



Preparing for 8-Member SSI (Preview)

John Franciscovich
z/VM Development and Service
francisj@us.ibm.com

Jacob Gagnon
z/VM Client Focus & Care
jpgagnon@us.ibm.com

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

Db2*	FlashCopy*	IBM eserver	OMEGAMON*	XIV*	z10 BC	zSecure
DirMaint	FlashSystem	IBM (logo)*	PR/SM	z13*	z10EC	zSeries*
DS8000*	GDPS*	IBM Z*	RACF*	z13s	z/Architecture*	z/VM*
ECKD	ibm.com	LinuxONE*	System z10*	z14	zEnterprise*	z Systems*
FICON*	IBM Cloud*	LinuxONE Emperor	System 390*	z15	zPDT	
		LinuxONE Rockhopper	WebSphere*	z16	z/OS*	

* Registered trademarks of IBM Corporation

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.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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.

Linux is a registered trademark of Linus Torvalds 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.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other product and service names might be trademarks of IBM or other 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.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Objectives

This presentation will help you to understand:

- The differences between 4- and 8-member capable SSI clusters
- How to chart your migration path from a 4- to an 8- member SSI cluster
- Steps required to add new members to your SSI cluster
 - Define and prepare resources for new members
 - Set up connectivity among cluster members
- Which system files need to be updated to expand your SSI cluster
 - System configuration file
 - Directory
 - Networking
 - etc.

Agenda

- SSI Overview/Refresher
- Orientation: key terms and concepts ([Driver education](#))
- Migration paths to become “SSI-8 Capable” ([Planning your route and starting your journey](#))
- Adding/Cloning new members into slots 5-8 ([Continuing your journey - the next leg](#))
- Going live – joining the new members to the SSI cluster ([Arriving at your destination](#))



Acknowledgments

Thanks to all who helped build this presentation:

- Alan Altmark
- Kay Blake
- Carol Everitt
- Glenda Ford
- Arielle Goldberg
- Timothy Greer
- Charlie Lawrence
- Mark Lorenc
- Lauren Maietti
- Cameron Miller
- Daniel Shekhtman

SSI Overview / Refresher

z/VM Single System Image (SSI)

- Multiple z/VM instances (members) may be included in a single system image (SSI) cluster
 - Same or different CECs
- Resources are shared among member z/VM systems and their guests
 - Managed as a single resource pool
- **Live Guest Relocation** provides virtual server mobility
 - Move Linux guests non-disruptively from one member of the cluster to another
 - Flexibility for planned outages for service and migration
- Simplified system management of a multi-z/VM environment
 - Concurrent installation of multiple-system cluster
 - Single maintenance stream
 - Reliable sharing of resources and data, managed by z/VM
- Allows growth of z/VM workloads
 - Distribution and balancing of resources and workloads
 - Move workload to where the resources are

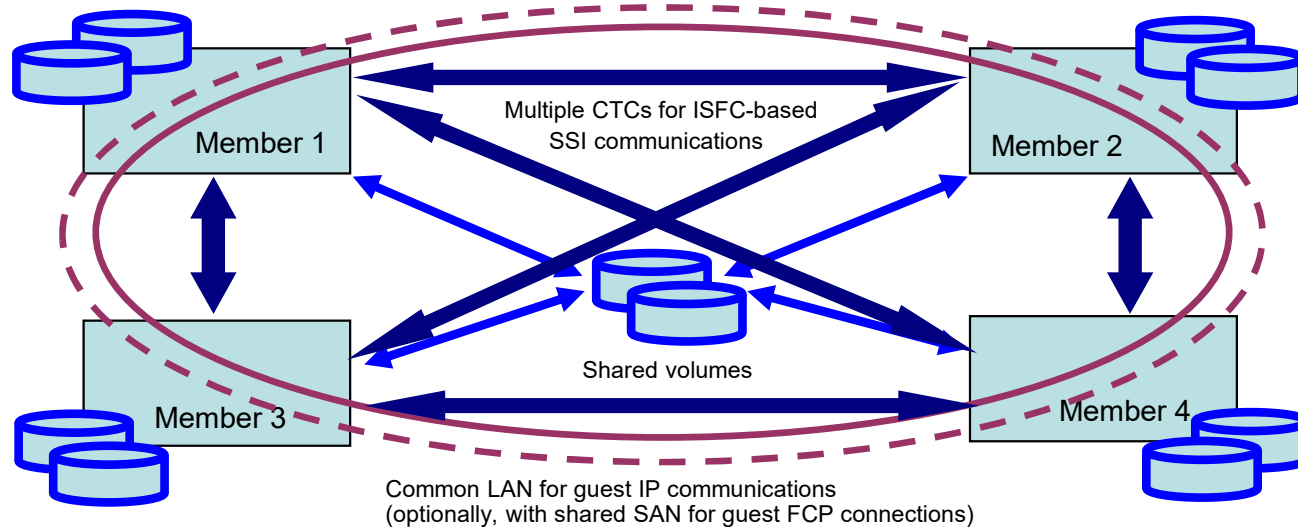
SSI Cluster Considerations

- ECKD DASD is required
 - Installation
 - RACF database
 - Common volume which contains information about cluster configuration and status
- Physical systems must be close enough to allow...
 - FICON CTC connections
 - Shared DASD
 - Common network and disk fabric connections
- Installation to SCSI devices is not supported
 - Guests may use SCSI devices
- Live Guest Relocation is supported only for Linux on IBM zSystems guests

