Planning and Migrating to z/VM Single System Image (SSI)

Session 11922

John Franciscovich
francisj@us.ibm.com

Emily Kate Hugenbruch
ekhugen@us.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

z/VM® z10™ z/Architecture® zEnterprise™

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Disclaimer

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "AS IS" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environments.

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice, and represent goals and objectives only. This is not a commitment to deliver the functions described herein

Topics

- z/VM 6.2 Installation Planning
- Planning and Configuring your SSI Cluster
- Migrating to SSI

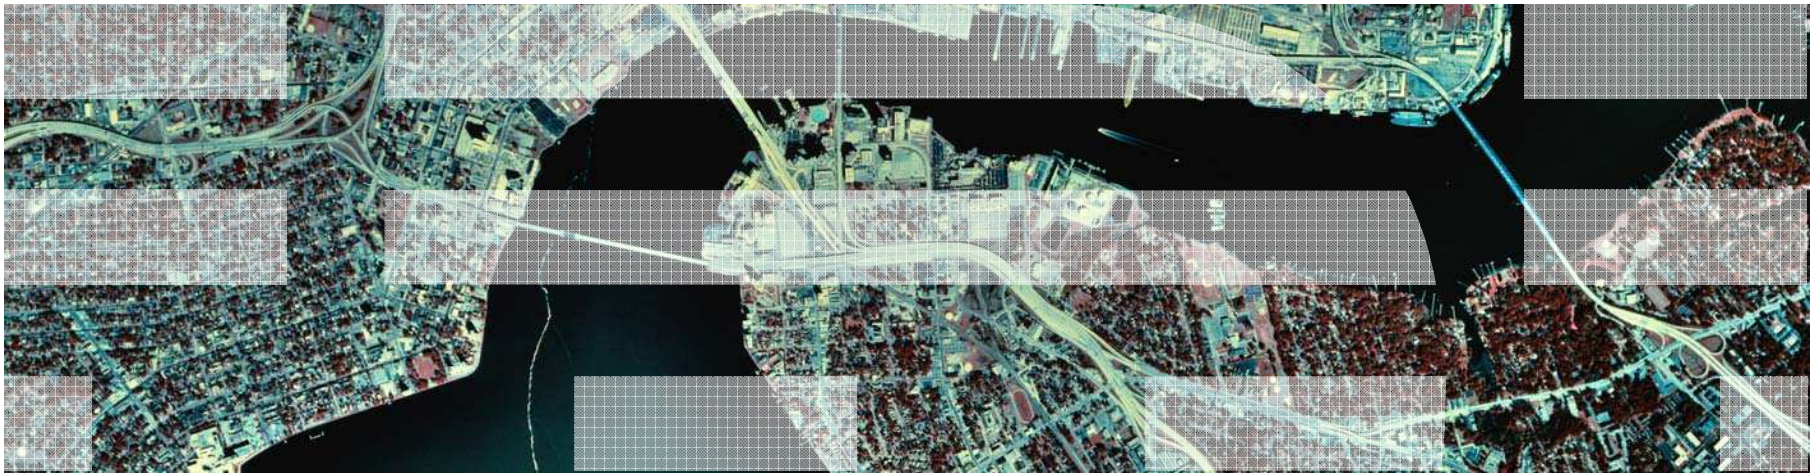

z/VM 6.2
Installation Planning

New MAINT Userids

MAINT	PMAINT	MAINT620
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, 550, 551 disks	Owns the service disks (e.g., 490, 493, 49D) and the CF2 parm disk
Use for work on a particular member, such as attaching devices, or relocating guests	Use for updating the system config, or for SSI-wide work, e.g., defining relocation domains	Use for applying 6.2.0 service. The CF2 parm disk contains 6.2.0 CPLOAD modules.

Minidisks for New MAINT Userids

Parm Disks *(Owner)*

- **CF0** ***(PMAINT)***
  - Common system configuration file
- **CF1** ***(MAINT)***
  - Production CPLOAD MODULE
- **CF2** ***(MAINT620)***
  - Used by SERVICE to hold test CPLOAD MODULE
- **CF3** ***(MAINT)***
  - Backup of CF1

Full Pack Minidisks

- **MAINT**
  - 122 M01S01
  - 123 M01RES
  - 124 M01W01
- **MAINT620**
  - 131 620RL1
  - 132 620RL2
  - 133 620RL3
- **PMAINT**
  - 141 VMCOM1
  - 142 VMCOM2

Minidisks for New MAINT Userids (by volume)

Cluster-Wide Volume (VMCOM1)

– PMAINT
- **CF0** - Common system configuration file
- **2CC** - Single source directory
- **41D** - VMSES/E production inventory disk
- **551** - SSI cluster common disk - contains utilities that must be at the highest level for all members of the SSI cluster, including

  CPFMTXA, DIRECTXA, DIRMAP, DISKMAP

Release Volumes (620RLn)

– MAINT620
- **490** - Test CMS system disk
- **493** - Test system tools disk
- **51D** - VMSES/E software inventory disk
- **CF2** – Test parm disk

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 directory
    - Shared system configuration file
    - Creates Persistent Data Record (PDR) on Common volume

