



Cloning z/VM 6.2 Systems - Everything You Need to Know!

Bruce Hayden
IBM Advanced Technical Sales Support

Thursday, August 15, 2013
Session Number 13400



Trademarks



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

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:

* AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

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.

Complete your sessions evaluation online at SHARE.org/BostonEval



System z Social Media Channels

- Top Facebook pages related to System z:

- IBM System z
- IBM Academic Initiative System z
- IBM Master the Mainframe Contest
- IBM Destination z
- Millennial Mainframer
- IBM Smarter Computing

- Top LinkedIn groups related to System z:

- System z Advocates
- SAP on System z
- IBM Mainframe- Unofficial Group
- IBM System z Events
- Mainframe Experts Network
- System z Linux
- Enterprise Systems
- Mainframe Security Gurus

- Twitter profiles related to System z:

- IBM System z
- IBM System z Events
- IBM DB2 on System z
- Millennial Mainframer
- Destination z
- IBM Smarter Computing

- YouTube accounts related to System z:

- IBM System z
- Destination z
- IBM Smarter Computing

- Top System z blogs to check out:

- Mainframe Insights
- Smarter Computing
- Millennial Mainframer
- Mainframe & Hybrid Computing
- The Mainframe Blog
- Mainframe Watch Belgium
- Mainframe Update
- Enterprise Systems Media Blog
- Dancing Dinosaur
- DB2 for z/OS
- IBM Destination z
- DB2utor



Agenda

- Background
- Changes in z/VM 6.2 and 6.3
- Tasks you still need to do
- “Easy” approach
- Cloning a non-SSI system
- Disk labels
- Cloning an SSI cluster

CP Utilities and Terminology

- Stand-Alone Program Loader, a.k.a. SAPL
 - The utility program that actually loads CP or another stand-alone program into your LPAR or virtual machine
- Stand-Alone Loader Creation Utility, a.k.a. SALIPL
 - The utility that writes the SAPL utility, along with your parameters and comments, on a DASD volume
 - You've probably used the module under CMS
 - There is also a stand-alone version
 - Can be loaded from a (virtual) reader, tape, or via SAPL!
- DASD Dump Restore utility a.k.a. DDR
 - Utility for copying disks, if you don't have or can't use flashcopy
 - It is both a CMS command and a stand-alone utility, loadable via SAPL
- Terminology – All these mean the same thing in this presentation
 - Node id, System id, Member name, System name

The way you used to do it .. (6.1 and older)

- DDR or Flashcopy the source disks to the target
- Define a 2nd level or other isolated environment to bring up the new VM
 - Or, somehow worry about duplicate labels
- IPL (override console if needed)
- Change labels if necessary
 - How to do this is discussed later
- Fix any other node identifiers, real addresses, etc.
 - These seem to hide all over the place.
- Move 2nd level system to LPAR (if that is your goal)
 - or, just use it!

What has changed?

- CP owned volumes can now have a system owner defined
 - Required on members in SSI clusters
 - Optional on non-SSI systems
 - but written by default during a non-SSI installation
 - Owner identifies the SSI cluster name and member name
 - SSI cluster name of “NOSSI” means the volume is not part of an SSI cluster
 - Special owner name of “NOSYS” means either:
 - SSI cluster volumes not tied to 1 SSI member
 - No ownership information on the volume
- The IPL parms in the Stand-alone Program Loader (SAPL) specify the real device address of the common volume
 - Used to find the SYSTEM CONFIG file

```
-----IPL PARAMETERS-----  
fn=SYSTEM ft=CONFIG pdnum=1 pdvol=7001
```

- VMSES/E (service commands) keep track of the system identifier

Cloning tasks that haven't changed

- Changing disk labels
 - The process is still the same
- System identifier may appear in several places
 - SYSTEM CONFIG
 - SYSTEM NETID S
 - DIRMAINT customization
 - TCPIP configuration
 - etc. (RSCS, RACF, names files, Perfkit config, <insert more here!>)
- Any reference to a real device
 - SYSTEM CONFIG
 - Directory
 - Execs (AUTOLOG1 or 2 system setup, others)
 - RACF (real device control, if used)
- Other **SYSTEM CONFIG** updates
 - Macprefix – important if using IP or “layer 2” network connections

