



IBM z/VM Single System Image and Live Guest Relocation Planning and Use

John Franciscovich

IBM z/VM Development, Endicott, NY

Alan Altmark

IBM Systems Lab Services, Endicott, NY



SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.



Topics

- Introduction
 - z/VM Single System Image Clusters
 - Live Guest Relocation (LGR)
- Planning for your SSI Cluster
- Configuring your SSI Cluster
- Relocation Domains
- Planning for Live Guest Relocation (LGR)

Introduction to SSI and LGR

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation

Multi-system Virtualization with z/VM Single System Image (SSI)

- VMSSI Feature of z/VM 6.2 and 6.3
- Up to 4 z/VM instances (members) in a single system image (SSI) cluster
 - Same or different CECs
- Provides a set of shared resources for the z/VM systems and their hosted virtual machines
 - Managed as a single resource pool
- **Live Guest Relocation** provides virtual server mobility
 - Move Linux virtual servers (guests) non-disruptively from one member of the cluster to another
- A single SSI cluster can have members running both z/VM 6.2 and 6.3

z/VM Single System Image (SSI) Cluster



- Common resource pool accessible from all members
 - Shared disks for system and virtual server data
 - Common network access
- All members of an SSI cluster are part of the same ISFC collection
- CP validates and manages all resource and data sharing
 - Uses ISFC messages that flow across channel-to-channel connections between members
 - No virtual servers required

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation



Benefits and Uses of z/VM SSI Clusters

- Horizontal growth of z/VM workloads
 - Increased control over server sprawl
 - Distribution and balancing of resources and workload
- Flexibility for planned outages for service and migration
 - z/VM or hardware
 - Less disruptive to virtual server workloads
- Workload testing
 - Different service/release levels
 - Various environments (stress, etc.)
 - New/changed workloads and applications can be tested before moving into production
- Simplified service management of a multi-z/VM environment
 - Concurrent installation of multiple-system cluster
 - Single maintenance stream
 - Reliable sharing of resources and data

Reliability and Integrity of Shared Data and Resources



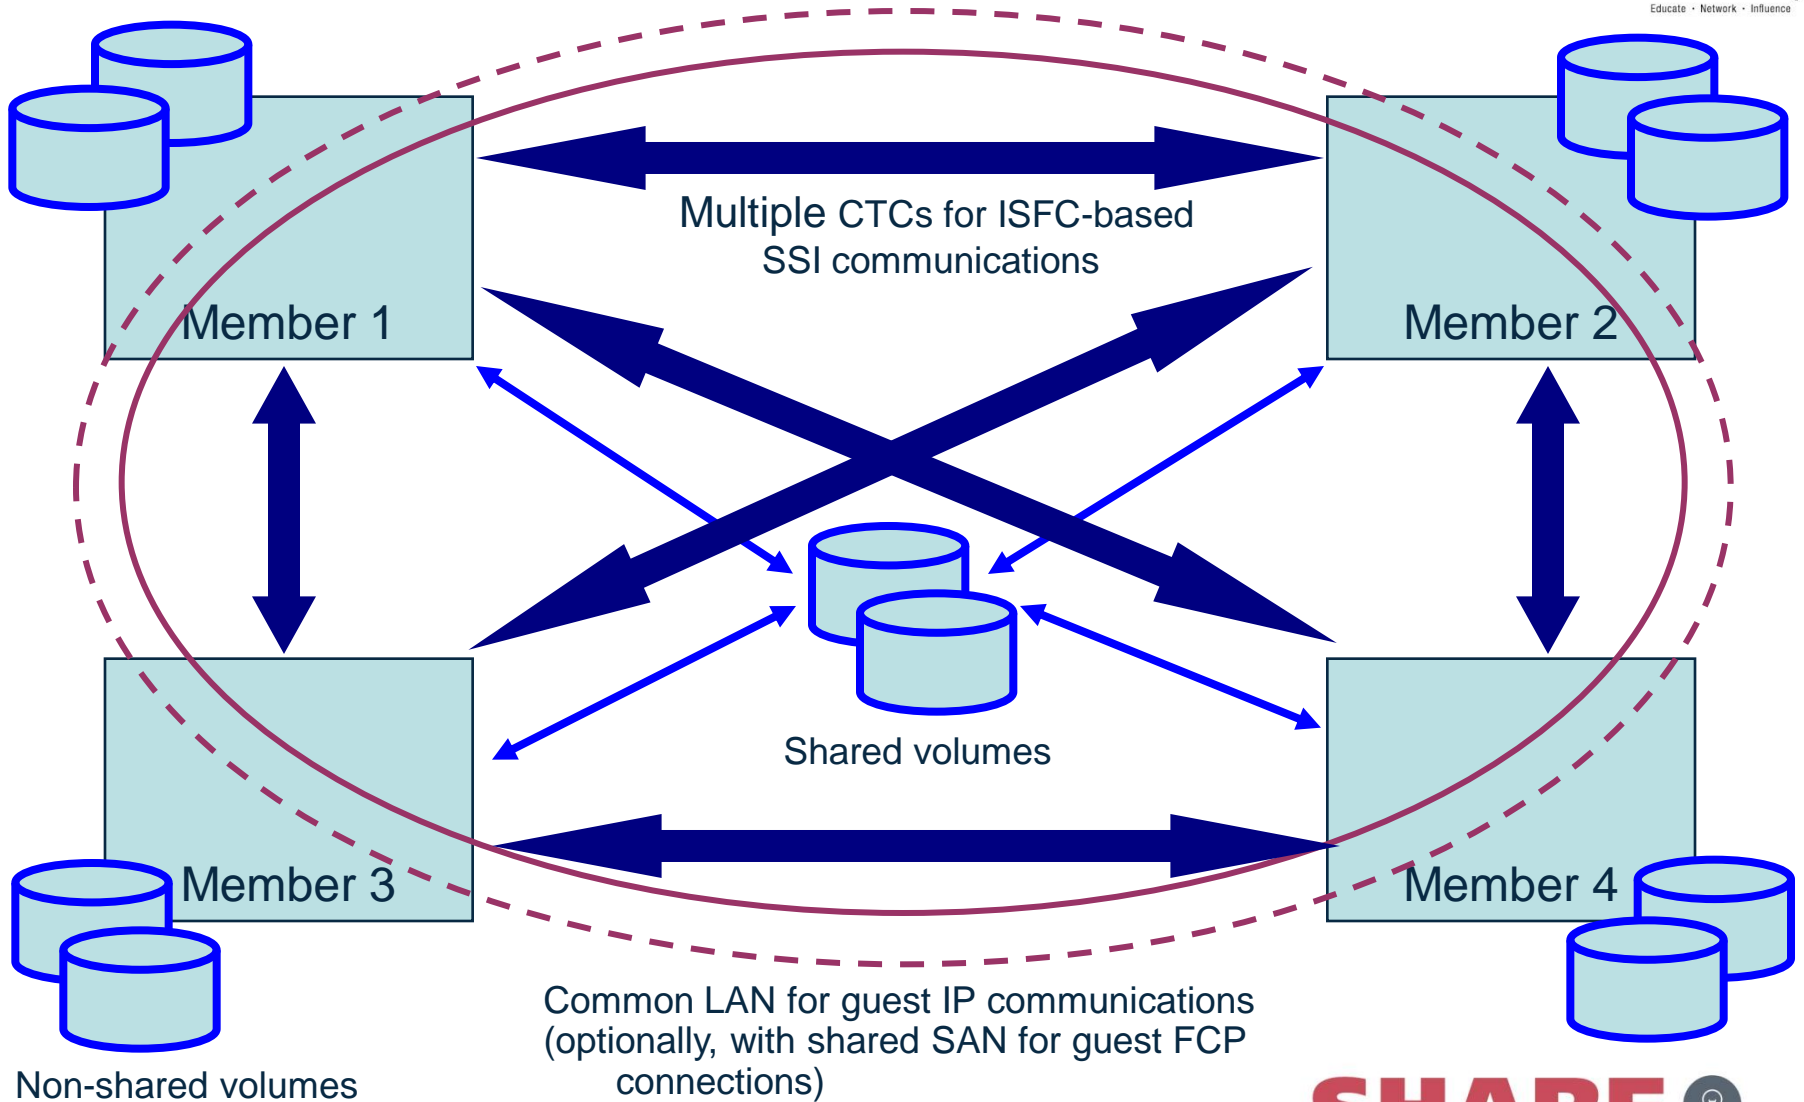
- Normal operating mode
 - All members communicating and sharing resources
 - Guests have access to same resources on all members
- Cluster-wide policing of resource access
 - Volume ownership marking
 - Coordinated minidisk link checking
 - Automatic minidisk cache management
 - Single logon enforcement
- Unexpected failure causes cluster to move to “SAFE” mode
 - Communications or catastrophic member failure
 - Existing running workloads continue to run
 - New access to shared resources is denied until failure is resolved
 - Most failures are resolved automatically
 - Manual intervention is rare