These have not changed for 8-member clusters

z/VM SSI Evolution

- z/VM 6.2
 - Introduced as a priced feature
 - Up to 4 member systems per SSI cluster
- z/VM 7.1
 - Included at no extra cost
- z/VM 7.3
 - Support is extended to allow SSI clusters with up to 8 members
 - Growth of existing clusters to allow for more workload
 - Increases flexibility for Live Guest Relocation and workload balancing



Orientation: Key Terms and Concepts

(Driver Education)



SSI Cluster Configurations

- **SSI-4 capable**
 - All SSI clusters in z/VM 6.2 - z/VM 7.2
 - Up to 4 members
 - Members are defined on SSI configuration statement, in any or all of slots 1-4
 - May continue to configure **SSI-4 capable** clusters in z/VM 7.3

- **SSI-8 capable**
 - Beginning with z/VM 7.3
 - Up to 8 members
 - New **SSI_CONTROLS** configuration statement (**SPOOL_MEMBERS 8**)
 - Members are defined on SSI configuration statement, in any or all of slots 1-8
 - Must be configured as **SSI-8 capable** to define members in any of slots 5-8
 - Regardless of how many members are defined

SSI Spool File Sharing and Limits

Things you should know

- **USERS**

- Full sharing of all spool files across cluster members
- Allowed half the number of spool files created on each member of an SSI-8 capable cluster as an SSI-4 capable cluster
- If a USER has too many files on a given member, excess files will be deleted during IPL
 - SFCOUNTR utility available to identify these users and process excess files ahead of time


- **IDENTITYs**

- No sharing of spool files across cluster members
- Allowed the same number of spool files in created on each member of an SSI-8 capable cluster as an SSI-4 capable cluster

SSI Spool File Sharing and Limits

The numbers


Non-SSI



SSI-4



SSI-8



SSI Spool File Sharing and Limits

SFCOUNTR Utility

- **Run the SFCOUNTR utility on all existing members to identify users with too many spool files for an SSI-8 capable cluster**
 - Included in z/VM 7.3 (PMAINT 551 disk)
 - Also available in the [z/VM Download Library](#)
 - User must have privilege class D and (C or E) to run SFCOUNTR
 - IDENTITYs in the SFCOUNTR output that have more than the file limit can be ignored

Sample Output:

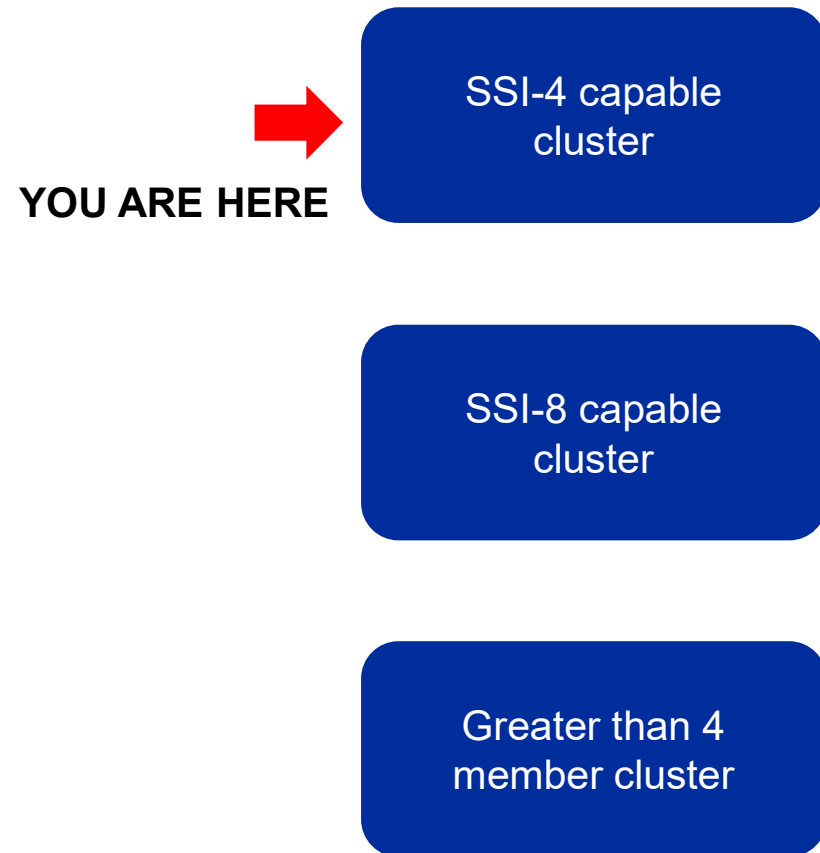
```
SFCOUNTR data collected by OPERATOR AT VM1  
TIME IS 11:32:58 EDT FRIDAY 07/02/21  
Threshold was set to 1124 files.
```

USER	FILES
====	=====
USERABC	2304
USERXYZ	1518
WHOEVER1	1125

Road Trip Check-in (1)

ARE YOU READY TO LEAVE? Before you can grow your cluster to more than 4 members, you need to have an SSI-8 capable cluster. There are a few routes to arriving at SSI-8 capable, so the first thing is to chart your path and set up your navigation.

STARTING YOUR JOURNEY: Once you've charted your path to SSI-8 capable, you are ready do your preparation and then follow the appropriate Upgrade or SSI Traditional Installation procedure in the *Installation Guide*.