How to make the task easier

- Some parts of z/VM can select their configuration based on system identifier, CPU id, or LPAR name.
 - SYSTEM CONFIG (CPU id or LPAR name to set system identifier)
 - SYSTEM NETID on 190 disk (select by CPU id)
 - TCPIP
 - Select profile based on System id
 - Qualify lines in TCPIP DATA based on System id
- SAPL allows you to select a different system configuration file when you load CP
 - “FN=” IPL parameter
- Customizing or creating customized copies of these files before you clone makes the job easier.

System configuration example

- Add qualifiers to statements or blocks of statements

```
System_Identifier LPAR LPARV1  VMSYSA
System_Identifier LPAR LPARV2  VMSYSB
System_Identifier_Default      ZVMV6R20
```

```
VMSYSA: System_Residence,
        Checkpoint  Valid M01R01  From CYL 21  For 9 ,
        Warmstart   Valid M01R01  From CYL 30  For 9
VMSYSB: System_Residence,
        Checkpoint  Valid M02R01  From CYL 21  For 9 ,
        Warmstart   Valid M02R01  From CYL 30  For 9
```

```
VMSYSA: Begin
        CP_Owned    Slot  10  M01S01
        CP_Owned    Slot  11  M01S02
VMSYSA: End
```

```
VMSYSB: Begin
        CP_Owned    Slot  10  M02S01
        CP_Owned    Slot  11  M02S02
VMSYSB: End
```

Do I really have to play by the new rules?

- Can I configure 6.2 or 6.3 so that it is as easy to clone as earlier releases? Is there a “short-cut”?
 - Yes, as long as your system is non-SSI
- Steps:
 - Move SYSTEM CONFIG from PMAINT CF0 to MAINT CF1
 - This is where it used to be
 - Note: Some service commands expect to find it on CF0!
 - Remove the IPL parms in SAPL. On MAINT6xx, run:
 - SALIPL 123 (EXTENT 1
 - Remove all ownership information from your CP owned volumes. Use CPFMTXA:
 - CPFMTXA 123 M01RES OWNER NONE
 - You must do this for all the CP owned volumes.
- Now you can copy your system using DDR or flashcopy

Short-cut method, continued

- When you Load or IPL the copied system:
 - May need to override console on SAPL screen
 - add CONS=xxxx as an IPL parameter
 - Or – define via loadparm as CONSxxxx
 - May want to specify a different config file on SAPL
 - add FN=name as an IPL parameter
 - ..assuming you created this file before the copy!
 - May want to force the IPL prompt so you can specify a “WARM NOAUTOLOG” start-up
 - add the IPL parameter PROMPT

SAPL screen with overrides

```
STAND ALONE PROGRAM LOADER: z/VM VERSION 6 RELEASE 2.0
DEVICE NUMBER: 7000      MINIDISK OFFSET: 39      EXTENT: 1
MODULE NAME:  CPLoad     LOAD ORIGIN: 1000
-----IPL PARAMETERS-----
FN=CLONE cons=4064 prompt_
-----COMMENTS-----
-----
9= FILELIST  10= LOAD  11= TOGGLE EXTENT/OFFSET
```

Short-cut method, continued

After the system is up:

- Fix the system id if the system didn't select the correct one automatically
 - Change SYSTEM CONFIG
 - Update TCPIP and other configurations
- Change labels, if needed
- One thing you still need to do:
 - Change the system id in the VMSES/E files
 - On MAINT6xx, run:
 - VMFUPDAT SYSPINV SYSTEM *new_sys_id old_sys_id*

But short-cuts aren't allowed!