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation



z/VM SSI Cluster



Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Cross-System Minidisk Management



- Minidisks can either be shared across all members or restricted to a single member
 - CP checks for conflicts throughout the cluster when a link is requested
- Automatic management of minidisk cache
 - Created and deleted as appropriate based on minidisk links throughout the cluster
- Virtual reserve/release for fullpack minidisks is supported across members
 - Only supported on one member at a time for non-fullpack minidisks
- Volumes can be shared with systems outside the SSI cluster
 - **SHARED YES** on RDEVICE statement or SET RDEVICE command
 - **Link conflicts must be managed manually**
 - SSI and XLINK are **not** compatible
 - Not eligible for minidisk cache
 - **Use with care**

Cross-System Spool

- Spool files are managed cooperatively and shared among all members of an SSI cluster
- Single-configuration virtual machines (most users) have a single logical view of all of their spool files
 - Access, manipulate, and transfer all files from any member where they are logged on
 - Regardless of which member they were created on
- Multiconfiguration virtual machines do not participate in cross-system spool
 - Each instance only has access to files created on the member where it is logged on
- All spool volumes in the SSI cluster are shared (R/W) by all members
 - Each member creates files on only the volumes that it owns
 - Each member can access and update files on all volumes

SLOT	VOL-ID	RDEV	TYPE	STATUS	SSIOWNER	SYSOWNER
10	M01S01	C4A8	OWN	ONLINE AND ATTACHED	CLUSTERA	VMSYS01
11	M02S01	C4B8	SHARE	ONLINE AND ATTACHED	CLUSTERA	VMSYS02
12	M01S02	C4A9	OWN	ONLINE AND ATTACHED	CLUSTERA	VMSYS01
13	M02S02	C4B9	SHARE	ONLINE AND ATTACHED	CLUSTERA	VMSYS02
14	M01S03	C4AA	DUMP	ONLINE AND ATTACHED	CLUSTERA	VMSYS01
15	M02S03	C4BA	DUMP	ONLINE AND ATTACHED	CLUSTERA	VMSYS02
16	-----	----	-----	RESERVED	-----	-----

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Cross-System CP Commands

- **AT** *command* can be used to issue most privileged commands on a different active member

AT sysname CMD cmdname

- **AT sysname** *operand* can be used to target virtual machines on different active member(s)
 - MESSAGE (MSG)
 - MSGNOH
 - SEND
 - SMSG
 - WARNING

MSG userid AT sysname

- Single-configuration virtual machines are usually found wherever they are logged on
- Multiconfiguration virtual machines require explicit targeting
- CMS TELL and SENDFILE commands require RSCS in order to communicate with multiconfiguration virtual machines on other members

Real Device Management

- Unique identification of real devices within an SSI cluster
 - Ensures that all members are using the same physical devices where required
- CP generates an equivalency identifier (EQID) for each disk volume and tape drive
 - Physical device has same EQID on all members
- EQID for network adapters (CTC, FCP, OSA, Hipersockets) must be defined by system administrator
 - Connected to same network/fabric
 - Conveying same access rights
- EQIDs used to select equivalent device for live guest relocation and to assure data integrity

Live Guest Relocation

- Relocate a running Linux virtual server (guest) from one member of an SSI cluster to another
 - Load balancing
 - Move workload off a member requiring maintenance
- **VMRELOCATE** command initiates and manages live guest relocations
 - Check status of relocations in progress
 - Cancel a relocation in progress
 - Relocations are **NOT** automatically done by the system
- Guests continue to run on source member while they are being relocated
 - Briefly quiesced
 - Resumed on destination member
- If a relocation fails or is cancelled, the guest continues to run on the source member
- Live Guest Relocation is supported only for Linux guests

Live Guest Relocation ...

- Relocation capacity and duration is determined by various factors including:
 - ISFC bandwidth
 - Guest size and activity
 - Destination system's memory size and availability
 - Load on destination system
- In order to be relocated, a guest must meet eligibility requirements, including:
 - The architecture and functional environment on destination member must be comparable
 - Relocation domains can be used define sets of members among which guests can relocate freely
 - Devices and resources used by the guest must be shared and available on the destination member
 - Is it "safe" to move the guest?
 - Eligibility checks are repeated multiple times during a relocation

Planning for your SSI Cluster

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation


SSI Cluster Requirements

- Shared and non-shared DASD
 - All volumes cabled to all members so that configuration problems can be fixed more easily.
 - 3390 volume required for the PDR
 - Shared 3390s required for RACF database (full-pack minidisks)
- 2-4 FICON paths between each pair of LPARs
 - 8-16 CTC devices (4 CTCs per path)
 - Provides direct ISFC links from each member to all other members
 - Point-to-point or FICON switch/director
- Compatible LAN and SAN connectivity
 - OSA access to the same LAN segments
 - FCP access to same ports and devices
- Shared or synchronized system configuration file for all members
- Shared or synchronized source user directory
- Capacity planning for each member of the SSI cluster
 - Ensure sufficient resources are available to contain shifting workload
 - Guests that will relocate
 - Guests that logon to different members

SSI Cluster Layout

1. How many members in your cluster?
2. Production configuration
 - How many CECs?
 - How many LPARS/CEC?
 - *Suggested configuration for 4-member cluster is 2 LPARs on each of 2 CECs*
3. Test configuration
 - VM guests?
 - LPARs?
 - Mixed?
4. Virtual server (guest) distribution
 - Each guest's "home" member?
 - Where can each guest be relocated?
 - *Distribute workload so each member has capacity to receive relocated guests*
 - CPU
 - Memory

- Each member of an SSI cluster must have a direct ISFC connection to every other member (logical link)




- Logical links are composed of 1-16 CTC connections
 - FICON channel paths
 - May be switched or unswitched
- CTCs distributed on two or more FICON channel paths between each pair of members
 - Avoids write collisions that affect link performance
 - Avoids severing logical link if one channel path is disconnected or damaged
- *Recommended practice:* Use same real device number for same CTC on each member
- *Recommended practice:* Plan and test out your CTCs BEFORE installing your SSI

CTC Connections – How Many Do I Need?