- **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

INSTPLAN - Select Installation Type

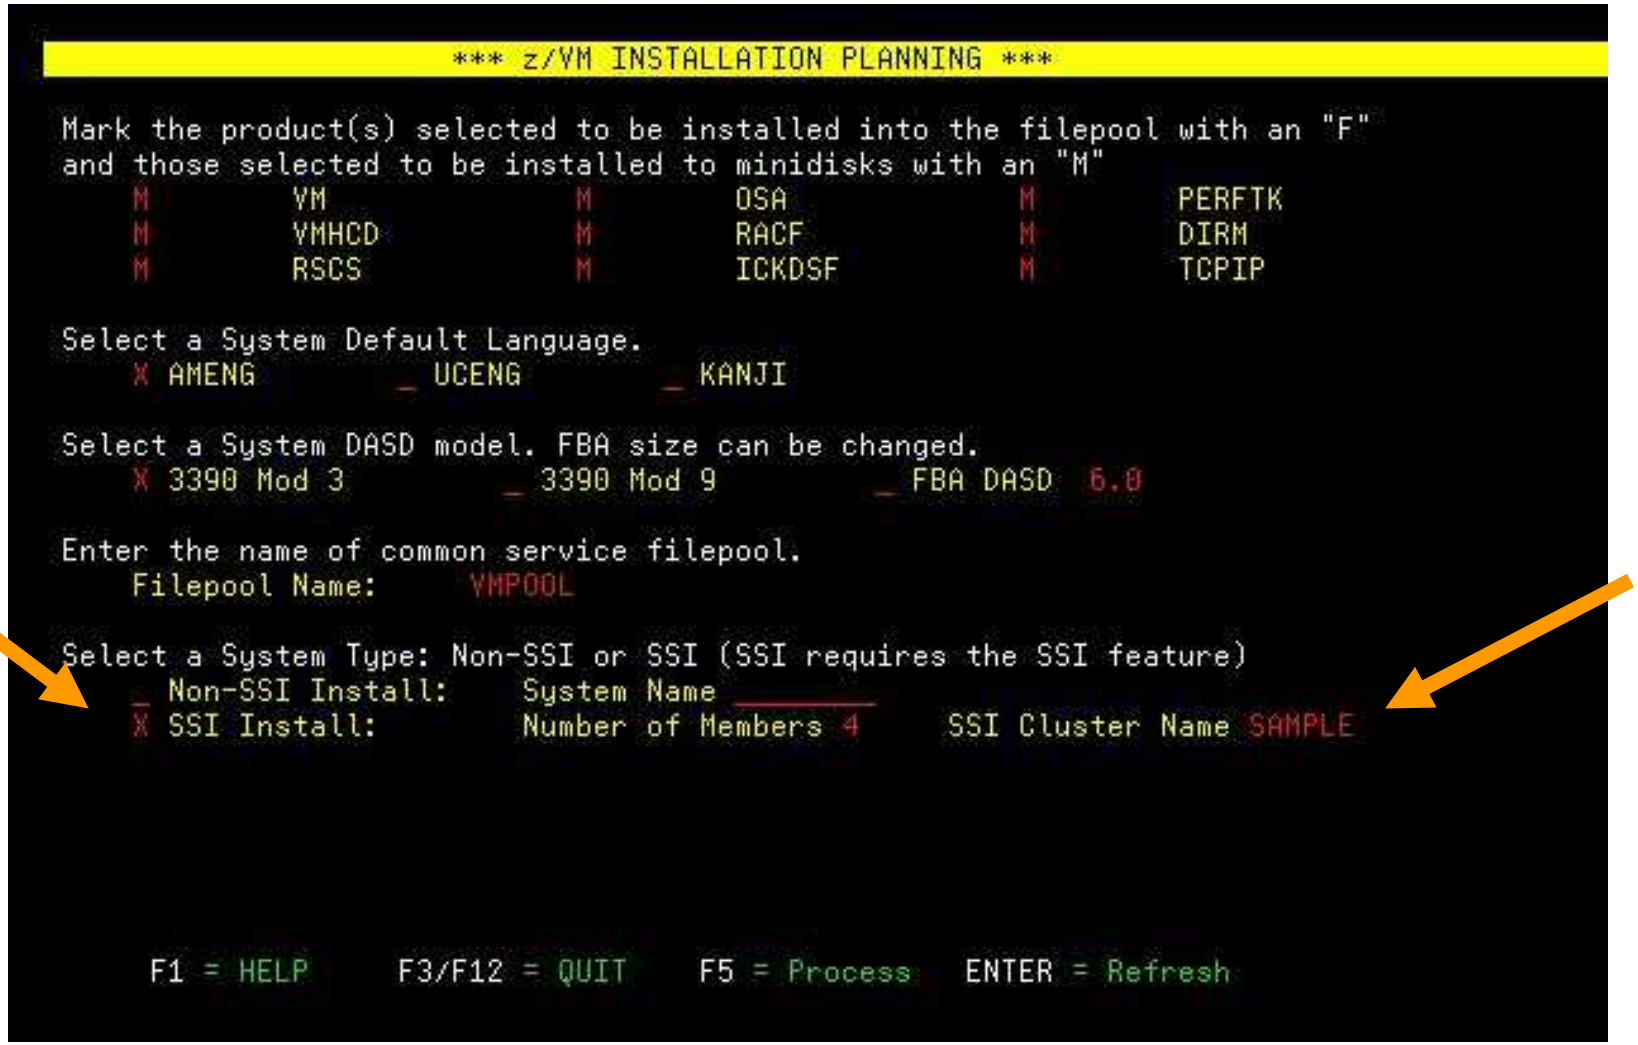

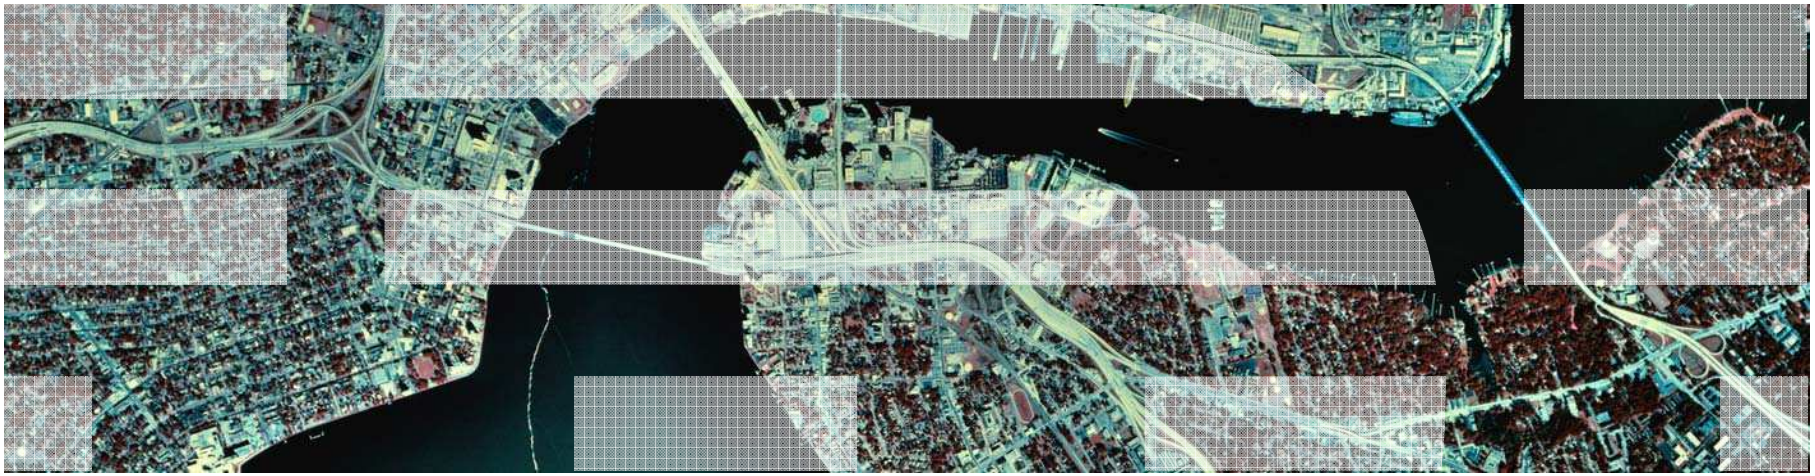

Planning and Configuring your SSI Cluster

SSI Cluster Requirements

- Servers must be IBM System z10 or later (z/VM Version 6)
- Shared and non-shared DASD
  - 3390 volume required for the PDR
  - All volumes should be cabled to all members
    - Makes non-shared disks accessible to other members to fix configuration problems
- LPARs
  - 1-16 FICON CTC devices between LPARs
    - Provide direct ISFC links from each member to all other members
  - FICON channels to shared DASD
  - OSA access to the same LAN segments
  - FCP access to same storage area networks (SANs) with same storage access rights
- Shared system configuration file for all members
- Shared source directory containing user definitions for all members
- 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 Topography

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

SSI Planning Worksheet

Table 4. Linux virtual server requirements for memory, processors, and devices (continued)

Linux server user ID	Memory	Virtual processors	DASD	Networking devices	Hardware feature or architecture	Member 1	Member 2	Member 3	Member 4
Maximum number of resident and relocated virtual servers:									
Maximum memory for normally resident and relocated virtual servers:									
Memory for z/VM:									
Total virtual memory requirement:									
Total real memory requirement (after considering overcommitment)[1]:									
Expanded storage estimate (Total real memory ×.25, but not more than 2 GB):									
Central storage estimate (Total real memory – expanded storage estimate):									
Number of real CPUs:									
DASD paging space (Total virtual memory × 2 or more):									

1. Total virtual memory should be no more than three times the total real memory.

Chapter 2. Planning for Linux virtual servers

SSI Cluster Planning

- CTC connections
- DASD
- Networks
- Cluster and member configuration
- Shared Source Directory

CTC Connections

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

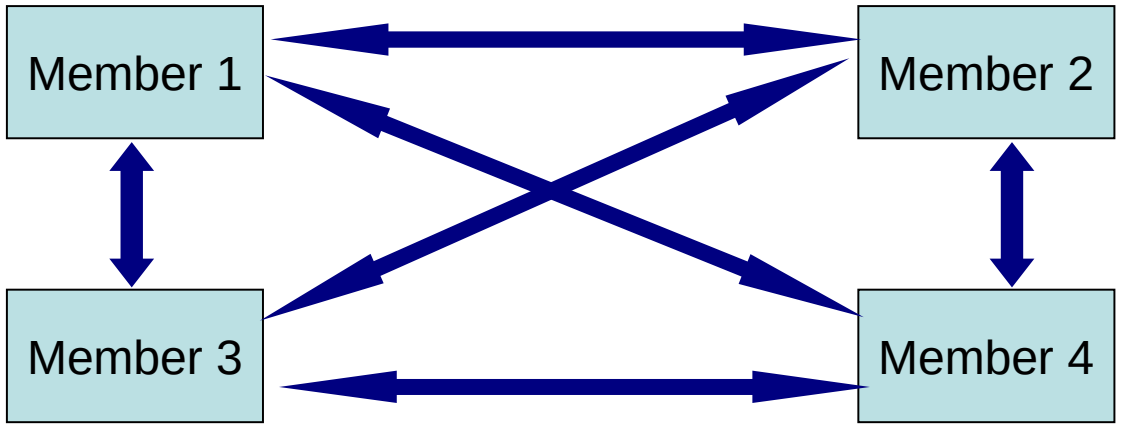

- Logical links are composed of 1-16 CTC connections
  - FICON channel paths
  - May be switched or unswitched
- Use multiple CTCs distributed on multiple FICON channel paths between each pair of members
  - Avoids write collisions that affect link performance
  - Avoids severing logical link if one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

CTC Connections – How Many Do I Need?

