An Introduction to SAN and Fixed Block Disk for ECKD Users

Samuel D. Cohen
IBM zSystems Consultant
Levi, Ray & Shoup, Inc
sam.cohen@lrs.com
(217) 862-9227

Overview

- z/OS uses primarily Variable-Block disk storage (aka Enhanced Count-Key-Data (ECKD)), while other zSystem operating systems can use ECKD and 512-byte Fixed Block (FB) disk.
- Accessing FB disk uses Fibre Channel instead of FICON
  - Same hardware
  - Different communication protocols
- This presentation will compare/contrast accessing ECKD disk vs. FB disk

Background: IOCDS

- I/O Configuration Dataset (IOCDS) defines pathing between zSystem and peripheral devices
- Path Management by I/O Subsystem (separate system within IBM zSystems
- Primary Entries in IOCDS
  - CHPID (hardware connection)
  - CNTLUNIT (control unit to receive I/O and route to device)
    - Connected to CHPID(s)
  - IOADDR (individual storage devices)
    - Connected to CNTLUNIT

Background: FCP

- Fibre Channel Protocol (FCP) attachment is handled differently from FICON attachment
  - FCP port is assigned a World-Wide Port Name (WWPN)
    - One WWPN per port
    - An FC Port is to a WWPN as an OSA Port is to a MACADDR
      - Same concept
    - IBM zSystems define WWPNs based on CPU Serial number and PCHID
    - Multiple subchannels available but every subchannel sees the same traffic because traffic is routed between WWPNs on both sides of the connection
  - No path management by I/O subsystem
  - Path management by zSystems operating system(s)

Background: NPIV

- How do you keep traffic different FCP subchannels from seeing traffic on all other subchannels?
- Virtualization!!
- N_Port ID Virtualization (NPIV) creates a virtual WWPN for each channel
  - Limited to 64 subchannels per FCP port in z15
- Using NPIV, traffic between an FCP subchannel and a disk subsystem will not be seen by any other FCP subchannel
  - Traffic could be seen at the disk subsystem channel interface unless it also uses NPIV

Background: SAN “Fabric”

- IBM zSystems cannot directly connect to FC HBAs
  - Must use a SAN switch that is certified for use with zSystems
- SAN provides the path management between FC-attached devices
- “Zoning” is the process of pairing these FC attachments
  - WWPNs are used in the zoning process
  - Not limited to a single point-to-point definition at each end
    - Can have 1:1, 1:many, many:1 or many:many
    - Pathing is managed by a multipath background process/started task/daemon in the host operating system
      - Configured by querying the SAN and devices attached at other end of the zone
- Usually want 2 separate fabrics for redundancy

ECKD Storage Devices

- 1 or more Hardware Bus Adapters (e.g. FICON channel)
- Pre-defined Logical Control Units (LCUs = CUADDR on IOCDS)
- Pre-defined Unit Addresses (0-255 per LCU)
  - Size of each logical disk is pre-defined
- IOCDS should have configuration statements matching pre-defined definitions in disk subsystem

FB Storage Devices

- 1 or more Hardware Bus Adapters (e.g. FC channel)
- Pre-defined Host Addresses (WWPNs)
- **No** Logical Control Units
- Pre-defined Logical Units (LUNs)
  - Size of each LUN is pre-defined
- SAN Zones providing the linkage between zSystems and the disk HBAs
- Disk subsystem definitions for the zSystems WWPNs that will be accepted and matched with local LUNs

Lost Yet?