- 4 CTC devices per per FICON chpid provides most efficient ISFC data transfer

6000 to
6003



6000 to
6003

- For large guests, relocation and quiesce times improve with more chpids


6020 to
6023



6020 to
6023

- Up to 4 chpid paths, with 4 CTCs each


6040 to
6043



6040 to
6043


- *Additional factors affect relocation and quiesce times*

6060 to
6063



6060 to
6063

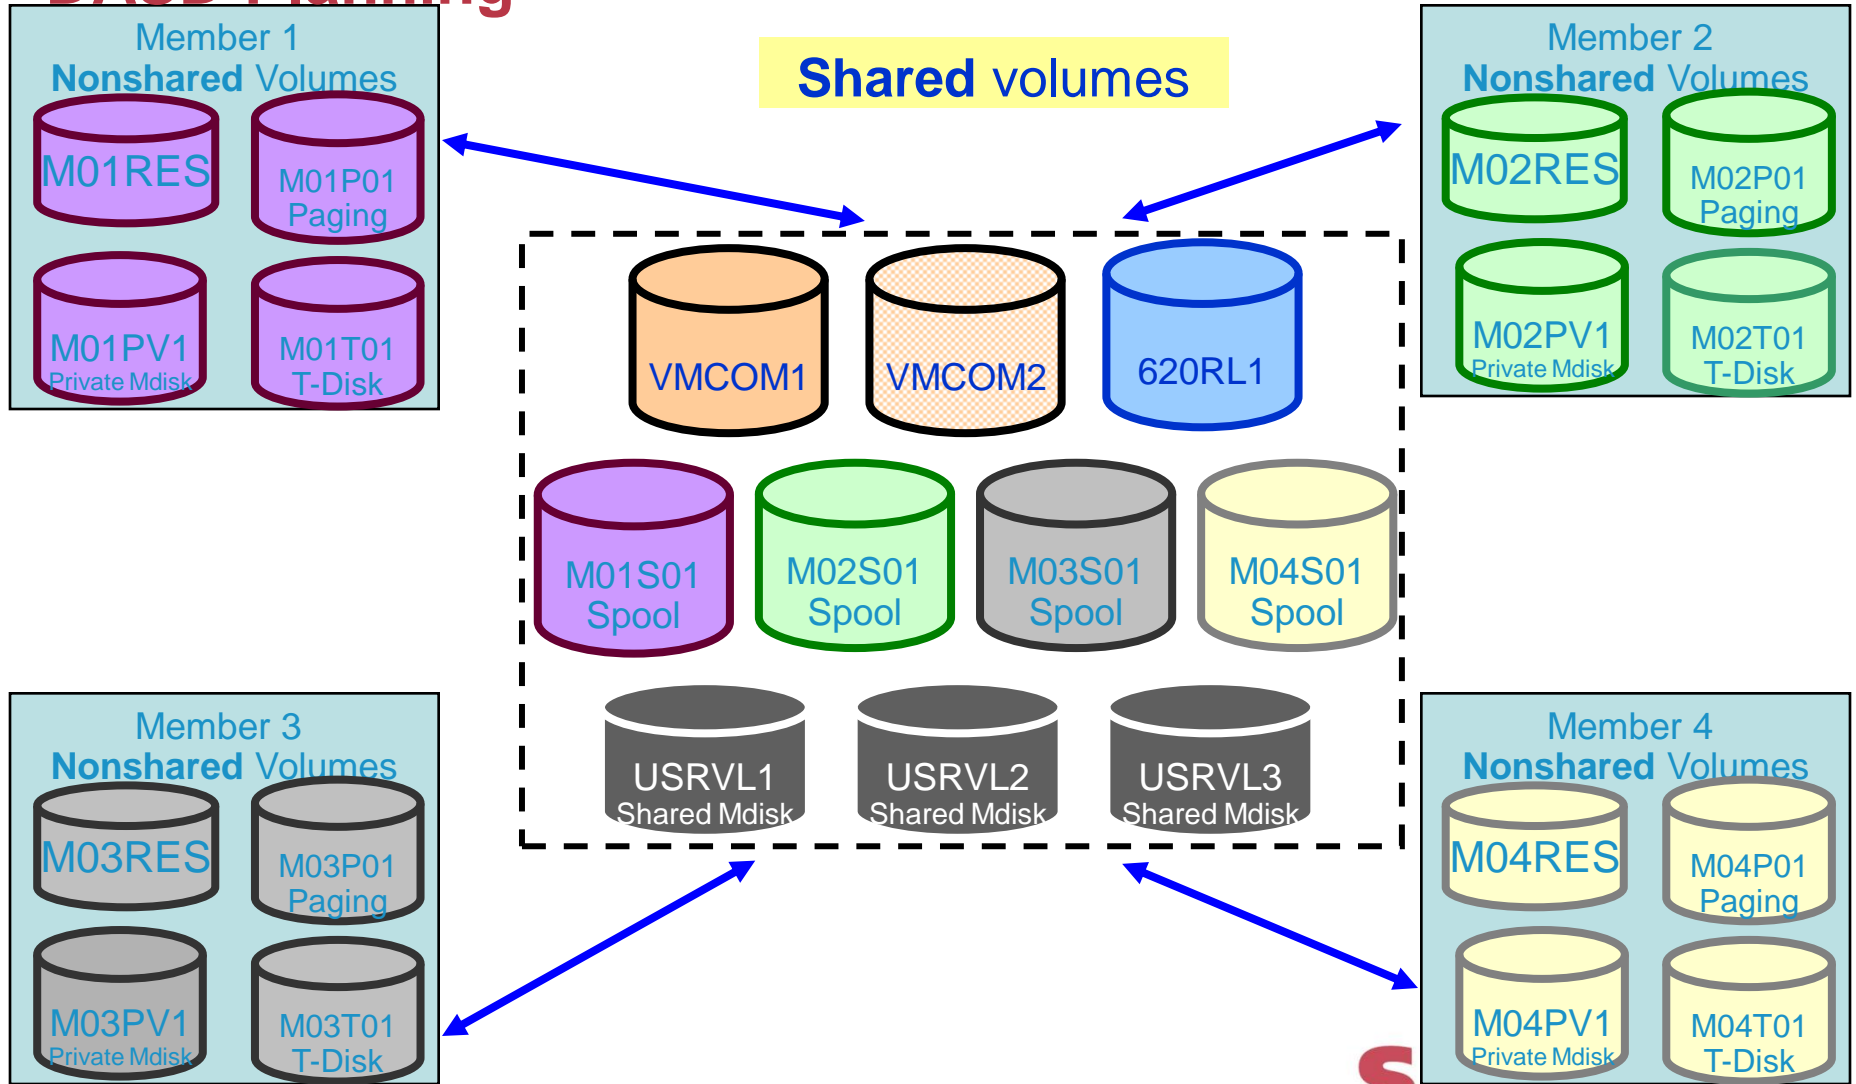
System Volumes and Minidisks – SSI Multiple Members



Complete your session evaluations online at www.ibm.com/eval

SHARI
in Orlando 2015

DASD Planning



Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Which Type of Installation Should I Choose?