- 4 CTC devices per per FICON chpid
  - provides most efficient ISFC data transfer
- For large guests, relocation and quiesce times improve with more chpids
  - Up to 4 chpid paths, with 4 CTCs each
    - *Additional factors affect relocation and quiesce times*

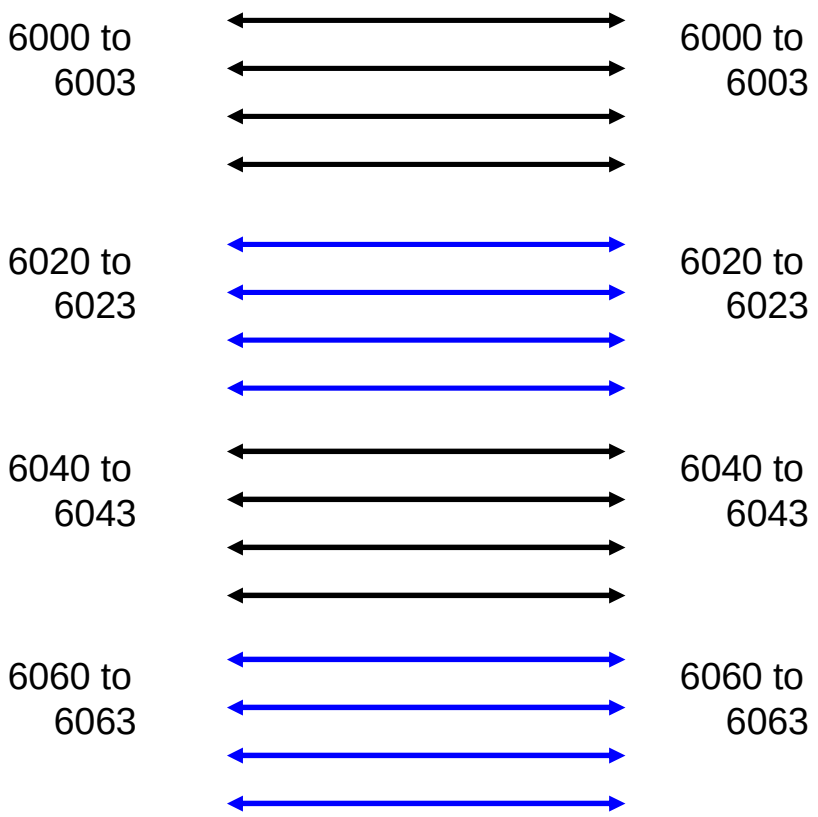

DASD Planning – Non-Shared and Shared System Volumes

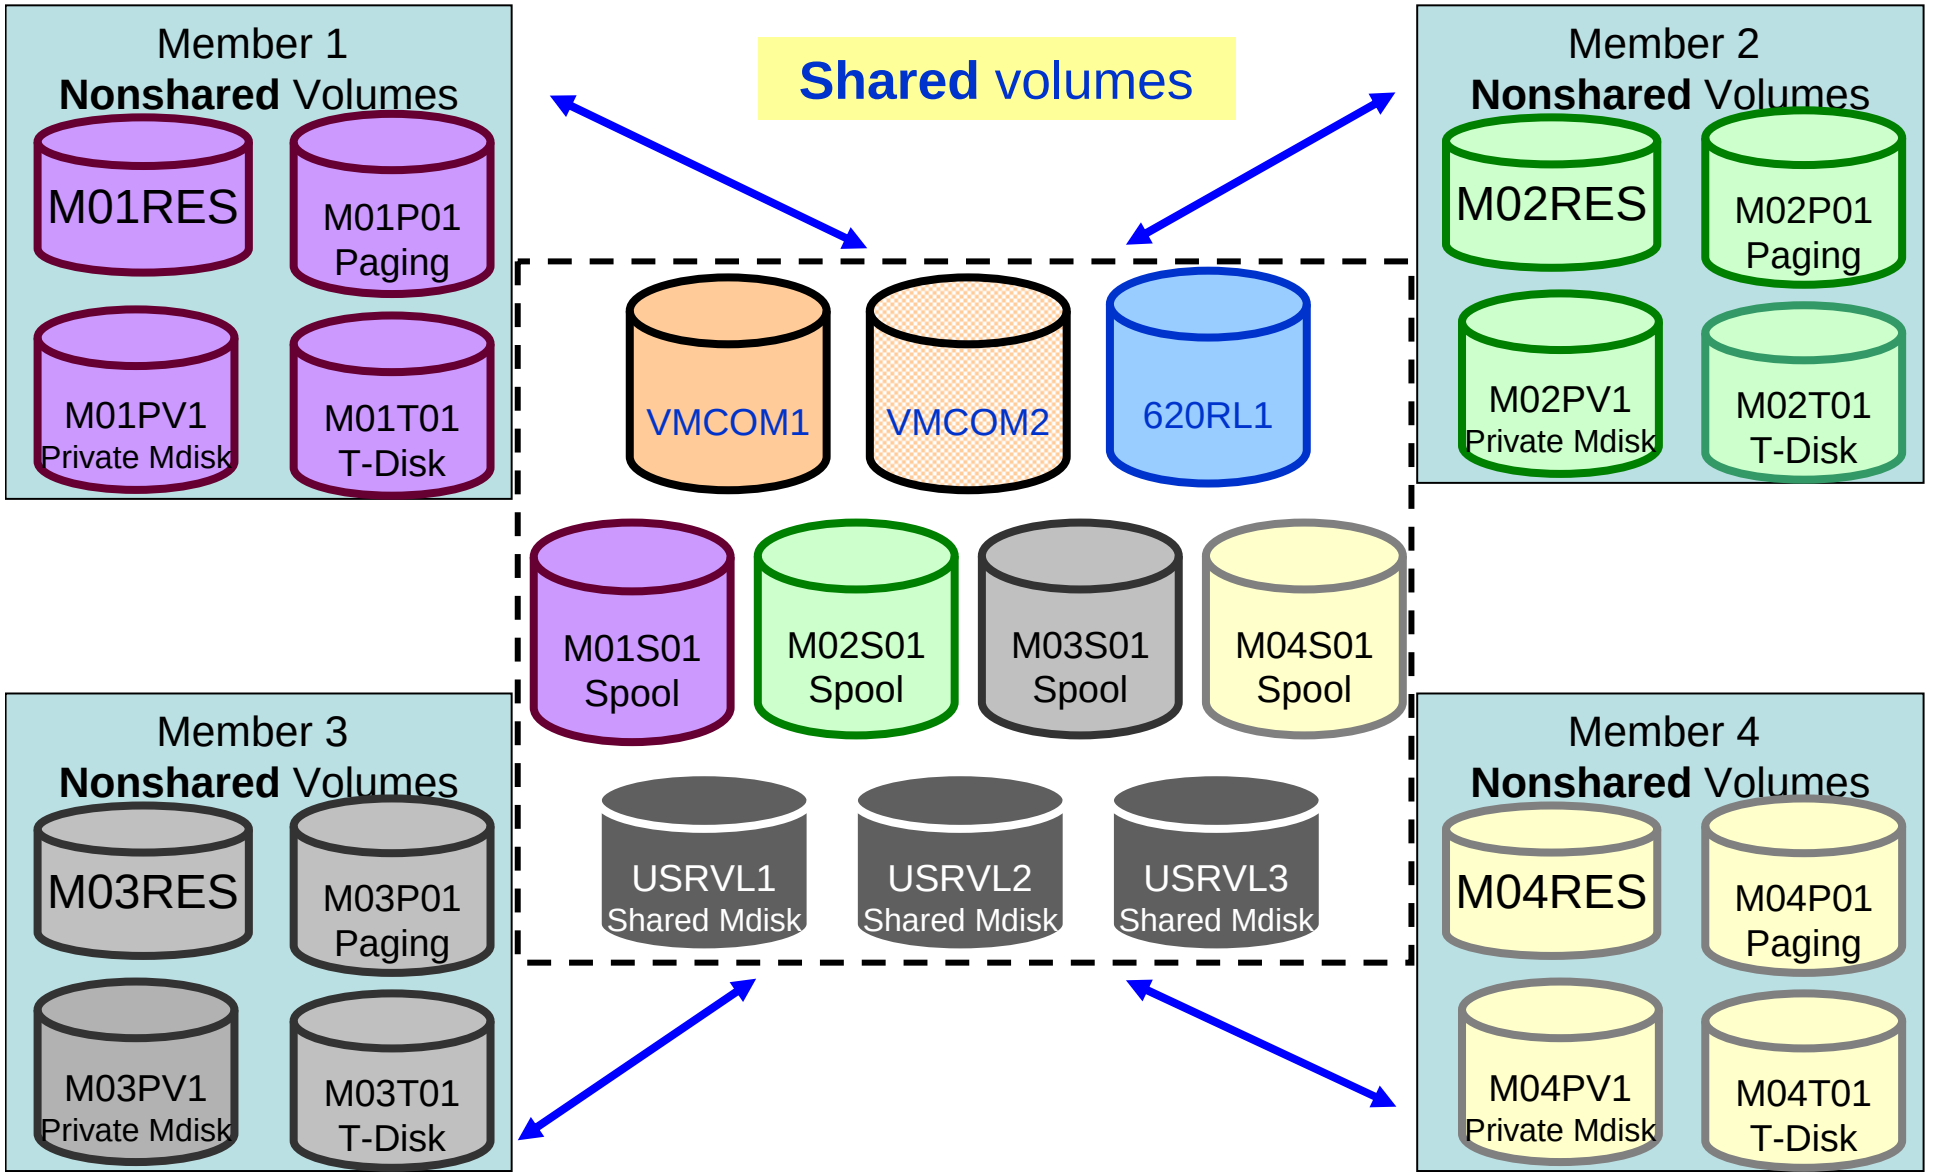

DASD Planning

- Decide which DASD volumes will be used for
  - Cluster-wide volume(s)
  - Release volumes
  - System volumes
    - Shared
    - Non-shared
  - User data (minidisks)
    - Shared
    - Non-shared
- Decide which member owns each CP-Owned volume

DASD Planning - 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

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

- Ownership information may also be used on non-SSI systems
  - System name but no cluster name
  - Default on non-SSI installs