Short-cuts

Clone a non-SSI 6.2 or 6.3 system

- OK, short-cuts not allowed...
- I'm assuming the new system identifier is changed, via:
 - Copied SYSTEM CONFIG with different identifier selected
 - Selected from CPU id or LPAR name
 - Fall back to System_Identifier_Default setting
- CP owned volume ownership must be changed after copying and before IPL
 - Otherwise you'll be most likely missing spool and page volumes!
 - But note! Only CP owned volumes need this
 - No allocation or ownership is read from user volumes
- How do you do this?
 - Use CPFMTXA from a z/VM 6.2+ sys that can write on the volumes
 - Use stand-alone ICKDSF

Using CPFMTXA

- CPFMTXA is a CP utility exec that invokes ICKDSF
 - Easier to use than ICKDSF
- Runs under CMS
- Options to format, allocate, relabel, and update owner
- Note: If you format cylinder 0, you wipe out the existing label, allocation, and ownership
 - Better to just change what you need
- Using CPFMTXA:
 - Syntax to update ownership:
 - CPFMTXA *vdev label* OWNER NOSSI.*newsysid*
 - The *label* must match the existing disk label
 - Syntax to write a new label:
 - CPFMTXA *vdev newlabel* LABEL

Using stand-alone ICKDSF

- Device Support Facilities (ICKDSF) can be loaded in the LPAR using the Stand-alone Program Loader (SAPL)
- Once it is up, you can change ownership information or volume labels
 - See example following this page
- Use this DSF command to change the owner
 - CPVOL OWNER UNIT(*rdev*) VFY(*label*) NOSSI SYSNAME(*sysid*)
- Use this DSF command to change a label:
 - CPVOL LABEL UNIT(*rdev*) VFY(*oldlabel*) VOLID(*newlabel*)
- Reply “U” to the prompt to actually make the update:
 - ICK003D REPLY U TO ALTER VOLUME *rdev* CONTENTS, ELSE T
 - U

Loading stand-alone ICKDSF

Specify ICKSADSF as the module name to load, press PF10

```
STAND ALONE PROGRAM LOADER: z/VM VERSION 6 RELEASE 2.0
DEVICE NUMBER:  7000      MINIDISK OFFSET:  39      EXTENT:  1
MODULE NAME:    ICKSADSF  LOAD ORIGIN:      1000

-----IPL PARAMETERS-----

-----COMMENTS-----

9= FILELIST  10= LOAD  11= TOGGLE EXTENT/OFFSET
```

Using stand-alone ICKDSF

Clear the screen, then enter ICKDSF commands

```
ICK005E DEFINE INPUT DEVICE, REPLY 'DDDD, CUU' OR 'CONSOLE'
ENTER INPUT/COMMAND:
CONSOLE ←
ICK006E DEFINE OUTPUT DEVICE, REPLY 'DDDD, CUU' OR 'CONSOLE'
ENTER INPUT/COMMAND:
CONSOLE ←
ICKDSF - SA/XA/ESA  DEVICE SUPPORT FACILITIES 17.0L          TIME: 19:59:54
          08/07/13    PAGE    1

ENTER INPUT/COMMAND:
  CPVOL OWNER UNIT(7001) VFY(M01P01) NOSSI SYSNAME(VMSYSB) ←
ICK00700I DEVICE INFORMATION FOR 7001 IS CURRENTLY AS FOLLOWS:
      PHYSICAL DEVICE = 3390
      STORAGE CONTROLLER = 3990
      STORAGE CONTROL DESCRIPTOR = E9
      DEVICE DESCRIPTOR = 0C
      ADDITIONAL DEVICE INFORMATION = 4A001F3C
      TRKS/CYL = 15, # PRIMARY CYLS = 100
ICK04000I DEVICE IS IN SIMPLEX STATE
ICK00703I DEVICE IS OPERATED AS A MINIDISK
ICK00091I 7001 NED=002107.900.IBM.75.00000000AKAZ1

-
CLEAR SCREEN WHEN READY
```

Other stand-alone utilities

- On your CP parm disk (where CPLOAD MODULE is located, MAINT CF1) there are 3 stand-alone utilities:
 - **ICKSADSF MODULE** (we just talked about)
 - **DDR MODULE**
 - Run the DASD Dump Restore program
 - Can be used to clone disks with z/VM down
 - **SALIPL MODULE**
 - Create or update SAPL on a CP formatted DASD
 - Can be used to update default IPL parameters
- All of them are loaded the same way using SAPL