Migration Paths to Become SSI-8 Capable

(Planning your route and starting your journey)



Charting your path

Start here: z/VM 7.1 or 7.2 members



Make each member SSI-8 capable one at a time (1 IPL per member)

Make the cluster SSI-8 capable all at once (2 IPLs per member)

Continue on: z/VM 7.3 members



Fresh install: z/VM 7.3

Up to 4 members


Fresh install: z/VM 7.3

Merge two z/VM 7.3 4-member clusters

Add members 5-8



Reach your destination: SSI-8 enabled cluster



Making all members SSI-8 Capable "at the end"

Upgrading an existing cluster

Start here: z/VM 7.1 or 7.2 members



Make each member SSI-8 capable one at a time (1 IPL per member)

Make the cluster SSI-8 capable all at once (2 IPLs per member)

Continue on: z/VM 7.3 members



Fresh install: z/VM 7.3



Up to 4 members

Fresh install: z/VM 7.3

Merge two z/VM 7.3 4-member clusters

Add members 5-8

Reach your destination: SSI-8 enabled cluster



Making all members SSI-8 Capable "at the end"

Things you should know

- This is the recommended upgrade path for an SSI cluster currently on z/VM 7.1 and/or 7.2
- Upgrade installation steps are documented in the *z/VM Installation Guide*
- Each member will be IPL'd twice
 1. IPL of z/VM 7.3 as part of the upgrade
 2. IPL to change each member to be SSI-8 capable
- After completing this upgrade, your cluster will:
 - Be SSI-8 capable
 - Still have only members 1-4 configured (as they were before)

Making all members SSI-8 Capable "at the end"

Upgrade steps

1. Run the **SFCOUNTR** utility with the default limit on all members
 - Process/prune spool files as needed
2. Follow the documented procedure to do an upgrade installation of **each member** to z/VM 7.3
 - IPL z/VM 7.3 as each member is upgraded

At this point, the cluster remains SSI-4 capable

3. When all members have been upgraded:
 - a) Add new **SSI_CONTROLS** statement to the (shared) system config file

SSI_CONTROLS SPOOL_MEMBERS 8

- b) IPL each member

The cluster is now SSI-8 capable; new members may be now be added in slots 5-8

Making all members SSI-8 Capable "at the end"

Upgrading an existing cluster

Start here: z/VM 7.1 or 7.2 members



Make each member SSI-8 capable one at a time (1 IPL per member)

Make the cluster SSI-8 capable all at once (2 IPLs per member)

Continue on: z/VM 7.3 members



Fresh install: z/VM 7.3



Up to 4 members

Fresh install: z/VM 7.3

Merge two z/VM 7.3 4-member clusters

Add members 5-8

Reach your destination: SSI-8 enabled cluster



Becoming SSI-8 Capable "As-you-go"

Things you should know

- This is an alternate upgrade path for an SSI cluster currently on z/VM 7.1 and/or 7.2
- Upgrade installation steps are documented in the *z/VM Installation Guide (7.3)*
- The PTF for compatibility APAR VM66462 must be applied to all z/VM 7.1 and 7.2 members before any z/VM 7.3 member is IPL'd with **SPOOL_MEMBERS 8**

Component	APAR	PTF	Available	RSU
CP	VM66462	z/VM 7.2 UM35787 z/VM 7.1 UM35786	March 30, 2021 June 29, 2021	2101

- Each member will be IPL'd once as part of its upgrade to z/VM 7.3
- After completing this upgrade, your cluster will:
 - Be SSI-8 capable
 - Still have only members 1-4 configured (as they were before)

Becoming SSI-8 Capable "As-you-go"

Upgrade steps


1. Ensure that the PTF for compatibility APAR VM66462 is applied to all z/VM 7.1 and 7.2 members
 - If not, the member being upgraded will not IPL
2. Run the **SFCOUNTR** utility with the default limit on all members
 - Process/prune spool files as needed
3. Follow the documented procedure to do an upgrade installation of the first member to z/VM 7.3
 - Add new **SSI_CONTROLS** statement to the (shared) system config file
 - » Add a record qualifier for the member being upgraded
MEMBER1: SSI_CONTROLS SPOOL_MEMBERS 8
 - IPL z/VM 7.3
4. Repeat Step 3 for each additional member
 - Add a record qualifier to the **SSI_CONTROLS** statement for each member being upgraded

The cluster is SSI-8 capable when completed for all members; new members may be added in slots 5-8

Note: After this process is completed for all members, the record qualifiers may be removed from the **SSI_CONTROLS** statement

Creating a new SSI-8 Capable cluster

Start here: z/VM 7.1 or 7.2 members



Make each member SSI-8 capable one at a time (1 IPL per member)

Make the cluster SSI-8 capable all at once (2 IPLs per member)

Continue on: z/VM 7.3 members



Fresh install: z/VM 7.3

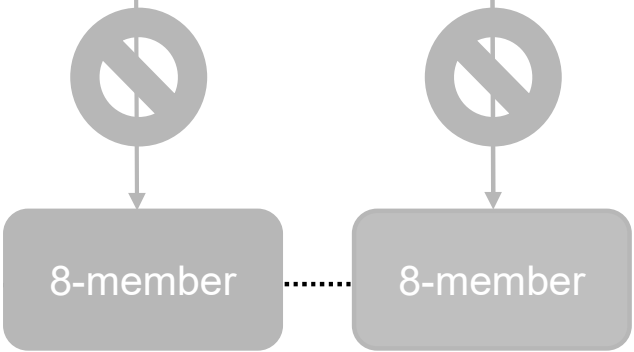
Up to 4 members

Fresh install: z/VM 7.3

Merge two z/VM 7.3 4-member clusters

Add members 5-8

Reach your destination: SSI-8 enabled cluster



Creating a new SSI-8 Capable cluster

z/VM 7.3 SSI Installation (Fresh Install)

- Up to 4 members may be installed (in slots 1-4)
- Follow the procedure in the *z/VM Installation Guide*
 - "SSI Traditional Installation Method" chapter
- Installed SSI cluster will be **SSI-8 capable**
 - Install process will add an **SSI_CONTROLS SPOOL_MEMBERS 8** statement to the system config file

Road Trip Check-in (2)

ARE YOU THERE YET? Whether you got here from a fresh install, or from an upgrade of an existing cluster, you are now SSI-8 capable. All members of your cluster should be running z/VM 7.3 or later.

