Managing and Monitoring Symmetrix for Linux on System z

Gail Riley
EMC Corporation

August 26, 2009
Session 9165

Objectives

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

- Describe the considerations needed for Linux on System z with EMC® Symmetrix® storage
- Define the drive emulation that best serves your Linux on System z environment
- Perform storage monitoring and management operations of a Symmetrix system in a Linux on System z environment

Agenda

- Discuss the basic tasks needed to configure Linux on System z to run with an EMC Symmetrix DMX™ system
- Discuss FBA and CKD relationship to z/VM and Linux
- Describe the processes available for managing Linux on System z with EMC Symmetrix DMX Storage

z/VM and Linux: Match Made in Cyber Heaven

- z/VM
  - Mature virtualization
  - Removing physical limitations dynamically
- Linux
  - Linux is an enterprise operating system based on UNIX standards
  - Innovative and built upon mainframe concepts
  - Open source
  - Community driven
- Linux on z/VM - brings the best of both worlds together
  - High-performance data access
  - Enables throughput benefits for Linux guest images and enhances overall system performance and scalability

Why Linux on z/VM

EMC² where information lives

- Consolidation
  - Workload - replaces the need for many application servers doing the same work which optimizes your given assets
  - Physically reduces the number of servers, footprint, power usage
- Administration
  - z/VM's support for scheduling, automation, performance monitoring and reporting and virtual machine management
- Manageability
  - Centralized management
  - Efficient use of existing resources
  - Dynamic data mobility
  - Accelerate application deployment (cloning)

Linux on System z Host Options

- What is Linux on System z?
  - Linux on System z is provided by Redhat and SUSE
- Linux on System z native in an LPAR (Logical Partition)
  - Single image
  - Dedicated system resources for a large application environment
- Linux on System z in a virtual machine (under z/VM)
  - More flexibility
    - Configuration
    - Easy expansion
  - Many virtual machines, easily cloned
  - Sharing of physical resources

"…160 Linux guests on 26 IFLs under 4 z/VM instances in production. In test we have 320 Linux guests on 17 IFLs under 5 z/VM instances."

Disk Terminology – Mainframe vs. Open Systems

Disk Devices are Used to Store Data but They Aren't Always Referred to by the Same Name

- Open Systems
  - LUN, Logical Unit Number - Open Systems storage device
- Symmetrix
  - SLV, Symmetrix Logical Volume - EMC Symmetrix storage device
- Mainframe
  - DASD, Direct Access Storage Device - Mainframe storage device
  - mdisk (minidisk) - z/VM storage abstraction (partial or full device)

Linux on System z Connectivity

- FICON – Fibre Connection
  - Follow-on from ESCON supporting full duplex data transfers enabling greater throughput
  - ESCON Multiple Image Facility (EMIF)
- FCP – Fibre Channel Protocol
  - Providing access to Open Systems SCSI FBA devices
  - New CHPID – FCP for IOCDS
  - Ability to save System z device numbers
- IO definitions
  - z/VM – IOCDS, IOCP description
    - CHPID (Channel Path Identifier Information)
    - CNTLUNIT (Control Unit Information)
    - IODEVICE (Device Information)
  - Linux on System z – sysfs
    - Device address

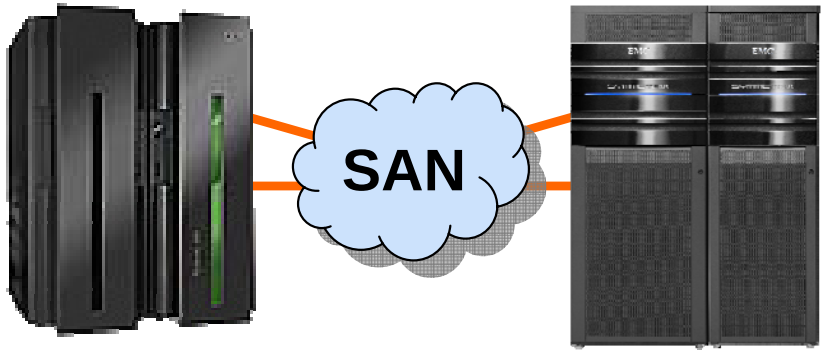

Linux on System z Disk Options

EMC² where information lives

- CKD (Count Key Data)
- Fibre Connection (FICON) Protocol
- Direct attached or dedicated
- Minidisks on Extended CKD (ECKD) via FICON channel
  - Full
  - Partial
- Manageability

- FBA (Fix Block Architecture)
- Fibre Channel Protocol (FCP)
- Direct attach or dedicated
  - SCSI via zFCP driver
- Minidisks on emulated FBA via FCP (9336)
  - Full
  - Partial
- Performance

Linux on System z Disk Options

EMC² where information lives

- ECKD – familiar to mainframe
  - HyperPAV support
  - Dedicated or attached to a virtual machine
    - IO handled by Linux
    - multipath support handled by System z
  - Minidisks – full or partial
    - IO handled by z/VM
    - Storage can be primarily managed and monitored from z/VM
    - Multipath support handled by System z
- FBA – familiar to open systems
  - Dedicated or attached to a virtual machine
    - IO handled by zFCP driver (SCSI) on Linux on System z
    - More like x86 open systems environment
    - Multipath handled by Linux on System z
  - Minidisks – full or partial (edev)
    - IO handled by z/VM
    - Storage can be primarily managed and monitored from z/VM
    - Multipath support handled by z/VM

Symmetrix Linux on System z I/O Support Connectivity Today

EMC² where information lives

	DMX-2	DMX-3, DMX-4			V-Max
Enginuity	**5671**	**5771**	**5772**	**5773**	**5874**
SUSE on System z	9,10	9,10,11	9,10,11	9,10,11	10,11
Red Hat on System z	4.5, 4.6, 4.7 5.0, 5.1, 5.2, 5.3	4.5, 4.6, 4.7 5.0, 5.1, 5.2, 5.3	4.5, 4.6, 4.7 5.0, 5.1, 5.2, 5.3	4.5, 4.6, 4.7 5.0, 5.1, 5.2, 5.3	4.5, 4.6, 4.7 5.0, 5.1, 5.2, 5.3

Symmetrix supports I/O from a Linux on System z host to CKD (FICON) or FBA (Fiber Channel)

Check eLab for latest support:
http://www.emc.com/products/interoperability/elab.htm

Symmetrix V-Max™ Series with Enginuity – 3X More Scalability for Hyper Consolidation

EMC² where information lives

Direct Matrix Architecture

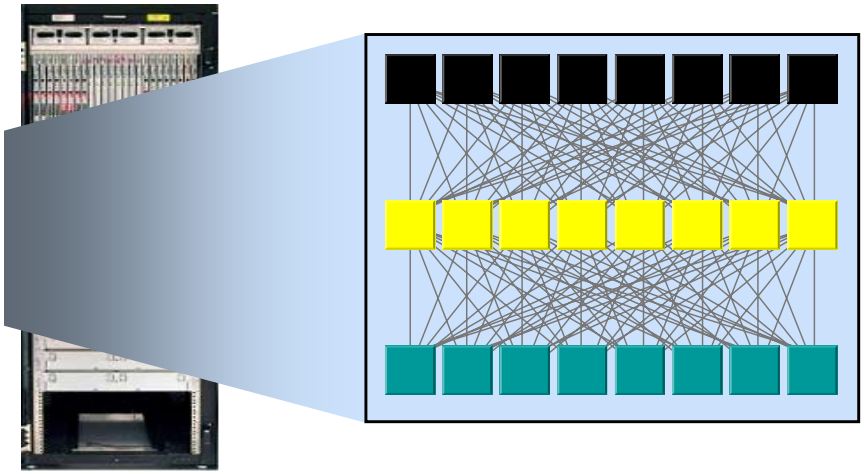

- Separate purpose-built directors
- 16 I/O directors w/4 slices each
- 64 total slices, 128 ports
- Up to 64 FE ports
- Up to 256 GB useable global memory
- Up to 585 TB useable storage capacity

Virtual Matrix Architecture

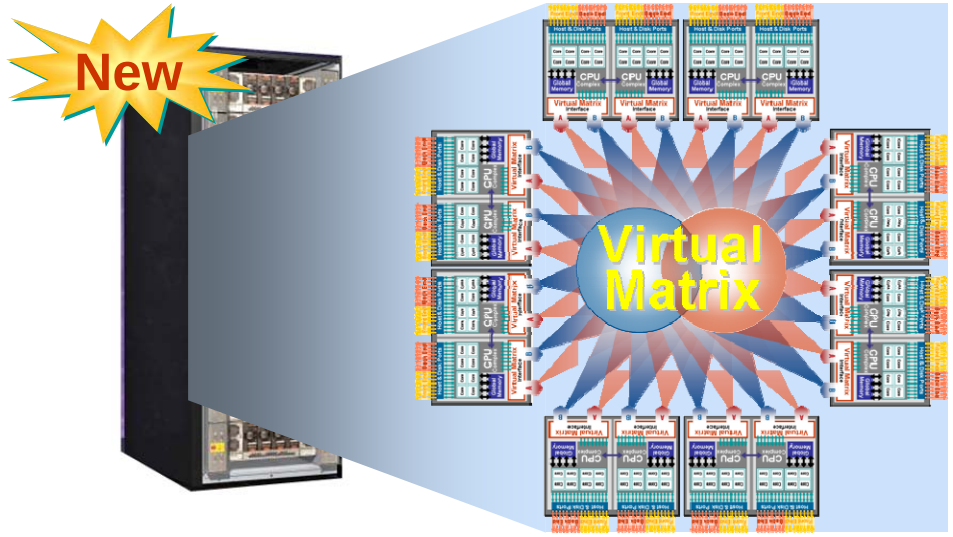

- Next High Availability Node
- One to eight V-Max engines (16 directors)
- Quad-core 2.3 GHz processors to provide more than twice the IOPS
- 128 total slices, 256 ports
- Up to 128 FE ports
- Up to 1 TB (512 GB useable) global mirrored memory
- Up to 2.1PB useable capacity

EMC Symmetrix Array-Specific FCP Settings

- The director bit settings can be modified by SYMCLI or Symmetrix Management Console (SMC)
- A qualified storage administrator should make these changes!
- When attaching to the Symmetrix array, the Front End Adapter (FA) settings for FCP should be

Director bit settings for Linux on System z:

PP	**Point-to-Point**
SPC2	**SPC2 SCSI Primary**
EAN	**Enable Auto Negotiation**
C	**Common Serial Number**
SC3	**SCSI 3 Interface**
UWN	**Unique Worldwide Name**

Linux on System z View of the Symmetrix System

- The Symmetrix system emulates disk drives
  - Open systems hosts see the Symmetrix system as one or more FBA SCSI disk drives
  - z/OS systems see the Symmetrix system as a Logical Control Unit and one or more CKD disk drives
  - z/VM and Linux on System z see both FBA and CKD devices
- Other than basic query data, the host has no knowledge of the Symmetrix internal configuration
- EMC provides the tools to provide visibility and control

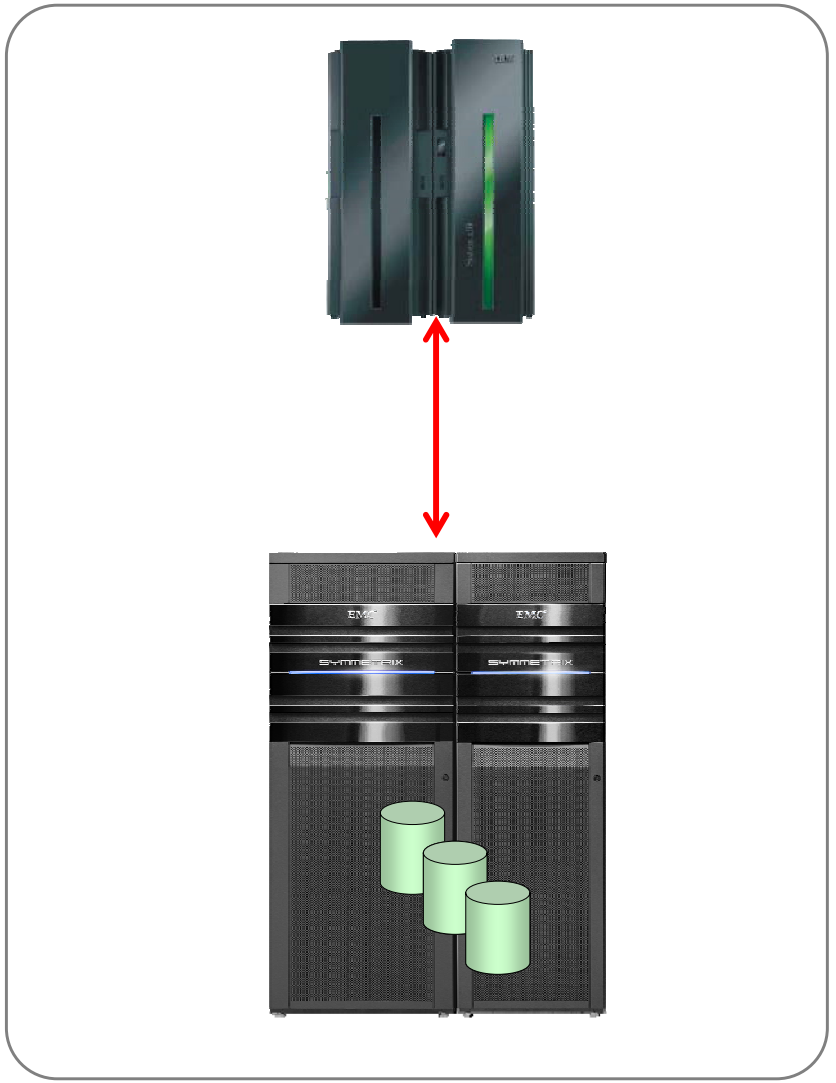

EMC Storage Management Options

- Solutions Enabler
  - On Linux on System z SLES 10
  - Open Systems hosts – Windows, Linux, UNIX
  - Used for general operations
- Mainframe Enablers
  - z/OS
  - Used for general operations and SRDF
- Symmetrix Management Console (SMC)
  - Windows, Linux, UNIX
  - Used for general and SRDF operations
- EMC z/OS Storage Manager (EzSM)
  - ISPF Panel Menu Interface
- EMC products for TPF
    - TimeFinder® Controls for TPF
    - SRDF® Controls for TPF
    - ResourcePak for TPF