Examples from a real system

- IOCDS source for an FCP channel

```
CHPID2D  CHPID PATH=(CSS(0),20),TYPE=FCP,PART=((PROD,TEST)),PCHID=100
CU6600   CNTLUNIT CUNUMBR=2000,PATH=20,UNIT=FCP
DEV2000  IODEVICE ADDRESS=(2000,64),CUNUMBR=(2000),UNIT=FCP
```

- z/VM WWPN Displays with NPIV active (2 LPARs)

```
q fcp wwpn 2000
FCP  2000      NPIV WWPN C05076D691800380
     CHPID 20  PERM WWPN C05076D691801141
     ATTACHED TO LNXUTILS

q fcp wwpn 2000
FCP  2000      NPIV WWPN C05076D691800400
     CHPID 20  PERM WWPN C05076D691801141
     FREE
```

Examples from a real system

- SAN Fabric Definitions
  - Aliases: Giving Names to WWPNs

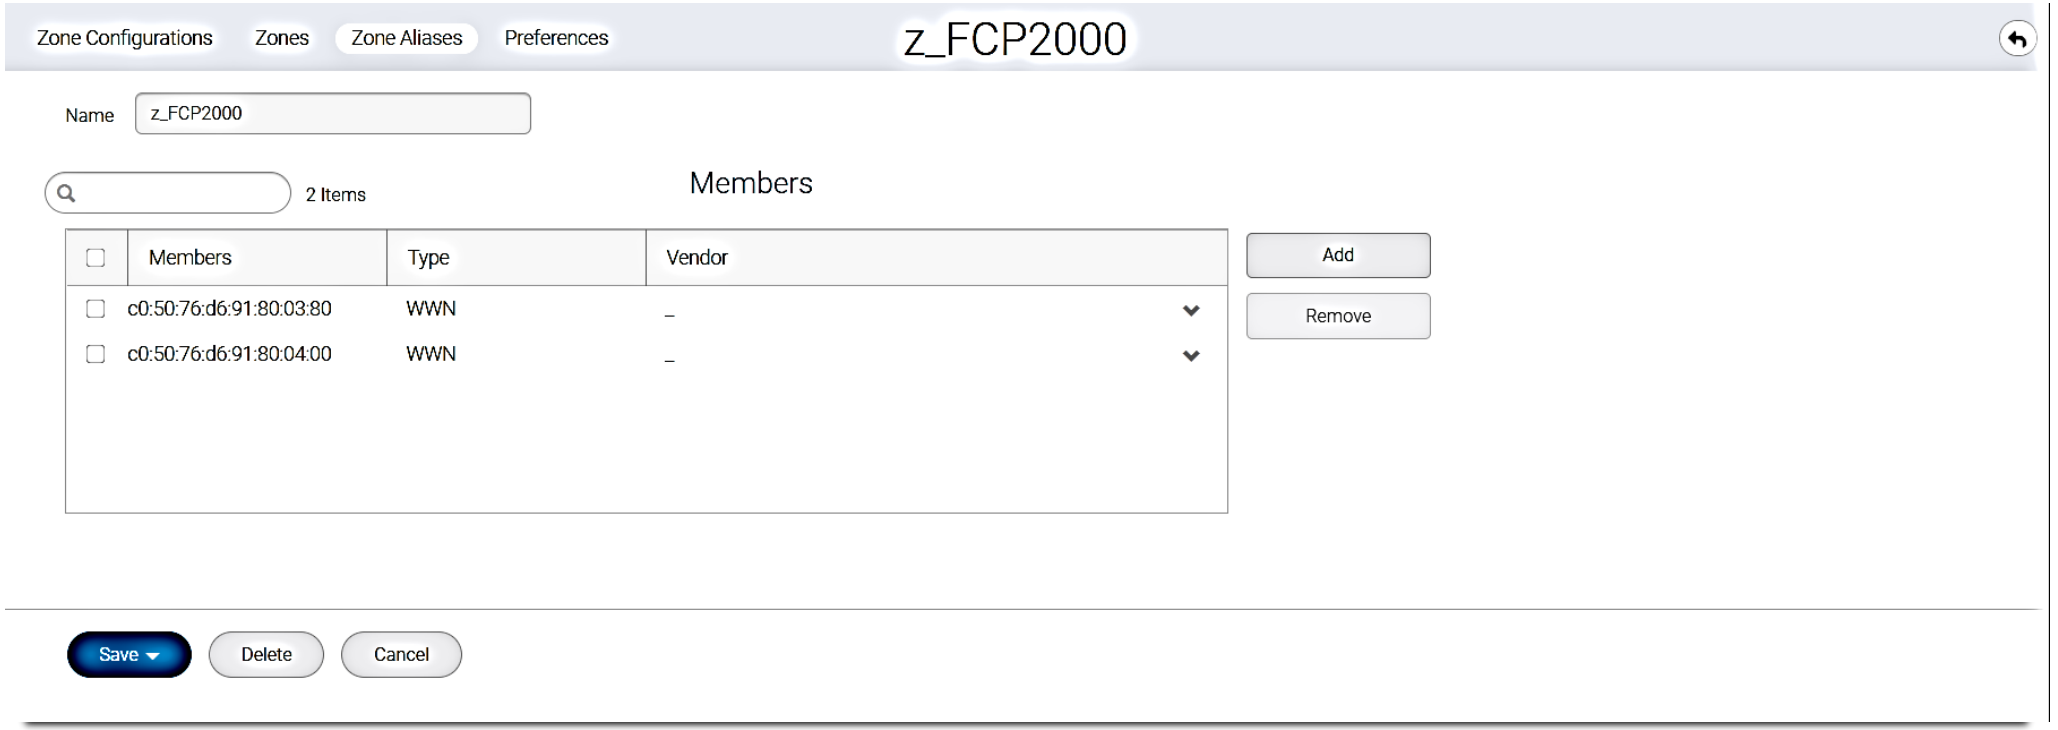

Examples from a real system

- SAN Fabric Definitions
  - Zones: Linking Aliases to create a path between connections

Zone Configurations | Zones | Zone Aliases | Preferences

z_FCP00_FS5030f

Name: z_FCP00_FS5030f

Type: Standard

4 Items

Members

Members	Type
FS5030f_node1_p1_NPIV	ALIAS
FS5030f_node2_p1_NPIV	ALIAS
z_FCP2000	ALIAS
z_FCP2200	ALIAS

Add

Remove

Save | Delete | Cancel

Examples from a real system

- Storage Subsystem
  - Host: Defining who can connect
    - Not needed with ECKD