- SSI Installation
 - Single installation for multiple z/VM images
 - Can also install a single system configured as an SSI member
 - Installed and configured as an SSI cluster
 - Single source directoryj
 - Shared system configuration file
 - Creates Persistent Data Record (PDR) on Common volume
 - SSI Installation to SCSI devices is not supported
 - Guests may use SCSI devices
- Non-SSI Installation
 - Single z/VM image
 - Can be converted to initial member of an SSI cluster later
 - Builds DASD layout, directory, and configuration file the same as SSI installation
- Both types of installation are different from previous releases of z/VM
 - Userids
 - Disks
 - Directory
 - System configuration file
- Review documented migration scenarios before deciding whether to do SSI or non-SSI install
 - CP Planning and Administration
 - SSI installation primarily for new or "from scratch" installs

Upgrade Installation

- New technique for upgrading from z/VM 6.2 to 6.3
 1. Install new release system as temporary second level guest of system being upgraded
 2. Move new code level to current system

- Current (to be upgraded) system can be either:
 - Non-SSI
 - Only member of single-member SSI cluster
 - Any member of a multi-member SSI cluster

- In a multi-member SSI cluster, only one member at a time is upgraded
 - Minimum impact to the cluster and other members
 - Can thoroughly test new release on one member before upgrading other members

Configuring your SSI Cluster

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation

Shared System Configuration File

- Resides on new shared parm disk
 - PMAINT CF0
- System_Identifier statements for each member

```
System_Identifier LPAR LP01 MEMBER1
System_Identifier LPAR LP02 MEMBER2
```

- System_Residence statements for each member

```
MEMBER1: System Residence,
Checkpoint Valid M01RES From CYL 21 For 9 ,
Warmstart Valid M01RES From CYL 30 For 9
MEMBER2: System Residence,
Checkpoint Valid M02RES From CYL 21 For 9 ,
Warmstart Valid M02RES From CYL 30 For 9
```

- Identify direct ISFC links between members
 - Must have at least one ISLINK from each member to every other member

```
MEMBER1: BEGIN
ACTIVATE ISLINK 912A /* Member 1 TO Member 2 */
ACTIVATE ISLINK 913A /* Member 1 TO Member 3 */
ACTIVATE ISLINK 914A /* Member 1 TO Member 4 */
MEMBER1: END
```

Shared System Configuration File...

- Define cluster configuration

```
SSI MYCLUSTR PDR_Volume VCOM1 ,  
Slot 1 MEMBER1,  
Slot 2 MEMBER2,  
Slot 3 Available,  
Slot 4 Available
```

- A system that is configured as a member of an SSI cluster joins the cluster during IPL
 - Verifies that its configuration is compatible with the cluster
 - Establishes communication with other members
- Members leave the SSI cluster when they shut down

DASD - Changes to the System Configuration File

- CP_OWNED list defines CP_OWNED volumes for **all** members:

```
/*
*****
/*
*****
MEMBER1: CP_Owned Slot 1 M01RES
MEMBER2: CP_Owned Slot 1 M02RES
*****
/*
*****
/*
*****
COMMON VOLUME
*****
/*
*****
CP_Owned Slot 5 VMCOM1
*****
/*
*****
```

- The User_Volume_List is now split between shared and private

```
/*
*****
/* Shared User Volumes
*****
User_Volume_List 620RL1 620RL2 USRVL1
*****
/*
*****
/* User volumes for local minidisks
*****
MEMBER1: User_Volume_List M01W01 M01PV1
MEMBER2: User_Volume_List M02W01 M02PV1
*****
/*
*****
```

DASD – CP Volume Ownership



- Each CP-owned volume in an SSI cluster will be marked with ownership information
 - Cluster name
 - System name of the owning member
 - The marking is created using CPFMTXA
- Ensures that one member does not allocate CP data on a volume owned by another member
 - Warm start, checkpoint, spool, paging, temporary disk, directory
- By the time AUTOLOG1 starts, it's too late!

DASD – CP Volume Ownership

- Link the full pack overlay for each CP-Owned volume
- Use CPFMTXA to add ownership information to each CP-Owned volume
 - Cluster name
 - System name of owning member

<u>Volume</u>	<u>Full Pack Overlay</u>	<u>Owner</u> <u>(CLUSTER.MEMBER)</u>
M01RES	MAINT 123	MYCLUSTER.MEMBER1
VMCOM1	PMAINT 141	MYCLUSTER.NOSYS
M01S01	MAINT 122	MYCLUSTER.MEMBER1
M01P01	\$PAGE\$ A01	MYCLUSTER.MEMBER1

Persistent Data Record (PDR)



- Cross-system serialization point on disk
 - Must be a shared 3390 volume (VMCOM1)
 - Created and viewed with new FORMSSI utility
- Contains information about member status
 - Used for health-checking
 - Ensures that a stalled or stopped member can be detected
- To query the PDR
 - **LINK** the fullpack overlay of VMCOM1, **PMaint 141**

Persistent Data Record (PDR)

formssi display 141

```
HCPPDF6618I Persistent Data Record on device 0141 (label VMCOM1) is for MYCLUSTR
HCPPDF6619I PDR state: Unlocked
HCPPDF6619I time stamp: 05/31/13 16:39:06
HCPPDF6619I cross-system timeouts: Enabled
HCPPDF6619I PDR slot 1 system: MEMBER1
HCPPDF6619I state: Joined
HCPPDF6619I time stamp: 05/31/13 16:40:03
HCPPDF6619I last change: MEMBER1
HCPPDF6619I PDR slot 2 system: MEMBER2
HCPPDF6619I state: Joined
HCPPDF6619I time stamp: 05/31/13 16:39:52
HCPPDF6619I last change: MEMBER2
HCPPDF6619I PDR slot 3 system: MEMBER3
HCPPDF6619I state: Down
HCPPDF6619I time stamp: 05/26/13 20:31:22
HCPPDF6619I last change: MEMBER3
HCPPDF6619I PDR slot 4 system: MEMBER4
HCPPDF6619I state: Joined
HCPPDF6619I time stamp: 05/31/13 16:39:06
HCPPDF6619I last change: MEMBER4
```


Shared Source Directory

- All user definitions in a single shared source directory
 - No system affinity (SYSAFFIN)!
- Run DIRECTXA on each member
 - Identical object directories on each member
- Single security context
 - Each user has same access rights and privileges on each member

Using a directory manager is strongly recommended!

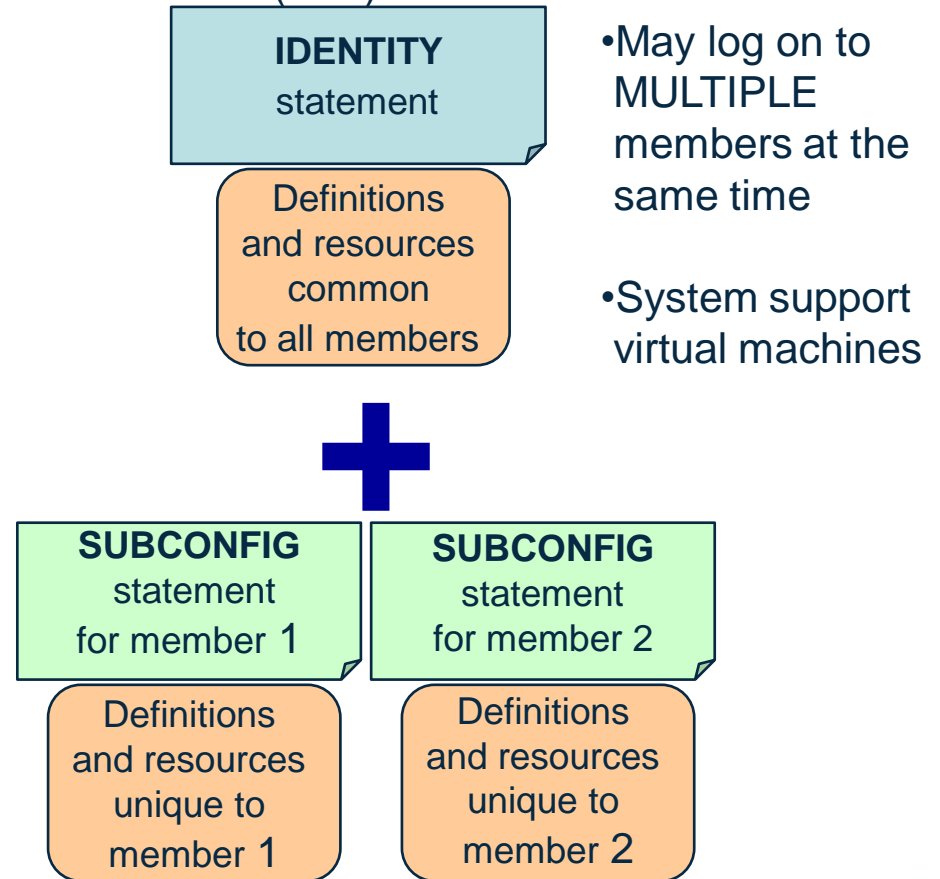
Virtual Machine Definitions

Single Configuration Virtual Machine (traditional)