DASD Planning – CP_OWNED List

Non-SSI

```
/*****************************************************************/
/*                      CP_Owned Volume Statements               */
/*****************************************************************/
/*                                                 SYSRES  VOLUME */
/*****************************************************************/

          CP_Owned   Slot   1  M01RES

/*****************************************************************/
/*                                                  COMMON VOLUME */
/*****************************************************************/

    CP_Owned   Slot   5  VMCOM1

/*****************************************************************/
/*                                           DUMP & SPOOL VOLUMES */
/*****************************************************************/

    CP_Owned   Slot  10  M01S01

/*****************************************************************/
/*                                           PAGE & TDISK VOLUMES */
/*****************************************************************/

           CP_Owned   Slot 255  M01P01
```

DASD Planning – CP_OWNED List

SSI

```
32 /*****************************************************************/
33 /*                      CP_Owned Volume Statements                 */
34 /*****************************************************************/
35 /*                                                 SYSRES  VOLUME  */
36 /*****************************************************************/
37
38 MEMBER1:  CP_Owned   Slot   1  M01RES
39
40 /*****************************************************************/
41 /*                                                  COMMON VOLUME  */
42 /*****************************************************************/
43
44     CP_Owned   Slot   5  VMCOM1
45
46 /*****************************************************************/
47 /*                                            DUMP & SPOOL VOLUMES */
48 /*****************************************************************/
49
50     CP_Owned   Slot  10  M01S01
51
52 /*****************************************************************/
53 /*                                            PAGE & TDISK VOLUMES */
54 /*****************************************************************/
55
56 MEMBER1:  CP_Owned   Slot 255  M01P01
```

DASD Planning – USER_VOLUME_LIST

Non-SSI

```
58 /*****************************************************************/
59 /*                          User_Volume_List                     */
60 /*****************************************************************/
61 /* These volumes contain the minidisks for your guests, as well as   */
62 /* the product disks for z/VM. Volumes that are not intended to hold  */
63 /* "local" minidisks, i.e., minidisks that would be unique to a       */
64 /* single system, should be kept on separate volumes.                 */
65
66 /*****************************************************************/
67 /* Shared User Volumes                                           */
68 /*****************************************************************/
69           User_Volume_List  620RL1 620RL2 USRVL1
70
71 /*****************************************************************/
72 /* User volumes for local minidisks                              */
73 /*****************************************************************/
74
75           User_Volume_List  M01W01
```

DASD Planning – USER_VOLUME_LIST

SSI

```
58 /*****************************************************************/
59 /*                          User_Volume_List                     */
60 /*****************************************************************/
61 /* These volumes contain the minidisks for your guests, as well as   */
62 /* the product disks for z/VM. Volumes that are not intended to hold  */
63 /* "local" minidisks, i.e., minidisks that would be unique to a       */
64 /* single system, should be kept on separate volumes.                 */
65
66 /*****************************************************************/
67 /* Shared User Volumes                                           */
68 /*****************************************************************/
69           User_Volume_List  620RL1 620RL2 USRVL1
70
71 /*****************************************************************/
72 /* User volumes for local minidisks                              */
73 /*****************************************************************/
74
75 MEMBER1:   User_Volume_List  M01W01 M01PV1
```

Networks in an SSI

- 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

- Updates to the main TCPIP stack configuration
  - *PROFILE TCPIP* now can have member-specific names like
    - *MEMBER1 TCPIP*
    - *MEMBER2 TCPIP*
  - *TCPIP DATA* file can be shared among SSI members, so you can add system qualifiers to statements like **HOSTNAME**

Networks in an SSI – 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 in an SSI – 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
```

Cluster and Member Configuration – SYSTEM_IDENTIFIER Statement

Non-SSI

```
 1 /*************************************************************************/
 2 /*              SYSTEM CONFIG FILE                                       */
 3 /*************************************************************************/
 4 /*                                                                       */
 5 /*  Refer to CP Planning and Administration for SYSTEM CONFIG rules      */
 6 /*                                                                       */
 7 /*  Warning - Always run CPSYNTAX after updating the SYSTEM CONFIG       */
 8 /*                                                                       */
 9 /*************************************************************************/
10
11 /*************************************************************************/
12 /*                    System_Identifier Information                      */
13 /*************************************************************************/
14
15 System_Identifier  *     *   MEMBER1
```

Cluster and Member Configuration – SYSTEM_IDENTIFIER Statement

SSI

```
 1 /*************************************************************************/
 2 /*              SYSTEM CONFIG FILE                                       */
 3 /*************************************************************************/
 4 /*                                                                       */
 5 /*  Refer to CP Planning and Administration for SYSTEM CONFIG rules      */
 6 /*                                                                       */
 7 /*  Warning - Always run CPSYNTAX after updating the SYSTEM CONFIG       */
 8 /*                                                                       */
 9 /*************************************************************************/
10
11 /*************************************************************************/
12 /*                    System_Identifier Information                      */
13 /*************************************************************************/
14
15 System_Identifier LPAR LP01 MEMBER1
```

Cluster and Member Configuration – SSI Statement

```
17 /*****************************************************************/
18 /*                      SSI Statement                            */
19 /*****************************************************************/
20
21 SSI MYCLUSTR PDR_Volume VMCOM1 ,
22 Slot 1 MEMBER1
23
```

Cluster and Member Configuration – SYSTEM_RESIDENCE Statement

Non-SSI

```
24 /*************************************************************************/
25 /*                  Checkpoint and Warmstart Information                 */
26 /*************************************************************************/
27
28 System_Residence,
29     Checkpoint  Volid M01RES    From CYL 21  For 9 ,
30     Warmstart   Volid M01RES    From CYL 30  For 9
31
```

Cluster and Member Configuration – SYSTEM_RESIDENCE Statement

SSI

```
24 /*****************************************************************/
25 /*                Checkpoint and Warmstart Information           */
26 /*****************************************************************/
27
28 MEMBER1:  System_Residence,
29     Checkpoint  Volid M01RES    From CYL 21  For 9 ,
30     Warmstart   Volid M01RES    From CYL 30  For 9
31
```

Cluster and Member Configuration – Additional Steps

- Enable the SSI feature

- If you're migrating from non-SSI to SSI, you'll want to enable the **PROMPT_AFTER_SHUTDOWN_REIPL** feature before you reIPL, so you can do a cold start

- Run CPSYNTAX

```
cpsyntax sysnew config (system member1
CONFIGURATION FILE PROCESSING COMPLETE -- NO ERRORS ENCOUNTERED.
```

Shared Source Directory – Virtual Machine Definition Types

Single Configuration Virtual Machine (traditional)

USER statement

Same definitions and resources on all members

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

Multiconfiguration Virtual Machine (new)

IDENTITY statement

Definitions and resources common to all members

- May log on to multiple members at the same time (known by IDENTITY name)
- System support virtual machines
- Service virtual machines

SUBCONFIG statement for member 1

Definitions and resources unique to member 1

SUBCONFIG statement for member 2

Definitions and resources unique to member 2

Shared Source Directory – Global and Local disks

- For each guest you're turning into a multiconfiguration virtual machine, decide which disks should be global and which should be local
  - You may want to split existing disks into global and local.

Global	Local
•All instances have access	•Only one instance has access
•Usually R/O	•Usually R/W
•EXECs	•Log files
•Control files	•Work files

Shared Source Directory - New Layout

- IBM-supplied directory will be significantly different than in previous releases
  - Both SSI and non-SSI installations
  - Directory for non-SSI installations will be in "SSI-ready" format
    - Facilitate future SSI deployment
- Many of the IBM-supplied userids will be defined as multiconfiguration virtual machines
- Determine if any of your guests should be defined as multiconfiguration virtual machines
  - Most will be single-configuration virtual machines
  - Userids defined on SYSTEM_USERIDS statements will usually be multiconfiguration virtual machines
- Merge your user definitions into the IBM-supplied directory

Shared Source Directory - Multiconfiguration Virtual Machine Definition

```
 IDENTITY MAINT    MAINTPAS     128M 1000M ABCDEFG

 BUILD ON SSIMEMB1 USING SUBCONFIG MAINT-1
 BUILD ON SSIMEMB2 USING SUBCONFIG MAINT-2
 BUILD ON SSIMEMB3 USING SUBCONFIG MAINT-3
 BUILD ON SSIMEMB4 USING SUBCONFIG MAINT-4

 CONSOLE 009 3215 T
 SPOOL 00C 2540 READER *
 SPOOL 00D 2540 PUNCH  A
 SPOOL 00E 1403 A
 LINK    USER1     2CC 2CC RR
 LINK    USER1     551 551 RR