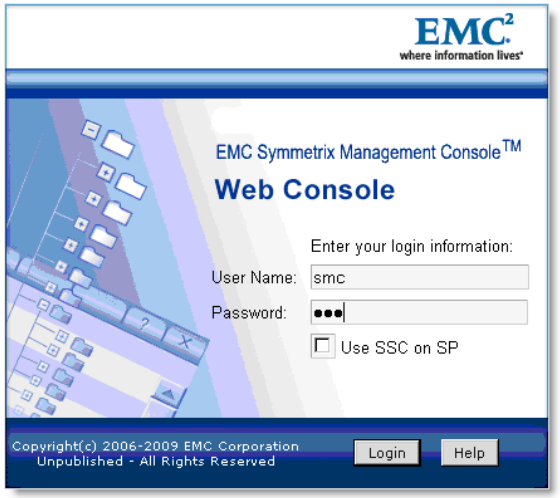

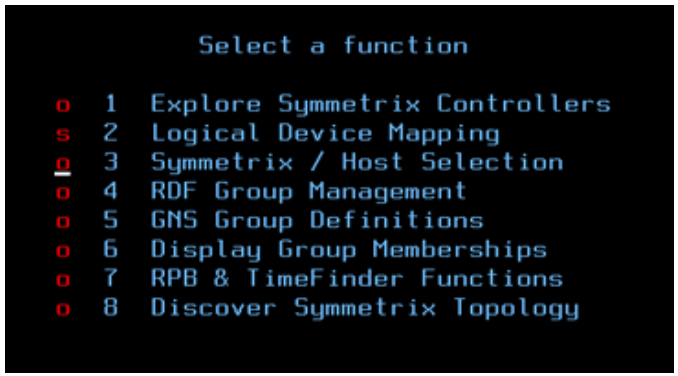

Symmetrix Storage – CKD or FBA

- Configure DASD – CKD
  - For Solutions Enabler support on SUSE gatekeeper devices need to be setup as rdev unsupported DASD
- Configure Linux on System z FCP (zfcp driver) Disks – FBA
  - Channel Number
  - WWPN – worldwide port name
- Create gatekeeper devices as dedicated communication path to the Symmetrix system from a management host
  - Allows Symmetrix management software to retrieve configuration and status information without interfering with normal Symmetrix I/O

FCP Connectivity Example

FCP Test Environment

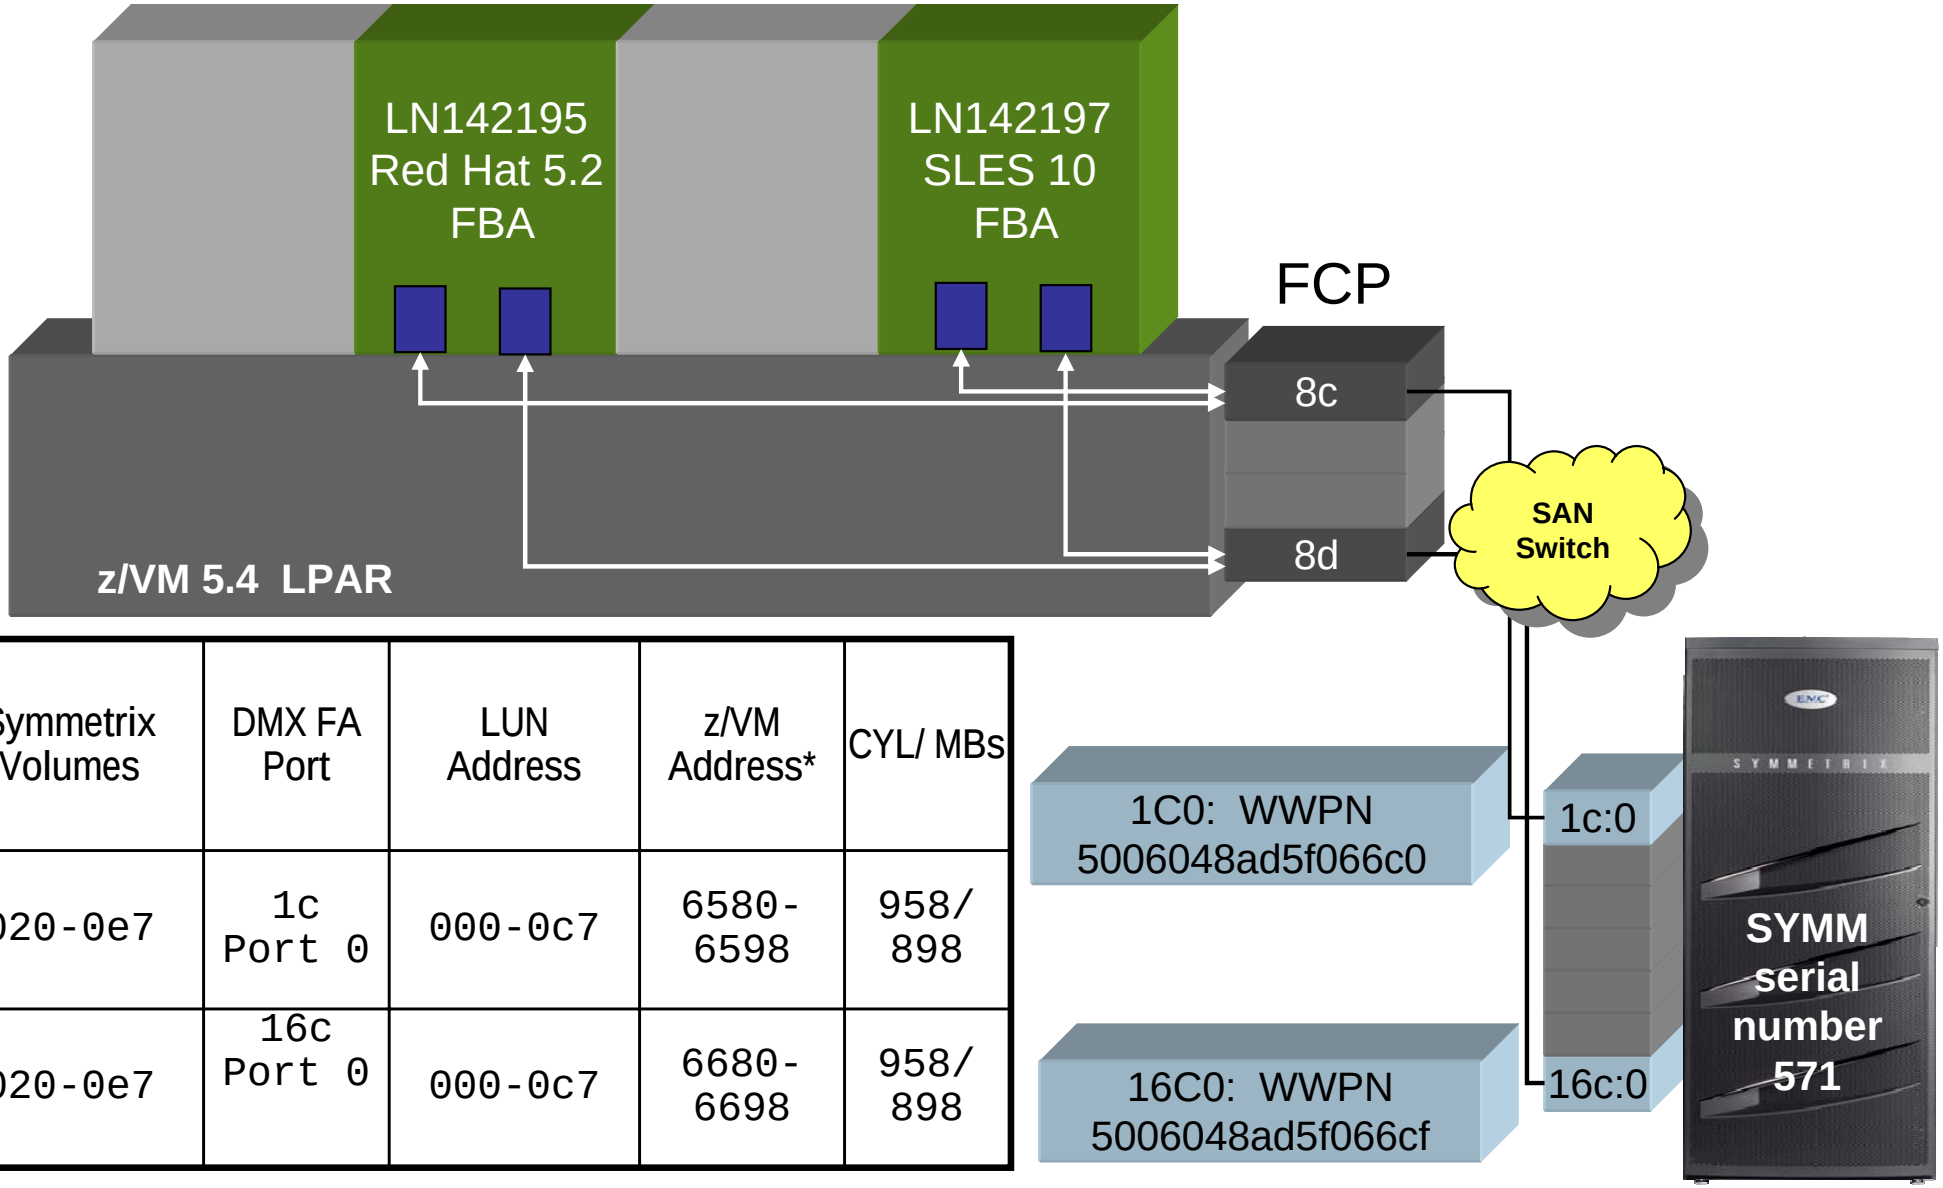

Symmetrix Volumes	DMX FA Port	LUN Address	z/VM Address*	CYL/ MBs
020-0e7	1c Port 0	000-0c7	6580-6598	958/898
020-0e7	16c Port 0	000-0c7	6680-6698	958/898

*Only one IODEVICE is required to address FCP LUNs

Linux on System z FBA Device Relationship Path

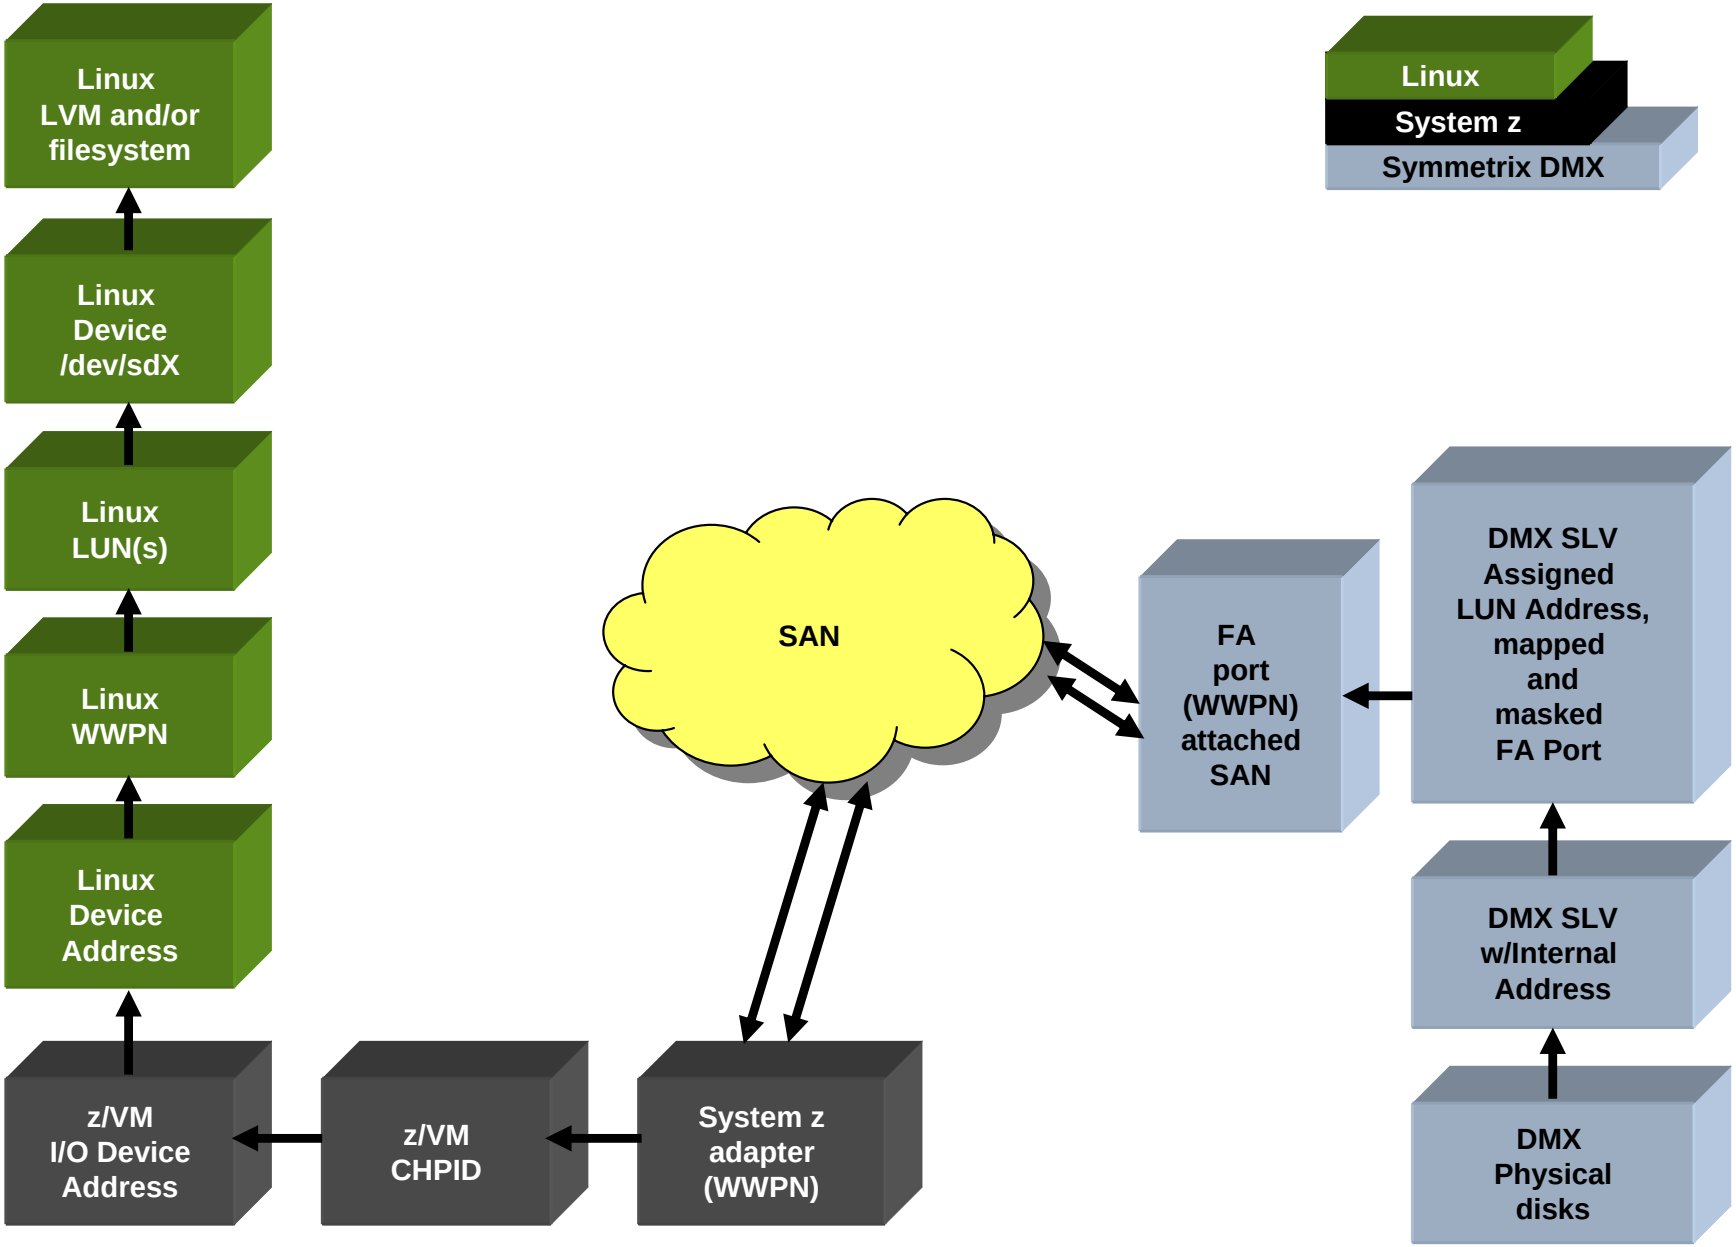

Linux on System z FBA Device Relationship Path

EMC² where information lives

- Symmetrix – FBA device created and mapped
  - Group of physical disks grouped to be a Symmetrix Logical Volume (SLV) and assigned an internal device address
  - SLV is assigned an external LUN address that the host will reference
  - SLV is mapped and masked out a Front Adapter port (FA, aka Director) which has a unique WWPN
  - FA port (WWPN) attached to a SAN switch
- SAN – switch connects the Symmetrix system to the mainframe
  - Symmetrix FA port (WWPN) is "zoned" together with the System z adapter (WWPN)
  - VM adapter (channel) is attached to SAN switch
- z/VM
  - Physical System z adapter assigned a CHPID as an FCP with WWPN
  - z/VM CHPID has IODEVICE associated to it
  - z/VM IODEVICE attached to Linux on System z as real or virtual device address
- Linux on System z
  - Assigned z/VM IODEVICE number recognized by Linux on System z via hotplug routine
  - Linux on System z device has a channel associated to it or control unit/target
  - Channel/control unit has a WWPN
  - IODEVICE address is associated with WWPN
  - LUN address(es) are assigned to a WWPN
  - LUN is associated to a Linux device - /dev/sdX
  - Linux device is associated to LVM and/or a Linux filesystem
- **NOTE: all WWPNs are unique**