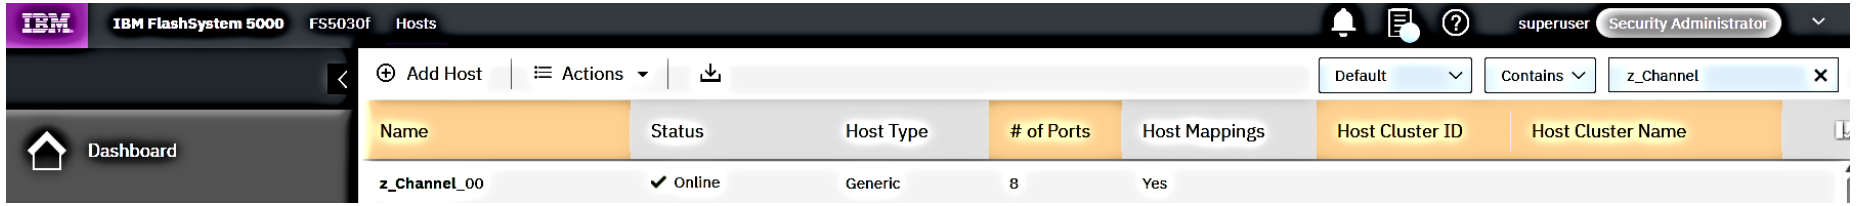

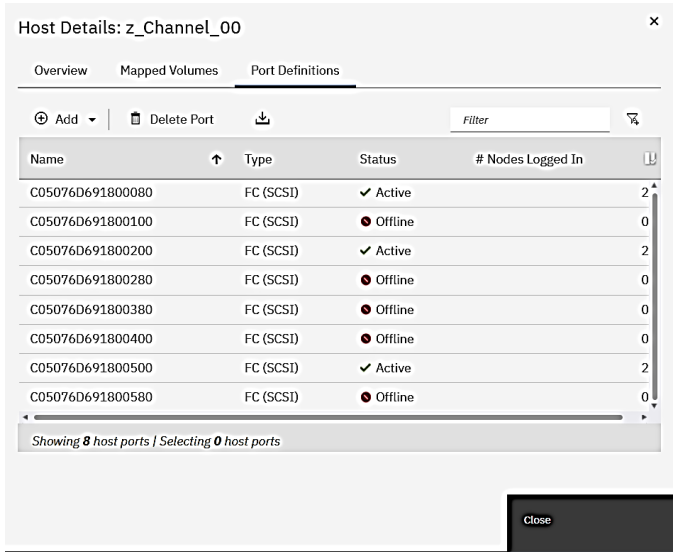

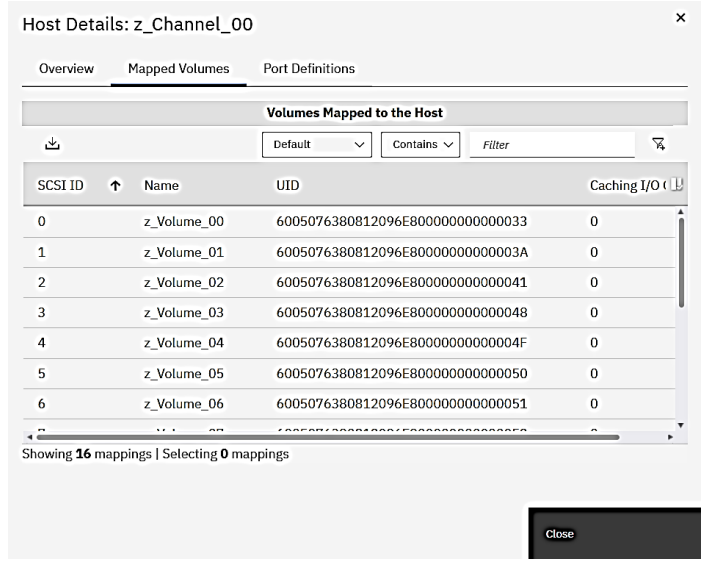

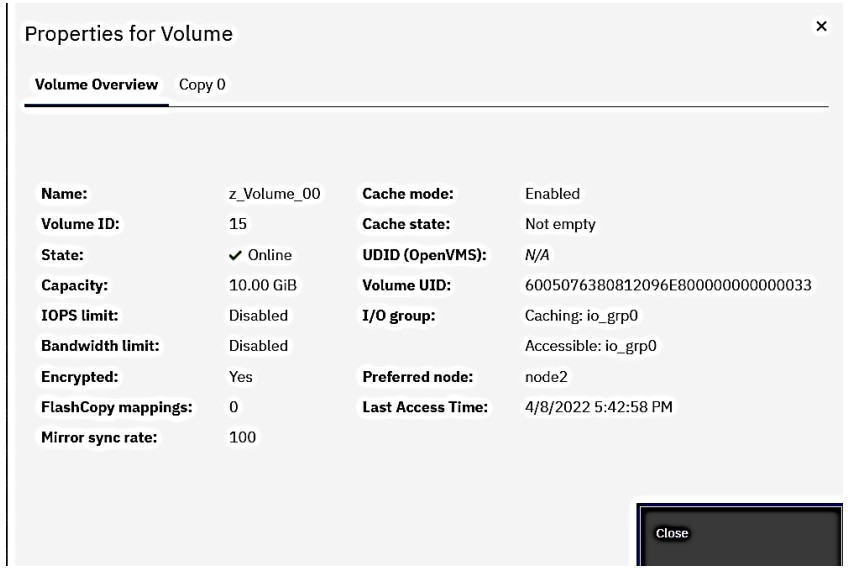

Questions?