HOW MUCH FURTHER? Now the real fun can begin! The next step is to break out *CP Planning and Administration*, and follow the steps to clone a new member into slots 5-8

YOU ARE HERE




SSI-4 capable cluster

SSI-8 capable cluster

Greater than 4 member cluster

Adding/Cloning New Members into Slots 5-8

(Continuing your journey - the next leg)



Adding Members 5-8 (n-8?)

Start here: z/VM 7.1 or 7.2 members



Make each member SSI-8 capable one at a time (1 IPL per member)

Make the cluster SSI-8 capable all at once (2 IPLs per member)

Continue on: z/VM 7.3 members



Fresh install: z/VM 7.3



Up to 4 members

Fresh install: z/VM 7.3

Merge two z/VM 7.3 4-member clusters

Add members 5-8

Reach your destination: SSI-8 enabled cluster



What Tasks to be Mindful of With SSI-8?

- ➡ Task 1: Prepare the CP-Owned Volumes for the Target Member
- ➡ Task 2: Create the TCP/IP Configuration for the Target Member
- ➡ Task 3: Update the Configuration Files for Other Service Virtual Machines and Servers
- ➡ Task 4: Customize the System Startup Virtual Machine
- ➡ Task 5: Copy the Source Volumes to the Target Volumes
- ➡ Task 6: Update the User Directory, or Task 6A: Update the User Directory Using DirMaint
- ➡ Task 7: Update the System Configuration File
- ➡ Task 8: Enable the Existing Members to Access the Target Member
- ➡ Task 9: IPL the Target Member
- ➡ Task 10: Start the ESM Service Virtual Machine
- ➡ Task 11: Update the VMSES/E System-Level Product Inventory Table
- ➡ Task 12: Update the CRR Server LU Name
- ➡ Task 13: Build the Saved Segments and Named Saved Systems
- ➡ Task 14: Start the Service Virtual Machines
- ➡ Task 15: Test and Verify

Legend	
➡	Business As Usual
➡	Handle With Caution (changes scale linearly with SSI-8)
➡	Handle With Care (considerable changes for SSI-8)

What Tasks to be Mindful of With SSI-8?

- ➔ Task 1: Prepare the CP-Owned Volumes for the Target Member
- ➔ Task 2: Create the TCP/IP Configuration for the Target Member
- ➔ Task 3: Update the Configuration Files for Other Service Virtual Machines and Servers
- ➔ Task 4: Customize the System Startup Virtual Machine
- ➔ Task 5: Copy the Source Volumes to the Target Volumes
- ➔ Task 6: Update the User Directory, or Task 6A: Update the User Directory Using DirMaint
- ➔ Task 7: Update the System Configuration File
- ➔ Task 8: Enable the Existing Members to Access the Target Member
- ➔ Task 9: IPL the Target Member
- ➔ Task 10: Start the ESM Service Virtual Machine
- ➔ Task 11: Update the VMSES/E System-Level Product Inventory Table
- ➔ Task 12: Update the CRR Server LU Name
- ➔ Task 13: Build the Saved Segments and Named Saved Systems
- ➔ Task 14: Start the Service Virtual Machines
- ➔ Task 15: Test and Verify

Legend	
➔	Business As Usual
➔	Handle With Caution (changes scale linearly with SSI-8)
➔	Handle With Care (considerable changes for SSI-8)

Business-As-Usual (cheat sheet) →

Task	Task Description	What Changes	Where it lives
1	Prepare CP-Owned Volumes of Target	CPFMTXA of System Page and Spool volumes	Task performed on MAINTvrm ID
5	Copy Source Volumes to Target	Flashcopy or DDR of System Page and Spool volumes	Task performed on MAINTvrm ID
9	IPL Target	N/A	Performed at HMC
10	Start the ESM Service Machine	Applicable ESM machine	N/A
11	Update VMSES/E Product Inventory Table	VM SYSPINV VM SYSSUF	Task performed on MAINTvrm ID
12	Update CRR Server LU Name	VMSEVR DMSPARMS	Task performed on VMSEVR ID
15	Test and Verify	N/A	N/A

* This is not a replacement for following the steps in CP Planning and Administration: "Adding a Member to a z/VM SSI Cluster by Cloning an Existing Member"

Preparation Makes the Ride Go Smoother

- Gather configuration information
- Complete your worksheets
- Think about naming/numbering conventions
- Use spreadsheet for easier reviewing / changes

Table 1. Traditional installation worksheet 1

Installation method (first-level or second-level): _____

In the **Install To** column, record an M if you will load the product to a minidisk or an F if you will load the product to the VMPSFS file pool.