Symmetrix to System z to Linux FBA Relationship

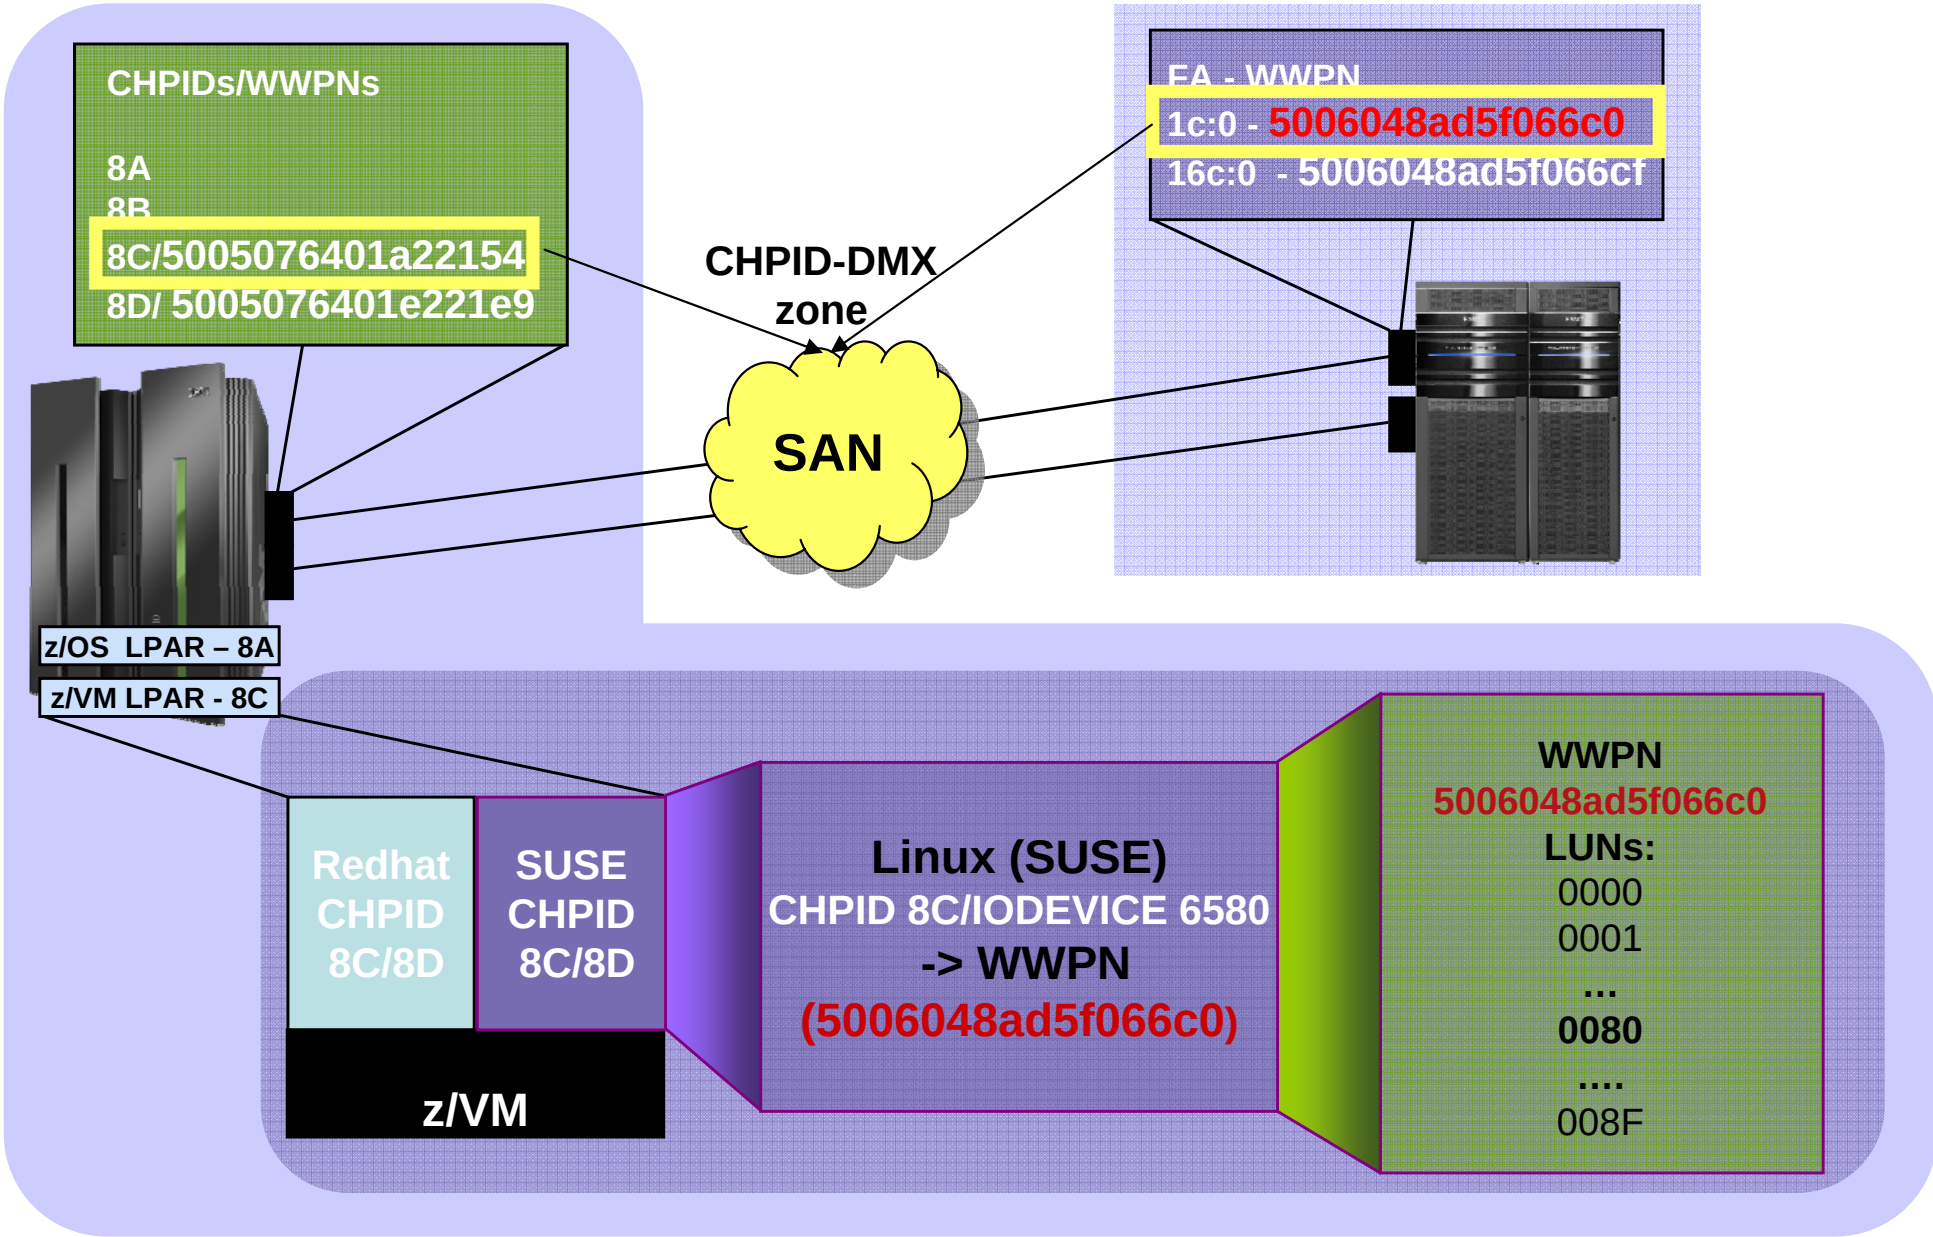

Steps to Adding FCP Devices to Linux on System z

EMC² where information lives

1. Setup hardware and make all physical connections between the Symmetrix array, System z, and SAN switch
2. Setup zoning on the switch to associate the appropriate components with the System z channel
3. On the Symmetrix define Symmetrix FBA LUNs, map and mask to the appropriate front end adapter port
4. Add I/O definitions to z/VM through IOCDS
5. Create z/VM virtual machine directory entries with assigned disk in the directory
6. Install Linux on System z in newly allocated Linux guest virtual machine
7. Add additional disk to the Linux on System z virtual machine via directory entry or CP attach command
   - Vary devices online
   - Associate WWPN with device address
   - Associate LUN(s) (SLV) with WWPN
8. Partition the Linux device, /dev/sdX
9. Add the Linux device (/dev/sdX) to LVM and/or create filesystem

z/VM: IOCDS Example for Symmetrix – FCP

```
CHPID PATH=(CSS(0),8C),SHARED,PARTITION=((0),(V109)),          *
                PCHID=392,TYPE=FCP
CHPID PATH=(CSS(0),8D),SHARED,PARTITION=((0),(V109)),          *
                PCHID=503,TYPE=FCP
CNTLUNIT CUNUMBR=0600,PATH=((CSS(0),8C)),UNIT=FCP
IODEVICE ADDRESS=(6580,25),CUNUMBR=(0600),UNIT=FCP
CNTLUNIT CUNUMBR=0640,PATH=((CSS(0),8D)),UNIT=FCP
IODEVICE ADDRESS=(6680,25),CUNUMBR=(0640),UNIT=FCP
```

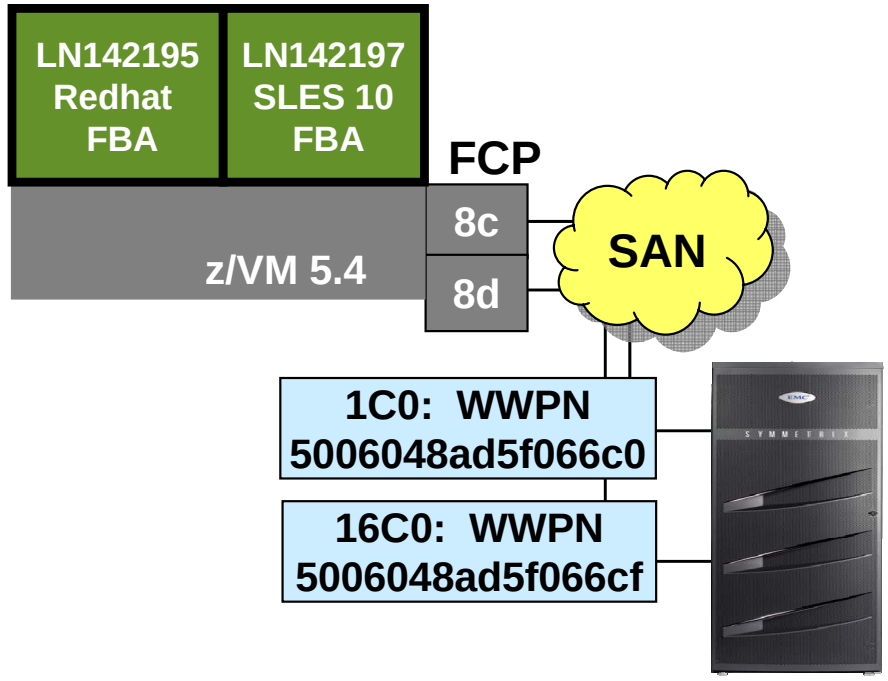

Example Linux on System z z/VM Directory Entry – FCP Devices


```
USER LN142197 LINUX 512M 1024M G
   INCLUDE LNXDFLT
   OPTION QUICKD
 DEDICATE 0191 6C03
 DEDICATE 6580 6580
 DEDICATE 6680 6680
```

z/VM View – Query FCP Devices

- View from z/VM of FCP devices attached to Linux on System z guest, ln142197

```
q fcp att ln142197

FCP  6580 ATTACHED TO LN142197 6580 CHPID 8C
```

