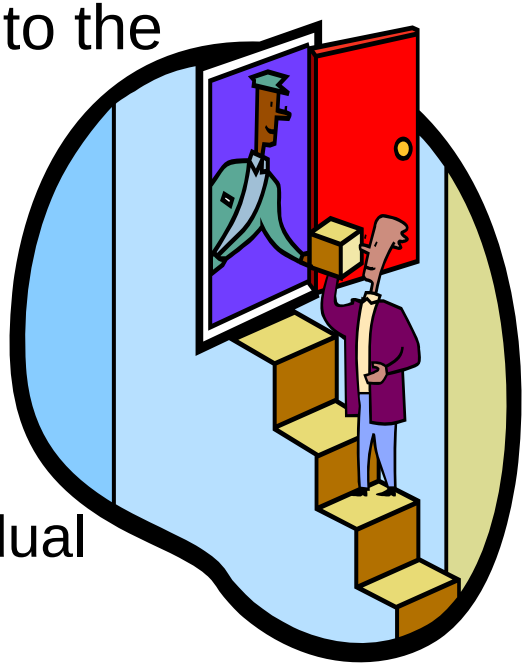

Replication Management

Replication Management from z/OS

- z/OS using EMC Software:
  - ResourcePak Base 5.8
  - SRDF 5.6
  - TimeFinder/Clone Snap 5.8
  - TimeFinder Mirror 5.6
  - TimeFinder Utilities 5.4
- Migrating to Mainframe Enablers 1Q2011

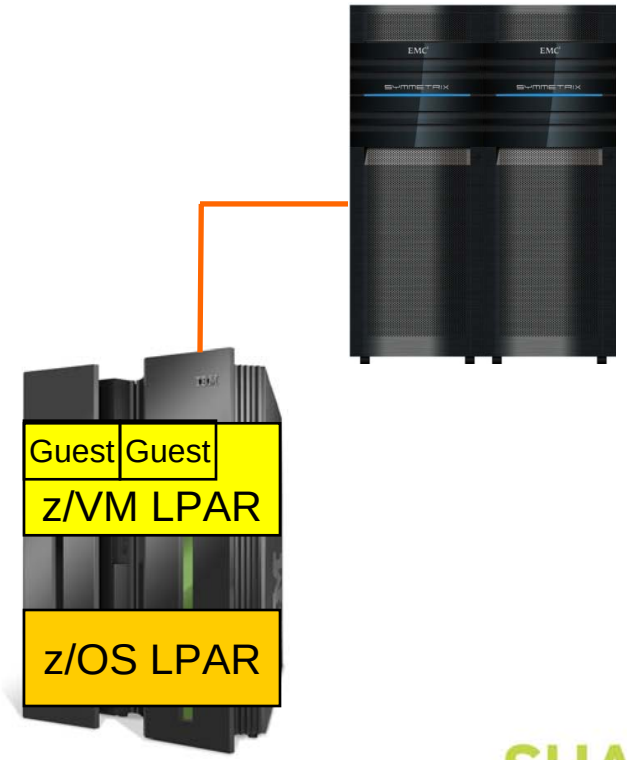

Future Plans

- Evaluating Linux next steps
- Planning upgrade to z/OS 1.12 Summer 2011
- EMC GDDR
- Mainframe Enablers
- Investigating other applications to move to Linux
- Symmetrix VMAX

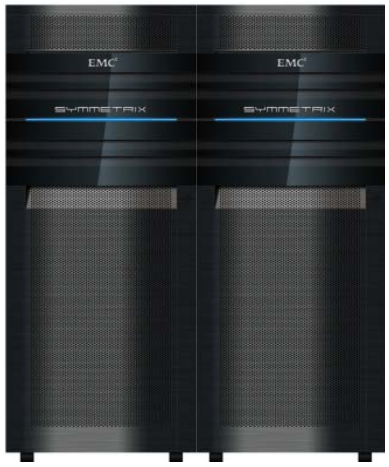

Related EMC Technical Documentation

- *White paper: Configuring EMC Symmetrix arrays for Linux-on-System z*
- *Linux on IBM System z: RHEL 5.x and SLES 10.x Installation and Configuration Guide, P/N 300-007-955, REV A01*
- *EMC Solutions Enabler Installation Guide*