IPL parameter “pdvol=*rdev*”

- The SYSTEM CONFIG file is now on PMAINT CF0
 - This is the first PARM disk on the common volume
 - This is not your IPL volume!
 - It is not CPACCESSED in SYSTEM CONFIG
- The stand-alone program loader, by default, specifies IPL parameters to find SYSTEM CONFIG
 - fn=SYSTEM ft=CONFIG pdnum=1 pdvol=*rdev*
 - You must correct this address – if you don't:
 - HCPARM6724W Requested parm disk volume is invalid.
 - CP entered; disabled wait PSW 00020000 00000000 00000000 00006724
 - You may also specify the console and “prompt”, if needed
- Use SALIPL to update (stand alone loader creation utility)
 - On MAINT6xx, enter:
SALIPL 123 (extent 1 iplparms fn=SYSTEM ft=CONFIG pdnum=1 pdvol=*newrdev*


Cloning non-SSI, after you IPL the clone

- Do the other normal updates
 - Update the system id in VMSES/E:
 - Logon to MAINT6xx
 - VMFUPDAT SYSPINV SYSTEM *new_sys_id old_sys_id*
 - Update the TCPIP configuration
 - If you IPLed with a renamed system config file, correct the name on the common parm disk
 - etc.

Update the CRR server LU name

- Note: This is not mentioned in the CP Planning Guide when you clone a system to create a new SSI member!
- The Coordinated Resource Recovery (CRR) server should have a unique LU name
 - This is the VMSEVR user id, filepool name is VMSYSR:
 - Install sets LU name to the system identifier
 - Can be any name, but must be unique if multiple systems are connected (such as an SSI cluster)
 - To fix: Log on to VMSEVR
 - Enter STOP NOBACKUP
 - Edit VMSEVR DMSPARMS
 - Change LU parameter and FILE the changes
 - Run FILESERV CRRLOG 306 307
 - IPL CMS and disconnect with #CP DISC
 - See the LU name by entering Q RESOURCE
 - Resource: .VMSYSA Type: Global Owning Userid: VMSEVR

Disk labels



What about duplicate disk labels?

If you've ever cloned a system, you may have seen this:

```
11:26:53 HCPZCO6718I Using parm disk 1 on volume VMCOM1 (device 1001) .
11:26:53 HCPZCO6718I Parm disk resides on cylinders 1 through 120.
11:26:53
11:26:53 HCPIIS954I DASD 2004 VOLID M01P01 IS A DUPLICATE OF DASD 1004
11:26:53 HCPIIS954I DASD 2003 VOLID M01S01 IS A DUPLICATE OF DASD 1003
11:26:53 HCPIIS954I DASD 2002 VOLID 630RL1 IS A DUPLICATE OF DASD 1002
11:26:53 HCPIIS954I DASD 2001 VOLID VMCOM1 IS A DUPLICATE OF DASD 1001
11:26:53 HCPIIS954I DASD 2000 VOLID M01RES IS A DUPLICATE OF DASD 1000
11:26:53 Start ((Warm|Force|COLD|CLEAN) (DRain) (DIisable) (NODIRect)
11:26:53          (NOAUTOlog)) or (SHUTDOWN)
```

- z/VM uses the labeled disk with the lower device address
- If the higher device is the correct one, don't start up your system!
- How do you fix this?

Avoiding duplicate labels

- Your choices:
 - Bring up new system second level so you can change labels
 - Remove devices with duplicate labels from the LPAR
 - Ignore devices with duplicate labels in your SYSTEM CONFIG
 - Devices statement, using either
 - *Offline_at_IPL*
 - *Sensed_but_offline*
 - **Sensed_but_offline** is preferred, because the device is sensed (CP will know it is a DASD) but left offline.
 - On z/VM 6.3, use the new support for specifying device numbers:
 - *CP_Owned Slot n valid RDEV rdev*
 - *User_Volume_RDEV valid RDEV rdev*
 - Example on the next page
 - Config changes require you to modify or copy your SYSTEM CONFIG before you clone the system.