Install To	Product	Install To	Product	Install To	Product
	VM		DBRM		ICKDSF
	PERFTK		RACF*		RSCS
	TCPIP		VMHCD		

Default system language: _____

DASD type: _____

Volume size: _____

Common service file pool name: _____

Installation Type: _____

Table 4. Traditional installation worksheet 4

IP address or host name: _____

User ID and password of FTP server: _____

Path name of DVD drive, USB flash drive, or F: _____

VM user ID and address of VM minidisk to upl: _____

Table 5. Traditional installation worksheet 5 (3390 SSI Only)

After installation is complete, SSI will be IPLed:

___ First-Level
___ Second-Level

SSI Member Name(s) / IPL LPAR Name(s) or User ID Name(s): _____

Slot Number	Member Name*	IPL LPAR/User ID
1		
2		
3		

Table 6. Traditional installation worksheet 6 (3390 SSI Only)

Volume Type	Default Label	New Label	Address
COMMON	VMCOM1		
RELVOL	720RL1		
RELVOL2*	720RL2		

Volume Type	Default Label	New Label	Address	Volume Type	Default Label	New Label
Member 1:			Member 2:			
RES	M01RES		RES	M02RES		
SPOOL	M01S01		SPOOL	M02S01		
PAGE	M01P01		PAGE	M02P01		
Member 3:			Member 4:			
RES	M03RES		RES	M04RES		
SPOOL	M03S01		SPOOL	M04S01		
PAGE	M03P01		PAGE	M04P01		

CTC device addresses:

From: Member 1	To: Member 2	From: Member 3	To: Member 4
To: Member 1	N/A	To: Member 1	
To: Member 2		To: Member 2	N/A
To: Member 3		To: Member 3	
To: Member 4		To: Member 4	
From: Member 3		From: Member 4	
To: Member 1		To: Member 1	
To: Member 2		To: Member 2	
To: Member 3	N/A	To: Member 3	
To: Member 4		To: Member 4	N/A

* Depending on your planning choices, you might or might not need this volume.
Note: You must not use any of IBM's default volume labels for a volume other than the volume for which it is originally defined.

Figure 1: Worksheets from Chapter 2 of z/VM Installation

Default Label	New Label	Address	Volumes accessible to all members:
VMCOM1			Mega V LCU 3E
720RL1			CE06-CE08 (3 mod 54's) SAD
720RL2			CE09-CE0C (4 mod 54's) Processing Dumps

Default Label	New Label	Address	Mega VI LCU 8
M01RES	C13A14	3A14	mod 9
M01S01	C13A15	3A15	mod 9
M01P01	C13A16	3A16	mod 9
	C13A17	3A17	mod 9
	C13A18	3A18	Alias
	C13A19	3A19	Alias
	C13A1A	3A1A	Alias
	C13A1B	3A1B	Alias
	C13A1C	3A1C	Alias
	C13A1D	3A1D	Alias

Default Label	New Label	Address	Mega VI LCU 8
M02RES	C23A00	3A00	mod 9
M02S01	C23A01	3A01	mod 9
M02P01	C23A02	3A02	mod 9
	C23A03	3A03	mod 9
	C23A04	3A04	mod 9
	C23A05	3A05	mod 9
	C23A06	3A06	Alias
	C23A07	3A07	Alias
	C23A08	3A08	Alias
	C23A09	3A09	Alias
	C23A0A	3A0A	Alias
	C23A0B	3A0B	Alias
	C23A0C	3A0C	Alias
	C23A0D	3A0D	Alias
	C23A0E	3A0E	Alias
	C23A0F	3A0F	Alias

Default Label	New Label	Address	Mega V LCU 4F
M03RES	C30370	D370	mod 9
M03S01	C30371	D371	mod 9
M03P01	C30372	D372	mod 9
	C30373	D373	mod 9
	C30374	D374	mod 9
	C30375	D375	mod 9
	C30376	D376	mod 9
	C30377	D377	Alias
	C30378	D378	Alias
	C30379	D379	Alias
	C3037A	D37A	Alias
	C3037B	D37B	Alias
	C3037C	D37C	Alias

What Tasks to be Mindful of With SSI-8?

- ➔ Task 1: Prepare the CP-Owned Volumes for the Target Member
- ➔ Task 2: Create the TCP/IP Configuration for the Target Member
- ➔ Task 3: Update the Configuration Files for Other Service Virtual Machines and Servers
- ➔ Task 4: Customize the System Startup Virtual Machine
- ➔ Task 5: Copy the Source Volumes to the Target Volumes
- ➔ Task 6: Update the User Directory”, or “Task 6A: Update the User Directory Using DirMaint
- ➔ Task 7: Update the System Configuration File
- ➔ Task 8: Enable the Existing Members to Access the Target Member
- ➔ Task 9: IPL the Target Member
- ➔ Task 10: Start the ESM Service Virtual Machine
- ➔ Task 11: Update the VMSES/E System-Level Product Inventory Table
- ➔ Task 12: Update the CRR Server LU Name
- ➔ Task 13: Build the Saved Segments and Named Saved Systems
- ➔ Task 14: Start the Service Virtual Machines
- ➔ Task 15: Test and Verify

Legend	
➔	Business As Usual
➔	Handle With Caution (changes scale linearly with SSI-8)
➔	Handle With Care (considerable changes for SSI-8)

Handle-With-Caution (Cheat Sheet) →

Task	Task Description	What Changes	Where it lives
2	Configure TCP/IP Stack	PROFILE TCPIP	TCPMAINT 198
		TCPIP DATA	TCPMAINT 592
3	Make Server Virtual Machines Multi-Nodal	Applicable Configuration Files	N/A
4	Customize Startup Machine	PROFILE EXEC (for AUTOLOG1 or AUTOLOG2)	AUTOLOG1 or AUTOLOG2 191 minidisk

* This is not a replacement for following the steps in CP Planning and Administration: "Adding a Member to a z/VM SSI Cluster by Cloning an Existing Member"

Know WHAT You're Changing

- These steps are generally meant to set up your services
- A misstep here could leave you without a key system service
- When your IDENTITYs enter their new home, make sure the lights are on!