Linux on System z FBA Device Relationship Path

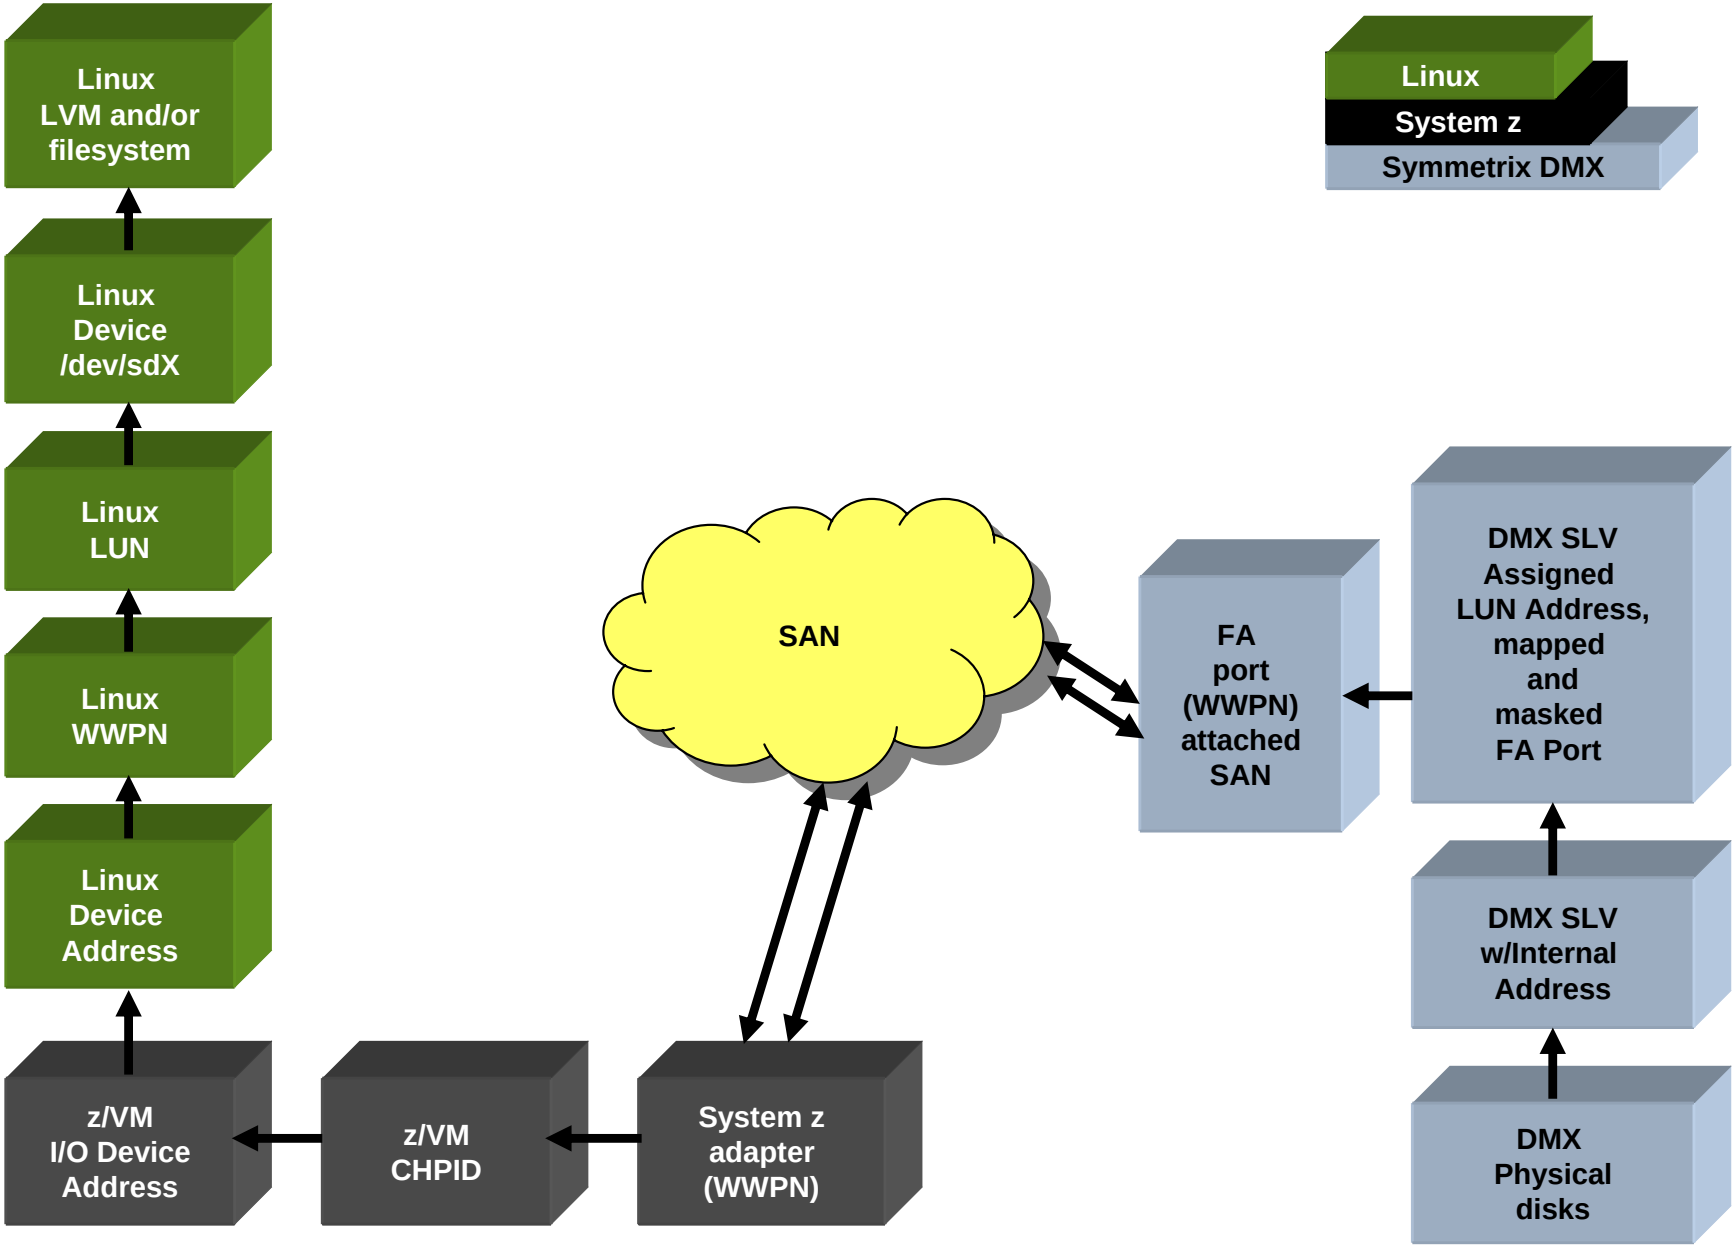

Adding FCP Devices from the Command Line

- Attach the FCP device to the virtual machine under z/VM
  - May be dedicated device in virtual machine's directory also

    attach 6580 ln142197
- Verify attachment of iodevice, from Linux on System z, via the vmcp command

vmcp q v all

```
FCP  6580 ON FCP    6580 CHPID 8C SUBCHANNEL = 0010
     6580                    TOKEN      = 000000001EDDE580
     6580 DEVTYPE FCP        CHPID 8C FCP
     6580 QDIO ACTIVE        QIOASSIST ACTIVE        QEBSM
     6580
     6580 INP + 01 IOCNT = 00000002  ADP = 128 PROG = 000 UNAVAIL = 000
     6580          BYTES = 0000000000000000
     6580 OUT + 01 IOCNT = 00000011  ADP = 000 PROG = 018 UNAVAIL = 110
     6580          BYTES = 0000000000010D40
```

Bring FCP Devices Online

- Bring the device online to Linux on System z with echo or chccwdev command

```
ln142197:/sys/bus/ccw/devices/0.0.6580 # cat online
0
ln142197:/sys/bus/ccw/devices/0.0.6580 # echo 1 > online

ln142197:/sys/bus/ccw/devices/0.0.6580 # cat online
1
```

```
ln142197:/sys/bus/ccw/devices # chccwdev --online 0.0.6580
Setting device 0.0.6580 online
```

Adding FCP Devices from the Command Line

- Message appears on the virtual machine console when device is online
- Shows establishment of communication between z/VM, Linux on System z and the FCP device

```
scsi4 : zfcp
zfcp: The adapter 0.0.6580 reported the following
  characteristics:
WWNN 0x5005076400c6cefe, WWPN 0x5005076401a22154, S_ID
  0x007b7813,
adapter version 0x4, LIC version 0x70b, FC link speed 2 Gb/s
zfcp: Switched fabric fibrechannel network detected at adapter
  0.0.6580.
```

- The above message is from z/VM adapter standpoint and not the Symmetrix system
  - WWNN, WWPN
  - FC link speed

FCP – Determine the WWPN

- Determine the WWPN from the SAN switch or from EMC Solutions Enabler (on another host)
- WWN information is used during addition of FCP LUNs and IPL from FCP device via *set loaddev* command

`symcfg -fa 1c list -sid 571 -v -P 0`

where -fa is the director you are using for FCP

where -sid is the Symmetrix identifier

where -P is the port being used on the director

```
Director Symbolic Number          : 01C
Director Numeric Number           : 33
Director Slot Number              :  1

Director Port: 0
  WWN Node Name                   : 5006048AD5F066C0
  WWN Port Name                   : 5006048AD5F066C0
```

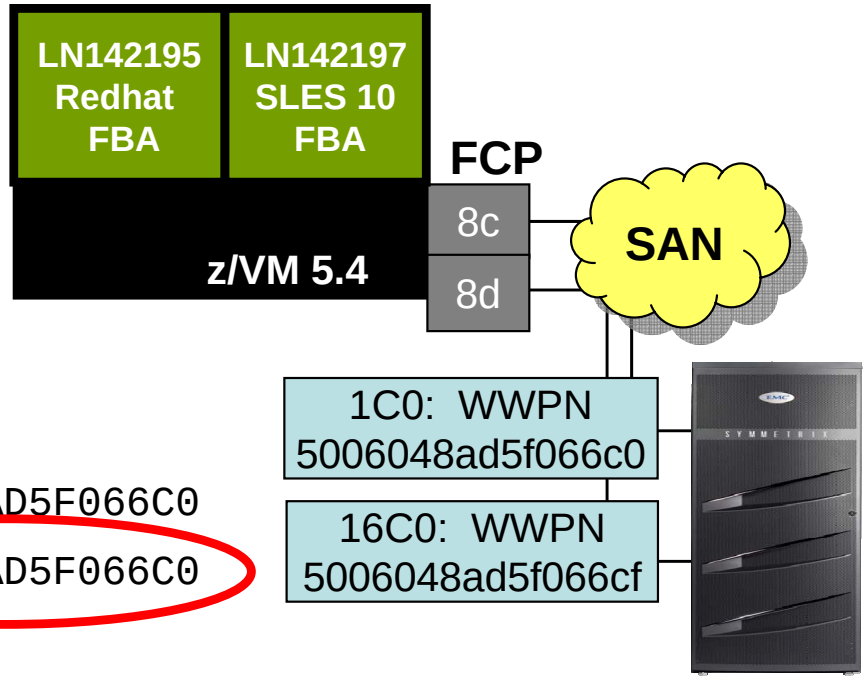

Adding FCP Devices from the Command Line

- Now add the Symmetrix system WWPN for the FCP device

```
echo 0x5006048ad5f066c0 >
/sys/bus/ccw/drivers/zfcp/0.0.6580/port_add
```

NOTE: port_add may not be necessary depending upon the Linux kernel version

- Manually add the LUN using the WWPN and LUN number
  - The LUN number is the first 4 bytes

```
echo 0x0080000000000000 >
/sys/bus/ccw/drivers/zfcp/0.0.6580/0x5006048ad5f066c0/unit_add
```

Adding FCP Devices – VM Console Message

- Virtual machine console message shows device assignment "sde"

```
 Vendor: EMC      Model: SYMMETRIX      Rev: 5773
  Type:   Direct-Access                 ANSI SCSI revision: 04
SCSI device sde: 1839360 512-byte hdwr sectors (942 MB)
sde: Write Protect is off
SCSI device sde: drive cache: write through
SCSI device sde: 1839360 512-byte hdwr sectors (942 MB)
sde: Write Protect is off
SCSI device sde: drive cache: write through
 sde: unknown partition table
sd 0:0:0:128: Attached scsi disk sde
sd 0:0:0:128: Attached scsi generic sg4 type 0
```

List New SCSI Device

- List SCSI device for *fdisk* command

```
ln142197:/sys/bus/ccw/devices/0.0.6580 # lsscsi -v

….
……
[0:0:0:128]    disk    EMC      SYMMETRIX        5773  /dev/sde
  dir: /sys/bus/scsi/devices/0:0:0:128
   [/sys/devices/css0/0.0.0001f/0.0.6580/host0/rport-0:0-
   0/target0:0:0/0:0:0:128]
```

/dev/.udev/db – Device Details

- You can also see the new device under udev

```
ln142197:/dev/.udev/db # cat block\@sde
N:sde
S:disk/by-id/scsi-360060480000190300571533030304130
S:disk/by-path/ccw-0.0.6580-zfcp-
  0x5006048ad5f066c0:0x0080000000000000
M:8:64
E:ID_VENDOR=EMC
E:ID_MODEL=SYMMETRIX
E:ID_REVISION=5773
E:ID_SERIAL=360060480000190300571533030304130
E:ID_TYPE=disk
E:ID_BUS=scsi
E:ID_PATH=ccw-0.0.6580-zfcp-
  0x5006048ad5f066c0:0x0080000000000000
```

Partition SCSI Device via fdisk

EMC² where information lives

```
ln142197:/ # fdisk /dev/sde
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF
    disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-1023, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-1023, default 1023):
Using default value 1023
```

Partition SCSI Device – continued

```
Command (m for help): p

Disk /dev/sde: 941 MB, 941752320 bytes
29 heads, 62 sectors/track, 1023 cylinders
Units = cylinders of 1798 * 512 = 920576 bytes

   Device Boot      Start         End      Blocks     Id   System
/dev/sde1               1        1023      919646     83   Linux

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Create a Filesystem

- Create a filesystem on the partitioned device