New z/VM 6.3 RDEV support

- My SYSTEM CONFIG statements:

```
CP_Owned      Slot      1  M01RES   RDEV 1000
CP_Owned      Slot      5  VMCOM1   RDEV 1001
CP_Owned      Slot     10  M01S01   RDEV 1003
CP_Owned      Slot    255  M01P01   RDEV 1004
User_Volume_RDEV      630RL1   RDEV 1002
```

- CP will still issue a message if it detects a duplicate label

- but you can now be sure which device it is using

```
HCPIIS291I DASD 2004 VOLID M01P01 DOES NOT MATCH DEVICE 1004 M01P01
HCPIIS291I DASD 2003 VOLID M01S01 DOES NOT MATCH DEVICE 1003 M01S01
HCPIIS291I DASD 2002 VOLID 630RL1 DOES NOT MATCH DEVICE 1002 630RL1
HCPIIS291I DASD 2001 VOLID VMCOM1 DOES NOT MATCH DEVICE 1001 VMCOM1
HCPIIS954I DASD 2000 VOLID M01RES IS A DUPLICATE OF DASD 1000
```

- The message 291I is confusing, but it says device 2xxx was found as a duplicate of the device CP is using (1xxx)
- 1000 is my IPL device. Duplicates of it have always been ignored.
 - This is why message 954I is issued instead of message 291I

z/VM 6.3 RDEV support, another test

- My SYSTEM CONFIG statements:

```
CP_Owned      Slot    1  M01RES  RDEV 1000
CP_Owned      Slot    5  VMCOM1  RDEV 1001
CP_Owned      Slot   10  M01S01  RDEV 1003
CP_Owned      Slot  255  M01P01  RDEV 1004
User_Volume_RDEV      630RL1  RDEV 1002
```

- Test with the duplicate volumes at 200 instead of 2000

- i.e. we must ignore the *lower* address this time

```
HCPIIS291I DASD 0204 VOLID M01P01 DOES NOT MATCH DEVICE 1004 M01P01
HCPIIS291I DASD 0203 VOLID M01S01 DOES NOT MATCH DEVICE 1003 M01S01
HCPIIS291I DASD 0202 VOLID 630RL1 DOES NOT MATCH DEVICE 1002 630RL1
HCPIIS291I DASD 0201 VOLID VMCOM1 DOES NOT MATCH DEVICE 1001 VMCOM1
HCPIIS954I DASD 0200 VOLID M01RES IS A DUPLICATE OF DASD 1000
```

- VM still has the correct volumes owned or attached to system:

```
Q DASD
DASD 1000 CP OWNED M01RES 90
DASD 1001 CP OWNED VMCOM1 12
DASD 1002 CP SYSTEM 630RL1 5
DASD 1003 CP OWNED M01S01 0
DASD 1004 CP OWNED M01P01 0
```