What Tasks to be Mindful of With SSI-8?

- ➔ Task 1: Prepare the CP-Owned Volumes for the Target Member
- ➔ Task 2: Create the TCP/IP Configuration for the Target Member
- ➔ Task 3: Update the Configuration Files for Other Service Virtual Machines and Servers
- ➔ Task 4: Customize the System Startup Virtual Machine
- ➔ Task 5: Copy the Source Volumes to the Target Volumes
- ➔ Task 6: Update the User Directory, or Task 6A: Update the User Directory Using DirMaint
- ➔ Task 7: Update the System Configuration File
- ➔ Task 8: Enable the Existing Members to Access the Target Member
- ➔ Task 9: IPL the Target Member
- ➔ Task 10: Start the ESM Service Virtual Machine
- ➔ Task 11: Update the VMSES/E System-Level Product Inventory Table
- ➔ Task 12: Update the CRR Server LU Name
- ➔ Task 13: Build the Saved Segments and Named Saved Systems
- ➔ Task 14: Start the Service Virtual Machines
- ➔ Task 15: Test and Verify

Legend	
➔	Business As Usual
➔	Handle With Caution (changes scale linearly with SSI-8)
➔	Handle With Care (considerable changes for SSI-8)

Handle-With-Care (Cheat Sheet)

Task	Task Description	What Changes	Where it lives
6	Update User Directory <ul style="list-style-type: none"> Use a directory manager to simplify this task 	DIRECTORY statement	USER DIRECT
		New subconfigs for IDENTITYs	
		New MDISK statements for system residence and member-specific volumes	
7	Update System Config	System_Identifier	System Config
		SSI Statement	
		System_Residence	
		CP_Owned Statements	
		User_Volume_List	
		ACTIVATE ISLINK Statements	
		Member-specific networking	
8	Enable access of existing members to new target	SET SSI SLOT command	Performed on current system's MAINTvrm ID
		ACTIVATE ISLINKs	
		Define new spool volumes	

* This is not a replacement for following the steps in CP Planning and Administration: "Adding a Member to a z/VM SSI Cluster by Cloning an Existing Member"

What's so Tedious? →

User Directory Changes

Add new system residence volume to DIRECTORY statement(s)

Add MDISKs for New DASD

SUBCONFIGs for IDENTITYs

DIRECTXA to create new object directory on each member

Adding ISFC Connectivity

Configure CTCs in IOCDs


Define Logical Links in SYSTEM CONFIG

Activate additional Links on each member

User Directory Key Points

General Information

- Changes can be made with a directory manager (highly recommended) or by editing USER DIRECT
- Add new system residence valid to DIRECTORY statement
 - Multiple DIRECTORY statements may be used (up to 8)
- IDENTITYs will need additional subconfigs for running on the new member systems
- Changes must be written to the object directory on ALL active members
 - A directory manager will take care of this



If Using DirMaint

- You'll need DIRMSATx and DATAMOVE servers for each new member you add.
- Make sure all satellite servers are active before adding a new member

Adding ISFC Connectivity

In a 4 Member SSI, each member needs 3 logical links.