SUBCONFIG MAINT-1
 MDISK 0191 3390 1000 20 MNTVL1 WR
 MDISK CF1  3390 100  20 M01RES RR
* END OF MAINT-1

SUBCONFIG MAINT-2
 MDISK 0191 3390 1000 20 MNTVL2 WR
 MDISK CF1  3390 100  20 M02RES RR
* END OF MAINT-2

SUBCONFIG MAINT-3
 MDISK 0191 3390 1000 20 MNTVL3 WR
 MDISK CF1  3390 100  20 M03RES RR
* END OF MAINT-3

SUBCONFIG MAINT-4
 MDISK 0191 3390 1000 20 MNTVL4 WR
 MDISK CF1  3390 100  20 M04RES RR
* END OF MAINT-4
```

These statements apply to all instances of MAINT on all members

These statements only apply to MAINT on member SSIMEMB1

These statements only apply to MAINT on member SSIMEMB2

These statements only apply to MAINT on member SSIMEMB3

These statements only apply to MAINT on member SSIMEMB4

Shared Source Directory – Multiconfiguration Virtual Machines

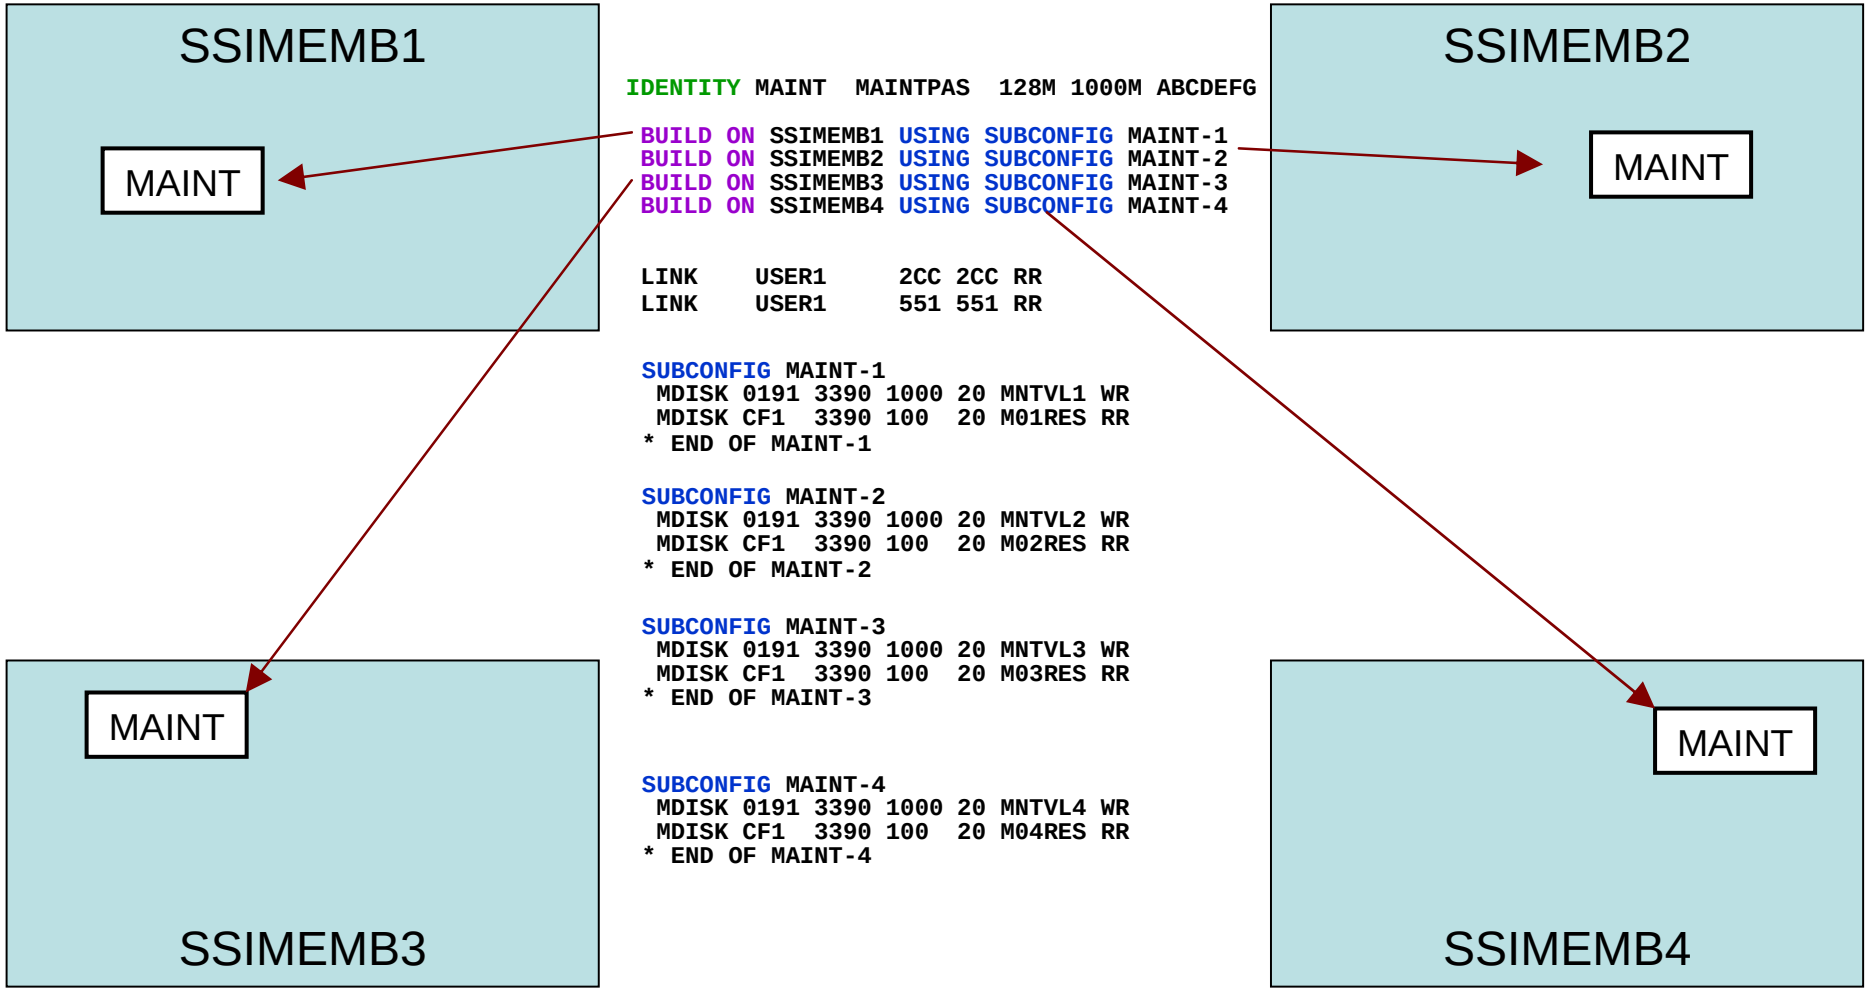

Shared Source Directory – Single Configuration Virtual Machines

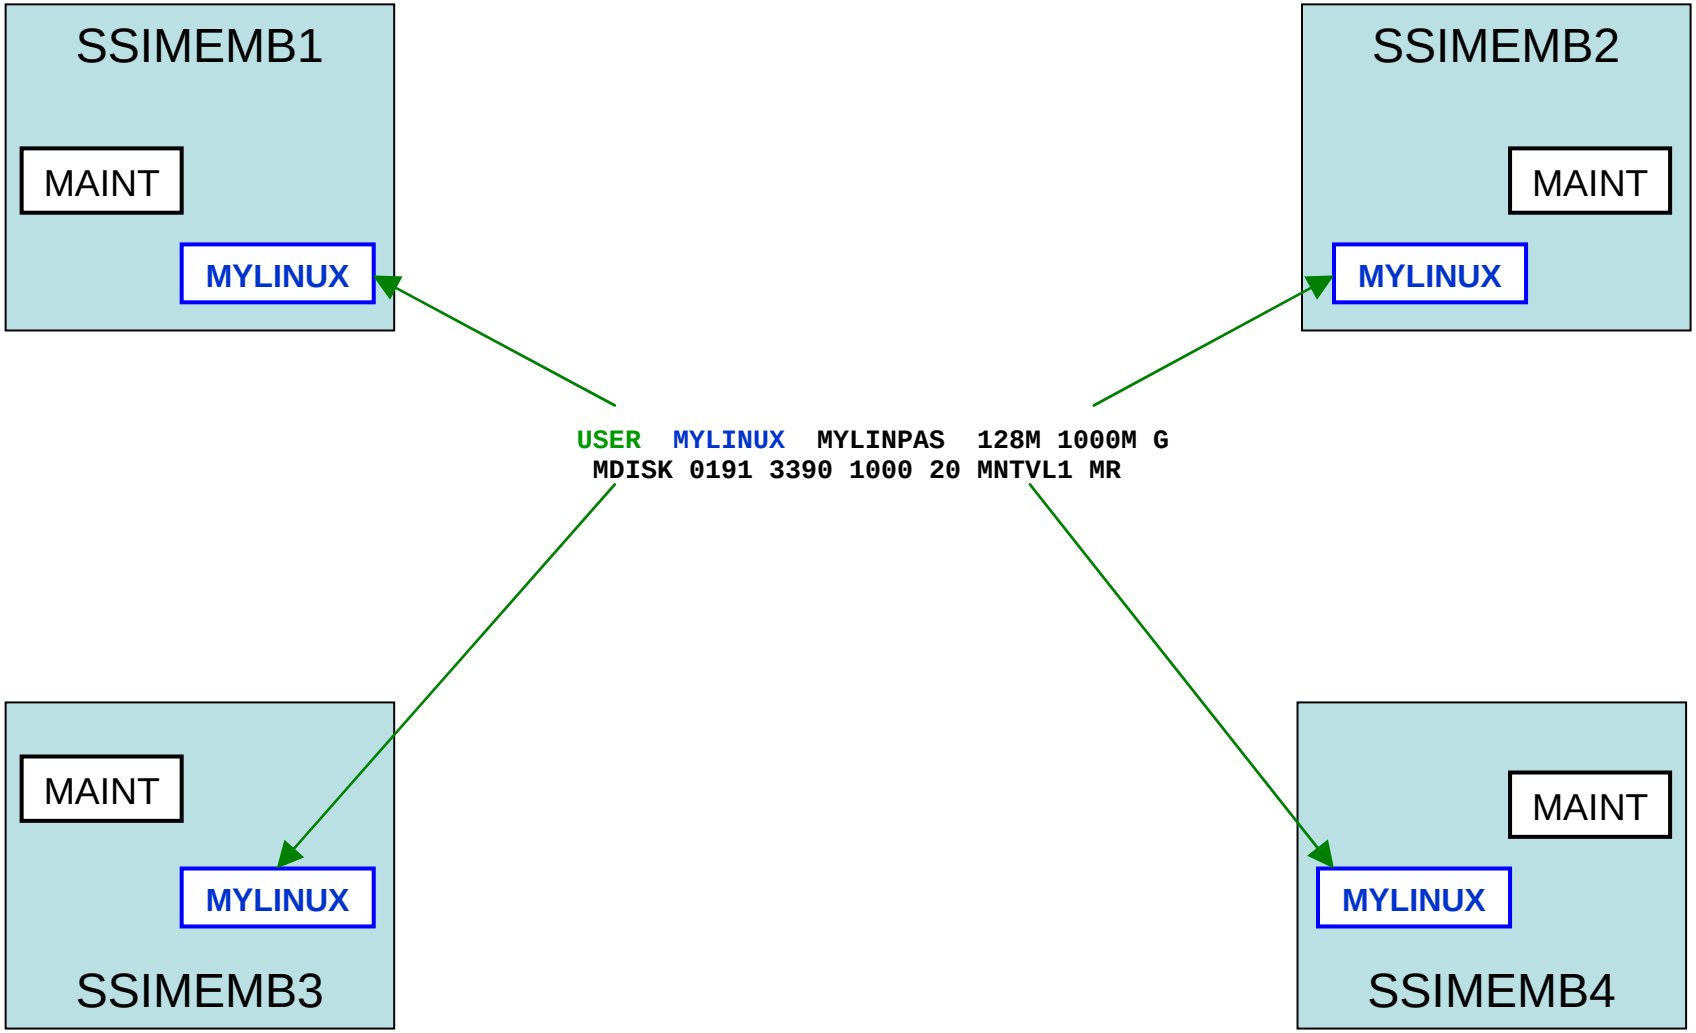

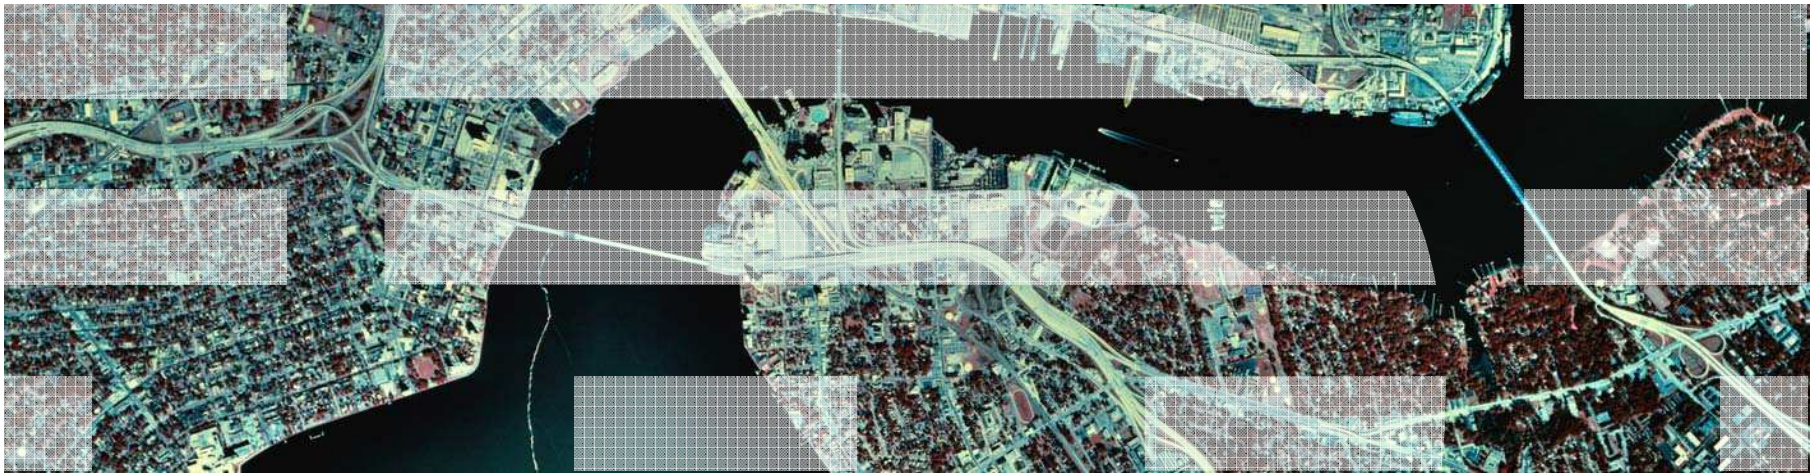

Migrating to SSI

Use Case Scenarios

- Migration procedures for existing z/VM environments
  - Documented in CP Planning and Administration
    - Converting a z/VM System to a Single-Member z/VM SSI Cluster
    - Adding a Member to a z/VM SSI Cluster by Cloning an Existing Member
    - Combining Two Non-SSI z/VM Systems to Create a z/VM SSI Cluster
    - Moving a Second-Level z/VM SSI Cluster to First-Level
    - Converting a CSE Complex to a z/VM SSI Cluster
    - Decommissioning a Member of a z/VM SSI Cluster
- Review documented procedures before deciding whether to do SSI or non-SSI install

Migrating from a Non-SSI 6.2.0 system to a Single Member SSI

1. Prepare the New DASD Volumes
2. Update the System Configuration File
3. Update the User Directory
4. Manage the User Spool Files
5. Prepare the CP-Owned Volumes
6. Create the PDR
7. Modify the Startup Parameters for the VMPSFS File Pool
8. Shut Down and Cold Start
9. Load the Spool Files
10. Change the User Directory to SSI-Enabled

"SSI-enable" the Shared Source Directory

Non-SSI