- May log on to any ONE member at a time
- General Workload
 - Guest Operating Systems
 - Service virtual machines requiring only one logon in the cluster

Multiconfiguration Virtual Machine (new)



- May log on to MULTIPLE members at the same time
- System support virtual machines

Global and Local Disks

Global

- All members have access
- Cross-system LINK protection
- Dependent on STABLE cluster

Local

- Only one member has access
- No cross-system LINK protection
- Cluster state not relevant
- Log and work files

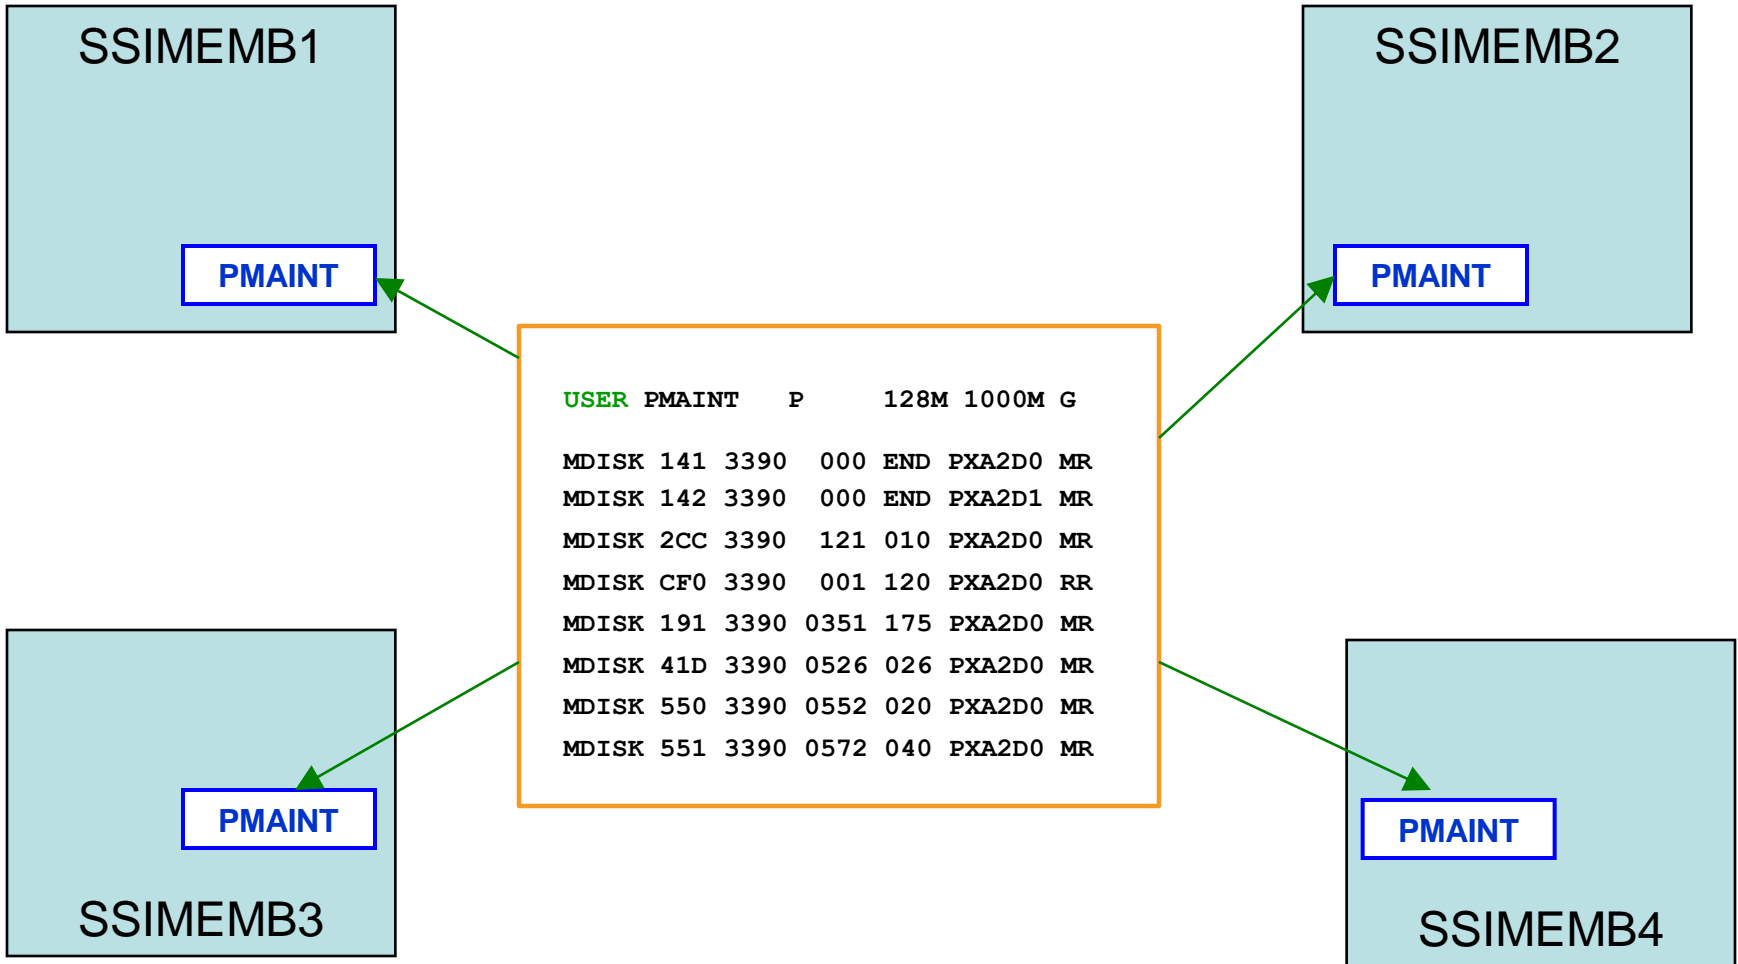
- Single configuration virtual machines typically have only **global** minidisks
- Multiconfiguration virtual machines typically have **local** minidisks
- OK to mix