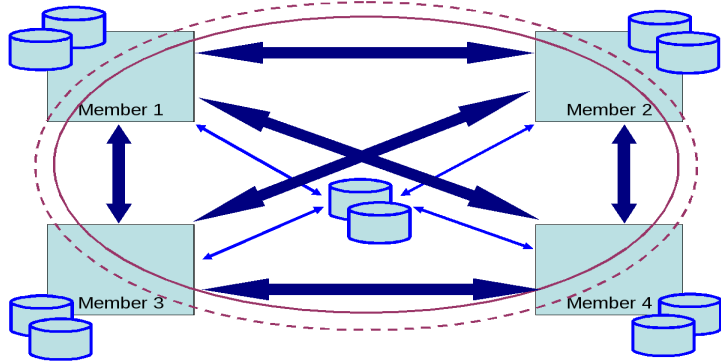


Changing volume labels

- Some things to know about DASD labels
 - Labels appear in SYSTEM CONFIG and the directory
 - SYSTEM CONFIG is only read when CP is initializing
 - Labels are only read when DASD is attached to the system
- So – updating labels can be done on a live system if you do it in a specific order
 - Suggestion:
 - Update a copy of your SYSTEM CONFIG and run CPSYNTAX on it
 - Update a copy of your USER DIRECT
 - *Change mdisk statements and the DIRECTORY statement*
 - *Use DIRECTXA with (EDIT option to syntax check!*
 - *Run DISKMAP EXEC or DIRMAP and check its output*
 - Change the real labels (more on this on next page)
 - Replace your SYSTEM CONFIG with the copy
 - Replace your USER DIRECT with the copy
 - DIRECTXA USER DIRECT
 - SHUTDOWN REIPL

Changing labels on a live system

- For example, after you've IPLed your cloned system
- The problem:
 - Any tool, such as CPFMTXA, operates on a virtual address
 - Devices in use by the system can't be attached to a user
 - You need a virtual device that lets you write on cylinder 0
- Solutions
 - Define a “0-END” minidisk or use an existing one
 - such as MAINT 122-124, MAINT6xx 131-133, PMAINT 141-142, etc.
 - Use DEFINE MDISK to dynamically create a minidisk
 - Requires the user id to have OPTION DEVMAINT
 - MAINT and MAINT6xx do not have this by default!
 - Use another system that can attach the volumes

Cloning an SSI cluster



Cloning an SSI cluster

- SSI cluster has member names and the SSI cluster name:
 - Written to the PDR (Persistent Data Record) on the common volume
 - In the ownership information on CP owned volumes
 - This is required for SSI!
 - Defined in SYSTEM CONFIG
 - Member names are in the user directory
- You must change this information before you IPL the cloned system the first time
 - so you really need to bring it up second level the first time or have another 6.2+ VM system available
 - Prep work with SYSTEM CONFIG and the directory before you clone is required!

Before you clone SSI ..

- On PMAINT CF0:
copy SYSTEM CONFIG to CLONE CONFIG
 - Edit CLONE CONFIG and change LPAR names, system names, SSI cluster name, macprefix, etc.
 - Change real addresses. Make sure the real CTC addresses on your ISLINK statements are correct for each LPAR and system!
 - You could even change labels if you're brave enough to change everything at once
 - Be sure to run CPSYNTAX for all LPAR names!
- Get a copy of the user directory and put it on your “work” VM system or the first level system.
 - a “work” system is another z/VM system that can write on SSI cluster's CP owned volumes
- Now you can clone your SSI
 - I recommend all systems are shut down first!

Update your CP owned disks on the clone

- Re-write the PDR on the cloned common volume
 - The FORMSSI command is found on PMAINT 551
 - ATTACH 1001 *
 - FORMSSI CREATE 1001 *newssi*
 - If all systems were down when you cloned disks, this works. If a system was up, you may need to add the FORCE option.
- Update ownership information on all CP owned volumes
 - ATTACH 1000 *
 - CPFMTXA 1000 M01RES OWNER *newssi.newsys*
 - Change labels now also if you're one of the brave ones
 - CPFMTXA 1000 M02RES LABEL
 - Note that the common volume is not owned by a system:
 - CPFMTXA 1001 VMCOM1 OWNER *newssi.NOSYS*
 - Also: all CP owned volumes must have the SSI cluster name.
 - You can't use NOSSI!

Update the IPL parms

- You could do this after you IPL
 - Just override the incorrect parms
 - Write out the changes after the system is up
- All RES disks will use the same parms
- It is easy to do now while you have disks attached:

SAL IPL 1000 (extent 1 iplparms FN=SYSTEM FT=CONFIG pdnum=1 pdvol=*rdev*)

- We will still have to override the IPL parms on our first IPL.