```
[root@ln142197 / ]# mke2fs -j /dev/sde1
mke2fs 1.39 (29-May-2006)
……

Writing inode tables: done
Creating journal (4096 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 32 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to
   override.
```

CKD Connectivity Example

CKD Test Environment

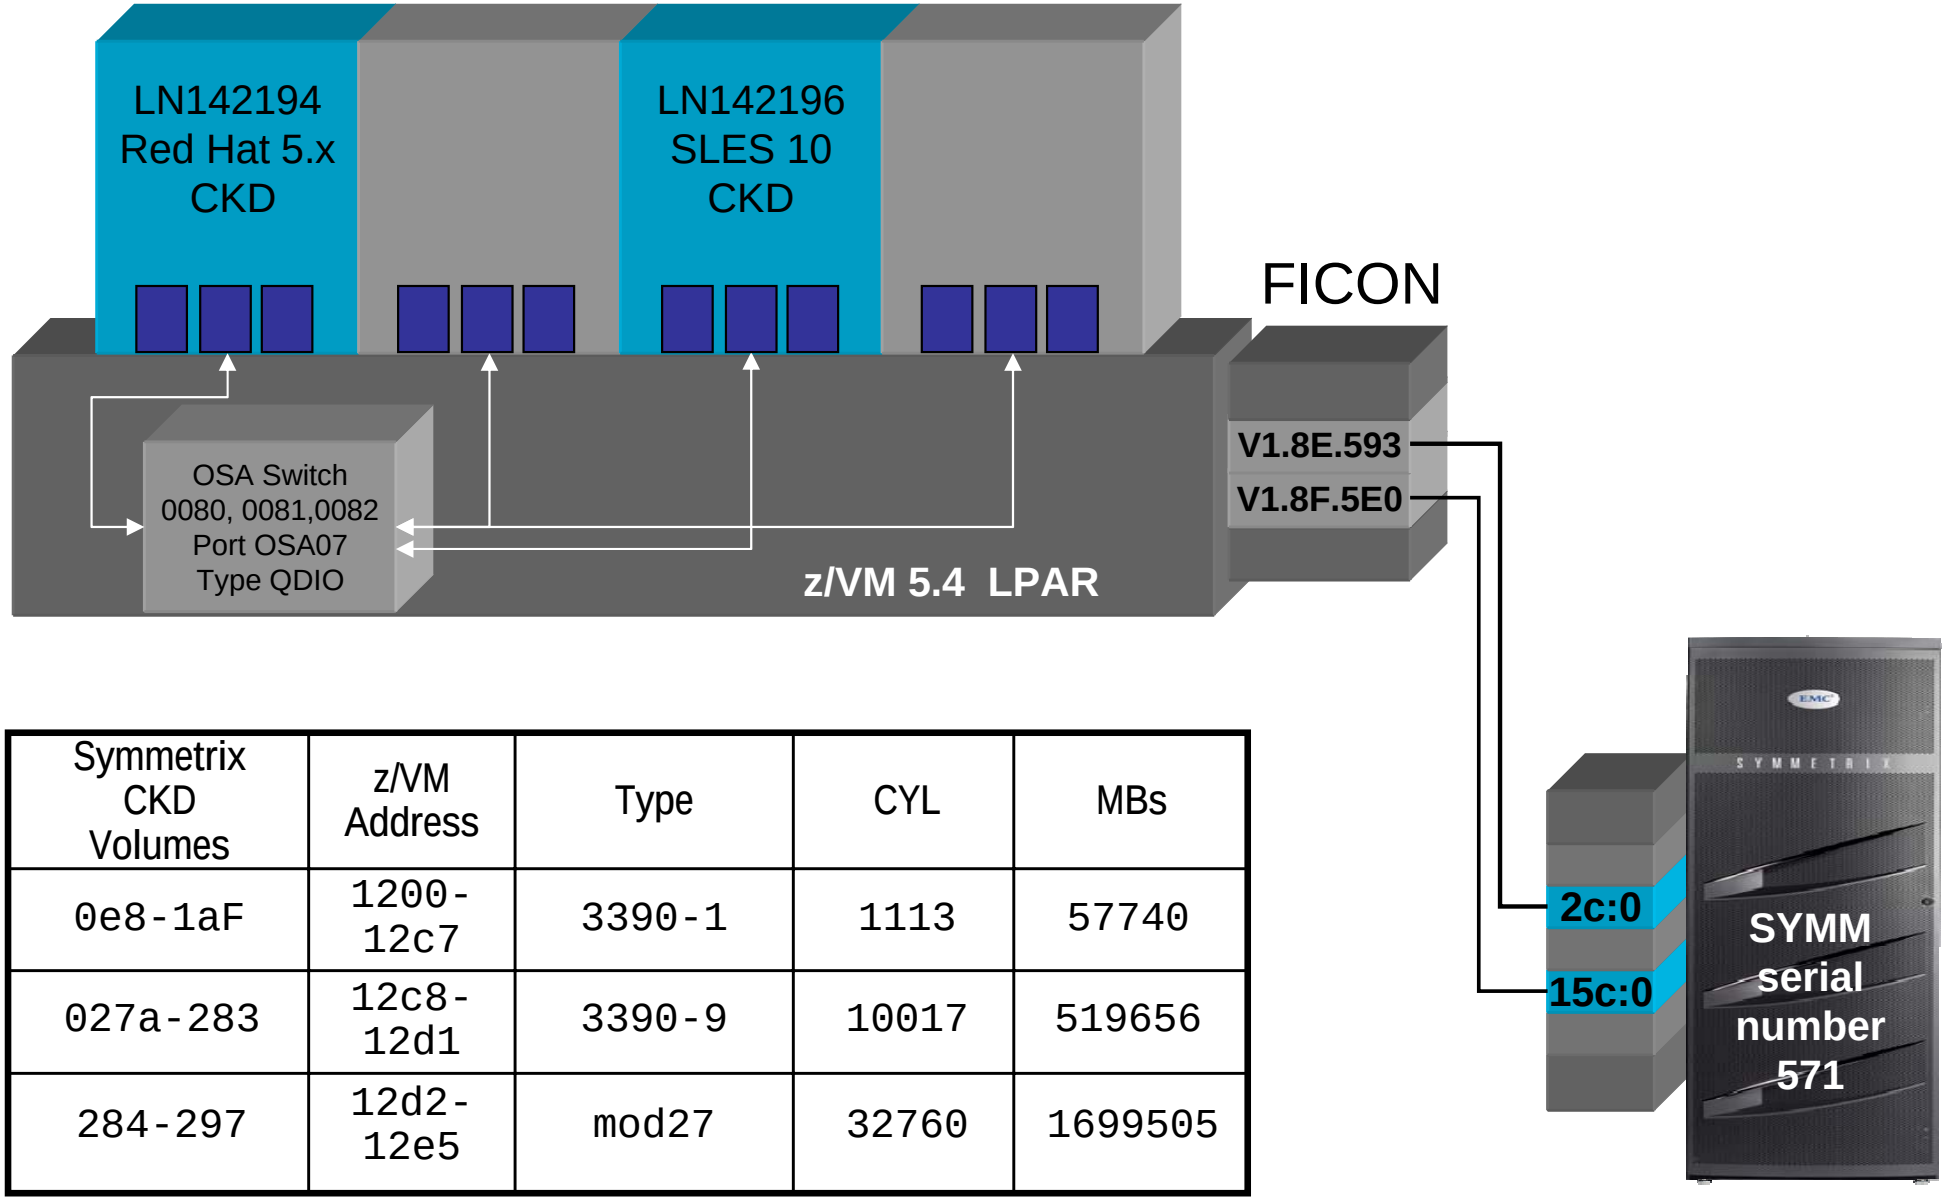

Symmetrix CKD Volumes	z/VM Address	Type	CYL	MBs
0e8-1aF	1200-12c7	3390-1	1113	57740
027a-283	12c8-12d1	3390-9	10017	519656
284-297	12d2-12e5	mod27	32760	1699505

IOCDS Example for Symmetrix – CKD

```
CHPID PATH=(CSS(0),8E),SHARED,                                          *
  PARTITION=((0),(V101,V102,V103,V104,V105,V106,V107,V108,              *
              V109)),PCHID=593,TYPE=FC
CHPID PATH=(CSS(0),8F),SHARED,                                          *
  PARTITION=((0),(V101,V102,V103,V104,V105,V106,V107,V108,              *
              V109)),PCHID=5E0,TYPE=FC


CNTLUNIT CUNUMBR=0680,PATH=((CSS(0),8E,8F)),                            *
       UNITADD=((00,256)),CUADD=0,UNIT=2105
IODEVICE ADDRESS=(6C00,250),CUNUMBR=(0680),STADET=Y,UNIT=3390
```

Linux on System z CKD Device Relationship Path

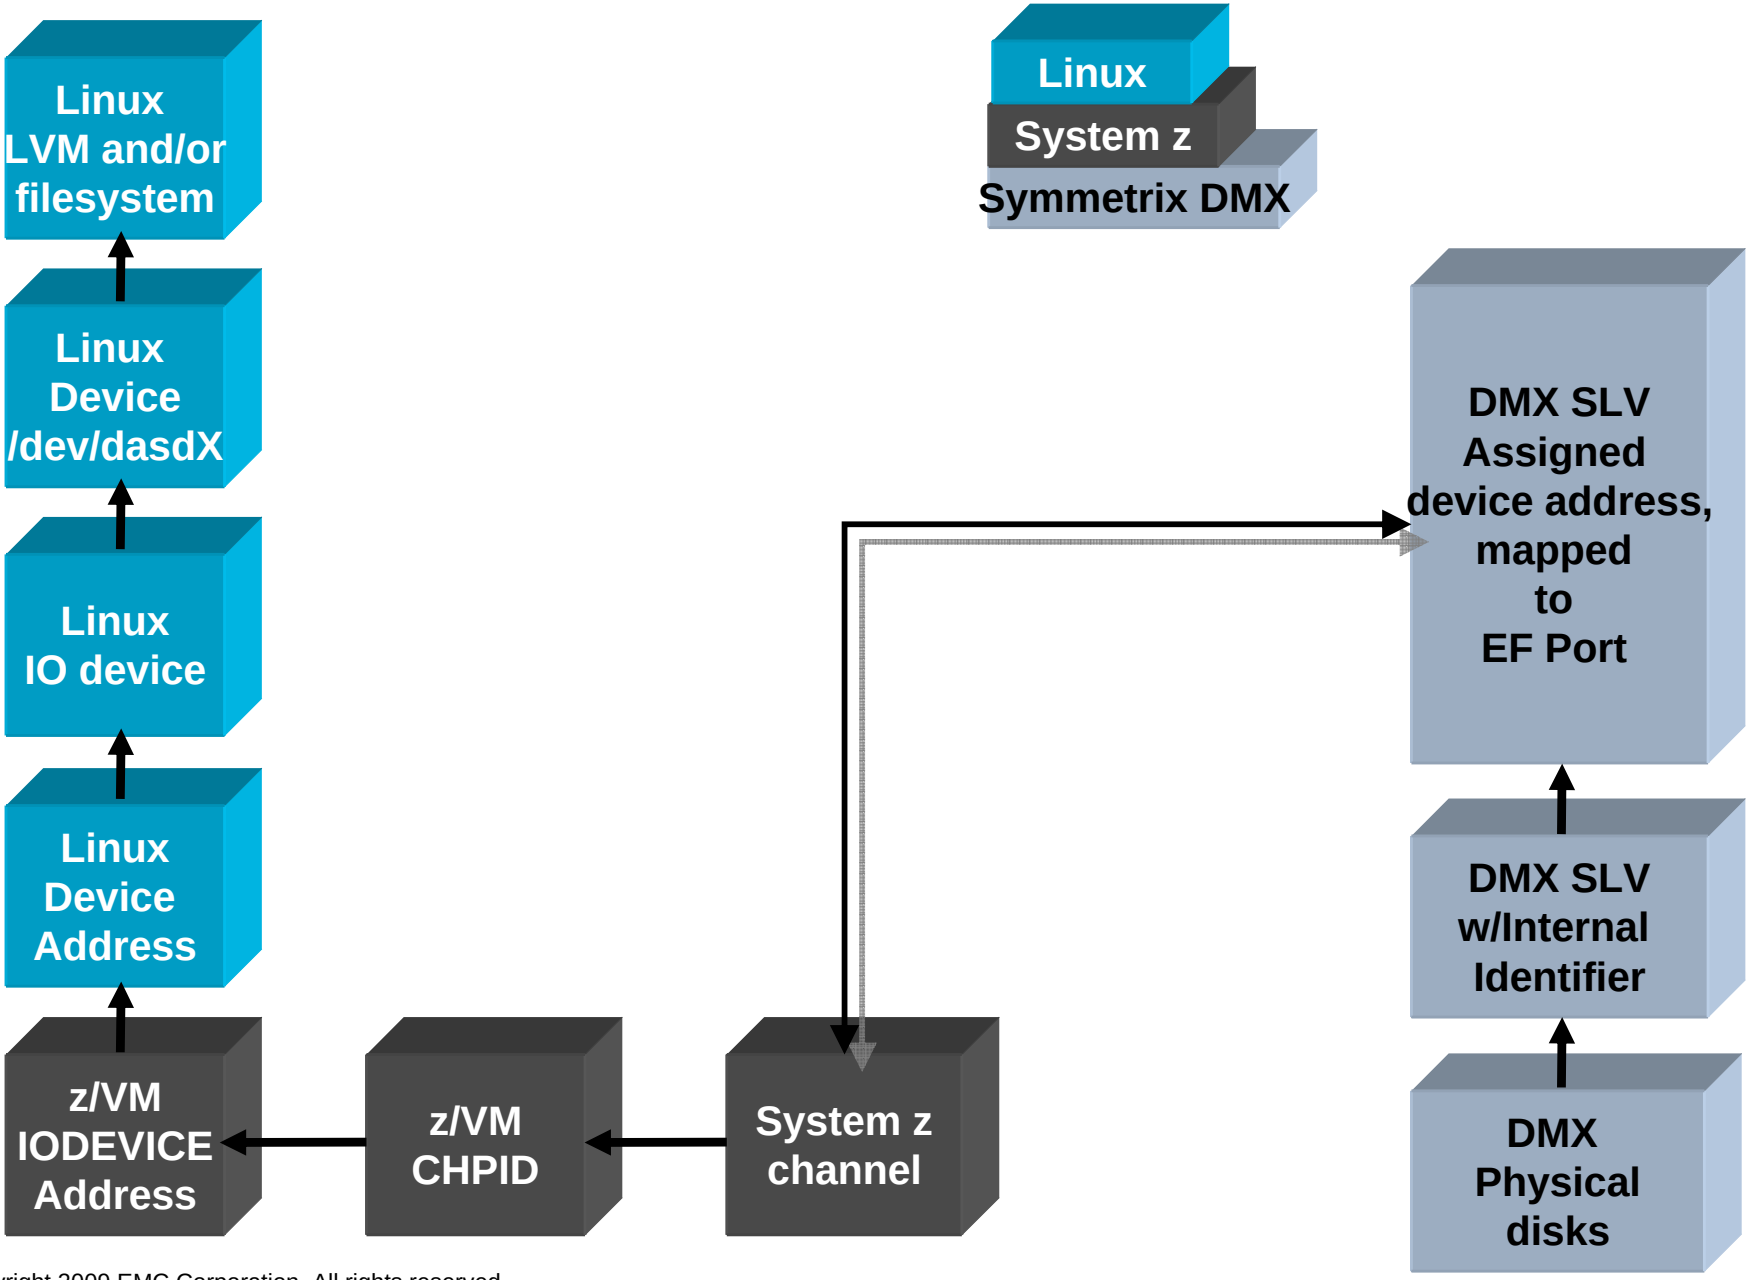

Linux on System z CKD Device Relationship Path

- Symmetrix – CKD device created and mapped
  - Group of physical disks grouped to be a Symmetrix Logical Volume (SLV) and assigned an internal device address
  - SLV is assigned an address
  - SLV is mapped out a Front Adapter port (FA, aka Director) which is FICON
  - FA port attached directly to mainframe, or optionally to a switch
- Optional switch connects the Symmetrix system to the mainframe
- z/VM
  - Physical System z channel adapter assigned a CHPID as FC
  - z/VM CHPID has IODEVICEs associated to it
  - z/VM IODEVICE attached to Linux on System z guest virtual machine as real or virtual device address
- Linux on System z
  - Assigned z/VM IODEVICE number recognized by Linux on System z via kernel hotplug routine
  - Linux on System z device has a subchannel associated to it (control unit)
  - Linux on System z device (LUN) is associated to a Linux device - /dev/dasdX
  - Linux device name, /dev/dasdX, is formatted via dasdfmt
  - Linux DASD is associated to LVM and/or a Linux filesystem

Steps to Adding ECKD Devices to Linux on System z

1. Setup hardware and make all physical connections
2. Add ECKD I/O definitions to z/VM via IOCP deck
3. Create z/VM directory entries with assigned disk dedicated to the virtual machine
4. Logon on to new virtual machine user
5. Install Linux on System z in new Linux guest virtual machine
6. Add additional disk to Linux on System z virtual machine via directory entry or CP attach command
7. Bring Linux on System z devices online
8. Format Linux on System z DASD devices using dasdfmt
9. Partition devices using fdasd
10. Add to LVM or create a filesystem directly on the partitioned device

z/VM View of CKD DASD

vmcp q v all

- DASD 6cca attached as virtual address 0100