```
 1 ************************************************************
 2 *  z/VM 6.2.0  SYSTEM DIRECTORY                            *
 3 ************************************************************
 4 *                                                          *
 5 *  THE ADDRESSES 123 THROUGH 133 ARE VIRTUAL ADDRESSES.    *
 6 *                                                          *
 7 *  NOTES:                                                  *
 8 *  REMEMBER THESE ARE ONLY VIRTUAL ADDRESSES NOT REAL      *
 9 *  ADDRESSES, SO THERE IS NO NEED TO CHANGE THEM TO MATCH  *
10 *  YOUR HARDWARE ADDRESSES.                                *
11 *                                                          *
12 *                                                          *
13 ************************************************************
14 *                                                          *
15 *  FOR A DESCRIPTION OF DIRECTORY STATEMENTS SEE:          *
16 *     VM ENTERPRISE SYSTEM ARCHITECTURE                    *
17 *     PLANNING AND ADMINISTRATION MANUAL.                  *
18 *                                                          *
19 ************************************************************
20 *
21 *
22 *

23 DIRECTORY 123 3390 M01RES
```

"SSI-enable" the Shared Source Directory

SSI

```
 1 ************************************************************
 2 *  z/VM 6.2.0  SYSTEM DIRECTORY                            *
 3 ************************************************************
 4 *                                                          *
 5 *  THE ADDRESSES 123 THROUGH 133 ARE VIRTUAL ADDRESSES.    *
 6 *                                                          *
 7 *  NOTES:                                                  *
 8 *  REMEMBER THESE ARE ONLY VIRTUAL ADDRESSES NOT REAL      *
 9 *  ADDRESSES, SO THERE IS NO NEED TO CHANGE THEM TO MATCH  *
10 *  YOUR HARDWARE ADDRESSES.                                *
11 *                                                          *
12 *                                                          *
13 ************************************************************
14 *                                                          *
15 *  FOR A DESCRIPTION OF DIRECTORY STATEMENTS SEE:          *
16 *     VM ENTERPRISE SYSTEM ARCHITECTURE                    *
17 *     PLANNING AND ADMINISTRATION MANUAL.                  *
18 *                                                          *
19 ************************************************************
20 *
21 *
22 *

23 DIRECTORY SSI 123 3390 M01RES M02RES M03RES M04RES
```

"SSI-enable" the Shared Source Directory

- Update the **BUILD** statements with the actual member name
  - Multiconfiguration virtual machines will have asterisks instead of machine names
- Run **DIRECTXA** to put the new directory into production

Non-SSI