New MAINT User IDs

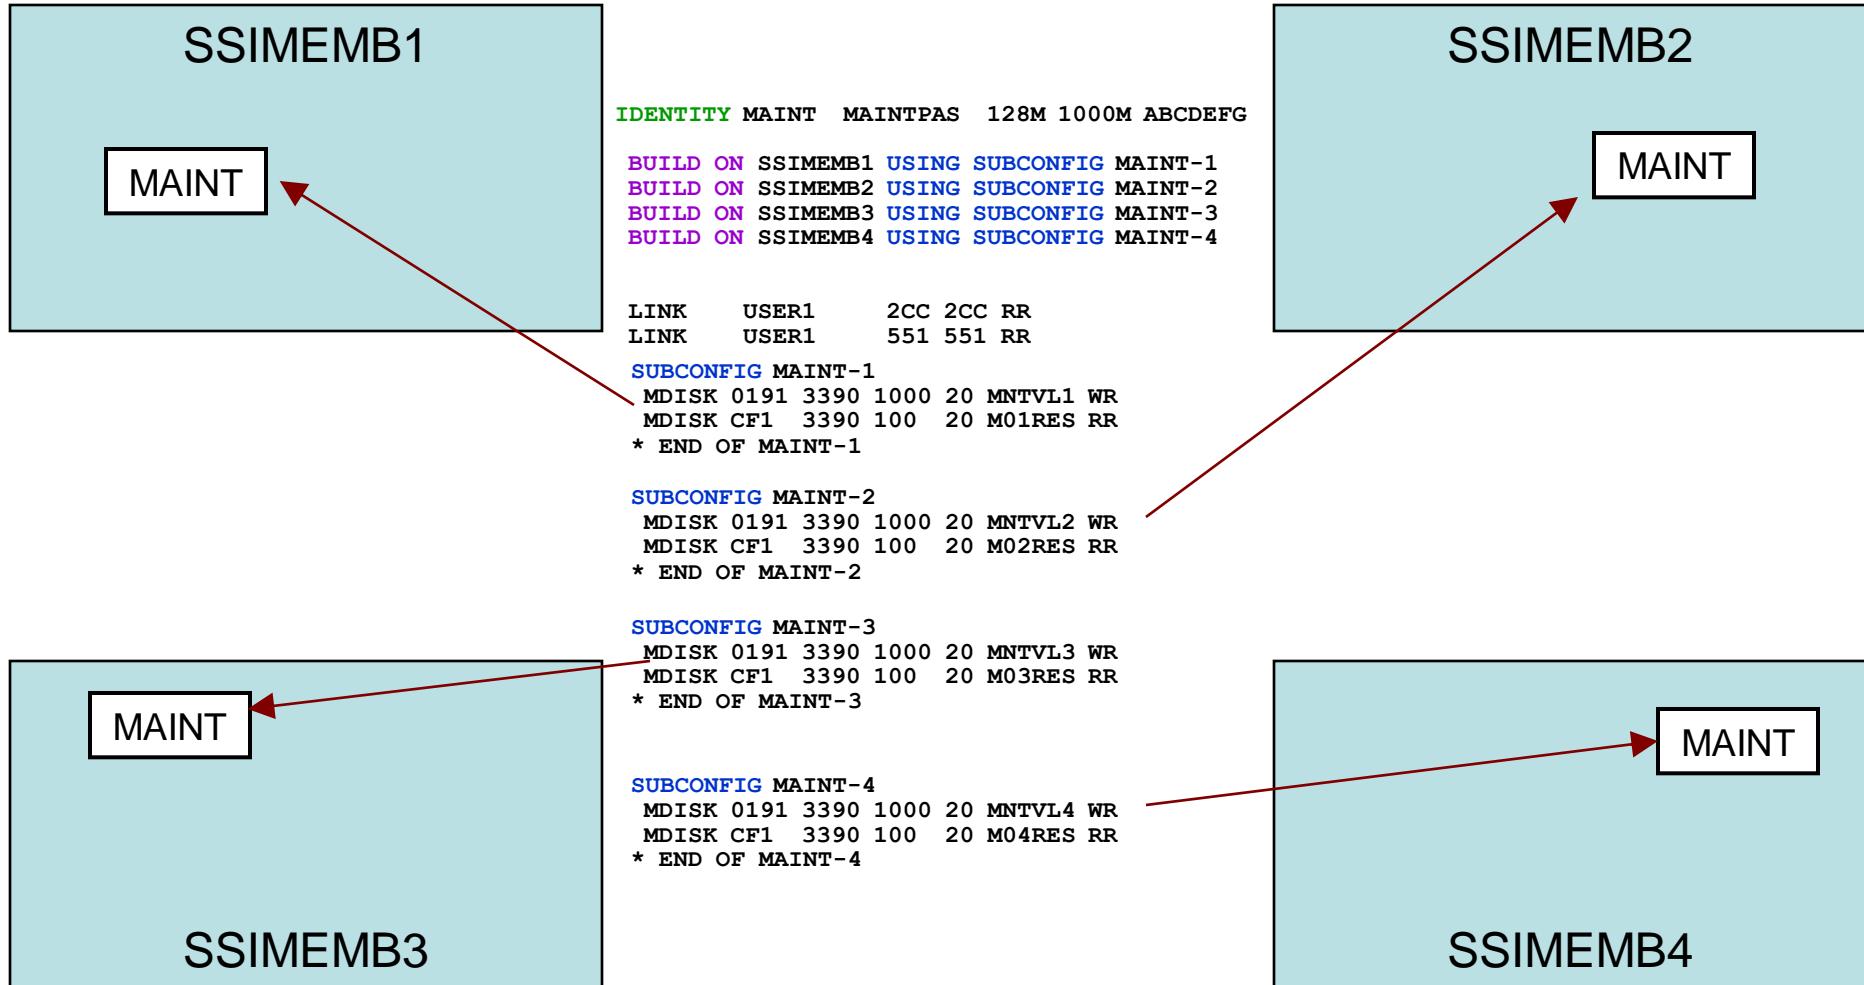
MAINT	PMAINT	MAINT630
Multi Configuration Virtual Machine	Single Configuration Virtual Machine	Single Configuration Virtual Machine
Owns CF1, CF3 parm disks, 190, 193, 19D, 19E, 401, 402, 990 CMS disks	Owns CF0 parm disk, 2CC directory disk, 550 and 551 utilities disks	Owns the service disks (e.g., 490, 493, 49D) and the CF2 backup parm disk
Use for work on a particular member, such as attaching devices, or relocating guests	Use for updating the SYSTEM CONFIG, or the source directory	Use for applying service

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Single Configuration Virtual Machines



Multiconfiguration Virtual Machines



Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Shared Source Directory - New Layout



- IBM-supplied directory is significantly different from prior releases
 - Both SSI and non-SSI installations
 - Directory for non-SSI installations will be in "SSI-ready" format to facilitate future SSI deployment
- Many of the IBM-supplied user IDs will be defined as multiconfiguration virtual machines
- Determine if any of your own guests should be defined as multiconfiguration virtual machines
 - Most will be single-configuration virtual machines
 - User IDs defined on `SYSTEM_USERIDS` statements must be multiconfiguration virtual machines
- Merge your user definitions into the IBM-supplied directory

Networks

- All members should have identical network connectivity
 - Connected to same physical LAN segments
 - Connected to same SAN fabric

- Assign equivalence identifiers (EQIDs) to all network devices
 - Devices assigned same EQID on each member must be
 - same type
 - have the same capabilities
 - have connectivity to the same destinations

Networks – Virtual Switches

- Define virtual switches with same name on each member
- For relocating guests:
 - Source and destination virtual switch guest NIC and port configurations must be equivalent
 - Port type
 - Authorizations (access, VLAN, promiscuous mode)
 - Source and destination virtual switches must be equivalent
 - Name and type
 - VLAN settings
 - Operational UPLINK port with matching EQID
 - Device and port numbers need not match, but connectivity to the same LAN segment is required