```
DASD 0100 ON DASD  6CCA R/W 0X6CCA SUBCHANNEL = 000F
DASD 0101 ON DASD  6CCB R/W 0X6CCB SUBCHANNEL = 0010
```

- DASD attached as real address 6c50

```
DEV  6C50 ON DEV   6C50 SUBCHANNEL = 000C
DEV  6C51 ON DEV   6C51 SUBCHANNEL = 000D
```

Bring CKD Devices Online and Format

- Bring the CKD device online

```
chccwdev --online 0.0.6c50
Setting device 0.0.6c50 online
Done
```

- Examine the status of the device via the *cat* command which will show a “1”, meaning it’s online

```
ln142196:/sys/bus/ccw/devices/0.0.6c50 # cat
    /sys/bus/ccw/devices/0.0.6c50/online
1
```

- Determine Linux device name for the dasdfmt via the *ls* command

```
ln142196:/sys/bus/ccw/devices/0.0.6c50 # ls
alias    bus      devtype     eer_enabled  readonly  use_diag
availability     cmb_enable  discipline   modalias  uevent
   vendor
block:dasdc      cutype      driver       online    uid
```

Format CKD Device – dasdfmt

```
ln142196:/sys/bus/ccw/devices/0.0.6c50 # dasdfmt -b 4096 -d
  cdl -l sym138  -v -f /dev/dasdc
Retrieving disk geometry...
Drive Geometry: 1113 Cylinders * 15 Heads =  16695 Tracks

I am going to format the device /dev/dasdc in the following
  way:
    Device number of device : 0x6c50
    Labelling device        : yes
    Disk label              : VOL1
    Disk identifier         : SYM138
    Extent start (trk no)   : 0
    Extent end (trk no)     : 16694
    Compatible Disk Layout  : yes
    Blocksize               : 4096

--->> ATTENTION! <<---
All data of that device will be lost.
Type "yes" to continue, no will leave the disk untouched: yes
```

ln142196 – Linux on System z View of CKD Devices

EMC² where information lives

```
ln142196:~ # lsdasd
0.0.0100(ECKD) at ( 94:  0) is dasda :active at blocksize 4096, 1803060 blocks,
0.0.0101(ECKD) at ( 94:  4) is dasdb :active at blocksize 4096, 1803060 blocks,
0.0.6c50(ECKD) at ( 94:  8) is dasdc :active at blocksize 4096, 200340 blocks,
0.0.6c51(ECKD) at ( 94: 12) is dasdd :active at blocksize 4096, 200340 blocks,
0.0.6c52(ECKD) at ( 94: 16) is dasde :active at blocksize 4096, 200340 blocks,
0.0.6c53(ECKD) at ( 94: 20) is dasdf :active at blocksize 4096, 200340 blocks,
0.0.6c54(ECKD) at ( 94: 24) is dasdg :active at blocksize 4096, 200340 blocks,
```

dasdfmt and CMS Format with Symmetrix (RAID5/6)

EMC² where information lives

- Linux dasdfmt performance
  - Linux fixes are complete – Red Hat and SuSE
    - Redhat: Bugzilla 486432 for RHEL 5.2, 486431 for RHEL 5.3
    - SUSE: Bugzilla 450989 for SUSE 10 SP2 , 477816 for SUSE 11
    - For more information see Primus cases
  - For minidisks, also require z/VM APAR VM64603
- CMS format performance
  - Two IBM APARs should be applied to z/VM
    - **VM64602: CMS FORMAT OF A MINIDISK IS MUCH SLOWER THAN CPFMTXA**

      CMS FORMAT of a minidisk or a full pack minidisk is slower than using CPFMTXA. This is due to the fact that the CCW chain created by CMS FORMAT does not set define extent byte 7 bit 5.
    - **VM64603: CCW TRANSLATION CHANGE FOR CMS FORMAT**

      VM's CCW fast translation fails to allow guest IO to a minidisk to use a performance feature related to writing out an entire track.

Multipathing

- DASD devices via ESCON or FICON attach
  - Multipathing is available via System z
  - Transparent at Linux on System z level
- FCP handled by the Linux v2.6.x kernel multipathing solution, DM-MPIO
- DM-MPIO requires packages
  - The device-mapper, udev, hotplug, udev, device-mapper-multipath
- Native Multipath Failover Based on DM-MPIO for v2.6.x Linux Kernel and EMC Storage Arrays, Red Hat Enterprise Linux 4 - U3, SuSE Linux Enterprise Server Configuration Guide
  - Available at http://powerlink.emc.com

EMC Storage Management Options

- Solutions Enabler
  - On Linux on System z SLES 10
  - Open Systems hosts – Windows, Linux, UNIX
  - Used for general operations
- Mainframe Enabler
  - z/OS
  - Used for general operations and SRDF
- Symmetrix Management Console (SMC)
  - Windows, Linux, UNIX
  - Used for general and SRDF operations
- EMC z/OS Storage Manager (EzSM)
  - ISPF Panel Menu Interface
- EMC products for TPF
    - TimeFinder Controls for TPF
    - SRDF Controls for TPF
    - ResourcePak for TPF

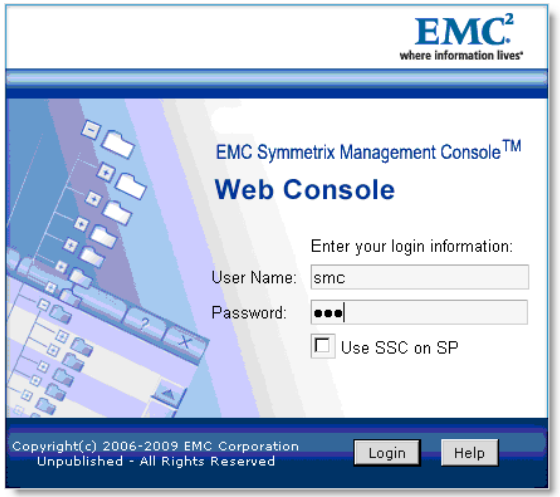

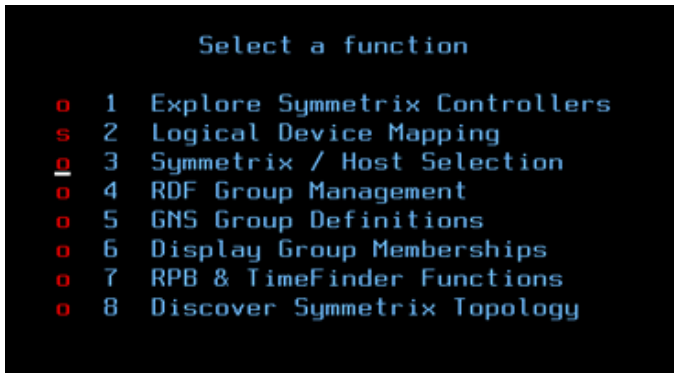

EMC Solutions Enabler Introduction

- Symmetrix Command Line Interface (SYMCLI)
- Provides a host with a comprehensive command set for managing a Symmetrix storage environment
  - Invoked from the host OS command line
  - Scripts that may provide further integration with OS and application
- Separate components licenses
- Security and access controls
  - Monitor only
  - Host-based and user-based controls
- Solution Enabler (SE) current support environment
  - Novell SUSE Linux Enterprise 10, including Service Pack 1 and 2

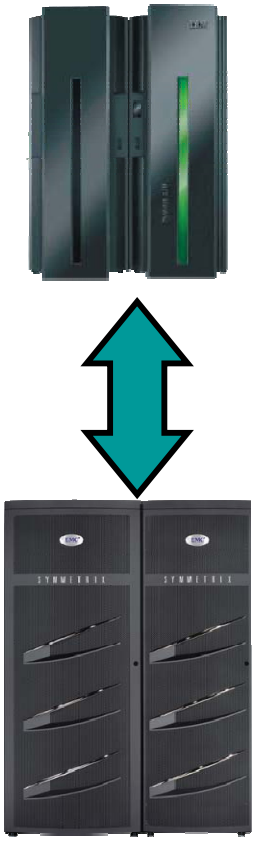

- ✓ **Detailed Configuration Information**
- ✓ **Status**
- ✓ **On-line Configuration Changes**
- ✓ **Performance**
- ✓ **Control**

Solutions Enabler Installation for Linux on System z

- Install Solutions Enabler from the rpm
- Install the Solutions Enabler kernel module via insmod
- Choose the correct module based upon your SUSE version.

  The modules can be found in one of the following directories
  - /opt/emc/SYMCLI/V6.5.2/ioctl/……/s390ioctl.ko
    - suse10
    - suse10sp1
    - suse10sp2
- Load the kernel module
  - `insmod   s390ioctl.ko`
- Verify the s390ioctl.ko module is loaded via lsmod command

```
ln142196:/opt/emc/SYMCLI/V6.5.2/ioctl # lsmod |grep s390
s390ioctl              26112  0
dasd_mod              127852  13 s390ioctl,dasd_eckd_mod
```

Solutions Enabler CKD Requirements

- If SUSE is running under z/VM, CKD gatekeeper devices must be defined as “unsupported” DASD via the rdev command

  set rdev XXXX type unsupported devclass dasd dps yes reserve_release yes

- Devices will be seen as type, CGK

```
ln142196:~ # syminq
     Device             Product                  Device
---------------  ----------------------------------------------
Name     Type   Vendor    ID           Rev  Ser Num      Cap (KB)
---------------- ----------------------------------------------
/dev/dasdc CGK    EMC     SYMMETRIX    5773 7100138340       5760
/dev/dasdd CGK    EMC     SYMMETRIX    5773 7100139340       5760
```

CKD sympd List

```
ln142196:~ # sympd list

Symmetrix ID: 000190300571

  Device Name           Directors                    Device
--------------------------- ------------ ----------------------------------

           Cap
Physical          Sym  SA :P DA :IT  Config      Attribute  Sts (MB)
--------------------------- ------------ ----------------------------------

/dev/dasdc        0138 02C:0 15B:D4  2-Way Mir   N/Grp'd    RW     6
/dev/dasdd        0139 02C:0 02A:D6  2-Way Mir   N/Grp'd    RW     6
/dev/dasde        013A 02C:0 15B:C5  2-Way Mir   N/Grp'd    RW   902
/dev/dasdf        013B 02C:0 15A:D5  2-Way Mir   N/Grp'd    RW   902
```

Solutions Enabler FBA Requirements

- There are no special requirements for SE for FCP devices on SUSE