Each provides connectivity to the other 3 members

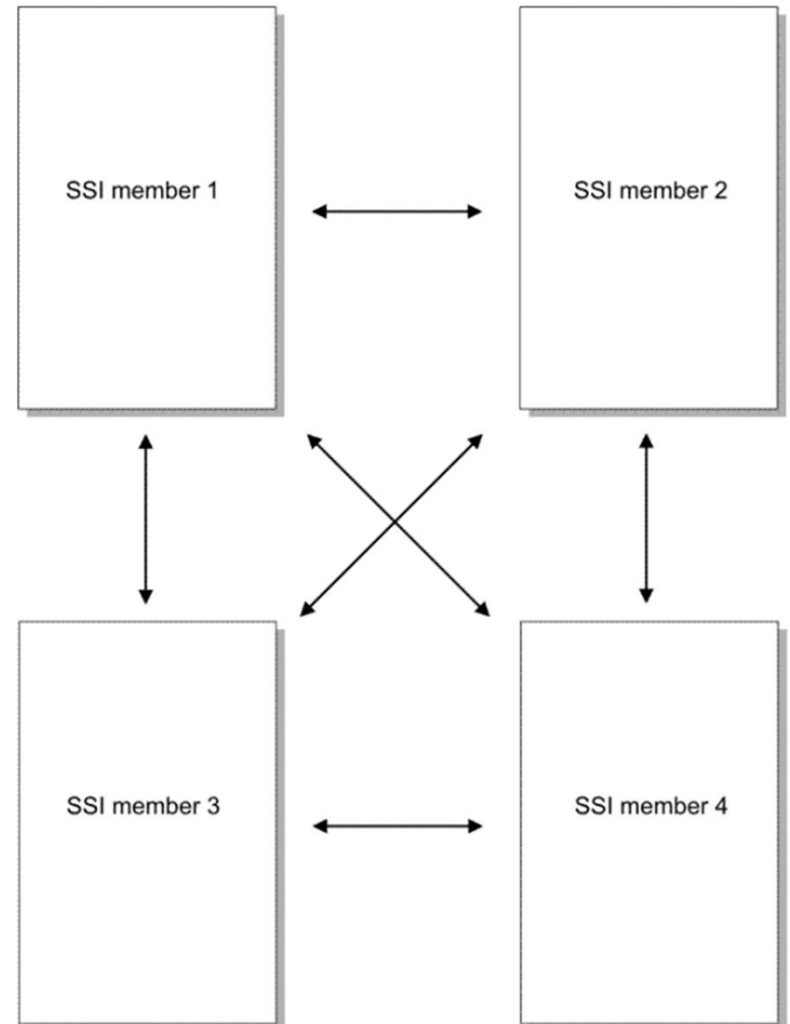
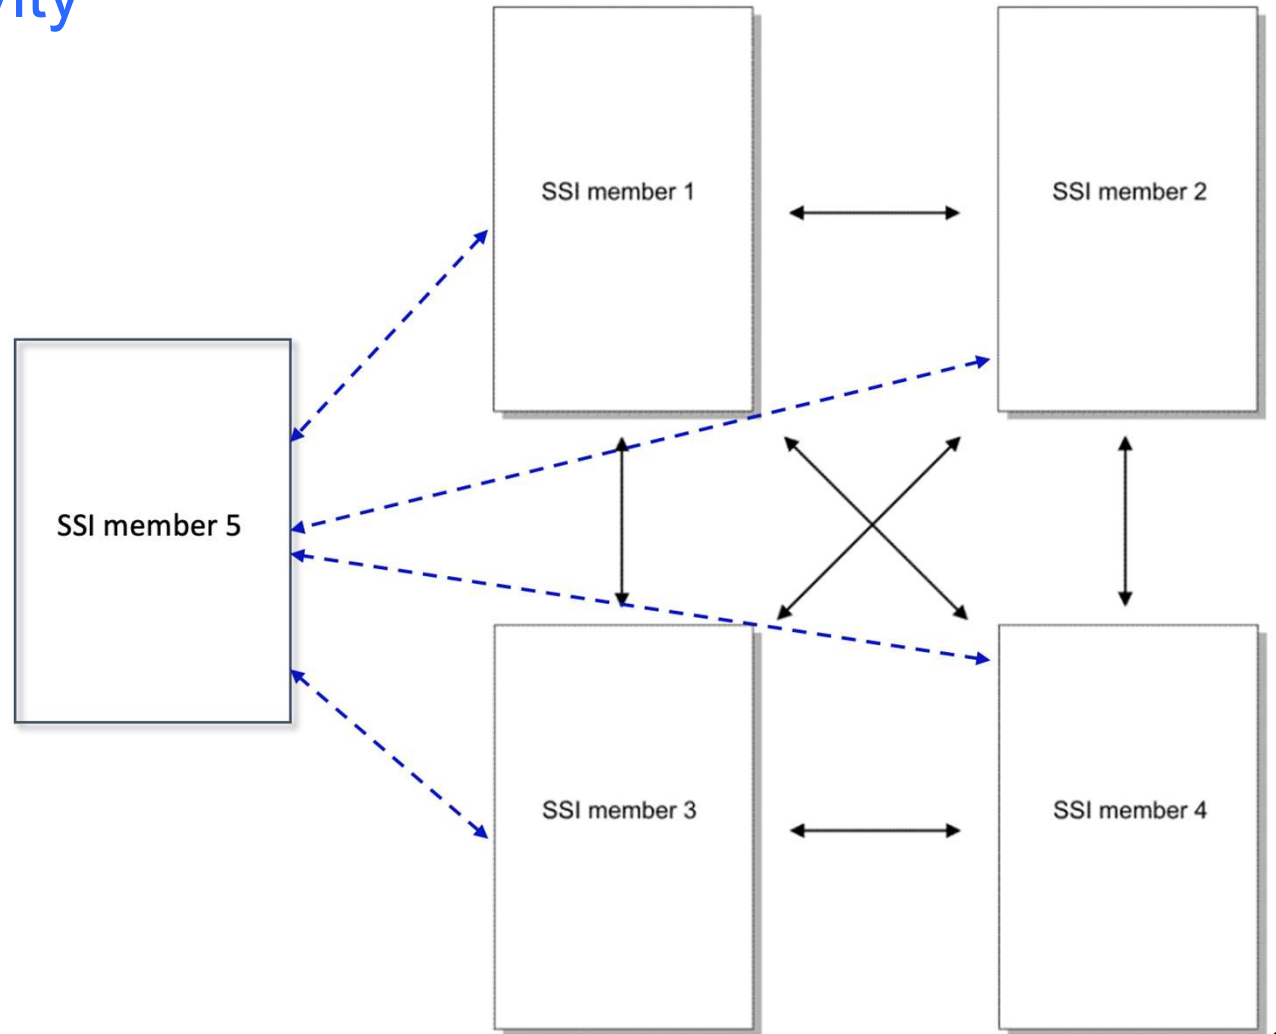


Figure 43. Fully Connected ISFC Topology

Adding ISFC Connectivity

In an (n) - Member SSI, each member needs (n-1) logical links.

Each provides connectivity to the other (n-1) members




Adding ISFC Connectivity (Diving Deeper)

- 3 Main Steps
 1. Configuration of physical CTCs
 2. Defining logical links in SYSTEM CONFIG for new or currently down cluster members
 3. Activating logical links for cluster members that are currently running

CAUTION: A logical link is required from EACH MEMBER to ALL OTHER MEMBERS. Missing one logical link definition will prevent a member from successfully joining the cluster.