Networks – MAC Addresses

- MAC address assignments are coordinated across an SSI cluster
 - VMLAN statement
 - MACPREFIX must be set to different value for each member
 - Default is 02-xx-xx where xx-xx is "system number" of member (e.g., 02-00-01 for member 1)
 - USERPREFIX must be set for SSI members
 - Must be identical for all members
 - Must not be equal to any member's MACPREFIX value
 - Default is 02-00-00
 - MACIDRANGE is ignored in an SSI cluster
 - Because MAC assignment is coordinated among members
 - Example:

```
VMSYS01: VMLAN MACPREFIX 021111 USERPREFIX 02AAAA
VMSYS02: VMLAN MACPREFIX 022222 USERPREFIX 02AAAA
VMSYS03: VMLAN MACPREFIX 023333 USERPREFIX 02AAAA
VMSYS04: VMLAN MACPREFIX 024444 USERPREFIX 02AAAA
```

Relocation Domains

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation

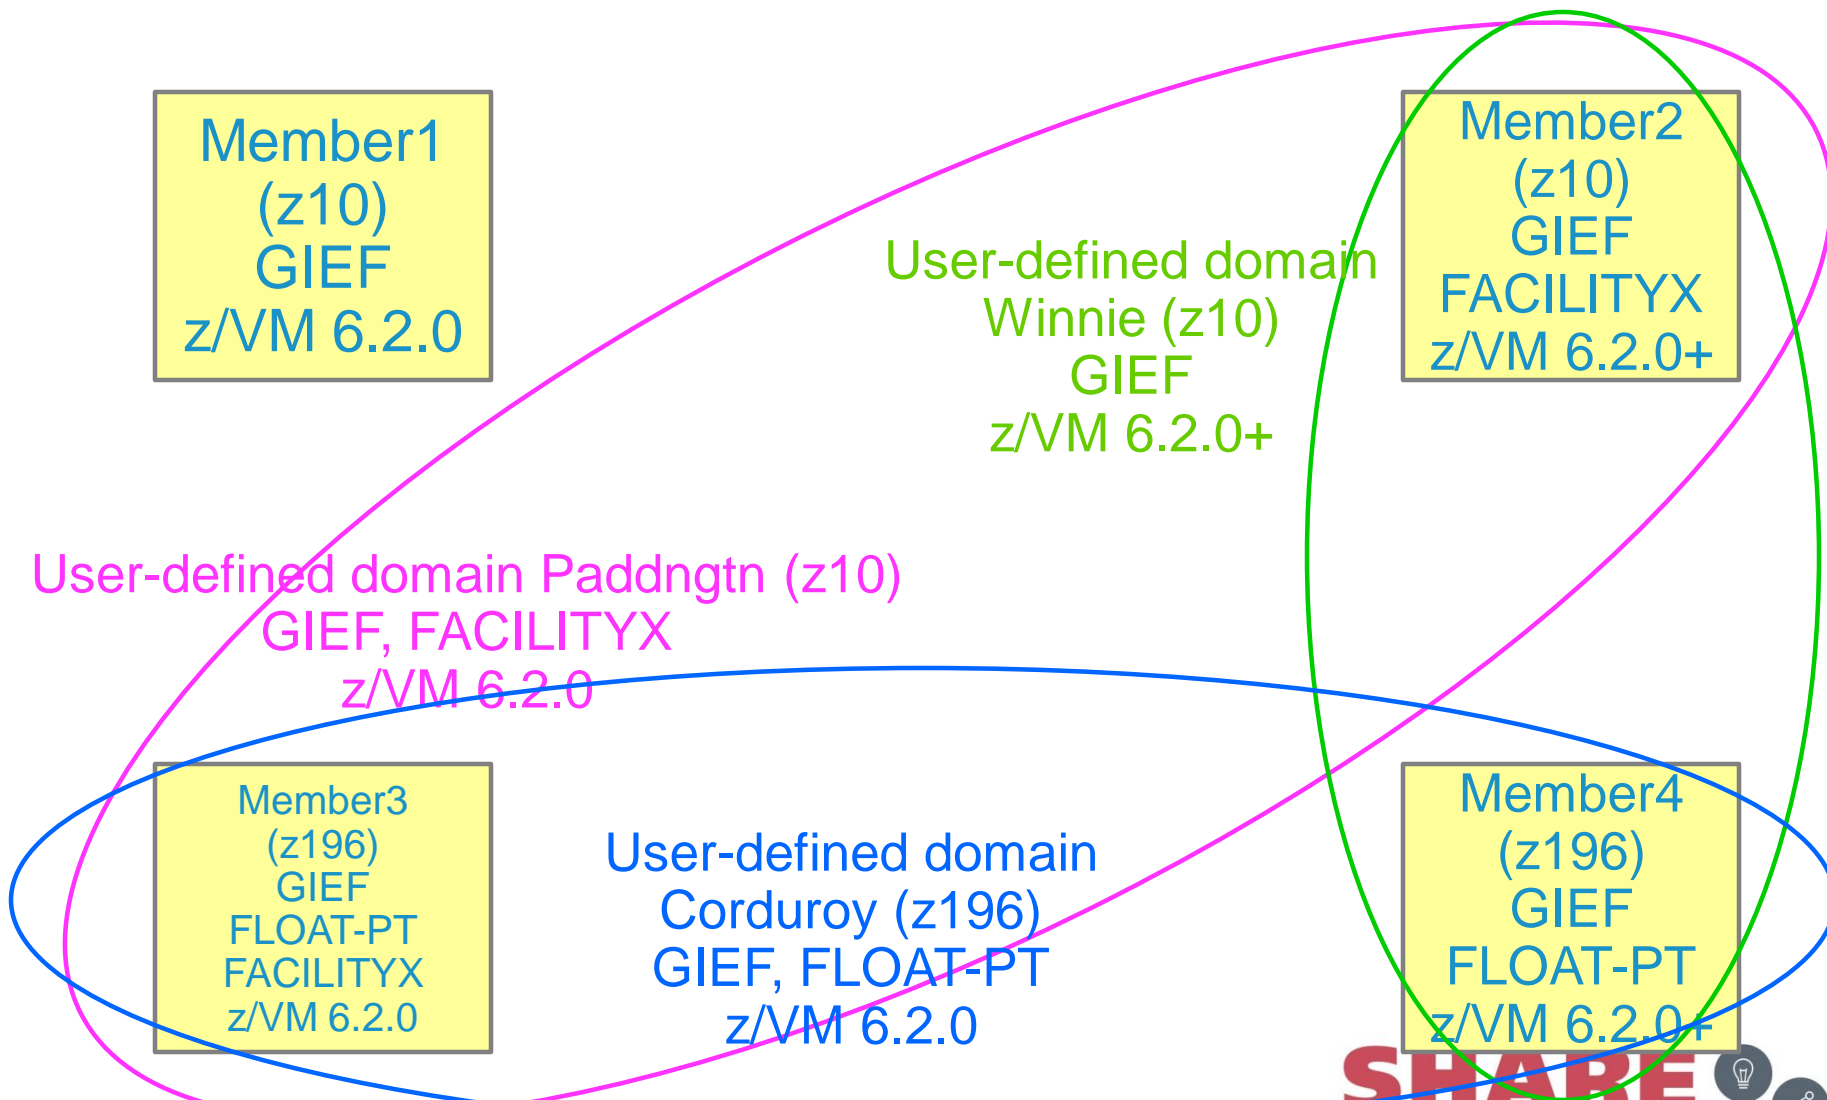
What is a Relocation Domain?