```
ln142197:/sys/bus/ccw/devices # syminq

    Device                Product                                Device
-------------------------------------------------------------------------
Name     Type   Vendor    ID            Rev  Ser Num       Cap (KB)
-------------------------------------------------------------------------
/dev/sda             EMC   SYMMETRIX    5773 71002C4000   17677440
/dev/sdb             EMC   SYMMETRIX    5773 71002C5000   17677440
/dev/sdc             EMC   SYMMETRIX    5773 71002C6000   17677440
/dev/sdd             EMC   SYMMETRIX    5773 71002C7000   17677440
/dev/sde             EMC   SYMMETRIX    5773 71000A0000     919680
```

Symmetrix Management Console (SMC)

- Device manager for the Symmetrix
  - Light-weight, web-based application
  - Intuitive browser interface
- Monitor, configuration and operational control of one or more arrays
  - Supports all configuration capabilities as Solutions Enabler CLI
  - Real-time reporting – No historical data
  - No SAN or Host monitoring and control
- Bundled with EMC Control Center Symmetrix Manager
- Requires Solutions Enabler server access

SMC Management Capabilities

Discovery

- Array and internals config
- Logical and physical devices

Monitoring

- Properties, Status, Alerts

Administration

- User-level security, logging, auditing
- Access-control management

Configuration

- Device creation, configuration, masking
- Supports Solutions Enabler code for Open Systems and Mainframe
- Symmetrix Priority Controls

Replication

- Configuration and control
- Discovery of objects and status
- SRDF monitoring

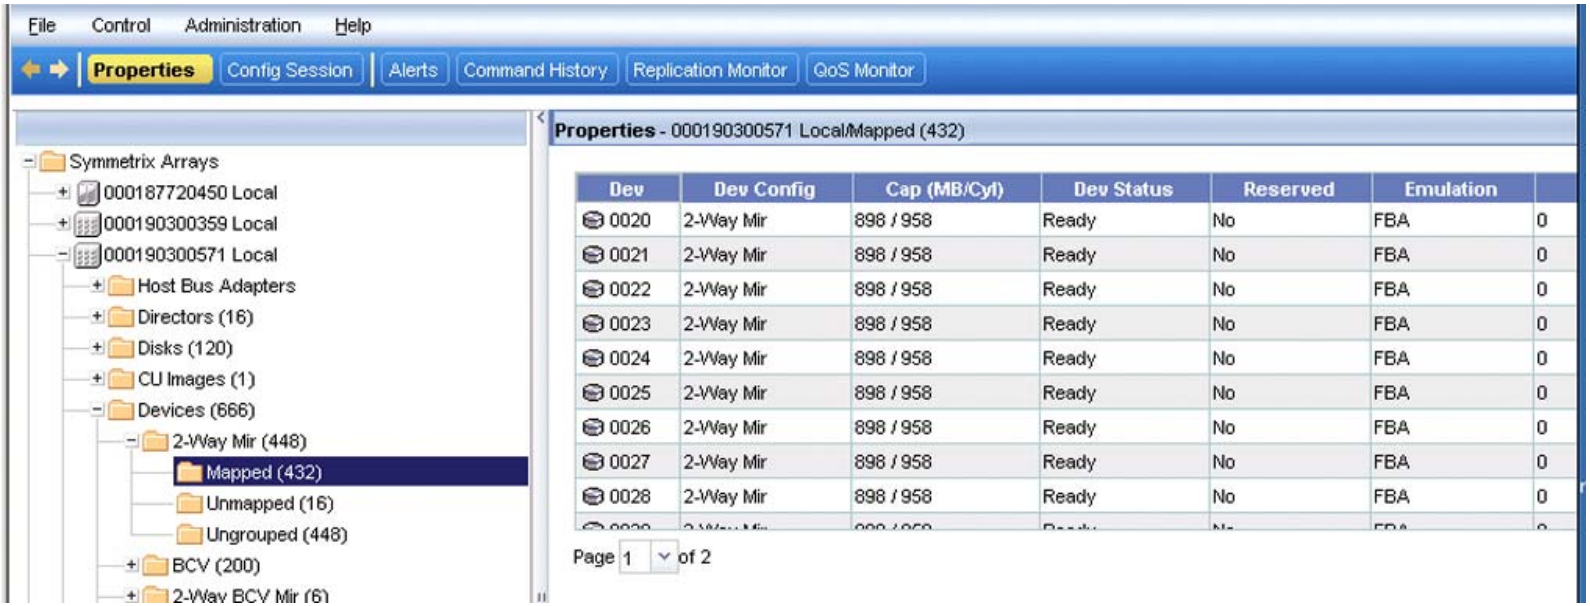

SMC – Director View

- View the Director bit settings for Linux on System z

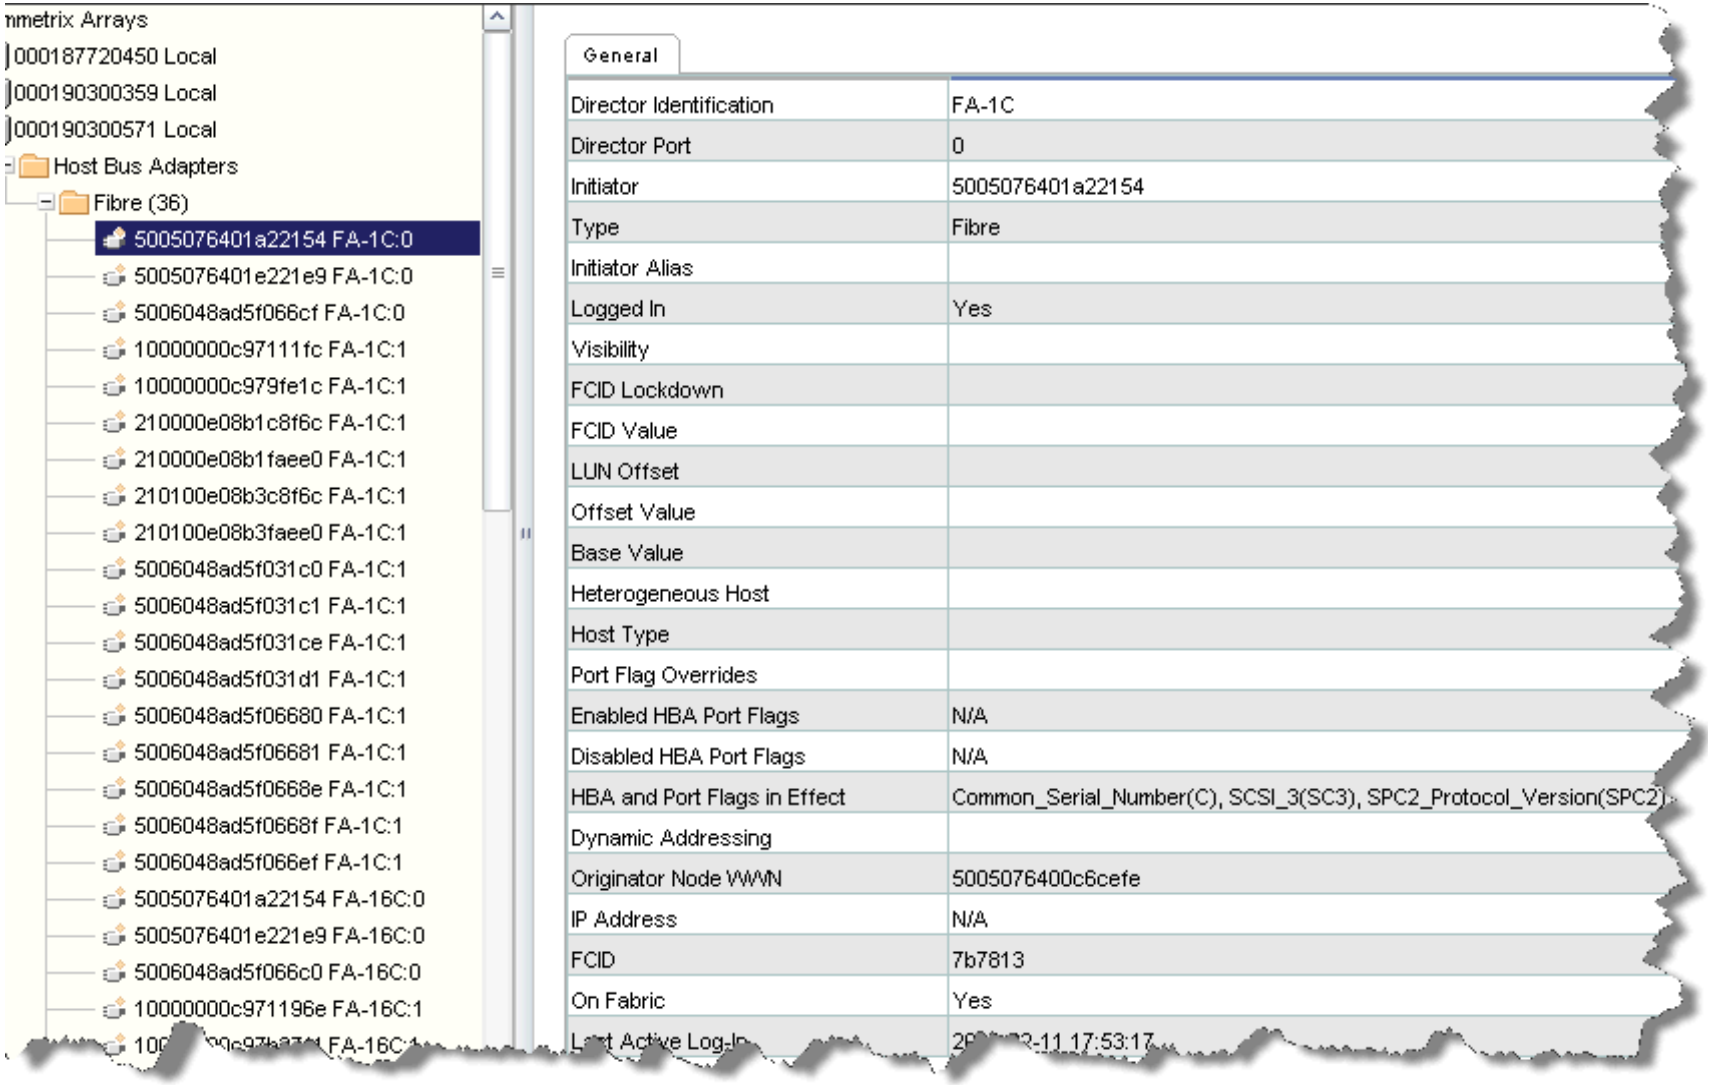

EMC TimeFinder Solutions