Directory updates

- You must have a copy of the source directory from the system you are cloning.
 - It doesn't matter if you're using Dirmaint or not
- Edit this file
 - Change names on the BUILD ON statements to the new names
 - If you are also changing labels now, change them
 - DIRECTORY statement
 - *DIRECTORY SSI 123 label1 label2 label3*
 - Minidisk statements
 - *MDISK 123 3390 0 END label1 MR*
- Check your work
 - Use DIRECTXA with the EDIT option to syntax check the directory
 - *DIRECTXA CLONE DIRECT A (EDIT*
 - Use a utility such as DISKMAP or DIRMAP to make sure you didn't incorrectly change a label
 - Look at the output to make sure only correct labels are there

Write out the updated directory

- Note that the directory must be written on each res disk
- The directory typically specifies address “123” to write the directory
 - This is on the DIRECTORY statement
 - The label of the RES disk must match a label on this statement
 - Make sure you don't update the wrong RES disk!
 - DET 123
 - ATTACH 1000 * 123
 - DIRECTXA CLONE DIRECT
 - DET 123
 - Do this for each member of the cloned SSI cluster
 - DIRECTXA will end with return code 5 – this is normal
 - The directory written out was not the active directory on your current system

IPL the first member of your SSI

- Specify a LOADPARM of your console
 - We need to override IPL parms
- Override the file name with CLONE
- Add parameters
 - “prompt” to force the IPL prompts
 - “cons=rdev” if needed to force the console

```
-----IPL PARAMETERS-----  
_FN=CLONE FT=CONFIG pdnum=1 pdvol=1001 prompt cons=4000
```

- Press PF10 to load

Initial start up

- Specify start-up options of FORCE NOAUTOLOG
 - WARM will give message
 - HCPWRM9200E System recovery failure; configuration error.
 - This may depend on if you changed labels or not.
- Go find your list of the usual things to fix on a cloned system
 - Get the directory source file you used to write the directory
 - Put on PMAINT 2CC or wherever you normally keep it
 - If using Dirmaint:
 - *Rename to USER INPUT*
 - *Put on DIRMAINT 1DF*
 - *ERASE USER DIRECT from that disk*
 - *Start up DIRMAINT*
 - Don't forget to:
 - Rename CLONE CONFIG on PMAINT CF0 to SYSTEM CONFIG
 - Run VMFUPDAT SYSPINV on MAINT6xx
 - Update SAPL if you didn't before

Final thoughts

- The service filepool (VMPSFS) is a global resource. If your cloned systems will be connected via ISFC, you may need to change this resource name.
- If you're using DIRMAINT and change labels, the EXTENT CONTROL file has to be updated also.
- If you're using Perfkit, some history files it writes use the system id as a file name.

The End

- **Thank you for listening!**

Session 13400

- **Contact information**

Bruce Hayden

bjhayden@us.ibm.com

IBM Plant #1

Endicott, NY



Additional information

- How to use Flashcopy on z/VM
- How to use DDR

How to use Flashcopy on z/VM

- FLASHCOPY is a CP command, but you specify virtual addresses to copy. So, the disks to be copied (from or to) must be either attached to your user id or a minidisk defined that allows reading the entire source disk and writing the entire target disk.
 - FLASHCOPY requires privilege class B
 - Your storage system must have the flashcopy feature in order for you to use this command. z/VM does not simulate the flashcopy function if the feature is not available.
- Syntax for copying an entire disk:
 - CP FLASHCOPY *vdev1* 0 END TO *vdev2* 0 END
 - *vdev1* is the source device
 - *vdev2* is the target device
 - You can add “LABEL *newlabel*” at the end of the command to also change the label on the target device.

How to use DDR

- The DDR utility can be run as a CMS command or started as a stand-alone program. Here are the DDR commands used to copy an entire disk.
 - SYSPRINT CONS
 - Send all program output to the console (a printer is the default!)
 - INPUT *dev1* DASD *valid1*
 - OUTPUT *dev2* DASD SCRTCH
 - Input device is *dev1* with label *valid1*
 - Output device is *dev2* and the current label is ignored
 - In a virtual machine, *dev1* and *dev2* are virtual addresses
 - As a stand-alone program in an LPAR, *dev1* and *dev2* are real addresses
 - COPY ALL
 - Copy the entire volume
 - You may enter additional INPUT, OUTPUT and COPY commands to DDR or press enter on a blank line to end DDR.