- A relocation domain defines a set of members of an SSI cluster among which virtual machines can relocate freely
- Relocation domains can be defined for business or technical reasons
- Regardless of differences in the facilities of the individual members, a domain has a common architectural level
 - This is the maximal common subset of all the members' facilities
- Several default domains are automatically defined by CP
 - Single member domains for each member in the SSI
 - An SSI domain that will have the features and facilities common to all members
- Defining your own domains is useful in a 3+ member cluster
 - In a 1 or 2 member cluster, all possible domains are defined by default
 - Defined via a SYSTEM CONFIG statement or dynamically by command

Architecture Fencing in Domains

- Guests cannot use facilities or features not included in the domain even if the member they are on has access to those features
 - We call this “fencing”
- Examples of commands/instructions with “fenced” responses:
 - **Q CPUID** -the model number will always reflect the virtual architecture level, the processor number is set at logon and not affected by relocation or relocation domain changes
 - **Diagnose x'00'** – will reflect the virtual CPLEVEL
 - **STFLE**

Relocation Domains



Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Assigning Relocation Domains

- Virtual machines may be assigned to a domain in their directory entry
 - Default for single configuration virtual machines is the SSI domain
 - Default for multiconfiguration virtual machines is their single member domain, which cannot be changed

- Virtual machines are assigned a virtual architecture level when they log on, according to what domain they are in

Planning for Live Guest Relocation

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation

General Guidelines for Relocating a Guest

- Make sure all resources used by the virtual machine are available on the destination member
 - Facilities (will be handled automatically if you are relocating within a domain)
 - Crypto cards
- Capacity for the virtual machine's memory and processor requirements
- Devices
 - Make sure that there really is an equivalent device on the destination
 - OSAs should be connected to the same LAN segment(s)
 - FCPs should have access to the same SAN fabric, WWPNs and LUNs
 - If possible, use the same device numbers to refer to equivalent devices
 - Equivalency ids (EQIDs) are defined for devices that need them
 - OSAs and FCPs
 - If connected to a VSWITCH, make sure the same VSWITCH is defined on the destination and the OSAs have been assigned the same set of EQIDs.
 - If the virtual machine has an FCP, make sure the “queue_if_no_path” option is specified in Linux

Guest Configuration for Live Guest Relocation

- In order to be eligible to relocate, a guest must be:
 - Defined as a single configuration virtual machine
 - Running in an ESA or XA virtual machine in ESA/390 or z/Architecture mode
 - Logged on and disconnected
 - Running only type CP or type IFL virtual processors
- **OPTION CHPIDVIRTUALIZATION ONE** should be specified in guest's directory entry
- If a guest is using a DCSS or NSS:
 - Identical NSS or DCSS must be available on the destination member
 - It cannot have the following types of page ranges
 - SW (shared write)
 - SC (shared with CP)
 - SN (shared with no data)

Guest Configuration for Live Guest Relocation...

- A guest can relocate if it has any of the following:
 - Private virtual disks in storage (created with DEFINE VFB-512 command)
 - An open console file
- A relocating guest can be using any of the following facilities:
 - Virtual machine time bomb (Diag x'288')
 - IUCV connections to *MSG and *MSGALL CP system services
 - Application monitor record (APPLDATA) collection
 - If guest buffer is not in a shared DCSS
 - Single Console Image Facility (SCIF)
 - Collaborative Memory Management Assist (CMMA)
- There are conditions that will prevent a guest from relocating
 - Documented in CP Planning and Administration - "Preparing for Live Guest Relocation in a z/VM SSI Cluster"

Starting and Managing a Live Guest Relocation

```
vmrelocate test lgrlin21 gdllcp3
01: User LGRLIN21 is eligible for relocation to GDLLCPX3
Ready; T=0.01/0.01 04:32:20
vmrelocate move lgrlin21 gdllcp3 asynchronous maxquiesce 10 maxtotal 10
01: Relocation of LGRLIN21 from GDLMPX4 to GDLLCPX3 started
Ready; T=0.01/0.01 04:32:25
vmrelocate status lgrlin21 details
01: User      From      To      By      Status      Elapsed
01: LGRLIN21  GDLMPX4  GDLLCPX3  CELESKEY  Final Mem Copy  00:00:07
01:
01: Options: ASYNCH IMMED NO
01: Max Total Time 10 sec
01: Max Quiesce Time 10 sec
01:
01: Total pages sent 560089 in 3 passes; 15881 pages sent in pass 4
Ready; T=0.01/0.01 04:32:32
User LGRLIN21 has been relocated from GDLMPX4 to GDLLCPX3
```

Starting and Managing a Live Guest Relocation...

```
at gdlldpx3 cmd vmrelocate move lgrlin21 gdlmcp4 async maxq 10 maxt 10
01: Relocation of LGRLIN21 from GDLLCPX3 to GDLMCPX4 started
Ready; T=0.01/0.01 04:35:56
vmrelocate cancel lgrlin21
01: Relocation of user LGRLIN21 from GDLLCPX3 to GDLMCPX4 has been canceled by CELESKEY on GDLMCPX4
Ready; T=0.01/0.01 04:35:58
query lgrlin21 at all
GDLLCPX3 : LGRLIN21 - DSC
Ready; T=0.01/0.01 04:36:09
at gdlldpx3 cmd vmrelocate move lgrlin21 gdlmcp4 async maxq 10 maxt 10
01: Relocation of LGRLIN21 from GDLLCPX3 to GDLMCPX4 started
Ready; T=0.01/0.01 04:36:17
vmrelocate status lgrlin21
01: No matching virtual machine relocation is currently in progress on this system
Ready; T=0.01/0.01 04:37:13
query lgrlin21 at all
01: GDLMCPX4 : LGRLIN21 - DSC
Ready; T=0.01/0.01 04:37:27
```

Summary

- An SSI cluster makes it easier to:
 - Manage and balance resources and workloads (move work to resources)
 - Schedule maintenance without disrupting key workloads
 - Test workloads in different environments
 - Operate and manage multiple z/VM images
 - Reliable sharing of resources and data
- Live Guest Relocation of Linux guests improves availability of key applications and workloads
 - Flexibility for planned outages
- Setting up an SSI cluster requires planning:
 - Migration from current environment
 - Configuration
 - Sharing resources and data
- Plan for extra
 - CPU capacity
 - Disk capacity
 - Memory
 - CTC connections

Reference Information

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© 2014, 2015 IBM Corporation

More Information

z/VM 6.3 resources

<http://www.vm.ibm.com/zvm630/>

<http://www.vm.ibm.com/events/>

z/VM Single System Image Overview

<http://www.vm.ibm.com/ssi/>

Live Virtual Classes for z/VM and Linux

<http://www.vm.ibm.com/education/lvc/>

Redbooks

- An Introduction to z/VM SSI and LGR

<http://publib-b.boulder.ibm.com/redpieces/abstracts/sg248006.html?Open>

- Using z/VM v 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)

<http://publib-b.boulder.ibm.com/abstracts/sg248039.html?Open>

- DB2 10 for Linux on System z Using z/VM v6.2, Single System Image Clusters and Live Guest Relocation

<http://www.redbooks.ibm.com/abstracts/sg248036.html?Open>

Whitepaper

- z/VM Migration: Migrating the User Directory and RACF Environment

<http://public.dhe.ibm.com/common/ssi/ecm/en/zsw03246usen/ZSW03246USEN.PDF>

Thanks!

John Franciscovich
IBM
z/VM Design and Development
Endicott, NY
francisj@us.ibm.com



Complete your session evaluations online at www.SHARE.org/Orlando-Eval