TimeFinder/Mirror

- Provides a full - volume mirror of the source

TimeFinder/Clone

- Provides a full - volume copy of the source

TimeFinder/Snap

- Provides a pointer - based snapshot copy of the source

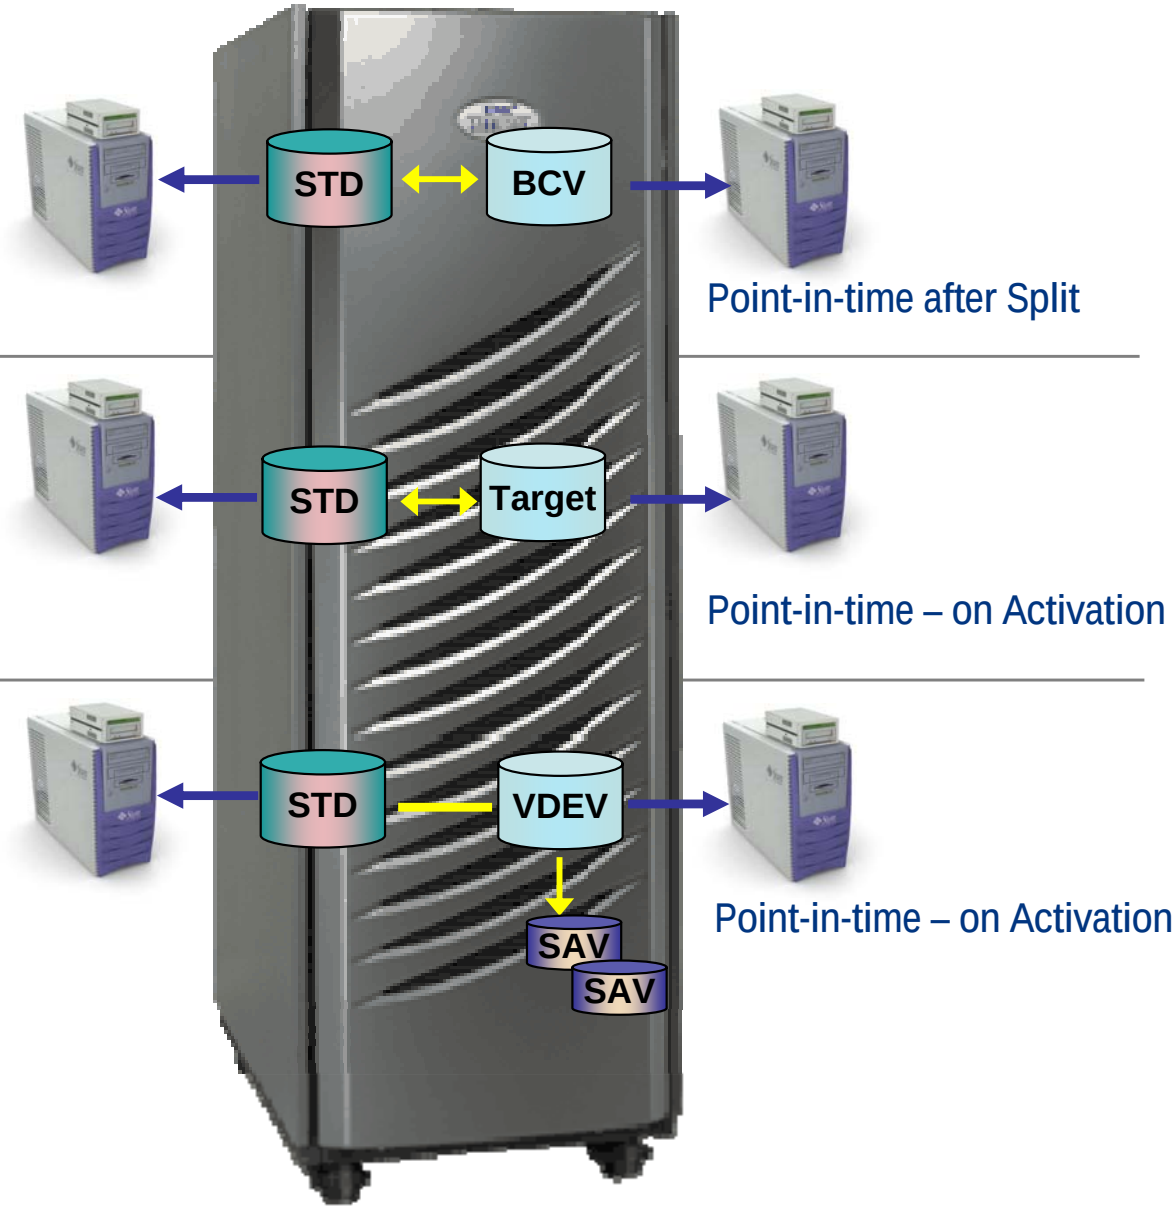

TimeFinder Functions Panel

- TimeFinder functions available through the EzSM panel interface

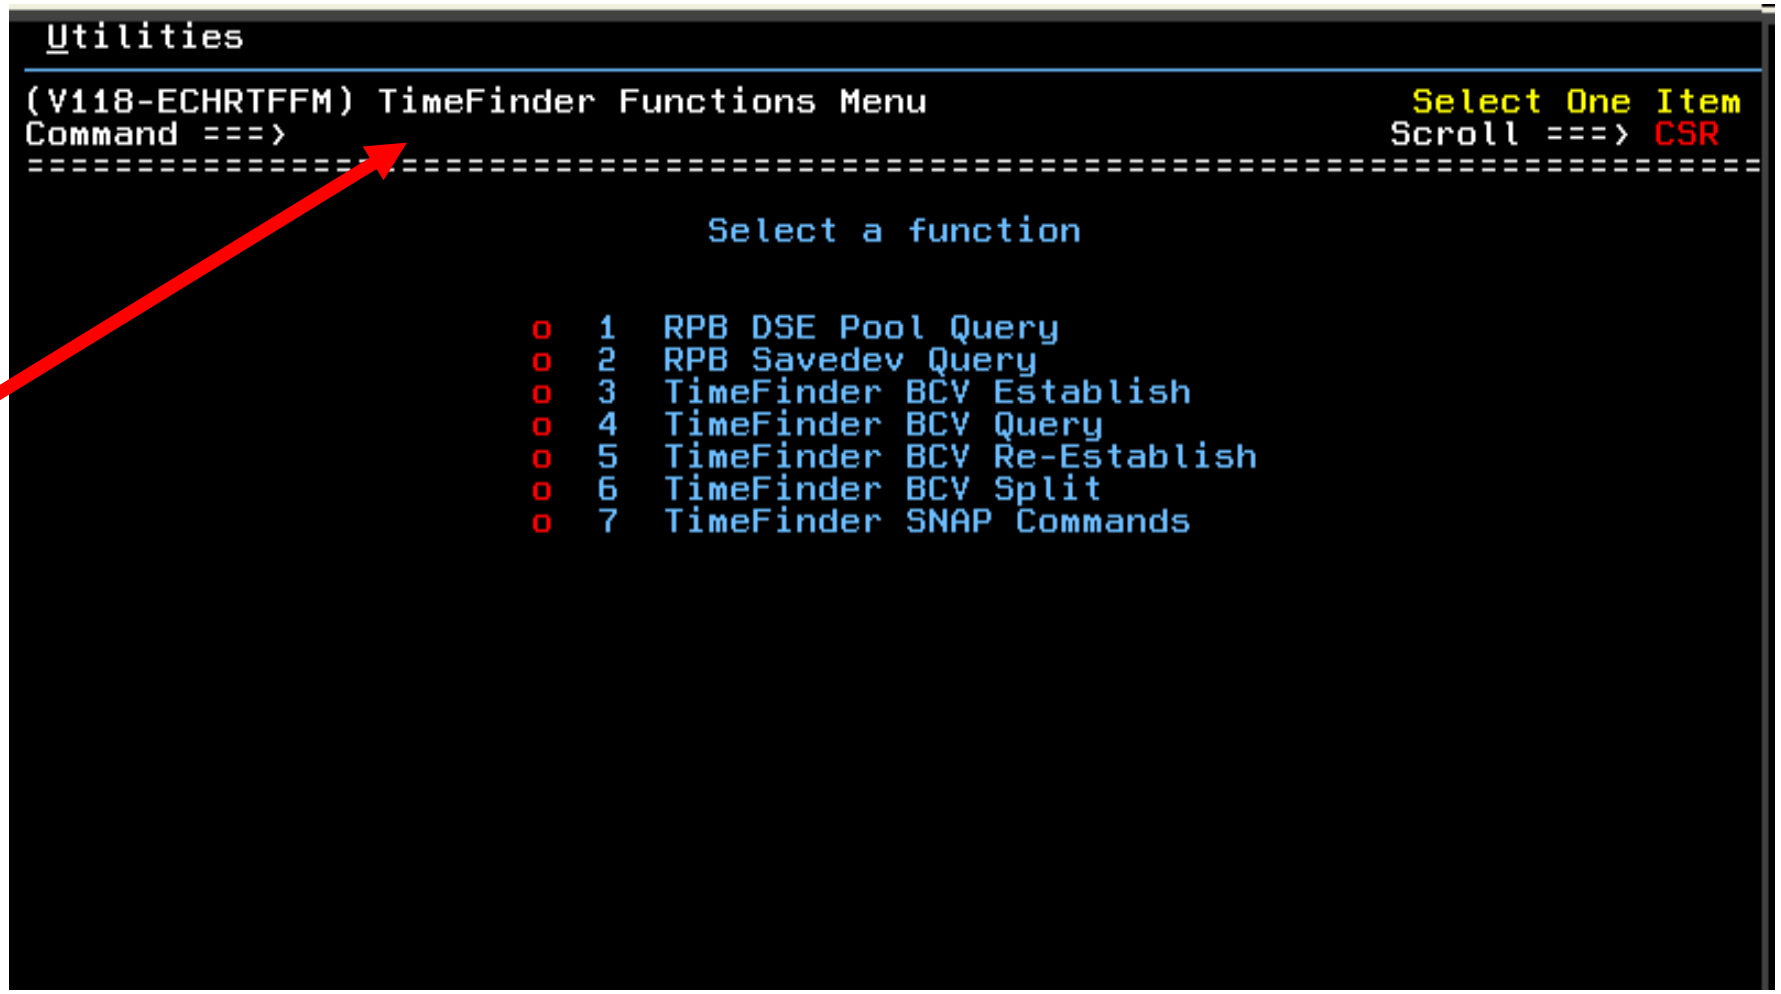

Synchronous and Asynchronous Options

SRDF/Synchronous

- RPO= 0 (No data exposure)
- Some performance impact
- Limited distance
- Mission Critical Applications

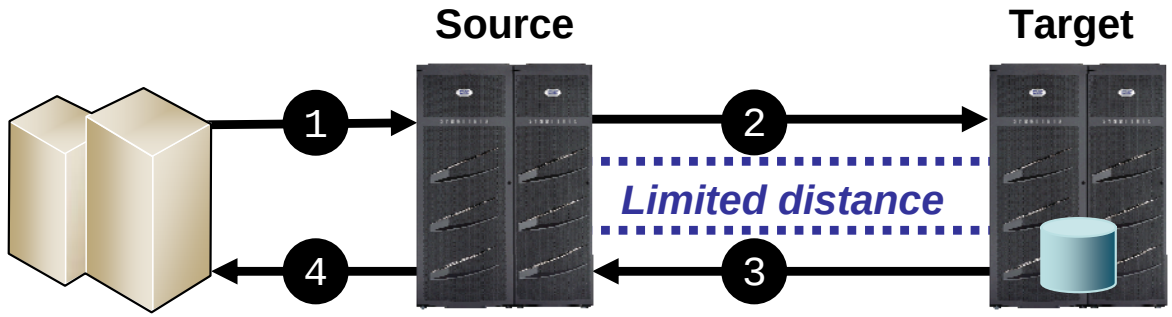

SRDF/Asynchronous

- RPO in Seconds
- Data transfer in Delta Sets
- No performance impact
- Unlimited distance

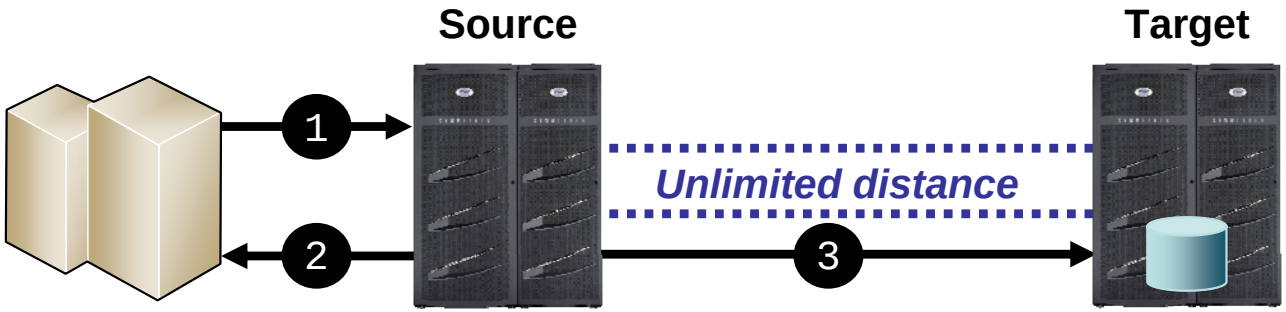

SRDF/Data Mobility

- Data Transfer between Symmetrix
- No performance impact
- Write ordering **NOT** maintained
- Unlimited distance

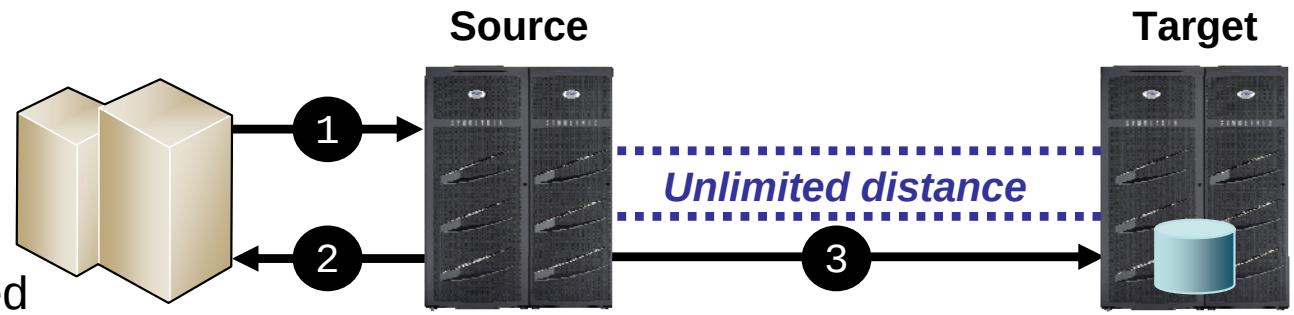

SRDF Host Component – EzSM Panel Interface

- Query, Configure, Display, Manage, CLI, Session Options

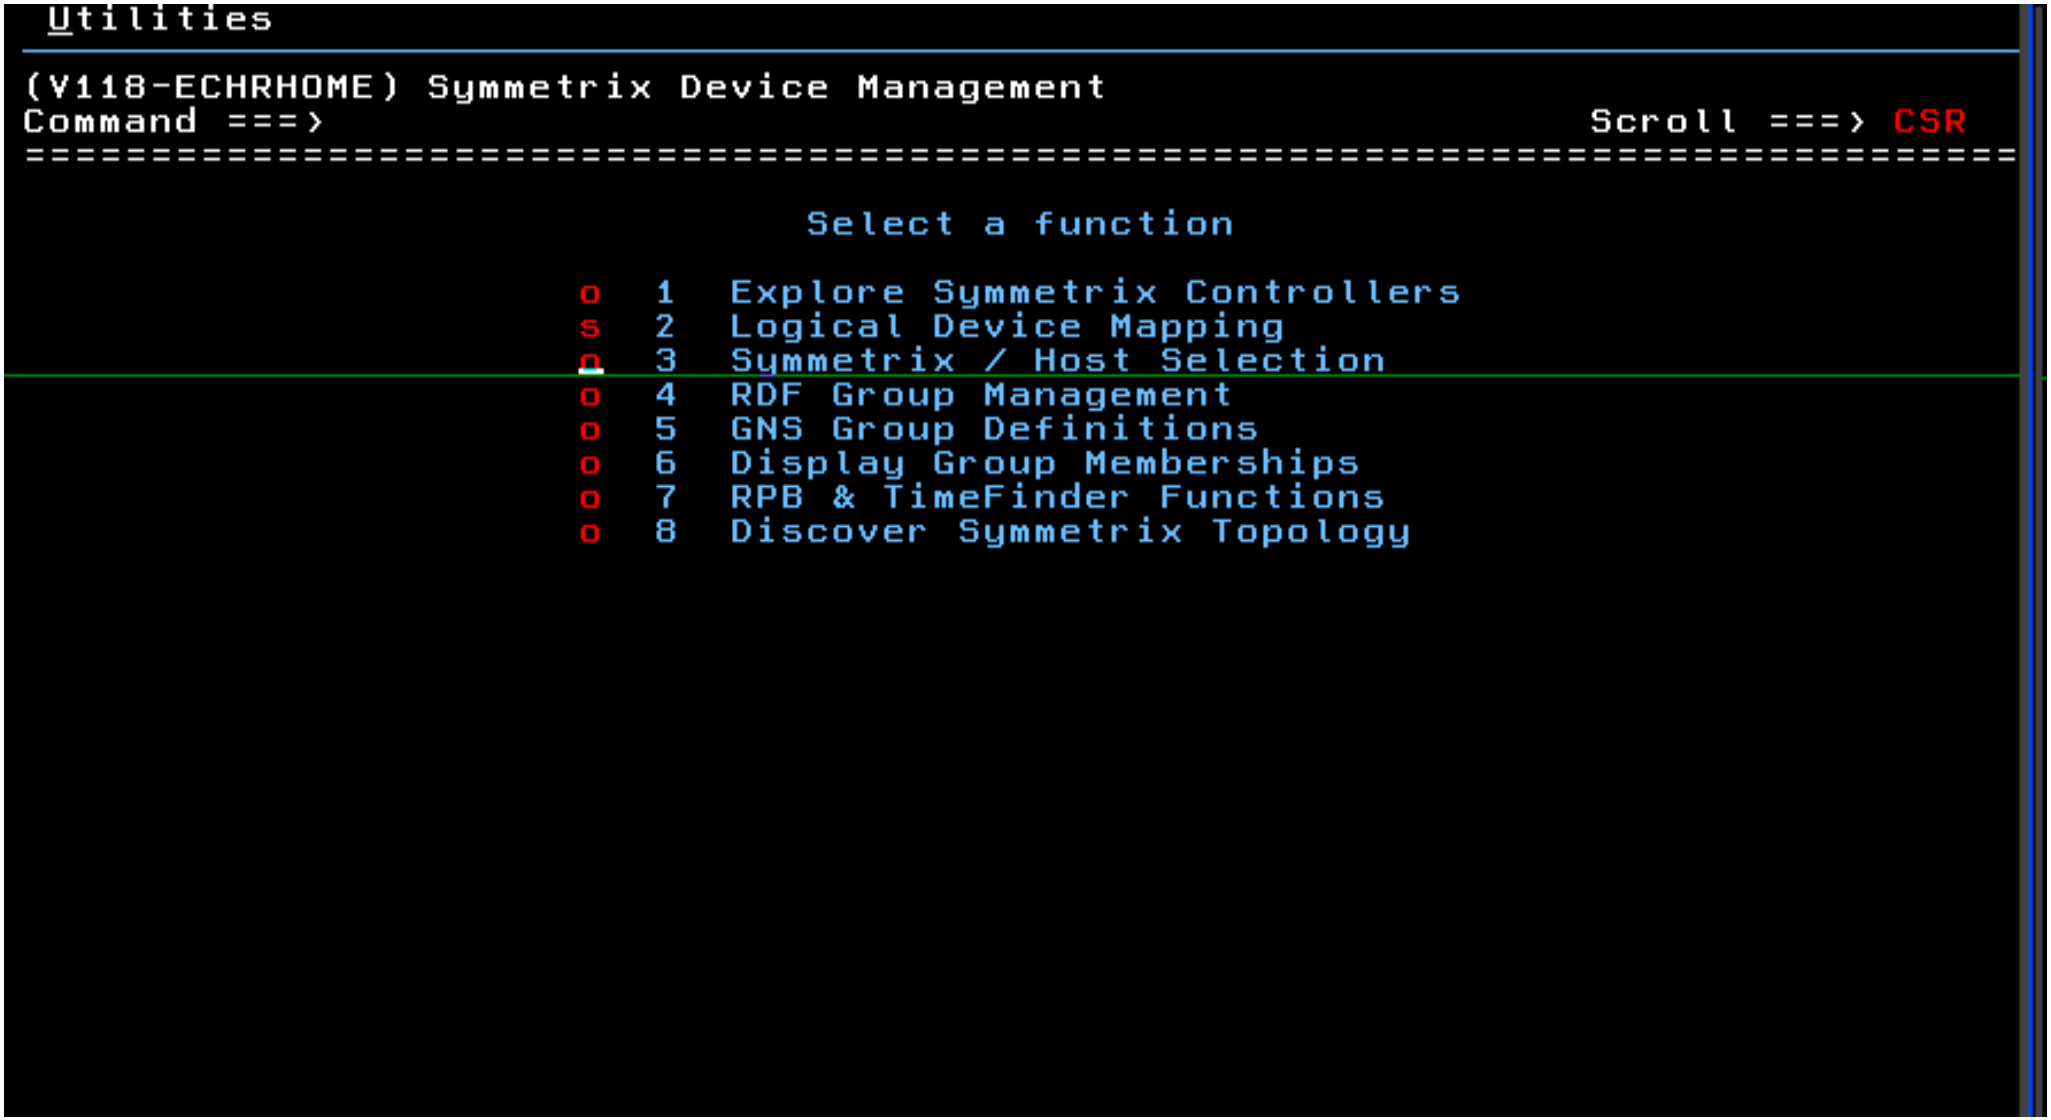

Summary and Questions

- Linux on System z Connectivity options – FICON or FCP
- Linux on System z Disk options – CKD and/or FBA
- Integration of Symmetrix devices
- Management of Symmetrix system
  - Solutions Enabler
  - Symmetrix Management Console (SMC)
  - Host Component on z/OS
  - EMC products for TPF

Related Technical Documentation

- *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, Version 6.5.2 Installation Guide P/N 300-002-289 REV A06*
- *Native Multipath Failover Based on DM-MPIO for v2.6x Linux Kernel and EMC Storage Arrays, Red Hat Enterprise Linux 4 - U3 and SuSE Linux Enterprise Server 9 - SP3, Configuration Guide*

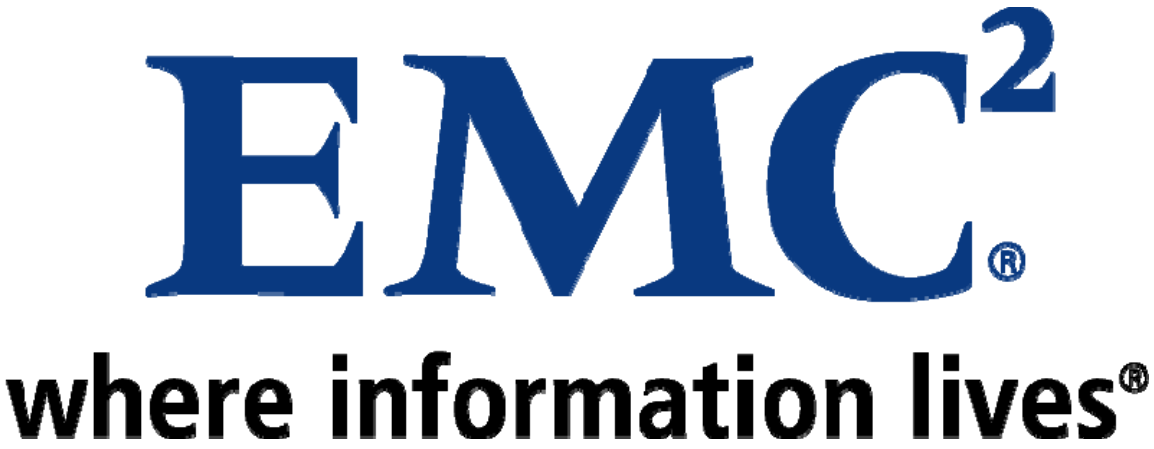
EMC²
where information lives®