```
165 IDENTITY MAINT MAINT 128M 1000M ABCDEFG
166  BUILD ON *        USING SUBCONFIG MAINT-1
```

"SSI-enable" the Shared Source Directory

- Update the **BUILD** statements with the actual member name
  - Multiconfiguration virtual machines will have asterisks instead of machine names
- Run **DIRECTXA** to put the new directory into production

SSI

```
165 IDENTITY MAINT MAINT 128M 1000M ABCDEFG
166  BUILD ON MEMBER1 USING SUBCONFIG MAINT-1
```

Create the Persistent Data Record (PDR)

- **LINK** the fullpack overlay of *VMCOM1*, **PMAINT** *141*

```
formssi create 141 myclustr
HCPPDF6613R Device 0141 label is VMCOM1 - continue (Yes/No)?
yes
HCPPDF6614I Persistent Data Record created on device 0141
Ready; T=0.01/0.01 14:35:48

formssi display 141
HCPPDF6618I Persistent Data Record on device 0141 (label VMCOM1) is for MYCLUSTR
HCPPDF6619I PDR                  state: Unlocked
HCPPDF6619I                time stamp: 09/23/11 14:35:48
HCPPDF6619I     cross-system timeouts: Enabled
Ready; T=0.01/0.01 14:35:54
```

Changes to the VMPSYS file pool

- In the *VMSERVP DMSPARMS* file, change the **LOCAL** startup parameter to **SSI**:

Non-SSI

```
00001 ADMIN MAINT MAINT620
AUTOLOG1
00002 NOBACKUP
00003 SAVESEGID CMSFILES
00004 LOCAL
00005 FILEPOOLID SERVPOOL
00006 USERS 100
```

Changes to the VMPSYS file pool

- In the *VMSERVP DMSPARMS* file, change the **LOCAL** startup parameter to **SSI**:

SSI

```
00001 ADMIN MAINT MAINT620
AUTOLOG1
00002 NOBACKUP
00003 SAVESEGID CMSFILES
00004 SSI
00005 FILEPOOLID SERVPOOL
00006 USERS 100
```

IPL the Single (First) Member of your SSI Cluster

```
20:12:47 HCPAAU2700I System gateway MEMBER1 identified.
20:12:47 HCPNET3010I Virtual machine network device configuration changes are permitted
20:12:47 HCPPLM1697I The state of SSI system MEMBER1 has changed from DOWN to JOINED
20:12:47 HCPPLM1698I The mode of the SSI cluster is STABLE
```

```
q ssi
16:57:39 SSI Name: MYCLUSTR
16:57:39 SSI Mode: Stable
16:57:39 Cross-System Timeouts: Enabled
16:57:39 SSI Persistent Data Record (PDR) device: VMCOM1 on E00A
16:57:39 SLOT SYSTEMID STATE     PDR HEARTBEAT        RECEIVED HEARTBEAT
16:57:39    1 MEMBER1  Joined    2011-10-13 16:57:17 2011-10-13 16:57:17
16:57:39    2 -------- Available
16:57:39    3 -------- Available
16:57:39    4 -------- Available
Ready; T=0.01/0.01 16:57:39
```

Adding a Second Member to Create a Two-member Cluster

1. Format the new member's volumes
2. Create the new member's services' configurations
3. Copy the member-specific volumes
4. Update the user directory
5. Update the shared system configuration
6. Enable the existing member to access the new member
7. IPL the new member
8. Update the Product Inventory Table
9. Build the saved segments
10. **XAUTOLOG AUTOLOG1** and check MEMBER2

Enable Existing Members to Accept the New Member

```
set ssi slot 2 member2
Ready; T=0.01/0.01 16:57:51
q ssi
16:57:53 SSI Name: MYCLUSTR
16:57:53 SSI Mode: Stable
16:57:53 Cross-System Timeouts: Enabled
16:57:53 SSI Persistent Data Record (PDR) device: VMCOM1 on E00A
16:57:53 SLOT SYSTEMID STATE      PDR HEARTBEAT        RECEIVED HEARTBEAT
16:57:53    1 MEMBER1  Joined     2011-10-13 16:57:47 2011-10-13 16:57:47
16:57:53    2 MEMBER2  Down (not IPLed)
16:57:53    3 -------- Available
16:57:53    4 -------- Available
Ready; T=0.01/0.01 16:57:53
```

```
activate islink 50 60 70
16:58:26 Link device 0050 activated.
16:58:26 Link device 0060 activated.
16:58:26 Link device 0070 activated.
Ready; T=0.01/0.01 16:58:26
```

Summary

- SSI is a new way to deploy z/VM images and resources
  - Benefit from clustering and virtual server mobility

- Planning and thought required
  - Capacity and equipment
  - Resource sharing
  - Virtual networks
  - Installation
    - SSI cluster configuration
    - Migrating from your current z/VM environment
  - User directory
    - Virtual machine (guest) definition and distribution
    - Live Guest Relocation

- New documentation to assist with
  - SSI Planning
  - Migrating to an SSI cluster

More Information

z/VM 6.2 resources
http://www.vm.ibm.com/zvm620/

z/VM Single System Image Overview
http://www.vm.ibm.com/ssi/

Redbook – An Introduction to z/VM SSI and LGR
http://publib-b.boulder.ibm.com/redpieces/abstracts/sg248006.html?Open

John Franciscovich
IBM
Endicott, NY
francisj@us.ibm.com

Session 11922

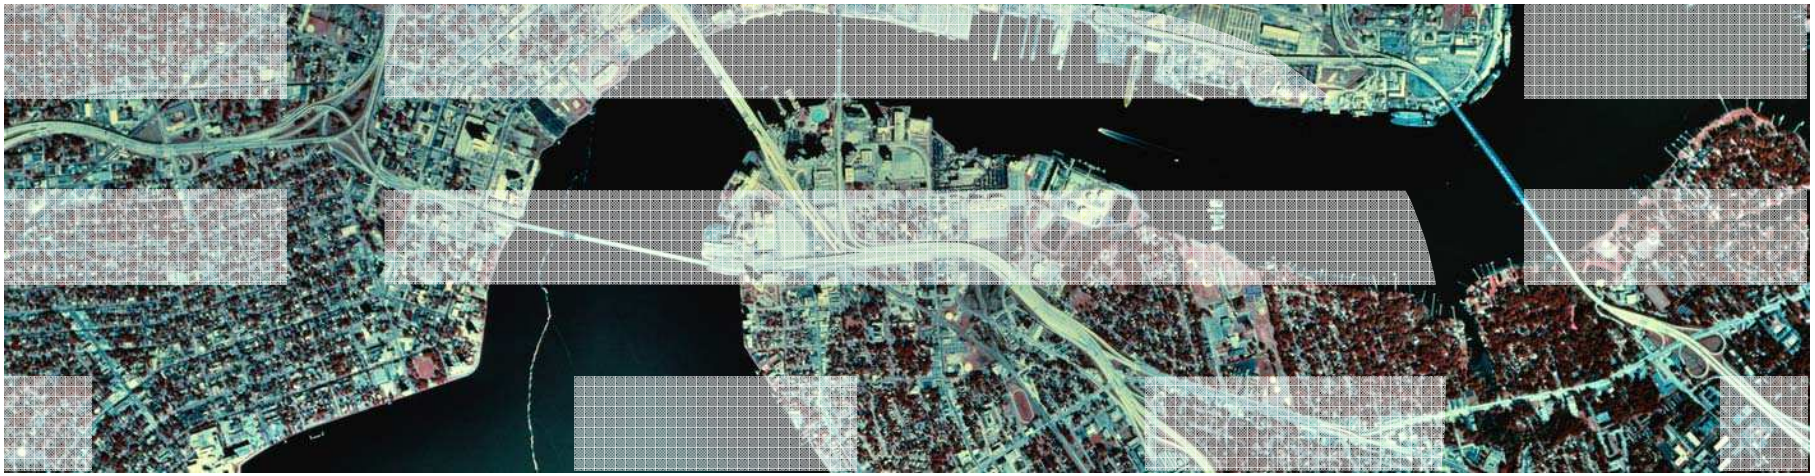

Additional Information

INSTPLAN - Select Installation Type

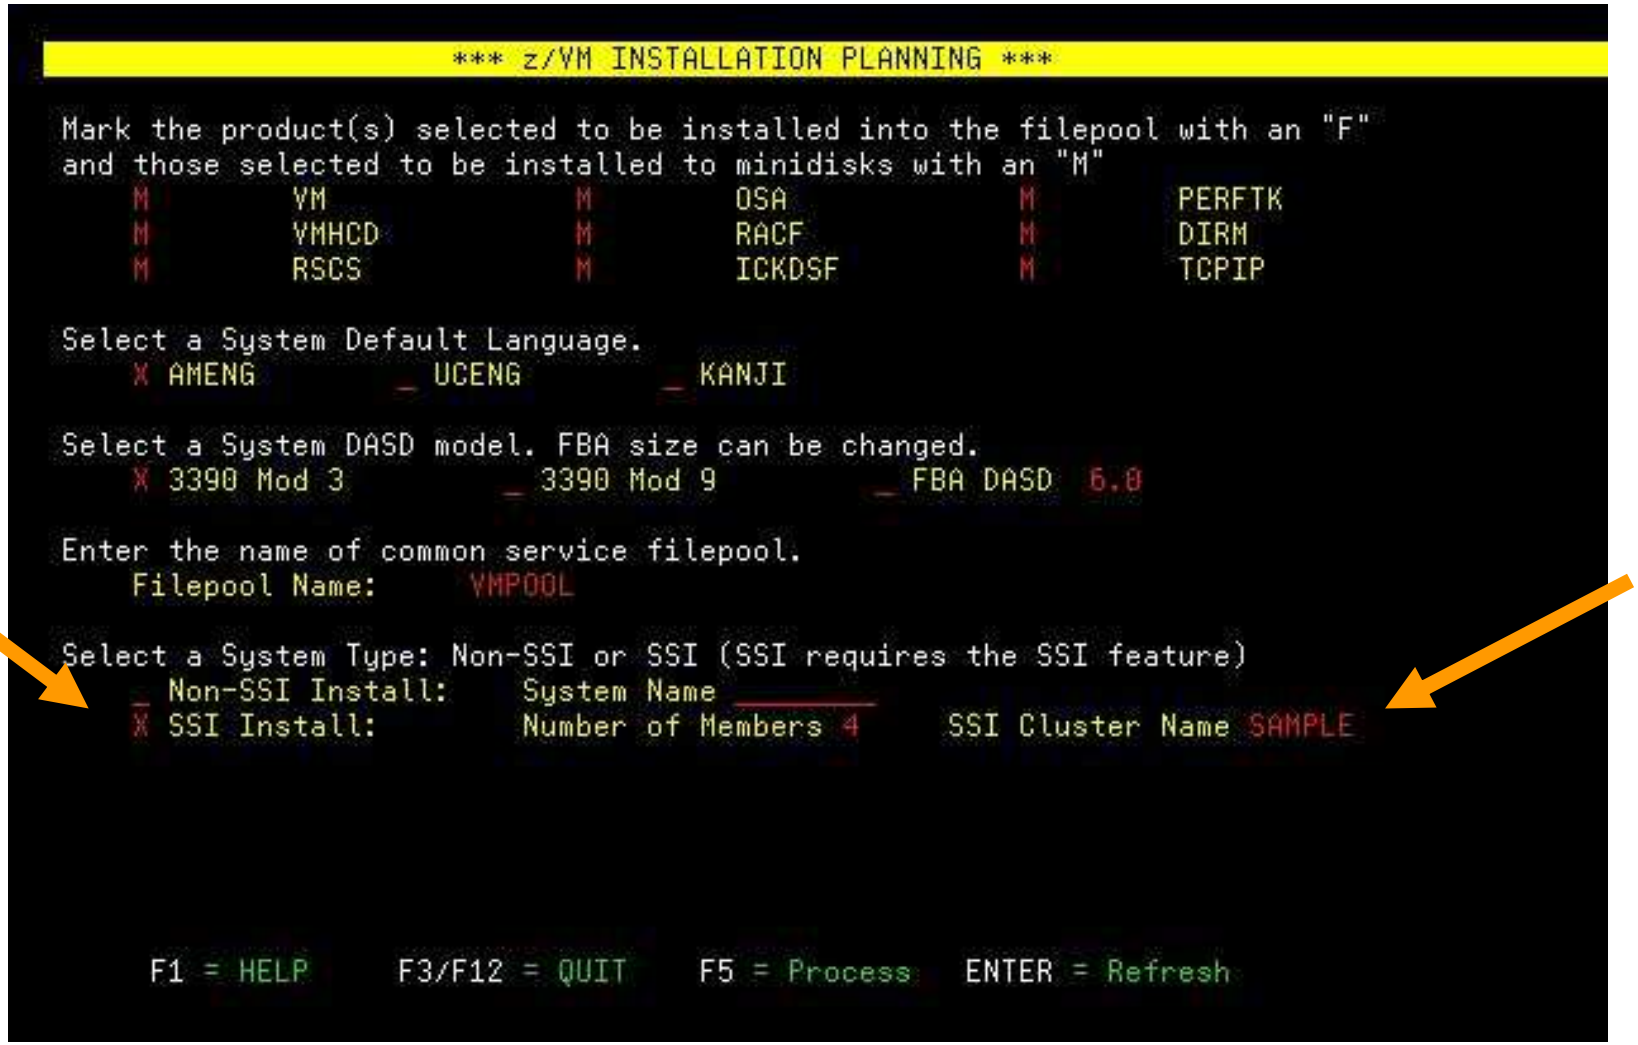

INSTPLAN - SSI Installation

Select first or second level and identify SSI member systems