A Note on Hardware

- The art of configuring a CTC network does not change with eight members
- Remember IBM Best Practices
 - 8 CTCs across 2 physical CHPIDS between systems (4 rdev's on each ISLINK) is recommended
 - 4 CTCs across 2 physical CHPIDS between systems (2 rdev's on each ISLINK) might be acceptable
 - The important thing is to have redundancy (multiple CHPIDS) and concurrency (multiple CTCs per CHPID)



Defining Logical Links in SYSTEM CONFIG

```
/*
*****
*/
/*          Activate ISLINK statements          */
/*
*****
*/

GDLTST8A:    ACTIVATE ISLINK 0850 0851  NODE GDLTST8B
GDLTST8A:    ACTIVATE ISLINK 0860 0861  NODE GDLTST8C
GDLTST8A:    ACTIVATE ISLINK 0870 0871  NODE GDLMST8D
GDLTST8A:    ACTIVATE ISLINK 0878 0879  NODE GDLMST8E
GDLTST8A:    ACTIVATE ISLINK 0880 0881  NODE GDLMST8F
GDLTST8A:    ACTIVATE ISLINK 0888 0889  NODE GDLAST8G
GDLTST8A:    ACTIVATE ISLINK 0890 0891  NODE GDLAST8H

GDLTST8B:    ACTIVATE ISLINK 0858 0859  NODE GDLTST8A
GDLTST8B:    ACTIVATE ISLINK 0898 0899  NODE GDLTST8C
GDLTST8B:    ACTIVATE ISLINK 08A8 08A9  NODE GDLMST8D
GDLTST8B:    ACTIVATE ISLINK 08B0 08B1  NODE GDLMST8E
GDLTST8B:    ACTIVATE ISLINK 08B8 08B9  NODE GDLMST8F
GDLTST8B:    ACTIVATE ISLINK 08C0 08C1  NODE GDLAST8G
GDLTST8B:    ACTIVATE ISLINK 08C8 08C9  NODE GDLAST8H
```

How Many ISLINK Statements Should You Have?

$$\text{ST_MEM} = \text{MEMBERS} - 1$$

$$\text{ST_TOT} = \text{ST_MEM} * \text{MEMBERS}$$

Where:

- **ST_MEM** = number of ACTIVATE ISLINK statements required for each member
- **ST_TOT** = total number of ACTIVATE ISLINK statements required in SYSTEM CONFIG
- **MEMBERS** = number of members in the cluster

Number Of Members	Number Of ACTIVATE ISLINK statements (per member)	Number of ACTIVATE ISLINK statements (total)
1	0	0
2	1	2
3	2	6
4	3	12
5	4	20
6	5	30
7	6	42
8	7	56

Enable Existing Members

- SYSTEM CONFIG statements cover new or down members joining the cluster, but you still need to define new slots and links on running members.
 - **SET SSI SLOT n** command
 - **ACTIVATE ISLINK** command

HINTS:

- Using EXECs to issue **ACTIVATE ISLINK**'s can help you desk check.
- Using the **AT** command can be helpful for issuing ALL required commands from ONE member.
 - From MEMBER1:
AT MEMBER2 CMD ACTIVATE ISLINK 0850 0857 NODE MEMBER3
- Spooling your console can aid debug if something goes wrong

What Could Go Wrong?

PROBLEM:

- Failure to define just one logical link can prevent successful IPL of a new or down member.

SOLUTION:

- Check ISLINK definitions
 - Active links with **QUERY ISLINK**
 - **ACTIVATE ISLINK** statements in SYSTEM CONFIG

Make it easier:

- CMS Pipelines, scripts, etc....can be helpful when desk checking

Going Live

(Arriving at your destination)



Joining New Member(s) to your SSI Cluster

Now that you've completed the cloning tasks, it's time to IPL each of your new members:

1. Verify that IPL completes successfully
2. Use the **QUERY SSI** command to verify that the new member has successfully joined the cluster
3. Verify that the resources you defined for each new member are available
 - ESM is active
 - All defined ISLINKs are active
 - Directory manager and/or satellites are active
 - Disk volumes are online
 - etc.

Road Trip Check-in (3)

YOU HAVE ARRIVED: You now have an SSI cluster with more than 4 members. Your road trip is complete. Congratulations on a safe journey!

TAKE ANOTHER LAP: Repeat the adding/cloning tasks until you have a full 8-member cluster. If you planned ahead during the first lap, additional laps will be easier.

SSI-4 capable cluster

SSI-8 capable cluster

Greater than 4 member cluster

YOU ARE HERE



Summary

- With z/VM 7.3, the benefits of a 4-member SSI cluster are extended for up to 8 members
 - Resources managed as a single resource pool
 - Virtual server mobility with Live Guest Relocation
 - Simplified system management of a multi-z/VM environment
 - Horizontal growth of z/VM workloads
- Careful planning is required for a successful road trip to more than 4 members
 - Plan naming and numbering conventions for 8 members even if you don't have immediate plans for 8 members
 - Use the worksheets
- Follow the documented procedures in order
 - Your chosen migration path to SSI-8 capable (upgrade or fresh installation)
 - Cloning an existing member to add members 5-8
- Enjoy the ride!



Thanks!

John Franciscovich
z/VM Development and Service
francisj@us.ibm.com

Jacob Gagnon
z/VM Client Focus & Care
jpgagnon@us.ibm.com