```
                    *** z/VM INSTALLATION PLANNING PANEL 3 ***


SSI Cluster Name:    SAMPLE

After installation is complete, the SSI cluster will be IPLed:

X    First-Level
_    Second-Level

SSI Member Name(s):

SLOT #      MEMBER NAME     IPL LPAR/USERID
======      ===========     ===============
  1          MEMBER1          LPAR1
  2          MEMBER2          LPAR2
  3          MEMBER3          LPAR3
  4          MEMBER4          LPAR4



     F1 = HELP    F3/F12 = QUIT    F5 = Process    ENTER = Refresh
```

INSTPLAN - SSI Installation (cont.)

Define CP-Owned and Release volumes for all members

```
                *** z/VM INSTALLATION VOLUME DEFINITION ***

          TYPE      LABEL      ADDRESS                    FORMAT (Y/N)
         =======   ========   =========                   ============
         COMMON    VMCOM1       2000                           Y
         COMMON2   VMCOM2       2001
         RELVOL    620RL1       2002
         RELVOL2   620RL2       2003

          TYPE      LABEL      ADDRESS          TYPE      LABEL      ADDRESS
         =======   ========   =========        =======   ========   =========
MEMBER1                                   MEMBER2
         RES       M01RES       3000              RES       M02RES       4000
         SPOOL     M01S01       3001              SPOOL     M02S01       4001
         PAGE      M01P01       3002              PAGE      M02P01       4002
         WORK      M01W01       3003              WORK      M02W01       4003
MEMBER3                                   MEMBER4
         RES       M03RES       5000              RES       M04RES       6000
         SPOOL     M03S01       5001              SPOOL     M04S01       6001
         PAGE      M03P01       5002              PAGE      M04P01       6002
         WORK      M03W01       5003              WORK      M04W01       6003


          F1 = HELP      F3/F12 = QUIT      F5 = Process      ENTER = Refresh
```

INSTPLAN - SSI Installation (cont.)

Define Common Volume and CTC Device addresses

```
                *** z/VM INSTALLATION FIRST-LEVEL CONFIGURATION ***

Real addresses for the common volume on each member LPAR:

  VOLUME     DASD       MEMBER1     MEMBER2     MEMBER3     MEMBER4
   TYPE      LABEL      ADDRESS     ADDRESS     ADDRESS     ADDRESS
  =======   ========   =========   =========   =========   =========
   COMMON    VMCOM1      2000        2000        2000        2000

CTC device addresses:

  From: MEMBER1                          From: MEMBER2
        To: MEMBER1      N/A                   To: MEMBER1  0100 0101
        To: MEMBER2  0100 0101                 To: MEMBER2      N/A
        To: MEMBER3  0300 0301                 To: MEMBER3  0310 0311
        To: MEMBER4  0400 0401                 To: MEMBER4  0410 0411

  From: MEMBER3                          From: MEMBER4
        To: MEMBER1  0300 0301                 To: MEMBER1  0400 0401
        To: MEMBER2  0310 0311                 To: MEMBER2  0410 0411
        To: MEMBER3      N/A                   To: MEMBER3  0320 0321
        To: MEMBER4  0320 0321                 To: MEMBER4      N/A

     F1 = HELP      F3/F12 = QUIT     F5 = Process     ENTER = Refresh
```

INSTPLAN - Non-SSI Installation

Identify CP-Owned and Release volumes

CTC Connections – Defining in the IOCP

```
*
CHPID PATH=(CSS(0,1),4A),PCHID=222,TYPE=FC,SHARED   SX*FC4 11/LG04/D3
CHPID PATH=(CSS(0,1),4E),PCHID=282,TYPE=FC,SHARED   SX*FC4 16/LG02/D3
**************************************************************
*** CHPID 4A SX FICON CTC                                  ***
**************************************************************
*
 CNTLUNIT CUNUMBR=0C00,PATH=((CSS(0),4A)),UNIT=FCTC,                  *
               UNITADD=((00,8)),CUADD=7
 IODEVICE ADDRESS=(0C00,8),CUNUMBR=(0C00),UNIT=FCTC,UNITADD=00,       *
               PART=((CSS(0),TEST7,TESTC))
*
**************************************************************
*** CHPID 4E SX FICON CTC                                  ***
**************************************************************
*
 CNTLUNIT CUNUMBR=0D00,PATH=((CSS(0),4E)),UNIT=FCTC,                  *
               UNITADD=((00,8)),CUADD=C
 IODEVICE ADDRESS=(0D00,8),CUNUMBR=(0D00),UNIT=FCTC,UNITADD=00,       *
               PART=((CSS(0),TEST7,TESTC))
```